

./ - ktap

- [Makefile](#)
- [_vimrc_local.vim](#)
- [genmobi.sh](#)
- [include/](#)
- [ktapvm.mod.c](#)
- [runtime/](#)
- [test/](#)
- [userspace/](#)
- [vim/](#)

Makefile - ktap

```
1
2 #
3 # Define NO_LIBELF if you do not want libelf dependency (e.g. cross-builds)
4 # (this will also disable resolve resolving symbols in DSO functionality)
5 #
6 # Define FFI if you want to compile ktap with FFI support. By default This
7 # toggle is off.
8 #
9 # Define amalg to enable amalgamation build, This compiles the ktapvm as
10 # one huge C file and allows GCC to generate faster and shorter code. Alas,
11 # this requires lots of memory during the build.
12 # Recommend to use amalgamation build as default.
13 amalg = 1
14
15 # Do not instrument the tracer itself:
16 ifdef CONFIG_FUNCTION_TRACER
17 ORIG_CFLAGS := $(KBUILD_CFLAGS)
18 KBUILD_CFLAGS = $(subst -pg,, $(ORIG_CFLAGS))
19 endif
20
21 all: mod ktap
22
23 INC = include
24 RUNTIME = runtime
25
26 FFIDIR = $(RUNTIME)/ffi
27 KTAP_LIBS = -lpthread
28
29 LIB_OBJS += $(RUNTIME)/lib_base.o $(RUNTIME)/lib_kdebug.o \
30             $(RUNTIME)/lib_timer.o $(RUNTIME)/lib_ansi.o \
31             $(RUNTIME)/lib_table.o $(RUNTIME)/lib_net.o
32
33 ifndef amalg
34 ifdef FFI
35 FFI_OBJS += $(FFIDIR)/ffi_call.o $(FFIDIR)/ffi_type.o $(FFIDIR)/ffi_symbol.o \
36             $(FFIDIR)/cdata.o $(FFIDIR)/ffi_util.o
37 RUNTIME_OBJS += $(FFI_OBJS)
38 LIB_OBJS += $(RUNTIME)/lib_ffi.o
39 endif
40 RUNTIME_OBJS += $(RUNTIME)/ktap.o $(RUNTIME)/kp_bcread.o $(RUNTIME)/kp_obj.o \
41                 $(RUNTIME)/kp_str.o $(RUNTIME)/kp_mempool.o \
42                 $(RUNTIME)/kp_tab.o $(RUNTIME)/kp_vm.o \
43                 $(RUNTIME)/kp_transport.o $(RUNTIME)/kp_events.o $(LIB_OBJS)
44 else
45 RUNTIME_OBJS += $(RUNTIME)/amalg.o
46 endif
47
48 ifdef FFI
49 ifeq ($(KBUILD_MODULES), 1)
50 ifdef CONFIG_X86_64
51 # call_x86_64.o is compiled from call_x86_64.S
52 RUNTIME_OBJS += $(FFIDIR)/call_x86_64.o
53 else
54 $(error ktap FFI only supports x86_64 for now!)
55 endif
56 endif
57
58
59 ccflags-y += -DCONFIG_KTAP_FFI
60 endif
61
62 obj-m += ktapvm.o
63 ktapvm-y := $(RUNTIME_OBJS)
64
65 KVERSION ?= $(shell uname -r)
66 KERNEL_SRC ?= /lib/modules/$(KVERSION)/build
67 PWD := $(shell pwd)
68 mod:
69     $(MAKE) -C $(KERNEL_SRC) M=$(PWD) modules
70
71 modules_install:
```

```

72     $(MAKE) -C $(KERNEL_SRC) M=$(PWD) modules_install
73
74     KTAPC_CFLAGS = -Wall -O2
75
76
77     # try-cc
78     # Usage: option = $(call try-cc, source-to-build, cc-options, msg)
79     ifneq ($(V),1)
80     TRY_CC_OUTPUT= > /dev/null 2>&1
81     endif
82     TRY_CC_MSG=echo "    CHK $(3)" 1>&2;
83
84     try-cc = $(shell sh -c
85         'TMP="/tmp/.$$$$";
86         $(TRY_CC_MSG)
87         echo "$(1)" |
88         $(CC) -x c - $(2) -o "$$TMP" $(TRY_CC_OUTPUT) && echo y;
89         rm -f "$$TMP"')
90
91
92     define SOURCE_LIBELF
93     #include <libelf.h>
94
95     int main(void)
96     {
97         Elf *elf = elf_begin(0, ELF_C_READ, 0);
98         return (long)elf;
99     }
100    endif
101
102    FLAGS_LIBELF = -lelf
103
104    ifndef NO_LIBELF
105        KTAPC_CFLAGS += -DNO_LIBELF
106    else
107        ifneq ($(call try-cc,$(SOURCE_LIBELF),$(FLAGS_LIBELF),libelf),y)
108            $(warning No libelf found, disables symbol resolving, please install elfutils-libelf-devel/libelf-dev);
109            NO_LIBELF := 1
110            KTAPC_CFLAGS += -DNO_LIBELF
111        else
112            KTAP_LIBS += -lelf
113        endif
114    endif
115
116    UDIR = userspace
117
118    $(UDIR)/kp_main.o: $(UDIR)/kp_main.c $(INC)/* KTAP-CFLAGS
119        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
120    $(UDIR)/kp_lex.o: $(UDIR)/kp_lex.c $(INC)/* KTAP-CFLAGS
121        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
122    $(UDIR)/kp_parse.o: $(UDIR)/kp_parse.c $(INC)/* KTAP-CFLAGS
123        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
124    $(UDIR)/kp_bcwrite.o: $(UDIR)/kp_bcwrite.c $(INC)/* KTAP-CFLAGS
125        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
126    $(UDIR)/kp_reader.o: $(UDIR)/kp_reader.c $(INC)/* KTAP-CFLAGS
127        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
128    $(UDIR)/kp_util.o: $(UDIR)/kp_util.c $(INC)/* KTAP-CFLAGS
129        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
130    $(UDIR)/kp_parse_events.o: $(UDIR)/kp_parse_events.c $(INC)/* KTAP-CFLAGS
131        $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
132    ifndef NO_LIBELF
133        $(UDIR)/kp_symbol.o: $(UDIR)/kp_symbol.c KTAP-CFLAGS
134            $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
135    endif
136    ifdef FFI
137        KTAPC_CFLAGS += -DCONFIG_KTAP_FFI
138        $(UDIR)/ffi_type.o: $(RUNTIME)/ffi/ffi_type.c $(INC)/* KTAP-CFLAGS
139            $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
140        $(UDIR)/ffi/cparser.o: $(UDIR)/ffi/cparser.c $(INC)/* KTAP-CFLAGS
141            $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
142        $(UDIR)/ffi/ctype.o: $(UDIR)/ffi/ctype.c $(INC)/* KTAP-CFLAGS
143            $(QUIET_CC)$(CC) $(DEBUGINFO_FLAG) $(KTAPC_CFLAGS) -o $@ -c $<
144    endif
145
146
147    KTAPOBJS =

```

```

148 KTAPOBJS += $(UDIR)/kp_main.o
149 KTAPOBJS += $(UDIR)/kp_lex.o
150 KTAPOBJS += $(UDIR)/kp_parse.o
151 KTAPOBJS += $(UDIR)/kp_bcwrite.o
152 KTAPOBJS += $(UDIR)/kp_reader.o
153 KTAPOBJS += $(UDIR)/kp_util.o
154 KTAPOBJS += $(UDIR)/kp_parse_events.o
155 ifndef NO_LIBELF
156 KTAPOBJS += $(UDIR)/kp_symbol.o
157 endif
158 ifdef FFI
159 KTAPOBJS += $(UDIR)/ffi_type.o
160 KTAPOBJS += $(UDIR)/ffi/cparser.o
161 KTAPOBJS += $(UDIR)/ffi/ctype.o
162 endif
163
164 ktap: $(KTAPOBJS) KTAP-CFLAGS
165     $(QUIET_LINK)$@ $(CC) $(KTAPC_CFLAGS) -o $@ $(KTAPOBJS) $(KTAP_LIBS)
166
167 KMISC := /lib/modules/$(KVERSION)/ktapvm/
168
169 install: mod ktap
170     make modules_install ktapvm.ko
171     install -c ktap /usr/bin/
172     mkdir -p ~/.vim/ftdetect
173     mkdir -p ~/.vim/syntax
174     cp vim/ftdetect/ktap.vim ~/.vim/ftdetect/
175     cp vim/syntax/ktap.vim ~/.vim/syntax/
176
177 load:
178     insmod ktapvm.ko
179
180 unload:
181     rmmod ktapvm
182
183 reload:
184     make unload; make load
185
186 test: FORCE
187     #start testing
188     prove -j4 -r test/
189
190 clean:
191     $(MAKE) -C $(KERNEL_SRC) M=$(PWD) clean
192     $(RM) ktap KTAP-CFLAGS
193
194
195 PHONY += FORCE
196 FORCE:
197
198 TRACK_FLAGS = KTAP
199 ifdef amalg
200 TRACK_FLAGS += AMALG
201 endif
202 ifdef FFI
203 TRACK_FLAGS += FFI
204 endif
205 ifdef NO_LIBELF
206 TRACK_FLAGS += NO_LIBELF
207 endif
208
209 KTAP-CFLAGS: FORCE
210     @FLAGS='$(TRACK_FLAGS)'; \
211     if test x"$$FLAGS" != x" `cat KTAP-CFLAGS 2>/dev/null` " ; then \
212         echo "$$FLAGS" >KTAP-CFLAGS; \
213     fi
214
215 #generate tags/etags/cscope index for editor.
216 define all_sources
217     (find . -name '*.[ch]' -print)
218 endef
219
220 .PHONY: tags
221 tags:
222     $(all_sources) | xargs ctags
223

```

```
224 .PHONY: etags
225 etags:
226     $(all_sources) | xargs etags
227
228 .PHONY: cscope
229 cscope:
230     $(all_sources) > cscope.files
231     cscope -k -b
```

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

_vimrc_local.vim - ktap

```
1 set noexpandtab
2 "filetype indent off
3 "filetype plugin off
4 "set nowrap autoindent nosmartindent nocindent indentexpr=
```

[One Level Up](#)

[Top Level](#)

genmobi.sh - ktap

```
1 #!/bin/bash
2
3 name=`pwd|perl -e '$d=<>;$d=~s{.*/}{}g;$d=~s/2$/g;print $d'`
4 src2html='src2html.pl -l -x -c -n'
5
6 echo "Generating HTMLs for $name..."
7 $src2html -t 4 --exclude='*/.*' --exclude='.*' ./ $name
8 #exit
9
10 echo "Generating .pdf file for $name..."
11 rm -rf *.pdf
12 ebook-convert html_out/index.html $name.pdf \
13     --override-profile-size \
14     --paper-size a4 \
15     --pdf-default-font-size 12 \
16     --pdf-mono-font-size 12 \
17     --margin-left 10 --margin-right 10 \
18     --margin-top 10 --margin-bottom 10 \
19     --page-breaks-before=/'
20 #exit
21
22 echo "Generating .epub file for $name..."
23 rm -rf *.epub
24 ebook-convert html_out/index.html $name.epub \
25     --no-default-epub-cover \
26     --publisher agentzh \
27     --output-profile ipad3 \
28     --title $name \
29     --language en --authors 'Jovi Zhangwei'
30
31 echo "Generating .mobi file for $name..."
32 rm -rf *.mobi
33 ebook-convert html_out/index.html $name.mobi \
34     --output-profile kindle_dx --no-inline-toc \
35     --publisher agentzh \
36     --title $name \
37     --language en --authors 'Jovi Zhangwei'
38
39 #cp -uv *.mobi ~/mobi/
40
```

[One Level Up](#)

[Top Level](#)

include/ - ktap

- [ktap_arch.h](#)
- [ktap_bc.h](#)
- [ktap_err.h](#)
- [ktap_errmsg.h](#)
- [ktap_ffi.h](#)
- [ktap_types.h](#)

[One Level Up](#)

[Top Level](#)

include/ktap_arch.h - ktap

Macros defined

- [KP_BE](#)
- [KP_BE](#)
- [KP_BE](#)
- [KP_BE](#)
- [KP_ENDIAN_SELECT](#)
- [KP_ENDIAN_SELECT](#)
- [KP_ENDIAN_SELECT](#)
- [KP_ENDIAN_SELECT](#)
- [KP_LE](#)
- [KP_LE](#)
- [KP_LE](#)
- [KP_LE](#)
- [__KTAP_ARCH](#)

Source code

```
1 #ifndef __KTAP_ARCH
2 #define __KTAP_ARCH__
3
4 #ifdef __KERNEL__
5 #include <linux/types.h>
6 #include <asm/byteorder.h>
7
8 #if defined(__LITTLE_ENDIAN)
9 #define KP_LE 1
10 #define KP_BE 0
11 #define KP_ENDIAN_SELECT(le, be) le
12 #elif defined(__BIG_ENDIAN)
13 #define KP_LE 0
14 #define KP_BE 1
15 #define KP_ENDIAN_SELECT(le, be) be
16 #endif
17
18 #else /* __KERNEL__ */
19
20 #if __BYTE_ORDER == __LITTLE_ENDIAN
21 #define KP_LE 1
22 #define KP_BE 0
23 #define KP_ENDIAN_SELECT(le, be) le
24 #elif __BYTE_ORDER == __BIG_ENDIAN
25 #define KP_LE 0
26 #define KP_BE 1
27 #define KP_ENDIAN_SELECT(le, be) be
28 #else
29 #error "could not determine byte order"
30 #endif
31
32 #endif
33 #endif
```

include/ktap_bc.h - ktap

Data types defined

- [BCIns](#)
- [BCLine](#)
- [BCMode](#)
- [BCOp](#)
- [BCPos](#)
- [BCReg](#)
- [MMS](#)

Functions defined

- [bc_isret](#)

Macros defined

- [BCBIAS_J](#)
- [BCDEF](#)
- [BCDUMP_F_BE](#)
- [BCDUMP_F_FFI](#)
- [BCDUMP_F_KNOWN](#)
- [BCDUMP_F_STRIP](#)
- [BCDUMP_HEAD1](#)
- [BCDUMP_HEAD2](#)
- [BCDUMP_HEAD3](#)
- [BCDUMP_VERSION](#)
- [BCENUM](#)
- [BCENUM](#)
- [BCINS_ABC](#)
- [BCINS_AD](#)
- [BCINS_AJ](#)
- [BCMAX_A](#)
- [BCMAX_B](#)
- [BCMAX_C](#)
- [BCMAX_D](#)

- [BCMODE](#)
- [BCMODE_FF](#)
- [BCM](#)
- [FF_next_N](#)
- [KP_STATIC_ASSERT](#)
- [NO_JMP](#)
- [NO_REG](#)
- [__KTAP_BC_H](#)
- [bc_a](#)
- [bc_b](#)
- [bc_c](#)
- [bc_d](#)
- [bc_j](#)
- [bc_op](#)
- [bcmode_a](#)
- [bcmode_b](#)
- [bcmode_c](#)
- [bcmode_d](#)
- [bcmode_hasd](#)
- [bcmode_mm](#)
- [kp_assert](#)
- [setbc_a](#)
- [setbc_b](#)
- [setbc_byte](#)
- [setbc_c](#)
- [setbc_d](#)
- [setbc_j](#)
- [setbc_op](#)

Source code

```

1 /*
2  * Bytecode instruction format.
3  *
4  * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
5  * Copyright (C) 2005-2014 Mike Pall.
6  */
7
8 #ifndef __KTAP_BC_H
9 #define __KTAP_BC_H

```

```

10
11 #include "../include/ktap_arch.h"
12
13 /*TODO*/
14 #define KP_STATIC_ASSERT(cond)
15 #define kp_assert(cond)
16
17 /* Types for handling bytecodes. */
18 typedef uint32_t BCIns; /* Bytecode instruction. */
19 typedef uint32_t BCPos; /* Bytecode position. */
20 typedef uint32_t BCReg; /* Bytecode register. */
21 typedef int32_t BCLine; /* Bytecode line number. */
22
23 /*
24 * Bytecode instruction format, 32 bit wide, fields of 8 or 16 bit:
25 *
26 * +-----+-----+-----+-----+
27 * | B | C | A | OP | Format ABC
28 * +-----+-----+-----+-----+
29 * |   D   | A | OP | Format AD
30 * +-----+-----+-----+-----+
31 * MSB                               LSB
32 *
33 * In-memory instructions are always stored in host byte order.
34 */
35
36 /* Operand ranges and related constants. */
37 #define BCMAX_A      0xff
38 #define BCMAX_B      0xff
39 #define BCMAX_C      0xff
40 #define BCMAX_D      0xffff
41 #define BCBIAS_J     0x8000
42 #define NO_REG       BCMAX_A
43 #define NO_JMP       (~(BCPos)0)
44
45 /* Macros to get instruction fields. */
46 #define bc_op(i)      ((BCOp)((i)&0xff))
47 #define bc_a(i)       ((BCReg)(((i)>>8)&0xff))
48 #define bc_b(i)       ((BCReg)((i)>>24))
49 #define bc_c(i)       ((BCReg)(((i)>>16)&0xff))
50 #define bc_d(i)       ((BCReg)((i)>>16))
51 #define bc_j(i)       ((ptrdiff_t)bc_d(i)-BCBIAS_J)
52
53 /* Macros to set instruction fields. */
54 #define setbc_byte(p, x, ofs) \
55     ((uint8_t *) (p))[KP_ENDIAN_SELECT(ofs, 3 - ofs)] = (uint8_t)(x)
56 #define setbc_op(p, x)      setbc_byte(p, (x), 0)
57 #define setbc_a(p, x)       setbc_byte(p, (x), 1)
58 #define setbc_b(p, x)       setbc_byte(p, (x), 3)
59 #define setbc_c(p, x)       setbc_byte(p, (x), 2)
60 #define setbc_d(p, x) \
61     ((uint16_t *) (p))[KP_ENDIAN_SELECT(1, 0)] = (uint16_t)(x)
62 #define setbc_j(p, x)       setbc_d(p, (BCPos)((int32_t)(x)+BCBIAS_J))
63
64 /* Macros to compose instructions. */
65 #define BCINS_ABC(o, a, b, c) \
66     (((BCIns)(o))|((BCIns)(a)<<8)|((BCIns)(b)<<24)|((BCIns)(c)<<16))
67 #define BCINS_AD(o, a, d) \
68     (((BCIns)(o))|((BCIns)(a)<<8)|((BCIns)(d)<<16))
69 #define BCINS_AJ(o, a, j)    BCINS_AD(o, a, (BCPos)((int32_t)(j)+BCBIAS_J))
70
71 /*
72 * Bytecode instruction definition. Order matters, see below.
73 *
74 * (name, filler, Amode, Bmode, Cmode or Dmode, metamethod)
75 *
76 * The opcode name suffixes specify the type for RB/RC or RD:
77 * V = variable slot
78 * S = string const
79 * N = number const
80 * P = primitive type (~itype)
81 * B = unsigned byte literal
82 * M = multiple args/results
83 */
84 #define BCDEF(_ ) \
85     /* Comparison ops. ORDER OPR. */ \

```

```

86  _(ISLT,   var,   ___,   var,   lt) \
87  _(ISGE,   var,   ___,   var,   lt) \
88  _(ISLE,   var,   ___,   var,   le) \
89  _(ISGT,   var,   ___,   var,   le) \
90  \
91  _(ISEQV,   var,   ___,   var,   eq) \
92  _(ISNEV,   var,   ___,   var,   eq) \
93  _(ISEQS,   var,   ___,   str,   eq) \
94  _(ISNES,   var,   ___,   str,   eq) \
95  _(ISEQN,   var,   ___,   num,   eq) \
96  _(ISNEN,   var,   ___,   num,   eq) \
97  _(ISEQP,   var,   ___,   pri,   eq) \
98  _(ISNEP,   var,   ___,   pri,   eq) \
99  \
100 /* Unary test and copy ops. */ \
101 _(ISTC,   dst,   ___,   var,   ___) \
102 _(ISFC,   dst,   ___,   var,   ___) \
103 _(IST,    ___,   ___,   var,   ___) \
104 _(ISF,    ___,   ___,   var,   ___) \
105 _(ISTYPE, var,   ___,   lit,   ___) \
106 _(ISNUM,  var,   ___,   lit,   ___) \
107 \
108 /* Unary ops. */ \
109 _(MOV,   dst,   ___,   var,   ___) \
110 _(NOT,   dst,   ___,   var,   ___) \
111 _(UNM,   dst,   ___,   var,   unm) \
112 \
113 /* Binary ops. ORDER OPR. VV last, POW must be next. */ \
114 _(ADDVN, dst,   var,   num,   add) \
115 _(SUBVN, dst,   var,   num,   sub) \
116 _(MULVN, dst,   var,   num,   mul) \
117 _(DIVVN, dst,   var,   num,   div) \
118 _(MODVN, dst,   var,   num,   mod) \
119 \
120 _(ADDNV, dst,   var,   num,   add) \
121 _(SUBNV, dst,   var,   num,   sub) \
122 _(MULNV, dst,   var,   num,   mul) \
123 _(DIVNV, dst,   var,   num,   div) \
124 _(MODNV, dst,   var,   num,   mod) \
125 \
126 _(ADDVV, dst,   var,   var,   add) \
127 _(SUBVV, dst,   var,   var,   sub) \
128 _(MULVV, dst,   var,   var,   mul) \
129 _(DIVVV, dst,   var,   var,   div) \
130 _(MODVV, dst,   var,   var,   mod) \
131 \
132 _(POW,   dst,   var,   var,   pow) \
133 _(CAT,   dst,   rbase, rbase, concat) \
134 \
135 /* Constant ops. */ \
136 _(KSTR,   dst,   ___,   str,   ___) \
137 _(KCDATA, dst,   ___,   cdata, ___) \
138 _(KSHORT, dst,   ___,   lits,  ___) \
139 _(KNUM,   dst,   ___,   num,   ___) \
140 _(KPRI,   dst,   ___,   pri,   ___) \
141 _(KNIL,   base,  ___,   base,  ___) \
142 \
143 /* Upvalue and function ops. */ \
144 _(UGET,   dst,   ___,   uv,   ___) \
145 _(USETV,  uv,   ___,   var,   ___) \
146 _(UINCV,  uv,   ___,   var,   ___) \
147 _(USETS,  uv,   ___,   str,   ___) \
148 _(USETN,  uv,   ___,   num,   ___) \
149 _(UINCN,  uv,   ___,   num,   ___) \
150 _(USETP,  uv,   ___,   pri,   ___) \
151 _(UCLO,   rbase, ___,   jump, ___) \
152 _(FNEW,   dst,   ___,   func, gc) \
153 \
154 /* Table ops. */ \
155 _(TNEW,   dst,   ___,   lit,   gc) \
156 _(TDUP,   dst,   ___,   tab,   gc) \
157 _(GGET,   dst,   ___,   str,   index) \
158 _(GSET,   var,   ___,   str,   newindex) \
159 _(GINC,   var,   ___,   str,   newindex) \
160 _(TGETV,  dst,   var,   var,   index) \
161 _(TGETS,  dst,   var,   str,   index) \

```

```

162  _(TGETB,  dst,  var,  lit,  index) \
163  _(TGETR,  dst,  var,  var,  index) \
164  _(TSETV,  var,  var,  var,  newindex) \
165  _(TINCV,  var,  var,  var,  newindex) \
166  _(TSETS,  var,  var,  str,  newindex) \
167  _(TINCS,  var,  var,  str,  newindex) \
168  _(TSETB,  var,  var,  lit,  newindex) \
169  _(TINCB,  var,  var,  lit,  newindex) \
170  _(TSETM,  base,  ___,  num,  newindex) \
171  _(TSETR,  var,  var,  var,  newindex) \
172  \
173  /* Calls and vararg handling. T = tail call. */ \
174  _(CALLM,  base,  lit,  lit,  call) \
175  _(CALL,   base,  lit,  lit,  call) \
176  _(CALLMT, base,  ___,  lit,  call) \
177  _(CALLT,  base,  ___,  lit,  call) \
178  _(ITERC,  base,  lit,  lit,  call) \
179  _(ITERN,  base,  lit,  lit,  call) \
180  _(VARG,   base,  lit,  lit,  ___) \
181  _(ISNEXT, base,  ___,  jump,  ___) \
182  \
183  /* Returns. */ \
184  _(RETM,   base,  ___,  lit,  ___) \
185  _(RET,    rbase, ___,  lit,  ___) \
186  _(RET0,   rbase, ___,  lit,  ___) \
187  _(RET1,   rbase, ___,  lit,  ___) \
188  \
189  /* Loops and branches. I/J = interp/JIT, I/C/L = init/call/loop. */ \
190  _(FORI,   base,  ___,  jump,  ___) \
191  _(JFORI,  base,  ___,  jump,  ___) \
192  \
193  _(FORL,   base,  ___,  jump,  ___) \
194  _(IFORL,  base,  ___,  jump,  ___) \
195  _(JFORL,  base,  ___,  lit,  ___) \
196  \
197  _(ITERL,  base,  ___,  jump,  ___) \
198  _(IITERL, base,  ___,  jump,  ___) \
199  _(JITERL, base,  ___,  lit,  ___) \
200  \
201  _(LOOP,   rbase, ___,  jump,  ___) \
202  _(ILOOP,  rbase, ___,  jump,  ___) \
203  _(JLOOP,  rbase, ___,  lit,  ___) \
204  \
205  _(JMP,    rbase, ___,  jump,  ___) \
206  \
207  /*Function headers. I/J = interp/JIT, F/V/C = fixarg/vararg/C func.*/ \
208  _(FUNCF,  rbase, ___,  ___,  ___) \
209  _(IFUNCF, rbase, ___,  ___,  ___) \
210  _(JFUNCF, rbase, ___,  lit,  ___) \
211  _(FUNCV,  rbase, ___,  ___,  ___) \
212  _(IFUNCV, rbase, ___,  ___,  ___) \
213  _(JFUNCV, rbase, ___,  lit,  ___) \
214  _(FUNCC,  rbase, ___,  ___,  ___) \
215  _(FUNCCW, rbase, ___,  ___,  ___) \
216  \
217  /* specific purpose bc. */ \
218  _(VARGN,  dst, ___,  lit,  ___) \
219  _(VARGSTR, dst, ___,  lit,  ___) \
220  _(VPROBENAME, dst, ___,  lit,  ___) \
221  _(VPID,    dst, ___,  lit,  ___) \
222  _(VTID,    dst, ___,  lit,  ___) \
223  _(VUID,    dst, ___,  lit,  ___) \
224  _(VCPU,    dst, ___,  lit,  ___) \
225  _(VEXECNAME, dst, ___,  lit,  ___) \
226  \
227  _(GFUNC,  dst, ___,  ___,  ___) /*load global C function*/
228
229 /* Bytecode opcode numbers. */
230 typedef enum {
231 #define BCENUM(name, ma, mb, mc, mt)    BC_##name,
232     BCDEF(BCENUM)
233 #undef BCENUM
234     BC__MAX
235 } BCOp;
236
237 KP\_STATIC\_ASSERT((int)BC_ISEQV+1 == (int)BC_ISNEV);

```

```

238 KP_STATIC_ASSERT(((int)BC_ISEQV^1) == (int)BC_ISNEV);
239 KP_STATIC_ASSERT(((int)BC_ISEQS^1) == (int)BC_ISNES);
240 KP_STATIC_ASSERT(((int)BC_ISEQN^1) == (int)BC_ISNEN);
241 KP_STATIC_ASSERT(((int)BC_ISEQP^1) == (int)BC_ISNEP);
242 KP_STATIC_ASSERT(((int)BC_ISLT^1) == (int)BC_ISGE);
243 KP_STATIC_ASSERT(((int)BC_ISLE^1) == (int)BC_ISGT);
244 KP_STATIC_ASSERT(((int)BC_ISLT^3) == (int)BC_ISGT);
245 KP_STATIC_ASSERT((int)BC_IST-(int)BC_ISTC == (int)BC_ISF-(int)BC_ISFC);
246 KP_STATIC_ASSERT((int)BC_CALLT-(int)BC_CALL == (int)BC_CALLMT-(int)BC_CALLM);
247 KP_STATIC_ASSERT((int)BC_CALLMT + 1 == (int)BC_CALLT);
248 KP_STATIC_ASSERT((int)BC_RET + 1 == (int)BC_RET);
249 KP_STATIC_ASSERT((int)BC_FORL + 1 == (int)BC_IFORL);
250 KP_STATIC_ASSERT((int)BC_FORL + 2 == (int)BC_JFORL);
251 KP_STATIC_ASSERT((int)BC_ITERL + 1 == (int)BC_IITERL);
252 KP_STATIC_ASSERT((int)BC_ITERL + 2 == (int)BC_JITERL);
253 KP_STATIC_ASSERT((int)BC_LOOP + 1 == (int)BC_ILOOP);
254 KP_STATIC_ASSERT((int)BC_LOOP + 2 == (int)BC_JLOOP);
255 KP_STATIC_ASSERT((int)BC_FUNC + 1 == (int)BC_IFUNC);
256 KP_STATIC_ASSERT((int)BC_FUNC + 2 == (int)BC_JFUNC);
257 KP_STATIC_ASSERT((int)BC_FUNCV + 1 == (int)BC_IFUNCV);
258 KP_STATIC_ASSERT((int)BC_FUNCV + 2 == (int)BC_JFUNCV);
259
260 /* This solves a circular dependency problem, change as needed. */
261 #define FF_next_N 4
262
263 /* Stack slots used by FORI/FORL, relative to operand A. */
264 enum {
265     FORL_IDX, FORL_STOP, FORL_STEP, FORL_EXT
266 };
267
268 /* Bytecode operand modes. ORDER BCMode */
269 typedef enum {
270     /* Mode A must be <= 7 */
271     BCMnone, BCMdst, BCMbase, BCMvar, BCMrbase, BCMuv,
272     BCMlit, BCMlits, BCMpri, BCMnum, BCMstr, BCMtab, BCMfunc,
273     BCMjump, BCMcdata,
274     BCM_max
275 } BCMode;
276
277 #define BCM___ BCMnone
278
279 #define bcmode_a(op) ((BCMode)(bc_mode[op] & 7))
280 #define bcmode_b(op) ((BCMode)((bc_mode[op] >> 3) & 15))
281 #define bcmode_c(op) ((BCMode)((bc_mode[op] >> 7) & 15))
282 #define bcmode_d(op) bcmode_c(op)
283 #define bcmode_hasd(op) ((bc_mode[op] & (15 << 3)) == (BCMnone << 3))
284 #define bcmode_mm(op) ((MMS)(bc_mode[op] >> 11))
285
286 #define BCMODE(name, ma, mb, mc, mm) \
287     (BCM##ma | (BCM##mb << 3) | (BCM##mc << 7)|(MM_##mm << 11)),
288 #define BCMODE_FF 0
289
290 static inline int bc_isret(BCOp op)
291 {
292     return (op == BC_RET || op == BC_RET0 ||
293            op == BC_RET1);
294 }
295
296 /*
297 * Metamethod definition
298 * Note ktap don't use any lua methmethod currently.
299 */
300 typedef enum {
301     MM_lt,
302     MM_le,
303     MM_eq,
304     MM_unm,
305     MM_add,
306     MM_sub,
307     MM_mul,
308     MM_div,
309     MM_mod,
310     MM_pow,
311     MM_concat,
312     MM_gc,
313     MM_index,

```

```

314     MM_newindex,
315     MM_call,
316     MM_MAX,
317     MM_____ = MM_MAX
318 } MMS;
319
320
321 /* -- Bytecode dump format ----- */
322
323 /*
324 ** dump = header proto+ 0U
325 ** header = ESC 'L' 'J' versionB flagsU [namelenU nameB*]
326 ** proto = lengthU pdata
327 ** pdata = phead bcinsw* uvdataH* kgc* knum* [debugB*]
328 ** phead = flagsB numparamsB framesizeB numuvB numkgcU numknU numbcU
329 **         [debuglenU [firstlineU numlineU]]
330 ** kgc = kgctypeU { ktab | (loU hiU) | (rloU rhiU iloU ihiU) | strB* }
331 ** knum = intU0 | (loU1 hiU)
332 ** ktab = narrayU nhashU karray* khash*
333 ** karray = ktabk
334 ** khash = ktabk ktabk
335 ** ktabk = ktabtypeU { intU | (loU hiU) | strB* }
336 **
337 ** B = 8 bit, H = 16 bit, W = 32 bit, U = ULEB128 of W, U0/U1 = ULEB128 of W+1
338 */
339
340 /* Bytecode dump header. */
341 #define BCDUMP_HEAD1      0x15
342 #define BCDUMP_HEAD2      0x22
343 #define BCDUMP_HEAD3      0x06
344
345 /* If you perform any kind of private modifications to the bytecode itself
346 ** or to the dump format, you must set BCDUMP_VERSION to 0x80 or higher.
347 */
348 #define BCDUMP_VERSION      1
349
350 /* Compatibility flags. */
351 #define BCDUMP_F_BE          0x01
352 #define BCDUMP_F_STRIP      0x02
353 #define BCDUMP_F_FFI        0x04
354
355 #define BCDUMP_F_KNOWN      (BCDUMP_F_FFI*2-1)
356
357 /* Type codes for the GC constants of a prototype. Plus length for strings. */
358 enum {
359     BCDUMP_KGC_CHILD, BCDUMP_KGC_TAB, BCDUMP_KGC_I64, BCDUMP_KGC_U64,
360     BCDUMP_KGC_COMPLEX, BCDUMP_KGC_STR
361 };
362
363 /* Type codes for the keys/values of a constant table. */
364 enum {
365     BCDUMP_KTAB_NIL, BCDUMP_KTAB_FALSE, BCDUMP_KTAB_TRUE,
366     BCDUMP_KTAB_INT, BCDUMP_KTAB_NUM, BCDUMP_KTAB_STR
367 };
368
369 #endif /* KTAP_BC_H */

```

[One Level Up](#)

[Top Level](#)

include/ktap_types.h - ktap

Data types defined

- [StkId](#)
- [cdata_number](#)
- [csymbol](#)
- [csymbol_id](#)
- [ffi_state](#)
- [ffi_state](#)
- [ktap_cdata](#)
- [ktap_cdata_t](#)
- [ktap_cfunction](#)
- [ktap_eventdesc](#)
- [ktap_eventdesc_t](#)
- [ktap_func](#)
- [ktap_func_t](#)
- [ktap_global_state](#)
- [ktap_global_state_t](#)
- [ktap_instr_t](#)
- [ktap_node_t](#)
- [ktap_node_t](#)
- [ktap_number](#)
- [ktap_obj](#)
- [ktap_obj_t](#)
- [ktap_option](#)
- [ktap_option_t](#)
- [ktap_proto](#)
- [ktap_proto_t](#)
- [ktap_rawobj](#)
- [ktap_rawobj](#)
- [ktap_state](#)
- [ktap_state_t](#)
- [ktap_stats](#)

- [ktap_stats_t](#)
- [ktap_str](#)
- [ktap_str_t](#)
- [ktap_tab](#)
- [ktap_tab_t](#)
- [ktap_upval](#)
- [ktap_upval_t](#)
- [ktap_val](#)
- [ktap_val_t](#)
- [ptrdiff_t](#)
- [u8](#)

Functions defined

- [set_cfunc](#)
- [set_eventstr](#)
- [set_func](#)
- [set_ip](#)
- [set_kstack](#)
- [set_number](#)
- [set_proto](#)
- [set_string](#)
- [set_table](#)

Macros defined

- [G](#)
- [GCHheader](#)
- [KP_MAX_ABITS](#)
- [KP_MAX_ASIZE](#)
- [KP_MAX_BCINS](#)
- [KP_MAX_CACHED_CFUNCTION](#)
- [KP_MAX_COLOSIZE](#)
- [KP_MAX_HBITS](#)
- [KP_MAX_LINE](#)
- [KP_MAX_LOCVAR](#)

- [KP_MAX_MEMPOOL_SIZE](#)
- [KP_MAX_SLOTS](#)
- [KP_MAX_STACK_DEPTH](#)
- [KP_MAX_STR](#)
- [KP_MAX_STRNUM](#)
- [KP_MAX_STRTAB](#)
- [KP_MAX_UPVAL](#)
- [KP_MAX_XLEVEL](#)
- [KTAP_AUTHOR](#)
- [KTAP_CMD_IOC_EXIT](#)
- [KTAP_CMD_IOC_RUN](#)
- [KTAP_CMD_IOC_VERSION](#)
- [KTAP_COPYRIGHT](#)
- [KTAP_ERROR](#)
- [KTAP_EXIT](#)
- [KTAP_QL](#)
- [KTAP_QS](#)
- [KTAP_RIDX_GLOBALS](#)
- [KTAP_RIDX_LAST](#)
- [KTAP_RUNNING](#)
- [KTAP_STATS](#)
- [KTAP_TCDATA](#)
- [KTAP_TCFUNC](#)
- [KTAP_TEVENTSTR](#)
- [KTAP_TFALSE](#)
- [KTAP_TFUNC](#)
- [KTAP_TKIP](#)
- [KTAP_TKSTACK](#)
- [KTAP_TLIGHTUD](#)
- [KTAP_TNIL](#)
- [KTAP_TNUM](#)
- [KTAP_TNUMX](#)
- [KTAP_TPROTO](#)

- [KTAP_TRACE_END](#)
- [KTAP_TSTR](#)
- [KTAP_TTAB](#)
- [KTAP_TTHREAD](#)
- [KTAP_TTRUE](#)
- [KTAP_TUDATA](#)
- [KTAP_TUIP](#)
- [KTAP_TUPVAL](#)
- [KTAP_VERSION](#)
- [KTAP_VERSION_MAJOR](#)
- [KTAP_VERSION_MINOR](#)
- [MYINT](#)
- [PROTO_CHILD](#)
- [PROTO_CLCOUNT](#)
- [PROTO_CLC_BITS](#)
- [PROTO_CLC_POLY](#)
- [PROTO_FFI](#)
- [PROTO_FIXUP_RETURN](#)
- [PROTO_HAS_RETURN](#)
- [PROTO_ILOOP](#)
- [PROTO_NOJIT](#)
- [PROTO_UV_IMMUTABLE](#)
- [PROTO_UV_LOCAL](#)
- [PROTO_VARARG](#)
- [VERSION](#)
- [_KTAP_TYPES_H](#)
- [boolvalue](#)
- [cdvalue](#)
- [clvalue](#)
- [fvalue](#)
- [gch](#)
- [gco2ts](#)
- [gco2uv](#)

- [gcvalue](#)
- [getstr](#)
- [hvalue](#)
- [incr_top](#)
- [is_bool](#)
- [is_btrace](#)
- [is_cdata](#)
- [is_cfunc](#)
- [is_eventstr](#)
- [is_false](#)
- [is_function](#)
- [is_kip](#)
- [is_nil](#)
- [is_number](#)
- [is_proto](#)
- [is_string](#)
- [is_table](#)
- [is_true](#)
- [itype](#)
- [nvalue](#)
- [obj2gco](#)
- [phvalue](#)
- [proto_bc](#)
- [proto_bcpos](#)
- [proto_chunkname](#)
- [proto_kgc](#)
- [proto_lineinfo](#)
- [proto_uv](#)
- [proto_uvinfo](#)
- [proto_varinfo](#)
- [ptvalue](#)
- [pvalue](#)
- [rawtsvalue](#)

- [set_bool](#)
- [set_cdata](#)
- [set_nil](#)
- [set_obj](#)
- [setitype](#)
- [svalue](#)
- [val](#)

Source code

```

1  /*
2  * ktap types definition.
3  *
4  * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
5  * Copyright (C) 2005-2014 Mike Pall.
6  * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
7  */
8
9  #ifndef __KTAP_TYPES_H__
10 #define __KTAP_TYPES_H__
11
12 #ifdef __KERNEL__
13 #include <linux/perf_event.h>
14 #else
15 typedef char u8;
16 #include <stdlib.h>
17 #include <stdio.h>
18 #include <string.h>
19 #include <stdint.h>
20 typedef int ptrdiff_t;
21 #endif
22
23 #include "../include/ktap_bc.h"
24
25 /* Various VM limits. */
26 #define KP_MAX_MEMPOOL_SIZE 10000 /* Max. mempool size(Kbytes). */
27 #define KP_MAX_STR 512 /* Max. string length. */
28 #define KP_MAX_STRNUM 9999 /* Max. string number. */
29
30 #define KP_MAX_STARTAB (1<<26) /* Max. string table size. */
31 #define KP_MAX_HBITS 26 /* Max. hash bits. */
32 #define KP_MAX_ABITS 28 /* Max. bits of array key. */
33 #define KP_MAX_ASIZE ((1<<(KP_MAX_ABITS-1))+1) /* Max. array part size. */
34 #define KP_MAX_COLOSIZE 16 /* Max. elems for colocated array. */
35
36 #define KP_MAX_LINE 1000 /* Max. source code line number. */
37 #define KP_MAX_XLEVEL 200 /* Max. syntactic nesting level. */
38 #define KP_MAX_BCINS (1<<26) /* Max. # of bytecode instructions. */
39 #define KP_MAX_SLOTS 250 /* Max. # of slots in a ktap func. */
40 #define KP_MAX_LOCVAR 200 /* Max. # of local variables. */
41 #define KP_MAX_UPVAL 60 /* Max. # of upvalues. */
42
43 #define KP_MAX_CACHED_CFUNCTION 128 /* Max. cached global cfunction */
44
45 #define KP_MAX_STACK_DEPTH 50 /* Max. stack depth */
46
47 /*
48 * The first argument type of kdebug.trace_by_id()
49 * The value is a userspace memory pointer.
50 * Maybe embed it into the trunk file in future.
51 */
52 typedef struct ktap_eventdesc {
53     int nr; /* the number of events id */
54     int *id_arr; /* events id array */
55     char *filter;
56 } ktap_eventdesc_t;

```

```

57
58
59 /* ktap option for each script */
60 typedef struct ktap_option {
61     char *trunk; /* __user */
62     int trunk_len;
63     int argc;
64     char **argv; /* __user */
65     int verbose;
66     int trace_pid;
67     int workload;
68     int trace_cpu;
69     int print_timestamp;
70     int quiet;
71     int dry_run;
72 } ktap_option_t;
73
74 /*
75 * Ioctls that can be done on a ktap fd:
76 * todo: use _IO macro in include/uapi/asm-generic/ioctl.h
77 */
78 #define KTAP_CMD_IOC_VERSION      ('$' + 0)
79 #define KTAP_CMD_IOC_RUN          ('$' + 1)
80 #define KTAP_CMD_IOC_EXIT        ('$' + 3)
81
82 #define KTAP_VERSION_MAJOR        "0"
83 #define KTAP_VERSION_MINOR        "4"
84
85 #define KTAP_VERSION              "ktap " KTAP_VERSION_MAJOR "." KTAP_VERSION_MINOR
86 #define KTAP_AUTHOR                "Jovi Zhangwei <jovi.zhangwei@gmail.com>"
87 #define KTAP_COPYRIGHT             KTAP_VERSION " Copyright (C) 2012-2014, " KTAP_AUTHOR
88
89 #define MYINT(s)                   (s[0] - '0')
90 #define VERSION                    (MYINT(KTAP_VERSION_MAJOR) * 16 + MYINT(KTAP_VERSION_MINOR))
91
92 typedef long ktap_number;
93 typedef int ktap_instr_t;
94 typedef union ktap_obj ktap_obj_t;
95
96 struct ktap_state;
97 typedef int (*ktap_cfunction) (struct ktap_state *ks);
98
99 /* ktap_val_t is basic value type in ktap stack, for reference all objects */
100 typedef struct ktap_val {
101     union {
102         ktap_obj_t *gc; /* collectable objects, str/tab/... */
103         void *p; /* light userdata */
104         ktap_cfunction f; /* light C functions */
105         ktap_number n; /* numbers */
106         struct {
107             uint16_t depth; /* stack depth */
108             uint16_t skip; /* skip stack entries */
109         } stack;
110     } val;
111     union {
112         int type; /* type for above val */
113         const unsigned int *pcr; /* Overlaps PC for ktap frames.*/
114     };
115 } ktap_val_t;
116
117 typedef ktap_val_t *StkId;
118
119 #define GCHheader ktap_obj_t *nextgc; u8 gct;
120
121 typedef struct ktap_str {
122     GCHheader;
123     u8 reserved; /* Used by lexer for fast lookup of reserved words. */
124     u8 extra;
125     unsigned int hash;
126     int len; /* number of characters in string */
127 } ktap_str_t;
128
129 typedef struct ktap_upval {
130     GCHheader;
131     uint8_t closed; /* Set if closed (i.e. uv->v == &uv->u.value). */
132     uint8_t immutable; /* Immutable value. */

```

```

133     union {
134         ktap_val_t tv; /* If closed: the value itself. */
135         struct { /* If open: double linked list, anchored at thread. */
136             struct ktap_upval *prev;
137             struct ktap_upval *next;
138         };
139     };
140     ktap_val_t *v; /* Points to stack slot (open) or above (closed). */
141 } ktap_upval_t;
142
143
144 typedef struct ktap_func {
145     GCHheader;
146     u8 nupvalues;
147     BCIns *pc;
148     struct ktap_proto *p;
149     struct ktap_upval *upvals[1]; /* list of upvalues */
150 } ktap_func_t;
151
152 typedef struct ktap_proto {
153     GCHheader;
154     uint8_t numparams; /* Number of parameters. */
155     uint8_t framesize; /* Fixed frame size. */
156     int sizebc; /* Number of bytecode instructions. */
157     ktap_obj_t *gclist;
158     void *k; /* Split constant array (points to the middle). */
159     void *uv; /* Upvalue list. local slot|0x8000 or parent uv idx. */
160     int sizekgc; /* Number of collectable constants. */
161     int sizekn; /* Number of lua_Number constants. */
162     int sizept; /* Total size including colocated arrays. */
163     uint8_t sizeuv; /* Number of upvalues. */
164     uint8_t flags; /* Miscellaneous flags (see below). */
165
166     /* --- The following fields are for debugging/tracebacks only --- */
167     ktap_str_t *chunkname; /* Chunk name this function was defined in. */
168     BCLine firstline; /* First line of the function definition. */
169     BCLine numline; /* Number of lines for the function definition. */
170     void *lineinfo; /* Compressed map from bytecode ins. to source line. */
171     void *uvinfo; /* Upvalue names. */
172     void *varinfo; /* Names and compressed extents of local variables. */
173 } ktap_proto_t;
174
175 /* Flags for prototype. */
176 #define PROTO_CHILD 0x01 /* Has child prototypes. */
177 #define PROTO_VARARG 0x02 /* Vararg function. */
178 #define PROTO_FFI 0x04 /* Uses BC_KCDATA for FFI datatypes. */
179 #define PROTO_NOJIT 0x08 /* JIT disabled for this function. */
180 #define PROTO_ILOOP 0x10 /* Patched bytecode with ILOOP etc. */
181 /* Only used during parsing. */
182 #define PROTO_HAS_RETURN 0x20 /* Already emitted a return. */
183 #define PROTO_FIXUP_RETURN 0x40 /* Need to fixup emitted returns. */
184 /* Top bits used for counting created closures. */
185 #define PROTO_CLC_COUNT 0x20 /* Base of saturating 3 bit counter. */
186 #define PROTO_CLC_BITS 3
187 #define PROTO_CLC_POLY (3*PROTO_CLC_COUNT) /* Polymorphic threshold. */
188
189 #define PROTO_UV_LOCAL 0x8000 /* Upvalue for local slot. */
190 #define PROTO_UV_IMMUTABLE 0x4000 /* Immutable upvalue. */
191
192 #define proto_kgc(pt, idx) (((ktap_obj_t *) (pt)->k)[idx])
193 #define proto_bc(pt) ((BCIns *) ((char *) (pt) + sizeof(ktap_proto_t)))
194 #define proto_bcpos(pt, pc) ((BCPos)((pc) - proto_bc(pt)))
195 #define proto_uv(pt) ((uint16_t *) (pt)->uv)
196
197 #define proto_chunkname(pt) ((pt)->chunkname)
198 #define proto_lineinfo(pt) ((const void *) (pt)->lineinfo)
199 #define proto_uvinfo(pt) ((const uint8_t *) (pt)->uvinfo)
200 #define proto_varinfo(pt) ((const uint8_t *) (pt)->varinfo)
201
202
203 typedef struct ktap_node_t {
204     ktap_val_t val; /* Value object. Must be first field. */
205     ktap_val_t key; /* Key object. */
206     struct ktap_node_t *next; /* hash chain */
207 } ktap_node_t;
208

```

```

209 /* ktap_tab */
210 typedef struct ktap_tab {
211     GCHheader;
212 #ifdef __KERNEL__
213     arch_spinlock_t lock;
214 #endif
215     ktap_val_t *array;    /* Array part. */
216     ktap_node_t *node;   /* Hash part. */
217     ktap_node_t *freetop; /* any free position is before this position */
218
219     uint32_t asize;      /* Size of array part (keys [0, asize-1]). */
220     uint32_t hmask;     /* log2 of size of `node' array */
221
222     uint32_t hnum;      /* number of all nodes */
223
224     ktap_obj_t *gclist;
225 } ktap_tab_t;
226
227 #ifdef CONFIG_KTAP_FFI
228 typedef int csymbol_id;
229 typedef uint64_t cdata_number;
230 typedef struct csymbol csymbol;
231
232 /* global ffi state maintained in each ktap vm instance */
233 typedef struct ffi_state {
234     ktap_tab_t *ctable;
235     int csym_nr;
236     csymbol *csym_arr;
237 } ffi_state;
238
239 /* instance of csymbol */
240 typedef struct ktap_cdata {
241     GCHheader;
242     csymbol_id id;
243     union {
244         cdata_number i;
245         struct {
246             void *addr;
247             int nmemb;    /* number of memory block */
248         } p;             /* pointer data */
249         void *rec;      /* struct member or union data */
250     } u;
251 } ktap_cdata_t;
252 #endif
253
254 typedef struct ktap_stats {
255     int mem_allocated;
256     int nr_mem_allocate;
257     int events_hits;
258     int events_missed;
259 } ktap_stats_t;
260
261 #define KTAP_STATS(ks)    this_cpu_ptr(G(ks)->stats)
262
263
264 #define KTAP_RUNNING      0 /* normal running state */
265 #define KTAP_TRACE_END    1 /* running in trace_end function */
266 #define KTAP_EXIT         2 /* normal exit, set when call exit() */
267 #define KTAP_ERROR        3 /* error state, called by kp\_error */
268
269 typedef struct ktap_global_state {
270     void *mempool;      /* string memory pool */
271     void *mp_freepos;   /* free position in memory pool */
272     int mp_size;        /* memory pool size */
273 #ifdef __KERNEL__
274     arch_spinlock_t mp_lock; /* mempool lock */
275 #endif
276
277     ktap_str_t **strhash; /* String hash table (hash chain anchors). */
278     int strmask;         /* String hash mask (size of hash table-1). */
279     int strnum;         /* Number of strings in hash table. */
280 #ifdef __KERNEL__
281     arch_spinlock_t str_lock; /* string operation lock */
282 #endif
283
284     ktap_val_t registry;

```

```

285 ktap_tab_t *gtab; /* global table contains cfunction and args */
286 ktap_obj_t *allgc; /* list of all collectable objects */
287 ktap_upval_t uvhead; /* head of list of all open upvalues */
288
289 struct ktap_state *mainthread; /*main state */
290 int state; /* status of ktapvm, KTAP_RUNNING, KTAP_TRACE_END, etc */
291 #ifdef __KERNEL__
292 /* reserved global percpu data */
293 void __percpu *percpu_state[PERF_NR_CONTEXTS];
294 void __percpu *percpu_print_buffer[PERF_NR_CONTEXTS];
295 void __percpu *percpu_temp_buffer[PERF_NR_CONTEXTS];
296
297 /* for recursion tracing check */
298 int __percpu *recursion_context[PERF_NR_CONTEXTS];
299
300 ktap_option_t *parm; /* ktap options */
301 pid_t trace_pid;
302 struct task_struct *trace_task;
303 cpumask_var_t cpumask;
304 struct ring_buffer *buffer;
305 struct dentry *trace_pipe_dentry;
306 struct task_struct *task;
307 int trace_enabled;
308 int wait_user; /* flag to indicat waiting user consume content */
309
310 struct list_head timers; /* timer list */
311 struct ktap_stats __percpu *stats; /* memory allocation stats */
312 struct list_head events_head; /* probe event list */
313
314 ktap_func_t *trace_end_closure; /* trace_end closure */
315 #ifdef CONFIG_KTAP_FFI
316 ffi_state ffig;
317 #endif
318
319 /* C function table for fast call */
320 int nr_builtin_cfunction;
321 ktap_cfunction gfunc_tbl[KP_MAX_CACHED_CFUNCTION];
322 #endif
323 } ktap_global_state_t;
324
325
326 typedef struct ktap_state {
327 ktap_global_state_t *g; /* global state */
328 int stop; /* don't enter tracing handler if stop is 1 */
329 StkId top; /* stack top */
330 StkId func; /* execute light C function */
331 StkId stack_last; /* last stack pointer */
332 StkId stack; /* ktap stack, percpu pre-reserved */
333 ktap_upval_t *openupval; /* opened upvals list */
334
335 #ifdef __KERNEL__
336 /* current fired event which allocated on stack */
337 struct ktap_event_data *current_event;
338 #endif
339 } ktap_state_t;
340
341 #define G(ks) (ks->g)
342
343 typedef struct ktap_rawobj {
344 GCHheader;
345 void *v;
346 } ktap_rawobj;
347
348 /*
349 * Union of all collectable objects
350 */
351 union ktap_obj {
352 struct { GCHheader } gch;
353 struct ktap_str ts;
354 struct ktap_func fn;
355 struct ktap_tab h;
356 struct ktap_proto pt;
357 struct ktap_upval uv;
358 struct ktap_state th; /* thread */
359 struct ktap_rawobj rawobj;
360 #ifdef CONFIG_KTAP_FFI

```

```

361     struct ktap_cdata cd;
362 #endif
363 };
364
365 #define gch(o)          (&(o)->gch)
366
367 /* macros to convert a ktap_obj_t into a specific value */
368 #define gco2ts(o)      (&((o)->ts))
369 #define gco2uv(o)      (&((o)->uv))
370 #define obj2gco(v)     ((ktap_obj_t *)(v))
371
372 /* predefined values in the registry */
373 #define KTAP_RIDX_GLOBALS    1
374 #define KTAP_RIDX_LAST      KTAP_RIDX_GLOBALS
375
376 /* ktap object types */
377 #define KTAP_TNIL           (~0u)
378 #define KTAP_TFALSE        (~1u)
379 #define KTAP_TTRUE         (~2u)
380 #define KTAP_TNUM          (~3u)
381 #define KTAP_TLIGHTUD      (~4u)
382 #define KTAP_TSTR          (~5u)
383 #define KTAP_TUPVAL        (~6u)
384 #define KTAP_TTHREAD       (~7u)
385 #define KTAP_TPROTO        (~8u)
386 #define KTAP_TFUNC         (~9u)
387 #define KTAP_TCFUNC        (~10u)
388 #define KTAP_TCDATA        (~11u)
389 #define KTAP_TTAB          (~12u)
390 #define KTAP_TUDATA        (~13u)
391
392 /* Specific types */
393 #define KTAP_TEVENTSTR      (~14u) /* argstr */
394 #define KTAP_TKSTACK        (~15u) /* stack(), not intern to string yet */
395 #define KTAP_TKIP           (~16u) /* kernel function ip address */
396 #define KTAP_TUIP           (~17u) /* userspace function ip address */
397
398 /* This is just the canonical number type used in some places. */
399 #define KTAP_TNUMX         (~18u)
400
401
402 #define itype(o)           ((o)->type)
403 #define setitype(o, t)    ((o)->type = (t))
404
405 #define val_(o)           ((o)->val)
406 #define gcvalue(o)       (val_(o).gc)
407
408 #define nvalue(o)         (val_(o).n)
409 #define boolvalue(o)     (KTAP_TFALSE - (o)->type)
410 #define hvalue(o)        (&val_(o).gc->h)
411 #define phvalue(o)       (&val_(o).gc->ph)
412 #define clvalue(o)       (&val_(o).gc->fn)
413 #define ptvalue(o)       (&val_(o).gc->pt)
414
415 #define getstr(ts)        (const char *)((ts) + 1)
416 #define rawtsvalue(o)    (&val_(o).gc->ts)
417 #define svalue(o)        getstr(rawtsvalue(o))
418
419 #define pvalue(o)        (&val_(o).p)
420 #define fvalue(o)        (val_(o).f)
421 #ifdef CONFIG_KTAP_FFI
422 #define cdvalue(o)       (&val_(o).gc->cd)
423 #endif
424
425 #define is_nil(o)         (itype(o) == KTAP_TNIL)
426 #define is_false(o)      (itype(o) == KTAP_TFALSE)
427 #define is_true(o)       (itype(o) == KTAP_TTRUE)
428 #define is_bool(o)      (is_false(o) || is_true(o))
429 #define is_string(o)     (itype(o) == KTAP_TSTR)
430 #define is_number(o)     (itype(o) == KTAP_TNUM)
431 #define is_table(o)      (itype(o) == KTAP_TTAB)
432 #define is_proto(o)      (itype(o) == KTAP_TPROTO)
433 #define is_function(o)   (itype(o) == KTAP_TFUNC)
434 #define is_cfunc(o)      (itype(o) == KTAP_TCFUNC)
435 #define is_eventstr(o)   (itype(o) == KTAP_TEVENTSTR)
436 #define is_kip(o)        (itype(o) == KTAP_TKIP)

```

```

437 #define is_btrace(o)          (itype(o) == KTAP_TBTRACE)
438 #ifndef CONFIG_KTAP_FFI
439 #define is_cdata(o)          (itype(o) == KTAP_TCDATA)
440 #endif
441
442
443 #define set_nil(o)            ((o)->type = KTAP_TNIL)
444 #define set_bool(o, x)       ((o)->type = KTAP_TFALSE-(uint32_t)(x))
445
446 static inline void set_number(ktap_val_t *o, ktap_number n)
447 {
448     setitype(o, KTAP_TNUM);
449     o->val.n = n;
450 }
451
452 static inline void set_string(ktap_val_t *o, const ktap_str_t *str)
453 {
454     setitype(o, KTAP_TSTR);
455     o->val.gc = (ktap_obj_t *)str;
456 }
457
458 static inline void set_table(ktap_val_t *o, ktap_tab_t *tab)
459 {
460     setitype(o, KTAP_TTAB);
461     o->val.gc = (ktap_obj_t *)tab;
462 }
463
464 static inline void set_proto(ktap_val_t *o, ktap_proto_t *pt)
465 {
466     setitype(o, KTAP_TPROTO);
467     o->val.gc = (ktap_obj_t *)pt;
468 }
469
470 static inline void set_kstack(ktap_val_t *o, uint16_t depth, uint16_t skip)
471 {
472     setitype(o, KTAP_TKSTACK);
473     o->val.stack.depth = depth;
474     o->val.stack.skip = skip;
475 }
476
477 static inline void set_func(ktap_val_t *o, ktap_func_t *fn)
478 {
479     setitype(o, KTAP_TFUNC);
480     o->val.gc = (ktap_obj_t *)fn;
481 }
482
483 static inline void set_cfunc(ktap_val_t *o, ktap_cfunction fn)
484 {
485     setitype(o, KTAP_TCFUNC);
486     o->val.f = fn;
487 }
488
489 static inline void set_eventstr(ktap_val_t *o)
490 {
491     setitype(o, KTAP_TEVENTSTR);
492 }
493
494 static inline void set_ip(ktap_val_t *o, unsigned long addr)
495 {
496     setitype(o, KTAP_TKIP);
497     o->val.n = addr;
498 }
499
500
501 #ifndef CONFIG_KTAP_FFI
502 #define set_cdata(o, x)        { setitype(o, KTAP_TCDATA); (o)->val.gc = x; }
503 #endif
504
505 #define set_obj(o1, o2)        { *(o1) = *(o2); }
506
507 #define incr_top(ks)           {ks->top++;}
508
509 /*
510 * KTAP\_QL describes how error messages quote program elements.
511 * CHANGE it if you want a different appearance.
512 */

```

```
513 #define KTAP_QL(x)      "" x ""
514 #define KTAP_QS        KTAP\_QL("s")
515
516 #endif /* \_KTAP\_TYPES\_H */
517
```

[One Level Up](#)

[Top Level](#)

userspace/kp_main.c - ktap

Global variables defined

- [dry_run](#)
- [dump_bytecode](#)
- [forks](#)
- [ktap_trunk_mem_size](#)
- [online_src](#)
- [output_filename](#)
- [print_timestamp](#)
- [quiet](#)
- [script_args_end](#)
- [script_args_start](#)
- [script_file](#)
- [trace_cpu](#)
- [trace_pid](#)
- [uparm](#)
- [verbose](#)
- [workload_argv](#)

Data types defined

- [binary_base](#)

Functions defined

- [compile](#)
- [fork_workload](#)
- [kp_writer](#)
- [main](#)
- [parse](#)
- [parse_option](#)
- [print_symbol](#)
- [run_ktapvm](#)
- [usage](#)

Macros defined

- [KTAPVM_PATH](#)
- [SIMPLE_ONE_LINER_FMT](#)
- [handle_error](#)

Source code

```

1  /*
2  * main.c - ktap compiler and loader entry
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <stdio.h>
23 #include <stdlib.h>
24 #include <sched.h>
25 #include <string.h>
26 #include <signal.h>
27 #include <stdarg.h>
28 #include <sys/mman.h>
29 #include <sys/stat.h>
30 #include <sys/ioctl.h>
31 #include <sys/types.h>
32 #include <unistd.h>
33 #include <fcntl.h>
34 #include <math.h>
35 #include <linux/errno.h>
36
37 #include "../include/ktap_types.h"
38 #include "kp_lex.h"
39 #include "kp_parse.h"
40 #include "kp_symbol.h"
41 #include "cparser.h"
42
43 static void usage(const char *msg_fmt, ...)
44 {
45     va_list ap;
46
47     va_start(ap, msg_fmt);
48     vfprintf(stderr, msg_fmt, ap);
49     va_end(ap);
50
51     fprintf(stderr,
52 "Usage: ktap [options] file [script args] -- cmd [args]\n"
53 "   or: ktap [options] -e one-liner -- cmd [args]\n"
54 "\n"
55 "Options and arguments:\n"
56 " -o file      : send script output to file, instead of stderr\n"
57 " -p pid       : specific tracing pid\n"
58 " -C cpu       : cpu to monitor in system-wide\n"
59 " -T           : show timestamp for event\n"
60 " -V           : show version\n"
61 " -v          : enable verbose mode\n"
62 " -q          : suppress start tracing message\n"
63 " -d          : dry run mode(register NULL callback to perf events)\n"
64 " -s          : simple event tracing\n"
65 " -b          : list byte codes\n"

```

```

66 " -le [glob]      : list pre-defined events in system\n"
67 #ifndef NO_LIBELF
68 " -lf DSO        : list available functions from DSO\n"
69 " -lm DSO        : list available sdt notes from DSO\n"
70 #endif
71 " file           : program read from script file\n"
72 " -- cmd [args]  : workload to tracing\n");
73
74     exit(EXIT_FAILURE);
75 }
76
77 #define handle_error(str) do { perror(str); exit(-1); } while(0)
78
79 ktap_option_t uparm;
80 static int ktap_trunk_mem_size = 1024;
81
82 static int kp_writer(const void* p, size_t sz, void* ud)
83 {
84     if (uparm.trunk_len + sz > ktap_trunk_mem_size) {
85         int new_size = (uparm.trunk_len + sz) * 2;
86         uparm.trunk = realloc(uparm.trunk, new_size);
87         ktap_trunk_mem_size = new_size;
88     }
89
90     memcpy(uparm.trunk + uparm.trunk_len, p, sz);
91     uparm.trunk_len += sz;
92
93     return 0;
94 }
95
96
97 static int forks;
98 static char **workload_argv;
99
100 static int fork_workload(int ktap_fd)
101 {
102     int pid;
103
104     pid = fork();
105     if (pid < 0)
106         handle_error("failed to fork");
107
108     if (pid > 0)
109         return pid;
110
111     signal(SIGTERM, SIG_DFL);
112
113     execvp("", workload_argv);
114
115     /*
116      * waiting ktapvm prepare all tracing event
117      * make it more robust in future.
118      */
119     pause();
120
121     execvp(workload_argv[0], workload_argv);
122
123     perror(workload_argv[0]);
124     exit(-1);
125
126     return -1;
127 }
128
129 #define KTAPVM_PATH "/sys/kernel/debug/ktap/ktapvm"
130
131 static char *output_filename;
132
133 static int run_ktapvm()
134 {
135     int ktapvm_fd, ktap_fd;
136     int ret;
137
138     ktapvm_fd = open(KTAPVM_PATH, O_RDONLY);
139     if (ktapvm_fd < 0)
140         handle_error("open " KTAPVM_PATH " failed");
141

```

```

142 ktap_fd = ioctl(ktapvm_fd, 0, NULL);
143 if (ktap_fd < 0)
144     handle\_error("ioctl ktapvm failed");
145
146 kp\_create\_reader(output\_filename);
147
148 if (forks) {
149     uparm.trace\_pid = fork\_workload(ktap_fd);
150     uparm.workload = 1;
151 }
152
153 ret = ioctl(ktap_fd, KTAP\_CMD\_IOC\_RUN, &uparm);
154 switch (ret) {
155 case -EPERM:
156 case -EACCES:
157     fprintf(stderr, "You may not have permission to run ktap\n");
158     break;
159 }
160
161 close(ktap_fd);
162 close(ktapvm_fd);
163
164 return ret;
165 }
166
167 int verbose;
168 static int quiet;
169 static int dry_run;
170 static int dump_bytecode;
171 static char oneline_src[1024];
172 static int trace_pid = -1;
173 static int trace_cpu = -1;
174 static int print_timestamp;
175
176 #define SIMPLE_ONE_LINER_FMT \
177     "trace %s { print(cpu(), tid(), execname(), argstr) }"
178
179 static const char *script_file;
180 static int script_args_start;
181 static int script_args_end;
182
183 #ifndef NO_LIBELF
184 struct binary_base
185 {
186     int type;
187     const char *binary;
188 };
189 static int print_symbol(const char *name, vaddr\_t addr, void *arg)
190 {
191     struct binary\_base *base = (struct binary\_base *)arg;
192     const char *type = base->type == FIND\_SYMBOL ?
193         "probe" : "sdt";
194
195     printf("%s:%s:%s\n", type, base->binary, name);
196     return 0;
197 }
198 #endif
199
200 static void parse_option(int argc, char **argv)
201 {
202     char pid[32] = {0};
203     char cpu_str[32] = {0};
204     char *next_arg;
205     int i, j;
206
207     for (i = 1; i < argc; i++) {
208         if (argv[i][0] != '-') {
209             script\_file = argv[i];
210             if (!script\_file)
211                 usage("");
212
213             script\_args\_start = i + 1;
214             script\_args\_end = argc;
215
216             for (j = i + 1; j < argc; j++) {
217                 if (argv[j][0] == '-' && argv[j][1] == '-')

```

```

218         goto found_cmd;
219     }
220
221     return;
222 }
223
224 if (argv[i][0] == '-' && argv[i][1] == '-') {
225     j = i;
226     goto found_cmd;
227 }
228
229 next_arg = argv[i + 1];
230
231 switch (argv[i][1]) {
232 case 'o':
233     output\_filename = malloc(strlen(next_arg) + 1);
234     if (!output\_filename)
235         return;
236
237     strncpy(output\_filename, next_arg, strlen(next_arg));
238     i++;
239     break;
240 case 'e':
241     strncpy(online\_src, next_arg, strlen(next_arg));
242     i++;
243     break;
244 case 'p':
245     strncpy(pid, next_arg, strlen(next_arg));
246     trace\_pid = atoi(pid);
247     i++;
248     break;
249 case 'C':
250     strncpy(cpu_str, next_arg, strlen(next_arg));
251     trace\_cpu = atoi(cpu_str);
252     i++;
253     break;
254 case 'T':
255     print\_timestamp = 1;
256     break;
257 case 'v':
258     verbose = 1;
259     break;
260 case 'q':
261     quiet = 1;
262     break;
263 case 'd':
264     dry\_run = 1;
265     break;
266 case 's':
267     sprintf(online\_src, SIMPLE ONE LINER FMT, next_arg);
268     i++;
269     break;
270 case 'b':
271     dump\_bytecode = 1;
272     break;
273 case 'l': /* list available events */
274     switch (argv[i][2]) {
275     case 'e': /* tracepoints */
276         list\_available\_events(next_arg);
277         exit(EXIT_SUCCESS);
278 #ifndef NO_LIBELF
279     case 'f': /* functions in DSO */
280     case 'm': /* static marks in DSO */ {
281         const char *binary = next_arg;
282         int type = argv[i][2] == 'f' ?
283             FIND\_SYMBOL : FIND\_STAPSDT\_NOTE;
284         struct binary\_base base = {
285             .type = type,
286             .binary = binary,
287         };
288         int ret;
289
290         ret = parse\_dso\_symbols(binary, type,
291             print\_symbol,
292             (void *)&base);
293         if (ret <= 0) {

```

```

294         fprintf(stderr,
295         "error: no symbols in binary %s\n",
296         binary);
297         exit(EXIT_FAILURE);
298     }
299     exit(EXIT_SUCCESS);
300 }
301 #endif
302     default:
303         exit(EXIT_FAILURE);
304     }
305     break;
306     case 'V':
307 #ifdef CONFIG_KTAP_FFI
308         usage("%s (with FFI)\n\n", KTAP\_VERSION);
309 #else
310         usage("%s\n\n", KTAP\_VERSION);
311 #endif
312     break;
313     case '?':
314     case 'h':
315         usage("");
316     break;
317     default:
318         usage("wrong argument\n");
319     break;
320 }
321 }
322
323     return;
324
325 found_cmd:
326     script\_args\_end = j;
327     forks = 1;
328     workload\_argv = &argv[j + 1];
329 }
330
331 static ktap\_proto\_t *parse(const char *chunkname, const char *src)
332 {
333     LexState ls;
334
335     ls.chunkarg = chunkname ? chunkname : "?";
336     kp\_lex\_init();
337     kp\_buf\_init(&ls.sb);
338     kp\_lex\_setup(&ls, src);
339     return kp\_parse(&ls);
340 }
341
342 static void compile(const char *input)
343 {
344     ktap\_proto\_t *pt;
345     char *buff;
346     struct stat sb;
347     int fdin;
348
349     kp\_str\_resize();
350
351     if (oneline\_src[0] != '\0') {
352         ffi\_cparser\_init();
353         pt = parse(input, oneline\_src);
354         goto dump;
355     }
356
357     fdin = open(input, O_RDONLY);
358     if (fdin < 0) {
359         fprintf(stderr, "open file %s failed\n", input);
360         exit(-1);
361     }
362
363     if (fstat(fdin, &sb) == -1)
364         handle\_error("fstat failed");
365
366     buff = mmap(NULL, sb.st_size, PROT_READ, MAP_PRIVATE, fdin, 0);
367     if (buff == MAP_FAILED)
368         handle\_error("mmap failed");
369

```

```

370     ffi\_cparser\_init();
371     pt = parse(input, buff);
372
373     munmap(buff, sb.st_size);
374     close(fdin);
375
376 dump:
377     if (dump\_bytecode) {
378 #ifndef CONFIG_KTAP_FFI
379         kp\_dump\_csymbols();
380 #endif
381         kp\_dump\_proto(pt);
382         exit(0);
383     }
384
385     /* bcwrite */
386     uparm.trunk = malloc(ktap\_trunk\_mem\_size);
387     if (!uparm.trunk)
388         handle\_error("malloc failed");
389
390     kp\_bcwrite(pt, kp\_writer, NULL, 0);
391     ffi\_cparser\_free();
392 }
393
394 int main(int argc, char **argv)
395 {
396     char **ktapvm_argv;
397     int new_index, i;
398     int ret;
399
400     if (argc == 1)
401         usage("");
402
403     parse\_option(argc, argv);
404
405     if (oneline\_src[0] != '\0')
406         script\_file = "(command line)";
407
408     compile(script\_file);
409
410     ktapvm_argv = (char **)malloc(sizeof(char *)*(script\_args\_end -
411         script\_args\_start + 1));
412     if (!ktapvm_argv) {
413         fprintf(stderr, "cannot allocate ktapvm_argv\n");
414         return -1;
415     }
416
417     ktapvm_argv[0] = malloc(strlen(script\_file) + 1);
418     if (!ktapvm_argv[0]) {
419         fprintf(stderr, "cannot allocate memory\n");
420         return -1;
421     }
422     strcpy(ktapvm_argv[0], script\_file);
423     ktapvm_argv[0][strlen(script\_file)] = '\0';
424
425     /* pass rest argv into ktapvm */
426     new_index = 1;
427     for (i = script\_args\_start; i < script\_args\_end; i++) {
428         ktapvm_argv[new_index] = malloc(strlen(argv[i]) + 1);
429         if (!ktapvm_argv[new_index]) {
430             fprintf(stderr, "cannot allocate memory\n");
431             free(ktapvm_argv);
432             return -1;
433         }
434         strcpy(ktapvm_argv[new_index], argv[i]);
435         ktapvm_argv[new_index][strlen(argv[i])] = '\0';
436         new_index++;
437     }
438
439     uparm.argv = ktapvm_argv;
440     uparm.argc = new_index;
441     uparm.verbose = verbose;
442     uparm.trace_pid = trace\_pid;
443     uparm.trace_cpu = trace\_cpu;
444     uparm.print_timestamp = print\_timestamp;
445     uparm.quiet = quiet;

```

```
446 uparm.dry\_run = dry\_run;  
447  
448 /* start running into kernel ktapvm */  
449 ret = run\_ktapvm();  
450  
451 cleanup\_event\_resources();  
452 return ret;  
453 }  
454  
455
```

[One Level Up](#)

[Top Level](#)

userspace/ - ktap

- [cparser.h](#)
- [ffi/](#)
- [kp_bcwrite.c](#)
- [kp_lex.c](#)
- [kp_lex.h](#)
- [kp_main.c](#)
- [kp_parse.c](#)
- [kp_parse.h](#)
- [kp_parse_events.c](#)
- [kp_reader.c](#)
- [kp_symbol.c](#)
- [kp_symbol.h](#)
- [kp_util.c](#)
- [kp_util.h](#)

userspace/cparser.h - ktap

Data types defined

- [cp_ctype](#)
- [parser](#)

Functions defined

- [cp_csymf_arg](#)
- [cp_csymf_ret](#)
- [ffi_cparser_free](#)
- [ffi_cparser_init](#)

Macros defined

- [ALIGNED_DEFAULT](#)
- [ALIGNOF](#)
- [DEFAULT_ALIGN_MASK](#)
- [FUNCTION_ALIGN_MASK](#)
- [IS_CHAR_UNSIGNED](#)
- [MAX_TYPE_NAME_LEN](#)
- [POINTER_BITS](#)
- [POINTER_MAX](#)
- [PTR_ALIGN_MASK](#)
- [__KTAP_CPARSER_H__](#)
- [ct_ffi_cs](#)

Source code

```
1 #ifndef \_\_KTAP\_CPARSER\_H\_\_
2 #define \_\_KTAP\_CPARSER\_H\_\_
3
4 /*
5  * Copyright (c) 2011 James R. McKaskill
6  *
7  * This software is licensed under the stock MIT license:
8  *
9  * Permission is hereby granted, free of charge, to any person obtaining a
10 * copy of this software and associated documentation files (the "Software"),
11 * to deal in the Software without restriction, including without limitation
12 * the rights to use, copy, modify, merge, publish, distribute, sublicense,
13 * and/or sell copies of the Software, and to permit persons to whom the
14 * Software is furnished to do so, subject to the following conditions:
15 *
16 * The above copyright notice and this permission notice shall be included in
17 * all copies or substantial portions of the Software.
18 *
19 * THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
```

```

20 * IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
21 * FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE
22 * AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER
23 * LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING
24 * FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
25 * DEALINGS IN THE SOFTWARE.
26 *
27 * -----
28 */
29
30 /*
31 * Adapted from luaffi commit: abc638c9341025580099dcf77795c4b320ba0e63
32 *
33 * Copyright (c) 2013 Yicheng Qin, Qingping Hou
34 */
35
36 #ifdef CONFIG_KTAP_FFI
37
38 #include <assert.h>
39 #include <stdlib.h>
40 #include <stdio.h>
41 #include <stdint.h>
42 #include <string.h>
43 #include <stdbool.h>
44
45 #include "../include/ktap_ffi.h"
46
47 #define PTR_ALIGN_MASK (sizeof(void*) - 1)
48 #define FUNCTION_ALIGN_MASK (sizeof(void (*)()) - 1)
49 #define DEFAULT_ALIGN_MASK 7
50
51 struct parser {
52     int line;
53     const char *next;
54     const char *prev;
55     unsigned align_mask;
56 };
57
58 enum {
59     C_CALL,
60     STD_CALL,
61     FAST_CALL,
62 };
63
64
65 #define MAX_TYPE_NAME_LEN CSYM_NAME_MAX_LEN
66
67 enum {
68     /* 0 - 4 */
69     INVALID_TYPE,
70     VOID_TYPE,
71     BOOL_TYPE,
72     INT8_TYPE,
73     INT16_TYPE,
74     /* 5 - 9 */
75     INT32_TYPE,
76     INT64_TYPE,
77     INTPTR_TYPE,
78     ENUM_TYPE,
79     UNION_TYPE,
80     /* 10 - 12 */
81     STRUCT_TYPE,
82     FUNCTION_TYPE,
83     FUNCTION_PTR_TYPE,
84 };
85
86
87 #define IS_CHAR_UNSIGNED (((char) -1) > 0)
88
89 #define POINTER_BITS 2
90 #define POINTER_MAX ((1 << POINTER_BITS) - 1)
91
92 #define ALIGNOF(S) ((int) ((char*) &S.v - (char*) &S - 1))
93
94
95 /* Note: if adding a new member that is associated with a struct/union

```

```

96 * definition then it needs to be copied over in ctype.c:set_defined for when
97 * we create types based off of the declaration alone.
98 *
99 * Since this is used as a header for every ctype and cdata, and we create a
100 * ton of them on the stack, we try and minimise its size.
101 */
102 struct cp_ctype {
103     size_t base_size; /* size of the base type in bytes */
104     int ffi_base_cs_id;
105     int ffi_cs_id; /* index for csymbol from ktap vm */
106     union {
107         /* valid if is_bitfield */
108         struct {
109             /* size of bitfield in bits */
110             unsigned bit_size : 7;
111             /* offset within the current byte between 0-63 */
112             unsigned bit_offset : 6;
113         };
114         /* Valid if is_array */
115         size_t array_size;
116         /* Valid for is_variable_struct or is_variable_array. If
117          * variable_size_known (only used for is_variable_struct)
118          * then this is the total increment otherwise this is the
119          * per element increment.
120          */
121         size_t variable_increment;
122     };
123     size_t offset;
124     /* as (align bytes - 1) eg 7 gives 8 byte alignment */
125     unsigned align_mask : 4;
126     /* number of dereferences to get to the base type
127      * including +1 for arrays */
128     unsigned pointers : POINTER_BITS;
129     /* const pointer mask, LSB is current pointer, +1 for the whether
130      * the base type is const */
131     unsigned const_mask : POINTER_MAX + 1;
132     unsigned type : 5; /* value given by type enum above */
133     unsigned is_reference : 1;
134     unsigned is_array : 1;
135     unsigned is_defined : 1;
136     unsigned is_null : 1;
137     unsigned has_member_name : 1;
138     unsigned calling_convention : 2;
139     unsigned has_var_arg : 1;
140     /* set for variable array types where we don't know
141      * the variable size yet */
142     unsigned is_variable_array : 1;
143     unsigned is_variable_struct : 1;
144     /* used for variable structs after we know the variable size */
145     unsigned variable_size_known : 1;
146     unsigned is_bitfield : 1;
147     unsigned has_bitfield : 1;
148     unsigned is_jitted : 1;
149     unsigned is_packed : 1;
150     unsigned is_unsigned : 1;
151 };
152
153 #define ALIGNED_DEFAULT (__alignof__(void* __attribute__((aligned))) - 1)
154
155 csymbol *cp_id_to_csym(int id);
156 #define ct_ffi_cs(ct) (cp_id_to_csym((ct)->ffi_cs_id))
157
158 size_t ctype_size(const struct cp_ctype* ct);
159 int ctype_stack_top();
160 int cp_ctype_init();
161 int cp_ctype_free();
162 struct cp_ctype *ctype_lookup_type(char *name);
163 void cp_ctype_dump_stack();
164 void cp_error(const char *err_msg_fmt, ...);
165 struct cp_ctype *cp_ctype_reg_type(char *name, struct cp_ctype *ct);
166
167 void cp_update_csym_in_ctype(struct cp_ctype *ct);
168 void cp_push_ctype_with_name(struct cp_ctype *ct, const char *name, int nlen);
169 void cp_push_ctype(struct cp_ctype *ct);
170 void cp_set_defined(struct cp_ctype *ct);
171

```

```
172 int cp\_symbol\_build\_func(struct cp\_ctype *type, const char *fname, int fn_size);
173 int cp\_symbol\_build\_record(const char *stname, int type, int start_top);
174 int cp\_symbol\_build\_fake\_record(const char *stname, int type);
175 int cp\_symbol\_build\_pointer(struct cp\_ctype *ct);
176
177 int ffi\_parse\_cdef(const char *s);
178 void ffi\_parse\_new(const char *s, struct cp\_ctype *ct);
179 int ffi\_lookup\_csymbol\_id\_by\_name(const char *s);
180
181 void ffi\_cparser\_init(void);
182 void ffi\_cparser\_free(void);
183
184
185 static inline csymbol *cp\_csymf\_ret(csymbol\_func *csf)
186 {
187     return cp\_id\_to\_csym(csf->ret_id);
188 }
189
190 static inline csymbol *cp\_csymf\_arg(csymbol\_func *csf, int idx)
191 {
192     return cp\_id\_to\_csym(csf->arg_ids[idx]);
193 }
194
195
196 #else
197 static inline void ffi\_cparser\_init(void)
198 {
199     return;
200 }
201 static inline void ffi\_cparser\_free(void)
202 {
203     return;
204 }
205 #endif /* CONFIG\_KTAP\_FFI */
206
207
208 #endif /* \_\_KTAP\_PARSER\_H */
```

[One Level Up](#)

[Top Level](#)

include/ktap_ffi.h - ktap

Data types defined

- [csymbol](#)
- [csymbol](#)
- [csymbol_func](#)
- [csymbol_func](#)
- [csymbol_struct](#)
- [csymbol_struct](#)
- [ffi_mode](#)
- [ffi_type](#)
- [struct_member](#)
- [struct_member](#)

Macros defined

- [CSYM_NAME_MAX_LEN](#)
- [NUM_FFI_TYPE](#)
- [__KTAP_FFI_H](#)
- [cd_csym](#)
- [cd_csym_id](#)
- [cd_int](#)
- [cd_ptr](#)
- [cd_ptr_nmemb](#)
- [cd_record](#)
- [cd_set_csym_id](#)
- [cd_struct](#)
- [cd_type](#)
- [cd_union](#)
- [csym_func](#)
- [csym_func_addr](#)
- [csym_func_arg](#)
- [csym_func_arg_ids](#)
- [csym_name](#)

- [csym_ptr_deref](#)
- [csym_ptr_deref_id](#)
- [csym_set_ptr_deref_id](#)
- [csym_struct](#)
- [csym_struct_mb](#)
- [csym_struct_mb_nr](#)
- [csym_type](#)
- [csymf_addr](#)
- [csymf_arg](#)
- [csymf_arg_id](#)
- [csymf_arg_ids](#)
- [csymf_arg_nr](#)
- [csymf_ret](#)
- [csymf_ret_id](#)
- [csymst_mb](#)
- [csymst_mb_csym](#)
- [csymst_mb_id](#)
- [csymst_mb_len](#)
- [csymst_mb_name](#)
- [csymst_mb_nr](#)
- [ffi_type_align](#)
- [ffi_type_name](#)
- [ffi_type_size](#)
- [id_to_csym](#)
- [max_csym_id](#)

Source code

```

1 #ifndef __KTAP_FFI_H
2 #define __KTAP_FFI_H
3
4 #ifdef CONFIG_KTAP_FFI
5
6 #include "../include/ktap_types.h"
7
8 /*
9  * Types design in FFI module
10  *
11  * ktap\_cdata\_t is an instance of csymbol, so it's a combination of csymbol
12  * and it's actual data in memory.
13  *
14  * csymbol structs are globally unique and readonly type that represent C
15  * types. For non scalar C types like struct and function, helper structs are
16  * used to store detailed information. See csymbol\_func and csymbol\_struct for

```

```

17  * more information.
18  */
19
20  typedef enum {
21      /* 0 - 4 */
22      FFI_VOID,
23      FFI_UINT8,
24      FFI_INT8,
25      FFI_UINT16,
26      FFI_INT16,
27      /* 5 - 9 */
28      FFI_UINT32,
29      FFI_INT32,
30      FFI_UINT64,
31      FFI_INT64,
32      FFI_PTR,
33      /* 10 - 12 */
34      FFI_FUNC,
35      FFI_STRUCT,
36      FFI_UNION,
37      FFI_UNKNOWN,
38  } ffi_type;
39  #define NUM_FFI_TYPE ((int)FFI_UNKNOWN)
40
41
42  /* following struct and macros are added for C typedef
43  * size and alignment calculation */
44  typedef struct {
45      size_t size;
46      size_t align;
47      const char *name;
48  } ffi_mode;
49  extern const ffi_mode const ffi_type_modes[];
50
51  #define ffi_type_size(t) (ffi_type_modes[t].size)
52  #define ffi_type_align(t) (ffi_type_modes[t].align)
53  #define ffi_type_name(t) (ffi_type_modes[t].name)
54
55
56  /* start of csymbol definition */
57  #define CSYM_NAME_MAX_LEN 64
58
59  typedef struct csymbol_func {
60      void *addr; /* function address */
61      csymbol_id ret_id; /* function return type */
62      int arg_nr; /* number of arguments */
63      csymbol_id *arg_ids; /* function argument types */
64      unsigned has_var_arg; /* is this a var arg function? */
65  } csymbol_func;
66
67  typedef struct struct_member {
68      char name[CSYM_NAME_MAX_LEN]; /* name for this struct member */
69      csymbol_id id; /* type for this struct member */
70      int len; /* length of the array, -1 for non-array */
71  } struct_member;
72
73  typedef struct csymbol_struct {
74      int memb_nr; /* number of members */
75      struct_member *members; /* array for each member definition */
76      size_t size; /* bytes used to store struct */
77      /* alignment of the struct, 0 indicates uninitialized */
78      size_t align;
79  } csymbol_struct;
80
81
82  /* wrapper struct for all C symbols */
83  typedef struct csymbol {
84      char name[CSYM_NAME_MAX_LEN]; /* name for this symbol */
85      ffi_type type; /* type for this symbol */
86      /* following members are used only for non scalar C types */
87      union {
88          csymbol_id p; /* pointer type */
89          csymbol_func f; /* C function type */
90          csymbol_struct st; /* struct/union type */
91          csymbol_id td; /* typedef type */
92      } u;

```

```

93 } csymbol;
94
95
96 /* helper macros for c symbols */
97 #define max_csym_id(ks) (G(ks)->ffis.csym_nr)
98
99 /* lookup csymbol address by its id */
100 inline csymbol *ffi_get_csym_by_id(ktap_state_t *ks, csymbol_id id);
101 #define id_to_csym(ks, id) (ffi_get_csym_by_id(ks, id))
102
103 /* helper macros for struct csymbol */
104 #define csym_type(cs) ((cs)->type)
105 #define csym_name(cs) ((cs)->name)
106
107 /*
108 * helper macros for pointer symbol
109 */
110 #define csym_ptr_deref_id(cs) ((cs)->u.p)
111 #define csym_set_ptr_deref_id(cs, id) ((cs)->u.p = (id))
112 /* following macro gets csymbol address */
113 #define csym_ptr_deref(ks, cs) (id_to_csym(ks, csym_ptr_deref_id(cs)))
114
115 /*
116 * helper macros for function symbol
117 * csym_* accepts csymbol type
118 * csymf_* accepts csymbol_func type
119 */
120 #define csymf_addr(csf) ((csf)->addr)
121 #define csymf_ret_id(csf) ((csf)->ret_id)
122 #define csymf_arg_nr(csf) ((csf)->arg_nr)
123 #define csymf_arg_ids(csf) ((csf)->arg_ids)
124 /* get csymbol id for the nth argument */
125 #define csymf_arg_id(csf, n) ((csf)->arg_ids[n])
126 #define csym_func(cs) (&((cs)->u.f))
127 #define csym_func_addr(cs) (csymf_addr(csym_func(cs)))
128 #define csym_func_arg_ids(cs) (csymf_arg_ids(csym_func(cs)))
129 /* following macros get csymbol address */
130 #define csymf_ret(ks, csf) (id_to_csym(ks, csymf_ret_id(csf)))
131 /* get csymbol address for the nth argument */
132 #define csymf_arg(ks, csf, n) (id_to_csym(ks, csymf_arg_id(csf, n)))
133 #define csym_func_arg(ks, cs, n) (csymf_arg(ks, csym_func(cs), n))
134
135 /*
136 * helper macros for struct symbol
137 * csym_* accepts csymbol type
138 * csymst_* accepts csymbol_struct type
139 */
140 #define csymst_mb_nr(csst) ((csst)->memb_nr)
141 #define csym_struct(cs) (&(cs)->u.st)
142 #define csym_struct_mb(cs, n) (csymst_mb(csym_struct(cs), n))
143 #define csym_struct_mb_nr(cs) (csymst_mb_nr(csym_struct(cs)))
144 /* following macro gets csymbol address for the nth struct member */
145 #define csymst_mb(csst, n) (&(csst)->members[n])
146 #define csymst_mb_name(csst, n) ((csst)->members[n].name)
147 #define csymst_mb_id(csst, n) ((csst)->members[n].id)
148 #define csymst_mb_len(csst, n) ((csst)->members[n].len)
149 #define csymst_mb_csym(ks, csst, n) (id_to_csym(ks, (csst)->members[n].id))
150
151
152 /*
153 * helper macros for ktap_cdata_t type
154 */
155 #define cd_csym_id(cd) ((cd)->id)
156 #define cd_set_csym_id(cd, id) (cd_csym_id(cd) = (id))
157 #define cd_csym(ks, cd) (id_to_csym(ks, cd_csym_id(cd)))
158 #define cd_type(ks, cd) (csym_type(cd_csym(ks, cd)))
159
160 #define cd_int(cd) ((cd)->u.i)
161 #define cd_ptr(cd) ((cd)->u.p.addr)
162 #define cd_ptr_nmemb(cd) ((cd)->u.p.nmemb)
163 #define cd_record(cd) ((cd)->u.rec)
164 #define cd_struct(cd) cd_record(cd)
165 #define cd_union(cd) cd_record(cd)
166
167 #ifdef __KERNEL__
168 size_t csym_size(ktap_state_t *ks, csymbol *sym);

```

```

169 size_t csym_align(ktap_state_t *ks, csymbol *sym);
170 size_t csym_record_mb_offset_by_name(ktap_state_t *ks,
171     csymbol *cs, const char *name);
172 struct member *csymst_mb_by_name(ktap_state_t *ks,
173     csymbol_struct *csst, const char *name);
174
175 void ffi_free_symbols(ktap_state_t *ks);
176 csymbol_id ffi_get_csym_id(ktap_state_t *ks, char *name);
177
178 ktap_cdata_t *kp_cdata_new(ktap_state_t *ks, csymbol_id id);
179 ktap_cdata_t *kp_cdata_new_ptr(ktap_state_t *ks, void *addr,
180     int nmemb, csymbol_id id, int to_allocate);
181 ktap_cdata_t *kp_cdata_new_record(ktap_state_t *ks, void *val, csymbol_id id);
182 void kp_cdata_dump(ktap_state_t *ks, ktap_cdata_t *cd);
183 int kp_cdata_type_match(ktap_state_t *ks, csymbol *cs, ktap_val_t *val);
184 void kp_cdata_unpack(ktap_state_t *ks, char *dst, csymbol *cs, ktap_val_t *val);
185 void kp_cdata_ptr_set(ktap_state_t *ks, ktap_cdata_t *cd,
186     ktap_val_t *key, ktap_val_t *val);
187 void kp_cdata_ptr_get(ktap_state_t *ks, ktap_cdata_t *cd,
188     ktap_val_t *key, ktap_val_t *val);
189 void kp_cdata_record_set(ktap_state_t *ks, ktap_cdata_t *cd,
190     ktap_val_t *key, ktap_val_t *val);
191 void kp_cdata_record_get(ktap_state_t *ks, ktap_cdata_t *cd,
192     ktap_val_t *key, ktap_val_t *val);
193
194 int ffi_call(ktap_state_t *ks, csymbol_func *cf);
195
196 #endif /* for CONFIG_FFI_KTAP */
197 #endif /* for __KERNEL__ */
198 #endif /* __KTAP_FFI_H */

```

[One Level Up](#)

[Top Level](#)

runtime/ffi/ffi_type.c - ktap

Global variables defined

- [ffi_type_modes](#)

Macros defined

- [CTYPE_MODE](#)
- [CTYPE_MODE_HELPER](#)
- [CTYPE_MODE_NAME](#)

Source code

```
1 #include "../../include/ktap_ffi.h"
2 #ifdef __KERNEL__
3 #include <linux/types.h>
4 #else
5 #include <stdint.h>
6 #include <stddef.h>
7 #endif
8
9 #define CTYPE_MODE_HELPER(name, type) \
10 struct _##name##_align { \
11     type t1; \
12     char c; \
13     type t2; \
14 };
15
16 #define CTYPE_MODE(name) \
17 { \
18     offsetof(struct _##name##_align, c), \
19     offsetof(struct _##name##_align, t2) - \
20     offsetof(struct _##name##_align, c), \
21     #name \
22 }
23
24 #define CTYPE_MODE_NAME(name) _##name##_mode
25
26 /* ffi_ctype_mode should be corresponded to ffi_ctype */
27 CTYPE_MODE_HELPER(uint8, uint8_t);
28 CTYPE_MODE_HELPER(int8, int8_t);
29 CTYPE_MODE_HELPER(uint16, uint16_t);
30 CTYPE_MODE_HELPER(int16, int16_t);
31 CTYPE_MODE_HELPER(uint32, uint32_t);
32 CTYPE_MODE_HELPER(int32, int32_t);
33 CTYPE_MODE_HELPER(uint64, uint64_t);
34 CTYPE_MODE_HELPER(int64, int64_t);
35 CTYPE_MODE_HELPER(pointer, void*);
36
37 const ffi_mode ffi_type_modes[NUM\_FFI\_TYPE+1] = {
38     {0, 1, "void"},
39     CTYPE_MODE(uint8),
40     CTYPE_MODE(int8),
41     CTYPE_MODE(uint16),
42     CTYPE_MODE(int16),
43     CTYPE_MODE(uint32),
44     CTYPE_MODE(int32),
45     CTYPE_MODE(uint64),
46     CTYPE_MODE(int64),
47     CTYPE_MODE(pointer),
48     {0, 1, "function"},
49     {0, 1, "struct"},
50     {0, 1, "union"},
51     {0, 1, "unknown"},
52 };
```

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

runtime/ffi/ - ktap

- [call_x86_64.S](#)
- [cdata.c](#)
- [ffi_call.c](#)
- [ffi_symbol.c](#)
- [ffi_type.c](#)
- [ffi_util.c](#)

[One Level Up](#)

[Top Level](#)

runtime/ - ktap

- [amalg.c](#)
- [ffi/](#)
- [kp_bcread.c](#)
- [kp_bcread.h](#)
- [kp_events.c](#)
- [kp_events.h](#)
- [kp_mempool.c](#)
- [kp_mempool.h](#)
- [kp_obj.c](#)
- [kp_obj.h](#)
- [kp_str.c](#)
- [kp_str.h](#)
- [kp_tab.c](#)
- [kp_tab.h](#)
- [kp_transport.c](#)
- [kp_transport.h](#)
- [kp_vm.c](#)
- [kp_vm.h](#)
- [ktap.c](#)
- [ktap.h](#)
- [lib_ansi.c](#)
- [lib_base.c](#)
- [lib_ffi.c](#)
- [lib_kdebug.c](#)
- [lib_net.c](#)
- [lib_table.c](#)
- [lib_timer.c](#)

runtime/amalg.c - ktap

```
1 /*
2  * kp_amalg.c - ktapvm kernel module amalgamation.
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include "ktap.c"
23 #include "kp_obj.c"
24 #include "kp_bcread.c"
25 #include "kp_str.c"
26 #include "kp_mempool.c"
27 #include "kp_tab.c"
28 #include "kp_transport.c"
29 #include "kp_vm.c"
30 #include "kp_events.c"
31 #include "lib_base.c"
32 #include "lib_ansi.c"
33 #include "lib_kdebug.c"
34 #include "lib_timer.c"
35 #include "lib_table.c"
36 #include "lib_net.c"
37
38 #ifdef CONFIG_KTAP_FFI
39 #include "ffi/ffi_call.c"
40 #include "ffi/ffi_type.c"
41 #include "ffi/ffi_symbol.c"
42 #include "ffi/cdata.c"
43 #include "ffi/ffi_util.c"
44 #include "lib_ffi.c"
45 #endif
```

runtime/kp_bcread.c - ktap

Data types defined

- [BCReadCtx](#)
- [BCReadCtx](#)

Functions defined

- [bcread_block](#)
- [bcread_byte](#)
- [bcread_bytecode](#)
- [bcread_dbg](#)
- [bcread_error](#)
- [bcread_header](#)
- [bcread_kgc](#)
- [bcread_knum](#)
- [bcread_ktab](#)
- [bcread_ktabk](#)
- [bcread_mem](#)
- [bcread_proto](#)
- [bcread_uint32](#)
- [bcread_uv](#)
- [bcread_varinfo](#)
- [bswap](#)
- [kp_bcread](#)

Macros defined

- [bcread_flags](#)
- [bcread_oldtop](#)
- [bcread_savetop](#)
- [bcread_swap](#)

Source code

```
1 /*
2  * Bytecode reader.
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
```

```

6 * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7 *
8 * Adapted from luajit and lua interpreter.
9 * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include "../include/ktap_types.h"
27 #include "../include/ktap_bc.h"
28 #include "../include/ktap_err.h"
29 #include "ktap.h"
30 #include "kp_obj.h"
31 #include "kp_str.h"
32 #include "kp_tab.h"
33
34
35 /* Context for bytecode reader. */
36 typedef struct BCReadCtx {
37     ktap_state_t *ks;
38     int flags;
39     char *start;
40     char *p;
41     char *pe;
42     ktap_str_t *chunkname;
43     ktap_val_t *savetop;
44 } BCReadCtx;
45
46
47 #define bcread_flags(ctx)    (ctx->flags)
48 #define bcread_swap(ctx) \
49     ((bcread_flags(ctx) & BCDUMP_F_BE) != KP_BE*BCDUMP_F_BE)
50 #define bcread_olddtop(ctx)    (ctx->savetop)
51 #define bcread_savetop(ctx)    (ctx->savetop = (ctx->ks->top);
52
53 static inline uint32_t bswap(uint32_t x)
54 {
55     return (uint32_t)__builtin_bswap32((int32_t)x);
56 }
57
58 /* -- Input buffer handling ----- */
59
60 /* Throw reader error. */
61 static void bcread_error(BCReadCtx *ctx, ErrMsg em)
62 {
63     kp_error(ctx->ks, "%s\n", err2msg(em));
64 }
65
66 /* Return memory block from buffer. */
67 static inline uint8_t *bcread_mem(BCReadCtx *ctx, int len)
68 {
69     uint8_t *p = (uint8_t *)ctx->p;
70     ctx->p += len;
71     kp_assert(ctx->p <= ctx->pe);
72     return p;
73 }
74
75 /* Copy memory block from buffer. */
76 static void bcread_block(BCReadCtx *ctx, void *q, int len)
77 {
78     memcpy(q, bcread_mem(ctx, len), len);
79 }
80
81 /* Read byte from buffer. */

```

```

82 static inline uint32_t bcread_byte(BCReadCtx *ctx)
83 {
84     kp_assert(ctx->p < ctx->pe);
85     return (uint32_t)(uint8_t)*ctx->p++;
86 }
87
88 /* Read ULEB128 value from buffer. */
89 static inline uint32_t bcread_uint32(BCReadCtx *ctx)
90 {
91     uint32_t v;
92     bcread_block(ctx, &v, sizeof(uint32_t));
93     kp_assert(ctx->p <= ctx->pe);
94     return v;
95 }
96
97 /* -- Bytecode reader ----- */
98
99 /* Read debug info of a prototype. */
100 static void bcread_dbg(BCReadCtx *ctx, ktap_proto_t *pt, int sizedbg)
101 {
102     void *lineinfo = (void *)proto_lineinfo(pt);
103
104     bcread_block(ctx, lineinfo, sizedbg);
105     /* Swap lineinfo if the endianness differs. */
106     if (bcread_swap(ctx) && pt->numline >= 256) {
107         int i, n = pt->sizebc-1;
108         if (pt->numline < 65536) {
109             uint16_t *p = (uint16_t *)lineinfo;
110             for (i = 0; i < n; i++)
111                 p[i] = (uint16_t)((p[i] >> 8)|(p[i] << 8));
112         } else {
113             uint32_t *p = (uint32_t *)lineinfo;
114             for (i = 0; i < n; i++)
115                 p[i] = bswap(p[i]);
116         }
117     }
118 }
119
120 /* Find pointer to varinfo. */
121 static const void *bcread_varinfo(ktap_proto_t *pt)
122 {
123     const uint8_t *p = proto_uvinfo(pt);
124     int n = pt->sizeuv;
125     if (n)
126         while (*p++ || --n);
127     return p;
128 }
129
130 /* Read a single constant key/value of a template table. */
131 static int bcread_ktab(BCReadCtx *ctx, ktap_val_t *o)
132 {
133     int tp = bcread_uint32(ctx);
134     if (tp >= BCDUMP_KTAB_STR) {
135         int len = tp - BCDUMP_KTAB_STR;
136         const char *p = (const char *)bcread_mem(ctx, len);
137         ktap_str_t *ts = kp_str_new(ctx->ks, p, len);
138         if (unlikely(!ts))
139             return -ENOMEM;
140
141         set_string(o, ts);
142     } else if (tp == BCDUMP_KTAB_NUM) {
143         set_number(o, *(ktap_number *)bcread_mem(ctx,
144             sizeof(ktap_number)));
145     } else {
146         kp_assert(tp <= BCDUMP_KTAB_TRUE);
147         setitype(o, ~tp);
148     }
149     return 0;
150 }
151
152 /* Read a template table. */
153 static ktap_tab_t *bcread_ktab(BCReadCtx *ctx)
154 {
155     int narray = bcread_uint32(ctx);
156     int nhash = bcread_uint32(ctx);
157

```

```

158 ktap_tab_t *t = kp_tab_new(ctx->ks, narray, hsize2hbits(nhash));
159 if (!t)
160     return NULL;
161
162 if (narray) { /* Read array entries. */
163     int i;
164     ktap_val_t *o = t->array;
165     for (i = 0; i < narray; i++, o++)
166         if (bcread_ktabk(ctx, o))
167             return NULL;
168 }
169 if (nhash) { /* Read hash entries. */
170     int i;
171     for (i = 0; i < nhash; i++) {
172         ktap_val_t key;
173         ktap_val_t val;
174         if (bcread_ktabk(ctx, &key))
175             return NULL;
176         kp_assert(!is_nil(&key));
177         if (bcread_ktabk(ctx, &val))
178             return NULL;
179         kp_tab_set(ctx->ks, t, &key, &val);
180     }
181 }
182 return t;
183 }
184
185 /* Read GC constants(string, table, child proto) of a prototype. */
186 static int bcread_kgc(BCReadCtx *ctx, ktap_proto_t *pt, int sizekgc)
187 {
188     ktap_obj_t **kr = (ktap_obj_t **)pt->k - (ptrdiff_t)sizekgc;
189     int i;
190
191     for (i = 0; i < sizekgc; i++, kr++) {
192         int tp = bcread_uint32(ctx);
193         if (tp >= BCDUMP_KGC_STR) {
194             int len = tp - BCDUMP_KGC_STR;
195             const char *p = (const char *)bcread_mem(ctx, len);
196             *kr = (ktap_obj_t *)kp_str_new(ctx->ks, p, len);
197             if (unlikely(!*kr))
198                 return -1;
199         } else if (tp == BCDUMP_KGC_TAB) {
200             *kr = (ktap_obj_t *)bcread_ktab(ctx);
201             if (unlikely(!*kr))
202                 return -1;
203         } else if (tp == BCDUMP_KGC_CHILD){
204             ktap_state_t *ks = ctx->ks;
205             if (ks->top <= bcread_oldtop(ctx)) {
206                 bcread_error(ctx, KP_ERR_BCBAD);
207                 return -1;
208             }
209             ks->top--;
210             *kr = (ktap_obj_t *)ptvalue(ks->top);
211         } else {
212             bcread_error(ctx, KP_ERR_BCBAD);
213             return -1;
214         }
215     }
216
217     return 0;
218 }
219
220 /* Read number constants of a prototype. */
221 static void bcread_knum(BCReadCtx *ctx, ktap_proto_t *pt, int sizekn)
222 {
223     int i;
224     ktap_val_t *o = pt->k;
225
226     for (i = 0; i < sizekn; i++, o++) {
227         set_number(o, *(ktap_number *)bcread_mem(ctx,
228             sizeof(ktap_number)));
229     }
230 }
231
232 /* Read bytecode instructions. */
233 static void bcread_bytecode(BCReadCtx *ctx, ktap_proto_t *pt, int sizebc)

```

```

234 {
235     BCIns *bc = proto_bc(pt);
236     bc[0] = BCINS_AD((pt->flags & PROTO_VARARG) ? BC_FUNCV : BC_FUNCF,
237         pt->framesize, 0);
238     bcread_block(ctx, bc+1, (sizebc-1)*(int)sizeof(BCIns));
239     /* Swap bytecode instructions if the endianness differs. */
240     if (bcread_swap(ctx)) {
241         int i;
242         for (i = 1; i < sizebc; i++) bc[i] = bswap(bc[i]);
243     }
244 }
245
246 /* Read upvalue refs. */
247 static void bcread_uv(BCReadCtx *ctx, ktap_proto_t *pt, int sizeuv)
248 {
249     if (sizeuv) {
250         uint16_t *uv = proto_uv(pt);
251         bcread_block(ctx, uv, sizeuv*2);
252         /* Swap upvalue refs if the endianness differs. */
253         if (bcread_swap(ctx)) {
254             int i;
255             for (i = 0; i < sizeuv; i++)
256                 uv[i] = (uint16_t)((uv[i] >> 8)|(uv[i] << 8));
257         }
258     }
259 }
260
261 /* Read a prototype. */
262 static ktap_proto_t *bcread_proto(BCReadCtx *ctx)
263 {
264     ktap_proto_t *pt;
265     int framesize, numparams, flags;
266     int sizeuv, sizekgc, sizekn, sizebc, sizept;
267     int ofsk, ofsuv, ofsdbg;
268     int sizedbg = 0;
269     BCLine firstline = 0, numline = 0;
270
271     /* Read prototype header. */
272     flags = bcread_byte(ctx);
273     numparams = bcread_byte(ctx);
274     framesize = bcread_byte(ctx);
275     sizeuv = bcread_byte(ctx);
276     sizekgc = bcread_uint32(ctx);
277     sizekn = bcread_uint32(ctx);
278     sizebc = bcread_uint32(ctx) + 1;
279     if (!(bcread_flags(ctx) & BCDUMP_F_STRIP)) {
280         sizedbg = bcread_uint32(ctx);
281         if (sizedbg) {
282             firstline = bcread_uint32(ctx);
283             numline = bcread_uint32(ctx);
284         }
285     }
286
287     /* Calculate total size of prototype including all colocated arrays. */
288     sizept = (int)sizeof(ktap_proto_t) + sizebc * (int)sizeof(BCIns) +
289         sizekgc * (int)sizeof(ktap_obj_t *);
290     sizept = (sizept + (int)sizeof(ktap_val_t)-1) &
291         ~((int)sizeof(ktap_val_t)-1);
292     ofsk = sizept; sizept += sizekn*(int)sizeof(ktap_val_t);
293     ofsuv = sizept; sizept += ((sizeuv+1)&-1)*2;
294     ofsdbg = sizept; sizept += sizedbg;
295
296     /* Allocate prototype object and initialize its fields. */
297     pt = (ktap_proto_t *)kp_obj_new(ctx->ks, (int)sizept);
298     pt->gct = ~KTAP_TPROTO;
299     pt->numparams = (uint8_t)numparams;
300     pt->framesize = (uint8_t)framesize;
301     pt->sizebc = sizebc;
302     pt->k = (char *)pt + ofsk;
303     pt->uv = (char *)pt + ofsuv;
304     pt->sizekgc = 0; /* Set to zero until fully initialized. */
305     pt->sizekn = sizekn;
306     pt->sizept = sizept;
307     pt->sizeuv = (uint8_t)sizeuv;
308     pt->flags = (uint8_t)flags;
309     pt->chunkname = ctx->chunkname;

```

```

310 /* Close potentially uninitialized gap between bc and kgc. */
311 *(uint32_t *)((char *)pt + ofsk - sizeof(ktap_obj_t *)*(sizekgc+1))
312         = 0;
313
314
315 /* Read bytecode instructions and upvalue refs. */
316 bcread_bytecode(ctx, pt, sizebc);
317 bcread_uv(ctx, pt, sizeuv);
318
319 /* Read constants. */
320 if (bcread_kgc(ctx, pt, sizekgc))
321     return NULL;
322 pt->sizekgc = sizekgc;
323 bcread_knum(ctx, pt, sizekn);
324
325 /* Read and initialize debug info. */
326 pt->firstline = firstline;
327 pt->numline = numline;
328 if (sizedbg) {
329     int sizeli = (sizebc-1) << (numline < 256 ? 0 :
330         numline < 65536 ? 1 : 2);
331     pt->lineinfo = (char *)pt + ofsdbg;
332     pt->uvinfo = (char *)pt + ofsdbg + sizeli;
333     bcread_dbg(ctx, pt, sizedbg);
334     pt->varinfo = (void *)bcread_varinfo(pt);
335 } else {
336     pt->lineinfo = NULL;
337     pt->uvinfo = NULL;
338     pt->varinfo = NULL;
339 }
340 return pt;
341 }
342
343 /* Read and check header of bytecode dump. */
344 static int bcread_header(BCReadCtx *ctx)
345 {
346     uint32_t flags;
347
348     if (bcread_byte(ctx) != BCDUMP_HEAD1 ||
349         bcread_byte(ctx) != BCDUMP_HEAD2 ||
350         bcread_byte(ctx) != BCDUMP_HEAD3 ||
351         bcread_byte(ctx) != BCDUMP_VERSION)
352         return -1;
353
354     bcread_flags(ctx) = flags = bcread_byte(ctx);
355
356     if ((flags & ~(BCDUMP_F_KNOWN)) != 0)
357         return -1;
358
359     if ((flags & BCDUMP_F_FFI)) {
360         return -1;
361     }
362
363     if ((flags & BCDUMP_F_STRIP)) {
364         ctx->chunkname = kp_str_newz(ctx->ks, "striped");
365     } else {
366         int len = bcread_uint32(ctx);
367         ctx->chunkname = kp_str_new(ctx->ks,
368             (const char *)bcread_mem(ctx, len), len);
369     }
370
371     if (unlikely(!ctx->chunkname))
372         return -1;
373
374     return 0;
375 }
376
377 /* Read a bytecode dump. */
378 ktap_proto_t *kp_bcread(ktap_state_t *ks, unsigned char *buff, int len)
379 {
380     BCReadCtx ctx;
381
382     ctx.ks = ks;
383     ctx.p = buff;
384     ctx.pe = buff + len;
385

```

```

386     ctx.start = buff;
387
388     bcread\_savetop(&ctx);
389     /* Check for a valid bytecode dump header. */
390     if (bcread\_header(&ctx)) {
391         bcread\_error(&ctx, KP_ERR_BCFMT);
392         return NULL;
393     }
394
395     for (;;) { /* Process all prototypes in the bytecode dump. */
396         ktag\_proto\_t *pt;
397         int len;
398         const char *startp;
399         /* Read length. */
400         if (ctx.p < ctx.pe && ctx.p[0] == 0) { /* Shortcut EOF. */
401             ctx.p++;
402             break;
403         }
404         len = bcread\_uint32(&ctx);
405         if (!len)
406             break; /* EOF */
407         startp = ctx.p;
408         pt = bcread\_proto(&ctx);
409         if (!pt)
410             return NULL;
411         if (ctx.p != startp + len) {
412             bcread\_error(&ctx, KP_ERR_BCBAD);
413             return NULL;
414         }
415         set\_proto(ks->top, pt);
416         incr\_top(ks);
417     }
418     if ((int32_t)(2*(uint32_t)(ctx.pe - ctx.p)) > 0 ||
419         ks->top-1 != bcread\_olddtop(&ctx)) {
420         bcread\_error(&ctx, KP_ERR_BCBAD);
421         return NULL;
422     }
423
424     /* Pop off last prototype. */
425     ks->top--;
426     return ptvalue(ks->top);
427 }
428

```

[One Level Up](#)

[Top Level](#)

include/ktap_err.h - ktap

Data types defined

- [ErrMsg](#)

Macros defined

- [ERRDEF](#)
- [__KTAP_ERR_H__](#)

Source code

```
1 #ifndef __KTAP_ERR_H__
2 #define __KTAP_ERR_H__
3
4 typedef enum {
5 #define ERRDEF(name, msg) \
6     KP_ERR_##name, KP_ERR_##name##_ = KP_ERR_##name + sizeof(msg)-1,
7 #include "ktap_errmsg.h"
8     KP_ERR__MAX
9 } ErrMsg;
10
11 #endif
```

runtime/ktap.h - ktap

Data types defined

- [ktap_libfunc](#)
- [ktap_libfunc_t](#)

Functions defined

- [get_recursion_context](#)
- [kp_this_cpu_print_buffer](#)
- [kp_this_cpu_state](#)
- [kp_this_cpu_temp_buffer](#)
- [put_recursion_context](#)
- [trace_get_context_bit](#)

Macros defined

- [SPRINT_SYMBOL](#)
- [SPRINT_SYMBOL](#)
- [__KTAP_H](#)
- [err2msg](#)
- [kp_arg](#)
- [kp_arg_check](#)
- [kp_arg_checkfunction](#)
- [kp_arg_checknumber](#)
- [kp_arg_checkoptnumber](#)
- [kp_arg_checkstring](#)
- [kp_arg_nr](#)
- [kp_error](#)
- [kp_puts](#)
- [kp_verbose_printf](#)
- [raw_cpu_ptr](#)

Source code

```
1 #ifndef __KTAP_H__
2 #define __KTAP_H__
3
4 #include <linux/version.h>
5 #include <linux/hardirq.h>
```

```

6 #include <linux/trace_seq.h>
7
8 #ifndef raw_cpu_ptr
9 #define raw_cpu_ptr __this_cpu_ptr
10 #endif
11
12 /* for built-in library C function register */
13 typedef struct ktap_libfunc {
14     const char *name; /* function name */
15     ktap_cfunction func; /* function pointer */
16 } ktap_libfunc_t;
17
18 long gettimeofday_ns(void); /* common helper function */
19 int kp_lib_init_base(ktap_state_t *ks);
20 int kp_lib_init_kdebug(ktap_state_t *ks);
21 int kp_lib_init_timer(ktap_state_t *ks);
22 int kp_lib_init_table(ktap_state_t *ks);
23 int kp_lib_init_ansi(ktap_state_t *ks);
24 int kp_lib_init_net(ktap_state_t *ks);
25 #ifdef CONFIG_KTAP_FFI
26 int kp_lib_init_ffi(ktap_state_t *ks);
27 #endif
28
29 void kp_exit_timers(ktap_state_t *ks);
30 void kp_freeupval(ktap_state_t *ks, ktap_upval_t *uv);
31
32 extern int (*kp_ftrace_profile_set_filter)(struct perf_event *event,
33     int event_id,
34     const char *filter_str);
35
36 extern struct syscall_metadata **syscalls_metadata;
37
38 /* get from kernel/trace/trace.h */
39 static __always_inline int trace_get_context_bit(void)
40 {
41     int bit;
42
43     if (in_interrupt()) {
44         if (in_nmi())
45             bit = 0;
46         else if (in_irq())
47             bit = 1;
48         else
49             bit = 2;
50     } else
51         bit = 3;
52
53     return bit;
54 }
55
56 static __always_inline int get_recursion_context(ktap_state_t *ks)
57 {
58     int rctx = trace_get_context_bit();
59     int *val = raw_cpu_ptr(G(ks)->recursion_context[rctx]);
60
61     if (*val)
62         return -1;
63
64     *val = true;
65     return rctx;
66 }
67
68 static inline void put_recursion_context(ktap_state_t *ks, int rctx)
69 {
70     int *val = raw_cpu_ptr(G(ks)->recursion_context[rctx]);
71     *val = false;
72 }
73
74 static inline void *kp_this_cpu_state(ktap_state_t *ks, int rctx)
75 {
76     return this_cpu_ptr(G(ks)->percpu_state[rctx]);
77 }
78
79 static inline void *kp_this_cpu_print_buffer(ktap_state_t *ks)
80 {
81     return this_cpu_ptr(G(ks)->percpu_print_buffer[trace_get_context_bit()]);

```

```

82 }
83
84 static inline void *kp_this_cpu_temp_buffer(ktap_state_t *ks)
85 {
86     return this_cpu_ptr(G(ks)->percpu_temp_buffer[trace_get_context_bit()]);
87 }
88
89 #define kp_verbose_printf(ks, ...) \
90     if (G(ks)->param->verbose) \
91         kp_printf(ks, "[verbose] " __VA_ARGS__);
92
93 /* argument operation macro */
94 #define kp_arg(ks, idx) ((ks)->func + (idx))
95 #define kp_arg_nr(ks) ((int)(ks->top - (ks->func + 1)))
96
97 #define kp_arg_check(ks, idx, type) \
98     do { \
99         if (unlikely(itype(kp_arg(ks, idx)) != type)) { \
100             kp_error(ks, "wrong type of argument %d\n", idx); \
101             return -1; \
102         } \
103     } while (0)
104
105 #define kp_arg_checkstring(ks, idx) \
106     ({ \
107         ktap_val_t *o = kp_arg(ks, idx); \
108         if (unlikely(!is_string(o))) { \
109             kp_error(ks, "wrong type of argument %d\n", idx); \
110             return -1; \
111         } \
112         svalue(o); \
113     })
114
115 #define kp_arg_checkfunction(ks, idx) \
116     ({ \
117         ktap_val_t *o = kp_arg(ks, idx); \
118         if (unlikely(!is_function(o))) { \
119             kp_error(ks, "wrong type of argument %d\n", idx); \
120             return -1; \
121         } \
122         clvalue(o); \
123     })
124
125 #define kp_arg_checknumber(ks, idx) \
126     ({ \
127         ktap_val_t *o = kp_arg(ks, idx); \
128         if (unlikely(!is_number(o))) { \
129             kp_error(ks, "wrong type of argument %d\n", idx); \
130             return -1; \
131         } \
132         nvalue(o); \
133     })
134
135 #define kp_arg_checkoptnumber(ks, idx, def) \
136     ({ \
137         ktap_number n; \
138         if (idx > kp_arg_nr(ks)) { \
139             n = def; \
140         } else { \
141             ktap_val_t *o = kp_arg(ks, idx); \
142             if (unlikely(!is_number(o))) { \
143                 kp_error(ks, "wrong type of argument %d\n", \
144                     idx); \
145                 return -1; \
146             } \
147             n = nvalue(o); \
148         } \
149         n; \
150     })
151
152 #define kp_error(ks, args...) \
153     do { \
154         kp_printf(ks, "error: "args); \
155         kp_vm_try_to_exit(ks); \
156         G(ks)->state = KTAP_ERROR; \
157     } while(0)

```

```

158
159
160 /* TODO: this code need to cleanup */
161 #if LINUX_VERSION_CODE > KERNEL_VERSION(3, 5, 0)
162 #define SPRINT_SYMBOL    sprint_symbol_no_offset
163 #else
164 #define SPRINT_SYMBOL    sprint_symbol
165 #endif
166
167 extern int kp_max_loop_count;
168
169 void kp_printf(ktap_state_t *ks, const char *fmt, ...);
170 void __kp_puts(ktap_state_t *ks, const char *str);
171 void __kp_bputs(ktap_state_t *ks, const char *str);
172
173 #define kp_puts(ks, str) ({
174     static const char *trace_printk_fmt
175         __attribute__((section("__trace_printk_fmt"))) =
176         __builtin_constant_p(str) ? str : NULL;
177
178     if (__builtin_constant_p(str))
179         __kp_bputs(ks, trace_printk_fmt);
180     else
181         __kp_puts(ks, str);
182 })
183
184 #define err2msg(em)    (kp_err_allmsg+(int)(em))
185 extern const char *kp_err_allmsg;
186
187 #endif /* KTAP_H */
188

```

[One Level Up](#)

[Top Level](#)

runtime/ktap.c - ktap

Global variables defined

- [exit_ktap](#)
- [init_ktap](#)
- [kp_dir_dentry](#)
- [kp_ftrace_profile_set_filter](#)
- [kp_max_loop_count](#)
- [ktap_fops](#)
- [ktapvm_fops](#)
- [syscalls_metadata](#)

Functions defined

- [exit_ktap](#)
- [gettimeofday_ns](#)
- [init_dummy_kernel_functions](#)
- [init_ktap](#)
- [ktap_ioctl](#)
- [ktap_main](#)
- [ktapvm_ioctl](#)
- [load_trunk](#)
- [print_version](#)

Source code

```
1 /*
2  * ktap.c - ktapvm kernel module main entry
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 /*
23 * this file is the first file to be compile, add CONFIG_ checking in here.
```

```

24 * See Requirements in doc/tutorial.md
25 */
26
27 #include <linux/version.h>
28 #if LINUX_VERSION_CODE < KERNEL_VERSION(3, 1, 0)
29 #error "Currently ktap don't support kernel older than 3.1"
30 #endif
31
32 #if !CONFIG_EVENT_TRACING
33 #error "Please enable CONFIG_EVENT_TRACING before compile ktap"
34 #endif
35
36 #if !CONFIG_PERF_EVENTS
37 #error "Please enable CONFIG_PERF_EVENTS before compile ktap"
38 #endif
39
40 #include <linux/module.h>
41 #include <linux/errno.h>
42 #include <linux/file.h>
43 #include <linux/slab.h>
44 #include <linux/fcntl.h>
45 #include <linux/sched.h>
46 #include <linux/poll.h>
47 #include <linux/anon_inodes.h>
48 #include <linux/debugfs.h>
49 #include <linux/vmalloc.h>
50 #include "../include/ktap_types.h"
51 #include "ktap.h"
52 #include "kp_bcread.h"
53 #include "kp_vm.h"
54
55 /* common helper function */
56 long gettimeofday_ns(void)
57 {
58     struct timespec now;
59
60     getnstimeofday(&now);
61     return now.tv_sec * NSEC_PER_SEC + now.tv_nsec;
62 }
63
64 static int load_trunk(ktap\_option\_t *parm, unsigned long **buff)
65 {
66     int ret;
67     unsigned long *vmstart;
68
69     vmstart = vmalloc(parm->trunk_len);
70     if (!vmstart)
71         return -ENOMEM;
72
73     ret = copy_from_user(vmstart, (void __user *)parm->trunk,
74                         parm->trunk_len);
75     if (ret < 0) {
76         vfree(vmstart);
77         return -EFAULT;
78     }
79
80     *buff = vmstart;
81     return 0;
82 }
83
84 static struct dentry *kp_dir_dentry;
85
86 /* Ktap Main Entry */
87 static int ktap_main(struct file *file, ktap\_option\_t *parm)
88 {
89     unsigned long *buff = NULL;
90     ktap\_state\_t *ks;
91     ktap\_proto\_t *pt;
92     long start_time, delta_time;
93     int ret;
94
95     start_time = gettimeofday\_ns();
96
97     ks = kp\_vm\_new\_state(parm, kp\_dir\_dentry);
98     if (unlikely(!ks))
99         return -ENOEXEC;

```

```

100 file->private_data = ks;
101
102
103 ret = load_trunk(parm, &buff);
104 if (ret) {
105     kp_error(ks, "cannot load file\n");
106     goto out;
107 }
108
109 pt = kp_bcread(ks, (unsigned char *)buff, parm->trunk_len);
110
111 vfree(buff);
112
113 if (pt) {
114     /* validate byte code */
115     if (kp_vm_validate_code(ks, pt, ks->stack))
116         goto out;
117
118     delta_time = (gettimeofday_ns() - start_time) / NSEC_PER_USEC;
119     kp_verbose_printf(ks, "booting time: %d (us)\n", delta_time);
120
121     /* enter vm */
122     kp_vm_call_proto(ks, pt);
123 }
124
125 out:
126     kp_vm_exit(ks);
127     return ret;
128 }
129
130
131 static void print_version(void)
132 {
133 }
134
135 static long ktap_ioctl(struct file *file, unsigned int cmd, unsigned long arg)
136 {
137     ktap_option_t parm;
138     int ret;
139
140     switch (cmd) {
141     case KTAP_CMD_IOC_VERSION:
142         print_version();
143         return 0;
144     case KTAP_CMD_IOC_RUN:
145         /*
146          * must be root to run ktap script (at least for now)
147          *
148          * TODO: check perf_paranoid sysctl and allow non-root user
149          * to use ktap for tracing process(like uprobe) ?
150          */
151         if (!capable(CAP_SYS_ADMIN))
152             return -EACCES;
153
154         ret = copy_from_user(&parm, (void __user *)arg,
155                             sizeof(ktap_option_t));
156         if (ret < 0)
157             return -EFAULT;
158
159         return ktap_main(file, &parm);
160     default:
161         return -EINVAL;
162     };
163
164     return 0;
165 }
166
167 static const struct file_operations ktap_fops = {
168     .llseek          = no_llseek,
169     .unlocked_ioctl = ktap_ioctl,
170 };
171
172 static long ktapvm_ioctl(struct file *file, unsigned int cmd, unsigned long arg)
173 {
174     int new_fd, err;
175     struct file *new_file;

```

```

176 new_fd = get_unused_fd();
177 if (new_fd < 0)
178     return new_fd;
179
180
181 new_file = anon_inode_getfile("[ktap]", &ktap_fops, NULL, O_RDWR);
182 if (IS_ERR(new_file)) {
183     err = PTR_ERR(new_file);
184     put_unused_fd(new_fd);
185     return err;
186 }
187
188 file->private_data = NULL;
189 fd_install(new_fd, new_file);
190 return new_fd;
191 }
192
193 static const struct file_operations ktapvm_fops = {
194     .owner = THIS_MODULE,
195     .unlocked_ioctl = ktapvm_ioctl,
196 };
197
198 int (*kp_ftrace_profile_set_filter)(struct perf_event *event, int event_id,
199     const char *filter_str);
200
201 struct syscall_metadata **syscalls_metadata;
202
203 /*TODO: kill this function in future */
204 static int __init init_dummy_kernel_functions(void)
205 {
206     unsigned long *addr;
207
208     /*
209      * ktap need symbol ftrace_profile_set_filter to set event filter,
210      * export it in future.
211     */
212 #ifdef CONFIG_PPC64
213     kp_ftrace_profile_set_filter =
214         (void *)kallsyms_lookup_name(".ftrace_profile_set_filter");
215 #else
216     kp_ftrace_profile_set_filter =
217         (void *)kallsyms_lookup_name("ftrace_profile_set_filter");
218 #endif
219 if (!kp_ftrace_profile_set_filter) {
220     pr_err("ktap: cannot lookup ftrace_profile_set_filter "
221         "in kallsyms\n");
222     return -1;
223 }
224
225 /* use syscalls_metadata for syscall event handling */
226 addr = (void *)kallsyms_lookup_name("syscalls_metadata");
227 if (!addr) {
228     pr_err("ktap: cannot lookup syscalls_metadata in kallsyms\n");
229     return -1;
230 }
231
232 syscalls_metadata = (struct syscall_metadata **)addr;
233 return 0;
234 }
235
236 static int __init init_ktap(void)
237 {
238     struct dentry *ktapvm_dentry;
239
240     if (init_dummy_kernel_functions())
241         return -1;
242
243     kp_dir_dentry = debugfs_create_dir("ktap", NULL);
244     if (!kp_dir_dentry) {
245         pr_err("ktap: debugfs_create_dir failed\n");
246         return -1;
247     }
248
249     ktapvm_dentry = debugfs_create_file("ktapvm", 0444, kp_dir_dentry, NULL,
250         &ktapvm_fops);
251

```

```
252     if (!ktapvm_dentry) {
253         pr_err("ktapvm: cannot create ktapvm file\n");
254         debugfs_remove_recursive(kp\_dir\_dentry);
255         return -1;
256     }
257
258     return 0;
259 }
260
261 static void __exit exit_ktap(void)
262 {
263     debugfs_remove_recursive(kp\_dir\_dentry);
264 }
265
266 module_init(init_ktap);
267 module_exit(exit_ktap);
268
269 MODULE_AUTHOR("Jovi Zhangwei <jovi.zhangwei@gmail.com>");
270 MODULE_DESCRIPTION("ktap");
271 MODULE_LICENSE("GPL");
272
273 int kp_max_loop_count = 100000;
274 module_param_named(max_loop_count, kp\_max\_loop\_count, int, S_IRUGO | S_IWUSR);
275 MODULE_PARM_DESC(max_loop_count, "max loop execution count");
276
```

[One Level Up](#)

[Top Level](#)

test/benchmark/sembench.c - ktap

Global variables defined

- [all_done](#)
- [allops](#)
- [futex_sem_ops](#)
- [ipc_sem_ops](#)
- [max_burns](#)
- [min_burns](#)
- [nanosleep_sem_ops](#)
- [num_threads](#)
- [semid_lookup](#)
- [thread_count](#)
- [timeout_test](#)
- [total_burns](#)
- [workers_started](#)
- [worklist](#)
- [worklist_mutex](#)

Data types defined

- [lockinfo](#)
- [sem_operations](#)
- [sem_wakeup_info](#)

Functions defined

- [cleanup_futex_sems](#)
- [cleanup_ipc_sems](#)
- [cleanup_nanosleep_sems](#)
- [futex](#)
- [futex_wake_some](#)
- [ipc_wake_some](#)
- [main](#)
- [nanosleep_wake_some](#)
- [print_usage](#)

- [setup_futex_sems](#)
- [setup_ipc_sems](#)
- [setup_nanosleep_sems](#)
- [smp_mb](#)
- [wait_futex_sem](#)
- [wait_ipc_sem](#)
- [wait_nanosleep_sem](#)
- [worker](#)
- [worklist_add](#)
- [worklist_rm](#)

Macros defined

- [FUTEX_CMP_REQUEUE](#)
- [FUTEX_FD](#)
- [FUTEX_REQUEUE](#)
- [FUTEX_WAIT](#)
- [FUTEX_WAKE](#)
- [FUTEX_WAKE_OP](#)
- [NUM_OPERATIONS](#)
- [SEMS_PERID](#)
- [VERSION](#)
- [_GNU_SOURCE](#)
- [_POSIX_C_SOURCE](#)

Source code

```

1 /*
2  * copyright Oracle 2007. Licensed under GPLv2
3  * To compile: gcc -Wall -o sembench sembench.c -lpthread
4  *
5  * usage: sembench -t thread count -w wakenum -r runtime -o op
6  * op can be: 0 (ipc sem) 1 (nanosleep) 2 (futexes)
7  *
8  * example:
9  *   sembench -t 1024 -w 512 -r 60 -o 2
10 * runs 1024 threads, waking up 512 at a time, running for 60 seconds using
11 * futex locking.
12 *
13 */
14 #define _GNU_SOURCE
15 #define _POSIX_C_SOURCE 199309
16 #include <fcntl.h>
17 #include <sched.h>
18 #include <stdio.h>
19 #include <stdlib.h>
20 #include <sys/sem.h>
21 #include <sys/ipc.h>

```

```

22 #include <sys/types.h>
23 #include <sys/mman.h>
24 #include <pthread.h>
25 #include <unistd.h>
26 #include <string.h>
27 #include <time.h>
28 #include <sys/time.h>
29 #include <sys/syscall.h>
30 #include <errno.h>
31
32 #define VERSION "0.2"
33
34 /* futexes have been around since 2.5.something, but it still seems I
35 * need to make my own syscall. Sigh.
36 */
37 #define FUTEX_WAIT          0
38 #define FUTEX_WAKE         1
39 #define FUTEX_FD           2
40 #define FUTEX_REQUEUE     3
41 #define FUTEX_CMP_REQUEUE 4
42 #define FUTEX_WAKE_OP     5
43 static inline int futex (int *uaddr, int op, int val,
44                         const struct timespec *timeout,
45                         int *uaddr2, int val3)
46 {
47     return syscall(__NR_futex, uaddr, op, val, timeout, uaddr2, val3);
48 }
49
50 static void smp_mb(void)
51 {
52     __sync_synchronize();
53 }
54
55 static int all_done = 0;
56 static int timeout_test = 0;
57
58 #define SEMS_PERID 250
59
60 struct sem_operations;
61
62 struct lockinfo {
63     unsigned long id;
64     unsigned long index;
65     int data;
66     pthread_t tid;
67     struct lockinfo *next;
68     struct sem_operations *ops;
69     unsigned long ready;
70 };
71
72 struct sem_wakeup_info {
73     int wakeup_count;
74     struct sembuf sb[SEMS_PERID];
75 };
76
77 struct sem_operations {
78     void (*wait)(struct lockinfo *l);
79     int (*wake)(struct sem_wakeup_info *wi, int num_semids, int num);
80     void (*setup)(struct sem_wakeup_info **wi, int num_semids);
81     void (*cleanup)(int num_semids);
82     char *name;
83 };
84
85 int *semid_lookup = NULL;
86
87 pthread_mutex_t worklist_mutex = PTHREAD_MUTEX_INITIALIZER;
88 static unsigned long total_burns = 0;
89 static unsigned long min_burns = ~0UL;
90 static unsigned long max_burns = 0;
91
92 /* currently running threads */
93 static int thread_count = 0;
94
95 struct lockinfo *worklist = NULL;
96 static int workers_started = 0;
97

```

```

98  /* total threads started */
99  static int num_threads = 2048;
100
101  static void worklist_add(struct lockinfo *l)
102  {
103      smp_mb();
104      l->ready = 1;
105  }
106
107  static struct lockinfo *worklist_rm(void)
108  {
109      static int last_index = 0;
110      int i;
111      struct lockinfo *l;
112
113      for (i = 0; i < num_threads; i++) {
114          int test = (last_index + i) % num_threads;
115
116          l = worklist + test;
117          smp_mb();
118          if (l->ready) {
119              l->ready = 0;
120              last_index = test;
121              return l;
122          }
123      }
124      return NULL;
125  }
126
127  /* ipc semaphore post& wait */
128  void wait_ipc_sem(struct lockinfo *l)
129  {
130      struct sembuf sb;
131      int ret;
132      struct timespec *tvp = NULL;
133      struct timespec tv = { 0, 1 };
134
135      sb.sem_num = l->index;
136      sb.sem_flg = 0;
137
138      sb.sem_op = -1;
139      l->data = 1;
140
141      if (timeout_test && (l->id % 5) == 0)
142          tvp = &tv;
143
144      worklist_add(l);
145      ret = semtimedop(semid_lookup[l->id], &sb, 1, tvp);
146
147      while(l->data != 0 && tvp) {
148          struct timespec tv2 = { 0, 500 };
149          nanosleep(&tv2, NULL);
150      }
151
152      if (l->data != 0) {
153          if (tvp)
154              return;
155          fprintf(stderr, "wakeup without data update\n");
156          exit(1);
157      }
158      if (ret) {
159          if (errno == EAGAIN && tvp)
160              return;
161          perror("semtimed op");
162          exit(1);
163      }
164  }
165
166  int ipc_wake_some(struct sem_wakeup_info *wi, int num_semids, int num)
167  {
168      int i;
169      int ret;
170      struct lockinfo *l;
171      int found = 0;
172
173      for (i = 0; i < num_semids; i++) {

```

```

174     wi[i].wakeup_count = 0;
175 }
176 while(num > 0) {
177     struct sembuf *sb;
178     l = worklist\_rm\(\);
179     if (!l)
180         break;
181     if (l->data != 1)
182         fprintf(stderr, "warning, lockinfo data was %d\n",
183             l->data);
184     l->data = 0;
185     sb = wi[l->id].sb + wi[l->id].wakeup_count;
186     sb->sem_num = l->index;
187     sb->sem_op = 1;
188     sb->sem_flg = IPC_NOWAIT;
189     wi[l->id].wakeup_count++;
190     found++;
191     num--;
192 }
193 if (!found)
194     return 0;
195 for (i = 0; i < num_semids; i++) {
196     int wakeup_total;
197     int cur;
198     int offset = 0;
199     if (!wi[i].wakeup_count)
200         continue;
201     wakeup_total = wi[i].wakeup_count;
202     while(wakeup_total > 0) {
203         cur = wakeup_total > 64 ? 64 : wakeup_total;
204         ret = semtimedop(semid\_lookup[i], wi[i].sb + offset,
205             cur, NULL);
206         if (ret) {
207             perror("semtimedop");
208             exit(1);
209         }
210         offset += cur;
211         wakeup_total -= cur;
212     }
213 }
214 return found;
215 }
216
217 void setup\_ipc\_sems(struct sem\_wakeup\_info **wi, int num_semids)
218 {
219     int i;
220     *wi = malloc(sizeof(**wi) * num_semids);
221     semid\_lookup = malloc(num_semids * sizeof(int));
222     for(i = 0; i < num_semids; i++) {
223         semid\_lookup[i] = semget(IPC_PRIVATE, SEMS\_PERID,
224             IPC_CREAT | 0777);
225         if (semid\_lookup[i] < 0) {
226             perror("semget");
227             exit(1);
228         }
229     }
230     sleep(10);
231 }
232
233 void cleanup\_ipc\_sems(int num)
234 {
235     int i;
236     for (i = 0; i < num; i++) {
237         semctl(semid\_lookup[i], 0, IPC_RMID);
238     }
239 }
240
241 struct sem\_operations ipc\_sem\_ops = {
242     .wait = wait\_ipc\_sem,
243     .wake = ipc\_wake\_some,
244     .setup = setup\_ipc\_sems,
245     .cleanup = cleanup\_ipc\_sems,
246     .name = "ipc sem operations",
247 };
248
249 /* futex post & wait */

```

```

250 void wait_futex_sem(struct lockinfo *l)
251 {
252     int ret;
253     l->data = 1;
254     worklist\_add(l);
255     while(l->data == 1) {
256         ret = futex(&l->data, FUTEX\_WAIT, 1, NULL, NULL, 0);
257         /*
258         if (ret && ret != EWOULDBLOCK) {
259             perror("futex wait");
260             exit(1);
261         }*/
262     }
263 }
264
265 int futex\_wake\_some(struct sem\_wakeup\_info *wi, int num_semids, int num)
266 {
267     int i;
268     int ret;
269     struct lockinfo *l;
270     int found = 0;
271
272     for (i = 0; i < num; i++) {
273         l = worklist\_rm();
274         if (!l)
275             break;
276         if (l->data != 1)
277             fprintf(stderr, "warning, lockinfo data was %d\n",
278                 l->data);
279         l->data = 0;
280         ret = futex(&l->data, FUTEX\_WAKE, 1, NULL, NULL, 0);
281         if (ret < 0) {
282             perror("futex wake");
283             exit(1);
284         }
285         found++;
286     }
287     return found;
288 }
289
290 void setup\_futex\_sems(struct sem\_wakeup\_info **wi, int num_semids)
291 {
292     return;
293 }
294
295 void cleanup\_futex\_sems(int num)
296 {
297     return;
298 }
299
300 struct sem\_operations futex\_sem\_ops = {
301     .wait = wait\_futex\_sem,
302     .wake = futex\_wake\_some,
303     .setup = setup\_futex\_sems,
304     .cleanup = cleanup\_futex\_sems,
305     .name = "futex sem operations",
306 };
307
308 /* nanosleep sems here */
309 void wait\_nanosleep\_sem(struct lockinfo *l)
310 {
311     int ret;
312     struct timespec tv = { 0, 1000000 };
313     int count = 0;
314
315     l->data = 1;
316     worklist\_add(l);
317     while(l->data) {
318         ret = nanosleep(&tv, NULL);
319         if (ret) {
320             perror("nanosleep");
321             exit(1);
322         }
323         count++;
324     }
325 }

```

```

326
327 int nanosleep_wake_some(struct sem\_wakeup\_info *wi, int num_semid, int num)
328 {
329     int i;
330     struct lockinfo *l;
331
332     for (i = 0; i < num; i++) {
333         l = worklist\_rm();
334         if (!l)
335             break;
336         if (l->data != 1)
337             fprintf(stderr, "warning, lockinfo data was %d\n",
338                 l->data);
339         l->data = 0;
340     }
341     return i;
342 }
343
344 void setup_nanosleep_sems(struct sem\_wakeup\_info **wi, int num_semid)
345 {
346     return;
347 }
348
349 void cleanup_nanosleep_sems(int num)
350 {
351     return;
352 }
353
354 struct sem\_operations nanosleep_sem_ops = {
355     .wait = wait\_nanosleep\_sem,
356     .wake = nanosleep\_wake\_some,
357     .setup = setup\_nanosleep\_sems,
358     .cleanup = cleanup\_nanosleep\_sems,
359     .name = "nano sleep sem operations",
360 };
361
362 void *worker(void *arg)
363 {
364     struct lockinfo *l = (struct lockinfo *)arg;
365     int burn_count = 0;
366     pthread_t tid = pthread_self();
367     size_t pagesize = getpagesize();
368     char *buf = malloc(pagesize);
369
370     if (!buf) {
371         perror("malloc");
372         exit(1);
373     }
374
375     l->tid = tid;
376     workers\_started = 1;
377     smp\_mb();
378
379     while(!all\_done) {
380         l->ops->wait(l);
381         if (all\_done)
382             break;
383         burn_count++;
384     }
385     pthread_mutex_lock(&worklist\_mutex);
386     total\_burns += burn_count;
387     if (burn_count < min\_burns)
388         min\_burns = burn_count;
389     if (burn_count > max\_burns)
390         max\_burns = burn_count;
391     thread\_count--;
392     pthread_mutex_unlock(&worklist\_mutex);
393     return (void *)0;
394 }
395
396 void print_usage(void)
397 {
398     printf("usage: sembench [-t threads] [-w wake incr] [-r runtime]);
399     printf("                [-o num] (0=ipc, 1=nanosleep, 2=futex)\n");
400     exit(1);
401 }

```

```

402
403 #define NUM_OPERATIONS 3
404 struct sem_operations *allops[NUM_OPERATIONS] = { &ipc_sem_ops,
405           &nanosleep_sem_ops,
406           &futex_sem_ops};
407
408 int main(int ac, char **av) {
409     int ret;
410     int i;
411     int semid = 0;
412     int sem_num = 0;
413     int burn_count = 0;
414     struct sem_wakeup_info *wi = NULL;
415     struct timeval start;
416     struct timeval now;
417     int num_semids = 0;
418     int wake_num = 256;
419     int run_secs = 30;
420     int pagesize = getpagesize();
421     char *buf = malloc(pagesize);
422     struct sem_operations *ops = allops[0];
423     cpu_set_t cpu_mask;
424     cpu_set_t target_mask;
425     int target_cpu = 0;
426     int max_cpu = -1;
427
428     if (!buf) {
429         perror("malloc");
430         exit(1);
431     }
432     for (i = 1; i < ac; i++) {
433         if (strcmp(av[i], "-t") == 0) {
434             if (i == ac - 1)
435                 print_usage();
436             num_threads = atoi(av[i+1]);
437             i++;
438         } else if (strcmp(av[i], "-w") == 0) {
439             if (i == ac - 1)
440                 print_usage();
441             wake_num = atoi(av[i+1]);
442             i++;
443         } else if (strcmp(av[i], "-r") == 0) {
444             if (i == ac - 1)
445                 print_usage();
446             run_secs = atoi(av[i+1]);
447             i++;
448         } else if (strcmp(av[i], "-o") == 0) {
449             int index;
450             if (i == ac - 1)
451                 print_usage();
452             index = atoi(av[i+1]);
453             if (index >= NUM_OPERATIONS) {
454                 fprintf(stderr, "invalid operations %d\n",
455                     index);
456                 exit(1);
457             }
458             ops = allops[index];
459             i++;
460         } else if (strcmp(av[i], "-T") == 0) {
461             timeout_test = 1;
462         } else if (strcmp(av[i], "-h") == 0) {
463             print_usage();
464         }
465     }
466     num_semids = (num_threads + SEMS_PERID - 1) / SEMS_PERID;
467     ops->setup(&wi, num_semids);
468
469     ret = sched_getaffinity(0, sizeof(cpu_set_t), &cpu_mask);
470     if (ret) {
471         perror("sched_getaffinity");
472         exit(1);
473     }
474     for (i = 0; i < CPU_SETSIZE; i++)
475         if (CPU_ISSET(i, &cpu_mask))
476             max_cpu = i;
477     if (max_cpu == -1) {

```

```

478     fprintf(stderr, "sched_getaffinity returned empty mask\n");
479     exit(1);
480 }
481
482 CPU_ZERO(&target_mask);
483
484 worklist = malloc(sizeof(*worklist) * num_threads);
485 memset(worklist, 0, sizeof(*worklist) * num_threads);
486
487 for (i = 0; i < num_threads; i++) {
488     struct lockinfo *l;
489     pthread_t tid;
490     thread_count++;
491     l = worklist + i;
492     if (!l) {
493         perror("malloc");
494         exit(1);
495     }
496     l->id = semid;
497     l->index = sem_num++;
498     l->ops = ops;
499     if (sem_num >= SEMS_PERID) {
500         semid++;
501         sem_num = 0;
502     }
503     ret = pthread_create(&tid, NULL, worker, (void *)l);
504     if (ret) {
505         perror("pthread_create");
506         exit(1);
507     }
508
509     while (!CPU_ISSET(target_cpu, &cpu_mask)) {
510         target_cpu++;
511         if (target_cpu > max_cpu)
512             target_cpu = 0;
513     }
514     CPU_SET(target_cpu, &target_mask);
515     ret = pthread_setaffinity_np(tid, sizeof(cpu_set_t),
516                                &target_mask);
517     CPU_CLR(target_cpu, &target_mask);
518     target_cpu++;
519
520     ret = pthread_detach(tid);
521     if (ret) {
522         perror("pthread_detach");
523         exit(1);
524     }
525 }
526 while(!workers_started) {
527     smp_mb();
528     usleep(200);
529 }
530 gettimeofday(&start, NULL);
531 //fprintf(stderr, "main loop going\n");
532 while(1) {
533     ops->wake(wi, num_semids, wake_num);
534     burn_count++;
535     gettimeofday(&now, NULL);
536     if (now.tv_sec - start.tv_sec >= run_secs)
537         break;
538 }
539 //fprintf(stderr, "all done\n");
540 all_done = 1;
541 while(thread_count > 0) {
542     ops->wake(wi, num_semids, wake_num);
543     usleep(200);
544 }
545 //printf("%d threads, waking %d at a time\n", num_threads, wake_num);
546 //printf("using %s\n", ops->name);
547 //printf("main thread burns: %d\n", burn_count);
548 //printf("worker burn count total %lu min %lu max %lu avg %lu\n",
549 //        total_burns, min_burns, max_burns, total_burns / num_threads);
550 printf("%d seconds: %lu worker burns per second\n",
551        (int)(now.tv_sec - start.tv_sec),
552        total_burns / (now.tv_sec - start.tv_sec));
553 ops->cleanup(num_semids);

```

```
554     return 0;  
555 }  
556
```

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

test/benchmark/ - ktap

- [cmp_neq.sh](#)
- [cmp_profile.sh](#)
- [cmp_table.sh](#)
- [sembench.c](#)
- [test.sh](#)

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

test/ - ktap

- [_vimrc_local.vim](#)
- [benchmark/](#)
- [ffi/](#)
- [ffi_test/](#)
- [lib/](#)
- [util/](#)

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

test/_vimrc_local.vim - ktap

```
1 set expandtab
2 set tw=10000
3 "filetype indent off
4 "filetype plugin off
5 "set nowrap autoindent nosmartindent nocindent indentexpr=
```

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

test/ffi/ - ktap

- [ktap_ffi_test.mod.c](#)

[One Level Up](#)

[Top Level](#)

test/ffi/ktap_ffi_test.mod.c - ktap

Global variables defined

- [__this_module](#)
- [__used](#)

Source code

```
1 #include <linux/module.h>
2 #include <linux/vermagic.h>
3 #include <linux/compiler.h>
4
5 MODULE_INFO(vermagic, VERMAGIC_STRING);
6
7 struct module __this_module
8 __attribute__((section(".gnu.linkonce.this_module"))) = {
9     .name = KBUILD_MODNAME,
10    .init = init_module,
11    #ifdef CONFIG_MODULE_UNLOAD
12    .exit = cleanup_module,
13    #endif
14    .arch = MODULE_ARCH_INIT,
15 };
16
17 static const char __module_depends[]
18 __used
19 __attribute__((section(".modinfo"))) =
20 "depends=";
21
```

[One Level Up](#)

[Top Level](#)

test/ffi_test/ - ktap

- [Makefile](#)
- [cparser_test.c](#)
- [ktap_ffi_test.c](#)

[One Level Up](#)

[Top Level](#)

test/ffi_test/Makefile - ktap

```
1 obj-m += ktap_ffi_test.o
2
3 all: funct_mod cparser_test
4
5 funct_mod:
6     make -C /lib/modules/$(shell uname -r)/build M=$(PWD) modules
7
8 INC=../../include
9 U_DIR=../../userspace
10 RUNTIME=../../runtime
11 U_FFI_DIR=$(U_DIR)/ffi
12 CPARSER_FILES=cparser.o ctype.o ffi_type.o
13 KTAPC_CFLAGS = -Wall -O2
14
15 cparser.o: $(U_FFI_DIR)/cparser.c $(INC)/*
16     $(QUIET_CC)$(CC) -DCONFIG_KTAP_FFI -o $@ -c $<
17
18 ctype.o: $(U_FFI_DIR)/ctype.c $(INC)/*
19     $(QUIET_CC)$(CC) -DCONFIG_KTAP_FFI -o $@ -c $<
20
21 ffi_type.o: $(RUNTIME)/ffi/ffi_type.c $(INC)/*
22     $(QUIET_CC)$(CC) -DCONFIG_KTAP_FFI -o $@ -c $<
23
24 cparser_test: cparser_test.c $(CPARSER_FILES) $(INC)/*
25     $(QUIET_CC)$(CC) -DCONFIG_KTAP_FFI -I$(INC) -I$(U_DIR) $(KTAPC_CFLAGS) \
26     -o $@ $< $(CPARSER_FILES)
27
28 load:
29     insmod ktap_ffi_test.ko
30
31 unload:
32     rmmod ktap_ffi_test
33
34 test: all
35     @echo "testing cparser:"
36     ./cparser_test
37     @echo "testing ffi module:"
38     rmmod ktap_ffi_test > /dev/null 2>&1 || true
39     insmod ktap_ffi_test.ko
40     ../../ktap ffi_test.kp
41     rmmod ktap_ffi_test.ko
42     @echo "[*] all ffi tests passed."
43
44 clean:
45     make -C /lib/modules/$(shell uname -r)/build M=$(PWD) clean
46     rm -rf cparser_test
```

test/ffi_test/cparser_test.c - ktap

Global variables defined

- [csym_state](#)

Functions defined

- [find_kernel_symbol](#)
- [lookup_csymbol_id_by_name](#)
- [main](#)
- [test_func_func_module](#)
- [test_func_sched_clock](#)
- [test_func_time_to_tm](#)
- [test_pointer_symbols](#)
- [test_struct_timespec](#)
- [test_var_arg_function](#)

Macros defined

- [DO_TEST](#)
- [assert_csym_arr_type](#)
- [assert_farg_type](#)
- [assert_fret_type](#)
- [cs_arr](#)
- [cs_arr_size](#)
- [cs_nr](#)

Source code

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <string.h>
4 #include <assert.h>
5
6 #include "ktap_types.h"
7 #include "ktap_opcodes.h"
8 #include "../../userspace/ktapc.h"
9 #include "cparser.h"
10
11 void ffi_cparser_init(void);
12 void ffi_cparser_free(void);
13 int ffi_parse_cdef(const char *s);
14
15 static cp\_csymbol\_state *csym_state;
16
17 #define cs\_nr (csym_state->cs_nr)
18 #define cs\_arr\_size (csym_state->cs_arr_size)
19 #define cs\_arr (csym_state->cs_arr)
```

```

20
21
22 #define DO_TEST(name) do {
23     ffi_cparser_init();
24     int ret;
25     printf("[*] start #name test... ");
26     ret = test_##name();
27     if (ret)
28         fprintf(stderr, "\n[!] #name test failed.\n");
29     else
30         printf(" passed.\n");
31     ffi_cparser_free();
32 } while (0)
33
34 #define assert_csym_arr_type(cs_arr, n, t) do {
35     csymbol *ncs;
36     ncs = &cs_arr[n];
37     assert(ncs->type == t);
38 } while (0)
39
40 #define assert_fret_type(fcs, t) do {
41     csymbol *ncs;
42     ncs = &cs_arr[fcs->ret_id];
43     assert(ncs->type == t);
44 } while (0)
45
46 #define assert_farg_type(fcs, n, t) do {
47     csymbol *ncs;
48     ncs = &cs_arr[fcs->arg_ids[n]];
49     assert(ncs->type == t);
50 } while (0)
51
52
53
54
55 /* mock find_kernel_symbol */
56 unsigned long find_kernel_symbol(const char *symbol)
57 {
58     return 0xdeadbeef;
59 }
60
61 int lookup_csymbol_id_by_name(char *name)
62 {
63     int i;
64
65     for (i = 0; i < cs_nr; i++) {
66         if (!strcmp(name, cs_arr[i].name)) {
67             return i;
68         }
69     }
70
71     return -1;
72 }
73
74 int test_func_sched_clock()
75 {
76     int idx;
77     csymbol *cs;
78     csymbol_func *fcs;
79
80     ffi_parse_cdef("unsigned long long sched_clock()");
81
82     csym_state = ctype_get_csym_state();
83     assert(cs_arr);
84
85     idx = lookup_csymbol_id_by_name("sched_clock");
86     assert(idx >= 0);
87     cs = &cs_arr[idx];
88     assert(cs->type == FFI_FUNC);
89
90     fcs = csym_func(cs);
91
92     /* check return type */
93     assert_fret_type(fcs, FFI_UINT64);
94
95     /* check arguments */

```

```

96     assert(fcs->arg_nr == 0);
97
98     return 0;
99 }
100
101 int test_func_funct_module()
102 {
103     int idx;
104     csymbol *cs;
105     csymbol_func *fcs;
106
107     ffi_parse_cdef("void funct_void();");
108     ffi_parse_cdef("int funct_int1(unsigned char a, char b, unsigned short c, "
109         "short d);");
110     ffi_parse_cdef("long long funct_int2(unsigned int a, int b, "
111         "unsigned long c, long d, unsigned long long e, "
112         "long long f, long long g);");
113     ffi_parse_cdef("void *funct_pointer1(char *a);");
114
115     csym_state = ctype_get_csym_state();
116     assert(cs_arr);
117
118     /* check funct_void function */
119     idx = lookup_csymbol_id_by_name("funct_void");
120     assert(idx >= 0);
121     cs = &cs_arr[idx];
122     assert(cs->type == FFI_FUNC);
123     fcs = csym_func(cs);
124
125     /* check return type */
126     assert_fret_type(fcs, FFI_VOID);
127
128     /* check arguments */
129     assert(fcs->arg_nr == 0);
130
131
132
133     /* check funct_int1 function */
134     idx = lookup_csymbol_id_by_name("funct_int1");
135     assert(idx >= 0);
136     cs = &cs_arr[idx];
137     assert(cs);
138     assert(cs->type == FFI_FUNC);
139     fcs = csym_func(cs);
140
141     /* check return type */
142     assert_fret_type(fcs, FFI_INT32);
143
144     /* check arguments */
145     assert(fcs->arg_nr == 4);
146     assert_farg_type(fcs, 0, FFI_UINT8);
147     assert_farg_type(fcs, 1, FFI_INT8);
148     assert_farg_type(fcs, 2, FFI_UINT16);
149     assert_farg_type(fcs, 3, FFI_INT16);
150
151
152
153     /* check funct_int2 function */
154     idx = lookup_csymbol_id_by_name("funct_int2");
155     assert(idx >= 0);
156     cs = &cs_arr[idx];
157     assert(cs);
158     assert(cs->type == FFI_FUNC);
159     fcs = csym_func(cs);
160
161     /* check return type */
162     assert_fret_type(fcs, FFI_INT64);
163
164     /* check arguments */
165     assert(fcs->arg_nr == 7);
166     assert_farg_type(fcs, 0, FFI_UINT32);
167     assert_farg_type(fcs, 1, FFI_INT32);
168     assert_farg_type(fcs, 2, FFI_UINT64);
169     assert_farg_type(fcs, 3, FFI_INT64);
170     assert_farg_type(fcs, 4, FFI_UINT64);
171     assert_farg_type(fcs, 5, FFI_INT64);

```

```

172 assert\_farg\_type(fcs, 6, FFI_INT64);
173
174
175
176 /* check funct_pointer1 function */
177 idx = lookup\_csymbol\_id\_by\_name("funct_pointer1");
178 assert(idx >= 0);
179 cs = &cs\_arr[idx];
180 assert(cs);
181 assert(cs->type == FFI_FUNC);
182 fcs = csym\_func(cs);
183
184 /* check return type */
185 assert\_fret\_type(fcs, FFI_PTR);
186
187 /* check arguments */
188 assert(fcs->arg_nr == 1);
189 assert\_farg\_type(fcs, 0, FFI_PTR);
190 /*@TODO check pointer dereference type 18.11 2013 (houqp)*/
191
192 return 0;
193 }
194
195 int test_struct_timespec()
196 {
197     int idx;
198     csymbol *cs;
199     csymbol\_struct *stcs;
200
201     ffi\_parse\_cdef("struct timespec { long ts_sec; long ts_nsec; };");
202
203     csym\_state = ctype\_get\_csym\_state();
204     assert(cs\_arr);
205
206     idx = lookup\_csymbol\_id\_by\_name("struct timespec");
207     assert(idx >= 0);
208     cs = &cs\_arr[idx];
209     assert(cs);
210     assert(cs->type == FFI_STRUCT);
211
212     stcs = csym\_struct(cs);
213     assert(stcs->memb_nr == 2);
214
215     return 0;
216 }
217
218 int test_func_time_to_tm()
219 {
220     int idx;
221     csymbol *cs, *arg_cs;
222     csymbol\_struct *stcs;
223     csymbol\_func *fcs;
224
225     ffi\_parse\_cdef("typedef long time_t;");
226     ffi\_parse\_cdef("struct tm { "
227         "int tm_sec;"
228         "int tm_min;"
229         "int tm_hour;"
230         "int tm_mday;"
231         "int tm_mon;"
232         "long tm_year;"
233         "int tm_wday;"
234         "int tm_yday;"
235         "};");
236     ffi\_parse\_cdef("void time_to_tm(time_t totalsecs, int offset, struct tm *result);");
237
238     csym\_state = ctype\_get\_csym\_state();
239     assert(cs\_arr);
240
241     idx = lookup\_csymbol\_id\_by\_name("struct tm");
242     assert(idx >= 0);
243     cs = cp\_id\_to\_csym(idx);
244     assert(cs);
245     assert(cs->type == FFI_STRUCT);
246
247     stcs = csym\_struct(cs);

```

```

248     assert(stcs->memb_nr == 8);
249
250
251     idx = lookup\_csymbol\_id\_by\_name("time_to_tm");
252     assert(idx >= 0);
253     cs = cp\_id\_to\_csym(idx);
254     assert(cs);
255     assert(cs->type == FFI_FUNC);
256
257     fcs = csym\_func(cs);
258     assert(csymf\_arg\_nr(fcs) == 3);
259     /* check first argument */
260     assert\_farg\_type(fcs, 0, FFI_INT64);
261
262     /* check second argument */
263     assert\_farg\_type(fcs, 1, FFI_INT32);
264     /* check third argument */
265     assert\_farg\_type(fcs, 2, FFI_PTR);
266     arg_cs = cp\_csymf\_arg(fcs, 2);
267     assert(!strcmp(csym\_name(arg_cs), "struct tm *"));
268     assert(csym\_ptr\_deref\_id(arg_cs) ==
269           lookup\_csymbol\_id\_by\_name("struct tm"));
270
271     return 0;
272 }
273
274 int test_pointer_symbols()
275 {
276     csymbol\_func *fcs_foo, *fcs_bar;
277
278     /* int pointer symbol should be resolved to the same id */
279     ffi\_parse\_cdef("void foo(int *a);");
280     ffi\_parse\_cdef("int *bar(void);");
281
282     csym\_state = ctype\_get\_csym\_state();
283     assert(cs\_arr);
284
285     fcs_foo = csym\_func(cp\_id\_to\_csym(lookup\_csymbol\_id\_by\_name("foo")));
286     fcs_bar = csym\_func(cp\_id\_to\_csym(lookup\_csymbol\_id\_by\_name("bar")));
287
288     assert(csymf\_arg\_ids(fcs_foo)[0] == csymf\_ret\_id(fcs_bar));
289     assert(cp\_csymf\_arg(fcs_foo, 0) == cp\_csymf\_ret(fcs_bar));
290
291     return 0;
292 }
293
294 int test_var_arg_function()
295 {
296     csymbol\_func *fcs;
297
298     ffi\_parse\_cdef("int printk(char *fmt, ...);");
299
300     fcs = csym\_func(cp\_id\_to\_csym(lookup\_csymbol\_id\_by\_name("printk")));
301
302     /* var arg function needs void * type argument type checking */
303     assert(lookup\_csymbol\_id\_by\_name("void *") >= 0);
304
305     assert\_fret\_type(fcs, FFI_INT32);
306     assert\_farg\_type(fcs, 0, FFI_PTR);
307     assert(fcs->has_var_arg);
308
309     return 0;
310 }
311
312 int main (int argc, char *argv[])
313 {
314     DO\_TEST(func_sched_clock);
315     DO\_TEST(func_funct_module);
316     DO\_TEST(struct_timespec);
317     DO\_TEST(func_time_to_tm);
318     DO\_TEST(pointer_symbols);
319     DO\_TEST(var_arg_function);
320
321     return 0;
322 }

```

[One Level Up](#)

[Top Level](#)

userspace/ffi/cparser.c - ktap

Global variables defined

- [tok1](#)
- [tok2](#)
- [tok3](#)

Data types defined

- [etoken](#)
- [test](#)
- [token](#)

Functions defined

- [calculate_constant](#)
- [calculate_constant1](#)
- [calculate_constant10](#)
- [calculate_constant11](#)
- [calculate_constant12](#)
- [calculate_constant13](#)
- [calculate_constant2](#)
- [calculate_constant3](#)
- [calculate_constant4](#)
- [calculate_constant5](#)
- [calculate_constant6](#)
- [calculate_constant7](#)
- [calculate_constant8](#)
- [calculate_constant9](#)
- [check_token](#)
- [cp_error](#)
- [cp_init_parser](#)
- [ffi_cparser_free](#)
- [ffi_cparser_init](#)
- [ffi_lookup_csymbol_id_by_name](#)
- [ffi_parse_cdef](#)

- [ffi_parse_new](#)
- [increase_ptr_deref_level](#)
- [instantiate_typedef](#)
- [max_bitfield_size](#)
- [next_token](#)
- [parse_argument](#)
- [parse_argument2](#)
- [parse_attribute](#)
- [parse_enum](#)
- [parse_function](#)
- [parse_function_arguments](#)
- [parse_record](#)
- [parse_root](#)
- [parse_struct](#)
- [parse_type](#)
- [parse_type_name](#)
- [parse_typedef](#)
- [put_back](#)
- [require_token](#)
- [set_type_name](#)

Macros defined

- [END](#)
- [INT_TYPE](#)
- [INT_TYPE](#)
- [IS_CONST](#)
- [IS_LITERAL](#)
- [IS_RESTRICT](#)
- [IS_VOLATILE](#)
- [LONG_TYPE](#)
- [LONG_TYPE](#)
- [PRAGMA_POP](#)
- [SHORT_TYPE](#)
- [SHORT_TYPE](#)

- [max](#)
- [min](#)

Source code

```

1 #include <stdarg.h>
2 #include "../cparser.h"
3
4 #define IS_CONST(tok) (IS_LITERAL(tok, "const") || IS_LITERAL(tok, "__const") \
5     || IS_LITERAL(tok, "__const_"))
6 #define IS_VOLATILE(tok) (IS_LITERAL(tok, "volatile") || \
7     IS_LITERAL(tok, "__volatile") || \
8     IS_LITERAL(tok, "__volatile_"))
9 #define IS_RESTRICT(tok) (IS_LITERAL(tok, "restrict") || \
10     IS_LITERAL(tok, "__restrict") || \
11     IS_LITERAL(tok, "__restrict_"))
12
13 #define max(a,b) ((a) < (b) ? (b) : (a))
14 #define min(a,b) ((a) < (b) ? (a) : (b))
15
16
17 enum etoken {
18     /* 0 - 3 */
19     TOK_NIL,
20     TOK_NUMBER,
21     TOK_STRING,
22     TOK_TOKEN,
23
24     /* the order of these values must match the token strings in lex.c */
25
26     /* 4 - 5 */
27     TOK_3_BEGIN,
28     TOK_VA_ARG,
29
30     /* 6 - 14 */
31     TOK_2_BEGIN,
32     TOK_LEFT_SHIFT, TOK_RIGHT_SHIFT, TOK_LOGICAL_AND, TOK_LOGICAL_OR,
33     TOK_LESS_EQUAL, TOK_GREATER_EQUAL, TOK_EQUAL, TOK_NOT_EQUAL,
34
35     /* 15 - 20 */
36     TOK_1_BEGIN,
37     TOK_OPEN_CURLY, TOK_CLOSE_CURLY, TOK_SEMICOLON, TOK_COMMA, TOK_COLON,
38     /* 21 - 30 */
39     TOK_ASSIGN, TOK_OPEN_PAREN, TOK_CLOSE_PAREN, TOK_OPEN_SQUARE, TOK_CLOSE_SQUARE,
40     TOK_DOT, TOK_AMPERSAND, TOK_LOGICAL_NOT, TOK_BITWISE_NOT, TOK_MINUS,
41     /* 31 - 40 */
42     TOK_PLUS, TOK_STAR, TOK_DIVIDE, TOK_MODULUS, TOK_LESS,
43     TOK_GREATER, TOK_BITWISE_XOR, TOK_BITWISE_OR, TOK_QUESTION, TOK_POUND,
44
45     /* 41 - 43 */
46     TOK_REFERENCE = TOK_AMPERSAND,
47     TOK_MULTIPLY = TOK_STAR,
48     TOK_BITWISE_AND = TOK_AMPERSAND,
49 };
50
51 struct token {
52     enum etoken type;
53     int64_t integer;
54     const char *str;
55     size_t size;
56 };
57
58 #define IS_LITERAL(TOK, STR) \
59     (((TOK).size == sizeof(STR) - 1) && \
60     0 == memcmp((TOK).str, STR, sizeof(STR) - 1))
61
62
63 static int parse_type_name(struct parser *P, char *type_name);
64 static void parse_argument(struct parser *P, struct cp_ctype *ct,
65     struct token *pname, struct parser *asmname);
66 static int parse_attribute(struct parser *P, struct token *tok,
67     struct cp_ctype *ct, struct parser *asmname);
68 static int parse_record(struct parser *P, struct cp_ctype *ct);

```

```

69 static void instantiate_typedef(struct parser *P, struct cp_ctype *tt,
70     const struct cp_ctype *ft);
71
72
73 /* the order of tokens _must_ match the order of the enum etoken enum */
74
75 static char tok3[][4] = {
76     "...", /* unused ">>=", "<<=", */
77 };
78
79 static char tok2[][3] = {
80     "<<", ">>", "&&", "||", "<=",
81     ">=", "==", "!=",
82     /* unused "+=", "-=", "*=", "/=", "%=", "&=", "^=",
83     * "|=", "++", "--", "->", "::", */
84 };
85
86 static char tok1[] = {
87     '{', '}', ';', ':', ',',
88     '=', '(', ')', '[', ']',
89     '.', '&', '!', '~', '-',
90     '+', '*', '/', '%', '<',
91     '>', '^', '|', '?', '#'
92 };
93
94
95 /* this function never returns, but it's an idiom to use it in C functions
96 * as return cp_error */
97 void cp_error(const char *err_msg_fmt, ...)
98 {
99     va_list ap;
100
101     fprintf(stderr, "cparser error:\n");
102
103     va_start(ap, err_msg_fmt);
104     vfprintf(stderr, err_msg_fmt, ap);
105     va_end(ap);
106
107     exit(EXIT_FAILURE);
108 }
109
110 static int set_type_name(char *dst, unsigned type, const char *src, int len)
111 {
112     int prefix_len;
113     char *prefix = NULL;
114
115     if (type == STRUCT_TYPE)
116         prefix = "struct ";
117     else if (type == UNION_TYPE)
118         prefix = "union ";
119     else
120         cp_error("Only set type name for struct or union\n");
121     prefix_len = sizeof(prefix);
122
123     if (len + prefix_len > MAX_TYPE_NAME_LEN)
124         return -1;
125
126     memset(dst, 0, MAX_TYPE_NAME_LEN);
127     strcpy(dst, prefix);
128     strncat(dst, src, len);
129
130     return 0;
131 }
132
133 static void increase_ptr_deref_level(struct parser *P, struct cp_ctype *ct)
134 {
135     if (ct->pointers == POINTER_MAX) {
136         cp_error("maximum number of pointer derefs reached - use a "
137             "struct to break up the pointers on line %d", P->line);
138     } else {
139         ct->pointers++;
140         ct->const_mask <<= 1;
141     }
142 }
143
144 static int next_token(struct parser *P, struct token *tok)

```

```

145 {
146     size_t i;
147     const char *s = P->next;
148
149     /* UTF8 BOM */
150     if (s[0] == '\xEF' && s[1] == '\xBB' && s[2] == '\xBF') {
151         s += 3;
152     }
153
154     /* consume whitespace and comments */
155     for (;;) {
156         /* consume whitespace */
157         while (*s == '\t' || *s == '\n' || *s == ' '
158             || *s == '\v' || *s == '\r') {
159             if (*s == '\n') {
160                 P->line++;
161             }
162             s++;
163         }
164
165         /* consume comments */
166         if (*s == '/' && *(s+1) == '/') {
167
168             s = strchr(s, '\n');
169             if (!s) {
170                 cp_error("non-terminated comment");
171             }
172
173         } else if (*s == '/' && *(s+1) == '*') {
174             s += 2;
175
176             for (;;) {
177                 if (s[0] == '\0') {
178                     cp_error("non-terminated comment");
179                 } else if (s[0] == '*' && s[1] == '/') {
180                     s += 2;
181                     break;
182                 } else if (s[0] == '\n') {
183                     P->line++;
184                 }
185                 s++;
186             }
187
188         } else if (*s == '\0') {
189             tok->type = TOK_NIL;
190             return 0;
191
192         } else {
193             break;
194         }
195     }
196
197     P->prev = s;
198
199     for (i = 0; i < sizeof(tok3) / sizeof(tok3[0]); i++) {
200         if (s[0] == tok3[i][0] && s[1] == tok3[i][1] && s[2] == tok3[i][2]) {
201             tok->type = (enum etoken) (TOK_3_BEGIN + 1 + i);
202             P->next = s + 3;
203             goto end;
204         }
205     }
206
207     for (i = 0; i < sizeof(tok2) / sizeof(tok2[0]); i++) {
208         if (s[0] == tok2[i][0] && s[1] == tok2[i][1]) {
209             tok->type = (enum etoken) (TOK_2_BEGIN + 1 + i);
210             P->next = s + 2;
211             goto end;
212         }
213     }
214
215     for (i = 0; i < sizeof(tok1) / sizeof(tok1[0]); i++) {
216         if (s[0] == tok1[i]) {
217             tok->type = (enum etoken) (TOK_1_BEGIN + 1 + i);
218             P->next = s + 1;
219             goto end;
220         }
221     }

```

```

221 }
222
223 if (*s == '.' || *s == '-' || ('0' <= *s && *s <= '9')) {
224     /* number */
225     tok->type = TOK_NUMBER;
226
227     /* split out the negative case so we get the full range of
228      * bits for unsigned (eg to support 0xFFFFFFFF where
229      * sizeof(long) == 4 */
230     if (*s == '-') {
231         tok->integer = strtol(s, (char**) &s, 0);
232     } else {
233         tok->integer = strtoul(s, (char**) &s, 0);
234     }
235
236     while (*s == 'u' || *s == 'U' || *s == 'l' || *s == 'L') {
237         s++;
238     }
239
240     P->next = s;
241     goto end;
242
243 } else if (*s == '\\' || *s == '\"') {
244     /* "\"" or "'" */
245     char quote = *s;
246     s++; /* jump over " */
247
248     tok->type = TOK_STRING;
249     tok->str = s;
250
251     while (*s != quote) {
252         if (*s == '\0' || (*s == '\\' && *(s+1) == '\0')) {
253             cp_error("string not finished\n");
254         }
255         if (*s == '\\') {
256             s++;
257         }
258         s++;
259     }
260
261     tok->size = s - tok->str;
262     s++; /* jump over " */
263     P->next = s;
264     goto end;
265
266 } else if (('a' <= *s && *s <= 'z') || ('A' <= *s && *s <= 'Z')
267           || *s == '_') {
268     /* tokens */
269     tok->type = TOK_TOKEN;
270     tok->str = s;
271
272     while (('a' <= *s && *s <= 'z') || ('A' <= *s && *s <= 'Z')
273           || *s == '_' || ('0' <= *s && *s <= '9')) {
274         s++;
275     }
276
277     tok->size = s - tok->str;
278     P->next = s;
279     goto end;
280 } else {
281     cp_error("invalid character %d", P->line);
282 }
283
284 end:
285     return 1;
286 }
287
288 static void require_token(struct parser *P, struct token *tok)
289 {
290     if (!next_token(P, tok)) {
291         cp_error("unexpected end");
292     }
293 }
294
295 static void check_token(struct parser *P, int type, const char *str,
296                       const char *err, ...)

```

```

297 {
298     va_list ap;
299     struct token tok;
300     if (!next_token(P, &tok) || tok.type != type
301         || (tok.type == TOK_TOKEN && (tok.size != strlen(str)
302             || memcmp(tok.str, str, tok.size) != 0))) {
303
304         va_start(ap, err);
305         vfprintf(stderr, err, ap);
306         va_end(ap);
307
308         exit(EXIT_FAILURE);
309     }
310 }
311
312 static void put_back(struct parser *P) {
313     P->next = P->prev;
314 }
315
316 int64_t calculate_constant(struct parser *P);
317
318 /* parses out the base type of a type expression in a function declaration,
319 * struct definition, typedef etc
320 *
321 * leaves the usr value of the type on the stack
322 */
323 int parse_type(struct parser *P, struct cp_ctype *ct)
324 {
325     struct token tok;
326
327     memset(ct, 0, sizeof(*ct));
328
329     require_token(P, &tok);
330
331     /* get function attributes before the return type */
332     while (parse_attribute(P, &tok, ct, NULL)) {
333         require_token(P, &tok);
334     }
335
336     /* get const/volatile before the base type */
337     for (;;) {
338         if (tok.type != TOK_TOKEN) {
339             cp_error("unexpected value before type name on line %d",
340                 P->line);
341             return 0;
342         } else if (IS_CONST(tok)) {
343             ct->const_mask = 1;
344             require_token(P, &tok);
345         } else if (IS_VOLATILE(tok) || IS_RESTRICT(tok)) {
346             /* ignored for now */
347             require_token(P, &tok);
348         } else {
349             break;
350         }
351     }
352
353     /* get base type */
354     if (tok.type != TOK_TOKEN) {
355         cp_error("unexpected value before type name on line %d", P->line);
356         return 0;
357     } else if (IS_LITERAL(tok, "struct")) {
358         ct->type = STRUCT_TYPE;
359         parse_record(P, ct);
360     } else if (IS_LITERAL(tok, "union")) {
361         ct->type = UNION_TYPE;
362         parse_record(P, ct);
363     } else if (IS_LITERAL(tok, "enum")) {
364         ct->type = ENUM_TYPE;
365         parse_record(P, ct);
366     } else {
367         put_back(P);
368
369         /* lookup type */
370         struct cp_ctype *lct;
371         char cur_type_name[MAX_TYPE_NAME_LEN];
372

```

```

373     memset(cur_type_name, 0, MAX_TYPE_NAME_LEN);
374     parse_type_name(P, cur_type_name);
375     lct = ctype_lookup_type(cur_type_name);
376     if (!lct)
377         cp_error("unknow type: \"%s\"\n", cur_type_name);
378
379     instantiate_typedef(P, ct, lct);
380 }
381
382 while (next_token(P, &tok)) {
383     if (tok.type != TOK_TOKEN) {
384         put_back(P);
385         break;
386     } else if (IS_CONST(tok) || IS_VOLATILE(tok)) {
387         /* ignore for now */
388     } else {
389         put_back(P);
390         break;
391     }
392 }
393
394 return 0;
395 }
396
397 enum test {TEST};
398
399 /* Parses an enum definition from after the open curly through to the close
400 * curly. Expects the user table to be on the top of the stack
401 */
402 static int parse_enum(struct parser *P, struct cp_ctype *type)
403 {
404     struct token tok;
405     int value = -1;
406
407     /*@TODO clean up this function when enum support is added*/
408     cp_error("TODO: enum not supported!\n");
409
410     for (;;) {
411         require_token(P, &tok);
412
413         if (tok.type == TOK_CLOSE_CURLY) {
414             break;
415         } else if (tok.type != TOK_TOKEN) {
416             cp_error("unexpected token in enum at line %d", P->line);
417             return 0;
418         }
419         require_token(P, &tok);
420
421         if (tok.type == TOK_COMMA || tok.type == TOK_CLOSE_CURLY) {
422             /* we have an auto calculated enum value */
423             value++;
424         } else if (tok.type == TOK_ASSIGN) {
425             /* we have an explicit enum value */
426             value = (int) calculate_constant(P);
427             require_token(P, &tok);
428         } else {
429             cp_error("unexpected token in enum at line %d", P->line);
430             return 0;
431         }
432
433         if (tok.type == TOK_CLOSE_CURLY) {
434             break;
435         } else if (tok.type != TOK_COMMA) {
436             cp_error("unexpected token in enum at line %d", P->line);
437             return 0;
438         }
439     }
440
441     type->base_size = sizeof(enum test);
442     type->align_mask = sizeof(enum test) - 1;
443
444     return 0;
445 }
446
447 /* Parses a struct from after the open curly through to the close curly. */
448 static int parse_struct(struct parser *P, const struct cp_ctype *ct)

```

```

449 {
450     struct token tok;
451
452     /* parse members */
453     for (;;) {
454         struct cp_ctype mbase;
455
456         /* see if we're at the end of the struct */
457         require_token(P, &tok);
458         if (tok.type == TOK_CLOSE_CURLY) {
459             break;
460         } else if (ct->is_variable_struct) {
461             cp_error("can't have members after a variable sized "
462                 "member on line %d", P->line);
463             return -1;
464         } else {
465             put_back(P);
466         }
467
468         /* members are of the form
469          * <base type> <arg>, <arg>, <arg>;
470          * eg struct foo bar, *bar2[2];
471          * mbase is 'struct foo'
472          * mtype is '' then '*[2]'
473          * mname is 'bar' then 'bar2'
474          */
475
476         parse_type(P, &mbase);
477
478         for (;;) {
479             struct token mname;
480             struct cp_ctype mt = mbase;
481
482             memset(&mname, 0, sizeof(mname));
483
484             if (ct->is_variable_struct) {
485                 cp_error("can't have members after a variable "
486                     "sized member on line %d", P->line);
487                 return -1;
488             }
489
490             parse_argument(P, &mt, &mname, NULL);
491
492             if (!mt.is_defined && (mt.pointers - mt.is_array) == 0) {
493                 cp_error("member type is undefined on line %d",
494                     P->line);
495                 return -1;
496             }
497
498             if (mt.type == VOID_TYPE
499                 && (mt.pointers - mt.is_array) == 0) {
500                 cp_error("member type can not be void on line %d",
501                     P->line);
502                 return -1;
503             }
504
505             mt.has_member_name = (mname.size > 0);
506             if (mt.has_member_name) {
507                 cp_push_ctype_with_name(&mt,
508                     mname.str, mname.size);
509             } else {
510                 cp_push_ctype(&mt);
511             }
512
513             require_token(P, &tok);
514             if (tok.type == TOK_SEMICOLON) {
515                 break;
516             } else if (tok.type != TOK_COMMA) {
517                 cp_error("unexpected token in struct "
518                     "definition on line %d", P->line);
519             }
520         }
521     }
522
523     return 0;
524 }

```

```

525
526 /* copy over attributes that could be specified before the typedef eg
527 * __attribute__((packed)) const type_t */
528 static void instantiate_typedef(struct parser *P, struct cp_ctype *tt,
529     const struct cp_ctype *ft)
530 {
531     struct cp_ctype pt = *tt;
532     *tt = *ft;
533
534     tt->const_mask |= pt.const_mask;
535     tt->is_packed = pt.is_packed;
536
537     if (tt->is_packed) {
538         tt->align_mask = 0;
539     } else {
540         /* Instantiate the typedef in the current packing. This may be
541         * further updated if a pointer is added or another alignment
542         * attribute is applied. If pt.align_mask is already non-zero
543         * than an increased alignment via __declspec(aligned(#)) has
544         * been set. */
545         tt->align_mask = max(min(P->align_mask, tt->align_mask),
546             pt.align_mask);
547     }
548 }
549
550 /* this parses a struct or union starting with the optional
551 * name before the opening brace
552 * leaves the type usr value on the stack */
553 static int parse_record(struct parser *P, struct cp_ctype *ct)
554 {
555     struct token tok;
556     char cur_type_name[MAX_TYPE_NAME_LEN];
557     int has_name;
558
559     require_token(P, &tok);
560
561     /* name is optional */
562     memset(cur_type_name, 0, MAX_TYPE_NAME_LEN);
563     if (tok.type == TOK_TOKEN) {
564         /* declaration */
565         struct cp_ctype *lct;
566
567         set_type_name(cur_type_name, ct->type, tok.str, tok.size);
568
569         /* lookup the name to see if we've seen this type before */
570         lct = ctype_lookup_type(cur_type_name);
571
572         if (!lct) {
573             /* new type, delay type registration to the end
574             * of this function */
575             ct->ffi_base_cs_id = ct->ffi_cs_id = -1;
576         } else {
577             /* get the existing declared type */
578             if (lct->type != ct->type) {
579                 cp_error("type '%s' previously declared as '%s'",
580                     cur_type_name,
581                     csym_name(ct ffi_cs(lct)));
582             }
583
584             instantiate_typedef(P, ct, lct);
585         }
586
587         /* if a name is given then we may be at the end of the string
588         * eg for ffi.new('struct foo') */
589         if (!next_token(P, &tok)) {
590             return 0;
591         }
592         has_name = 1;
593     } else {
594         char anon_name[MAX_TYPE_NAME_LEN];
595         /* create a new unnamed record */
596         sprintf(anon_name, "%d line", P->line);
597         set_type_name(cur_type_name, ct->type,
598             anon_name, strlen(anon_name));
599         ct->ffi_base_cs_id = ct->ffi_cs_id = -1;
600         has_name = 0;

```

```

601 }
602
603 if (tok.type != TOK_OPEN_CURLY) {
604     /* this may just be a declaration or use of the type as an
605      * argument or member */
606     put_back(P);
607     if (!has_name)
608         cp_error("noname record type declaration\n");
609
610     /* build symbol for vm */
611     ct->ffi_base_cs_id =
612         cp_symbol_build_fake_record(cur_type_name, ct->type);
613     ct->ffi_cs_id = ct->ffi_base_cs_id;
614     return 0;
615 }
616
617 if (ct->is_defined) {
618     cp_error("redefinition in line %d", P->line);
619     return 0;
620 }
621
622 if (ct->type == ENUM_TYPE) {
623     parse_enum(P, ct);
624     cp_set_defined(ct);
625 } else {
626     int start_top = ctype_stack_top();
627     /* we do a two stage parse, where we parse the content first
628      * and build up the temp user table. We then iterate over that
629      * to calculate the offsets and fill out ct_usr. This is so we
630      * can handle out of order members (eg vtable) and attributes
631      * specified at the end of the struct. */
632     parse_struct(P, ct);
633     cp_set_defined(ct);
634     /* build symbol for vm */
635     ct->ffi_base_cs_id = cp_symbol_build_record(
636         cur_type_name, ct->type, start_top);
637     ct->ffi_cs_id = ct->ffi_base_cs_id;
638     /* save cp_ctype for parser */
639     cp_ctype_reg_type(cur_type_name, ct);
640 }
641
642 return 0;
643 }
644
645 /* parses single or multi work built in types, and pushes it onto the stack */
646 static int parse_type_name(struct parser *P, char *type_name)
647 {
648     struct token tok;
649     int flags = 0;
650
651     enum {
652         UNSIGNED = 0x01,
653         SIGNED = 0x02,
654         LONG = 0x04,
655         SHORT = 0x08,
656         INT = 0x10,
657         CHAR = 0x20,
658         LONG_LONG = 0x40,
659         INT8 = 0x80,
660         INT16 = 0x100,
661         INT32 = 0x200,
662         INT64 = 0x400,
663     };
664
665     require_token(P, &tok);
666
667     /* we have to manually decode the builtin types since they can take up
668      * more than one token */
669     for (;;) {
670         if (tok.type != TOK_TOKEN) {
671             break;
672         } else if (IS_LITERAL(tok, "unsigned")) {
673             flags |= UNSIGNED;
674         } else if (IS_LITERAL(tok, "signed")) {
675             flags |= SIGNED;
676         } else if (IS_LITERAL(tok, "short")) {

```

```

677     flags |= SHORT;
678 } else if (IS_LITERAL(tok, "char")) {
679     flags |= CHAR;
680 } else if (IS_LITERAL(tok, "long")) {
681     flags |= (flags & LONG) ? LONG_LONG : LONG;
682 } else if (IS_LITERAL(tok, "int")) {
683     flags |= INT;
684 } else if (IS_LITERAL(tok, "__int8")) {
685     flags |= INT8;
686 } else if (IS_LITERAL(tok, "__int16")) {
687     flags |= INT16;
688 } else if (IS_LITERAL(tok, "__int32")) {
689     flags |= INT32;
690 } else if (IS_LITERAL(tok, "__int64")) {
691     flags |= INT64;
692 } else if (IS_LITERAL(tok, "register")) {
693     /* ignore */
694 } else {
695     break;
696 }
697
698 if (!next_token(P, &tok)) {
699     break;
700 }
701 }
702
703 if (flags) {
704     put_back(P);
705 }
706
707 if (flags & CHAR) {
708     if (flags & SIGNED) {
709         strcpy(type_name, "int8_t");
710     } else if (flags & UNSIGNED) {
711         strcpy(type_name, "uint8_t");
712     } else {
713         if (((char) -1) > 0) {
714             strcpy(type_name, "uint8_t");
715         } else {
716             strcpy(type_name, "int8_t");
717         }
718     }
719 } else if (flags & INT8) {
720     strcpy(type_name, (flags & UNSIGNED) ? "uint8_t" : "int8_t");
721 } else if (flags & INT16) {
722     strcpy(type_name, (flags & UNSIGNED) ? "uint16_t" : "int16_t");
723 } else if (flags & INT32) {
724     strcpy(type_name, (flags & UNSIGNED) ? "uint32_t" : "int32_t");
725 } else if (flags & INT64) {
726     strcpy(type_name, (flags & UNSIGNED) ? "uint64_t" : "int64_t");
727 } else if (flags & LONG_LONG) {
728     strcpy(type_name, (flags & UNSIGNED) ? "uint64_t" : "int64_t");
729 } else if (flags & SHORT) {
730 #define SHORT_TYPE(u) (sizeof(short) == sizeof(int64_t) ? \
731     u "int64_t" : sizeof(short) == sizeof(int32_t) ? \
732     u "int32_t" : u "int16_t")
733     if (flags & UNSIGNED) {
734         strcpy(type_name, SHORT_TYPE("u"));
735     } else {
736         strcpy(type_name, SHORT_TYPE(""));
737     }
738 #undef SHORT_TYPE
739 } else if (flags & LONG) {
740 #define LONG_TYPE(u) (sizeof(long) == sizeof(int64_t) ? \
741     u "int64_t" : u "int32_t")
742     if (flags & UNSIGNED) {
743         strcpy(type_name, LONG_TYPE("u"));
744     } else {
745         strcpy(type_name, LONG_TYPE(""));
746     }
747 #undef LONG_TYPE
748 } else if (flags) {
749 #define INT_TYPE(u) (sizeof(int) == sizeof(int64_t) ? \
750     u "int64_t" : sizeof(int) == sizeof(int32_t) ? \
751     u "int32_t" : u "int16_t")
752     if (flags & UNSIGNED) {

```



```

829         "on line %d", P->line);
830     }
831
832     switch (tok->integer) {
833     case 1: align = 0; break;
834     case 2: align = 1; break;
835     case 4: align = 3; break;
836     case 8: align = 7; break;
837     case 16: align = 15; break;
838     default:
839         cp_error("unsupported align "
840             "size on line %d",
841             P->line);
842     }
843
844     check_token(P, TOK_CLOSE_PAREN, NULL,
845         "expected align(#) on line %d",
846         P->line);
847     break;
848
849 default:
850     cp_error("expected align(#) on line %d",
851         P->line);
852 }
853
854 /* __attribute__(aligned(#)) is only supposed
855  * to increase alignment */
856 ct->align_mask = max(align, ct->align_mask);
857
858 } else if (IS_LITERAL(*tok, "packed")
859     || IS_LITERAL(*tok, "__packed__")) {
860     ct->align_mask = 0;
861     ct->is_packed = 1;
862
863 } else if (IS_LITERAL(*tok, "mode")
864     || IS_LITERAL(*tok, "__mode__")) {
865
866     check_token(P, TOK_OPEN_PAREN, NULL,
867         "expected mode(MODE) on line %d",
868         P->line);
869
870     require_token(P, tok);
871     if (tok->type != TOK_TOKEN) {
872         cp_error("expected mode(MODE) on line %d",
873             P->line);
874     }
875
876
877     struct {char ch; uint16_t v;} a16;
878     struct {char ch; uint32_t v;} a32;
879     struct {char ch; uint64_t v;} a64;
880
881     if (IS_LITERAL(*tok, "QI")
882         || IS_LITERAL(*tok, "__QI__")
883         || IS_LITERAL(*tok, "byte")
884         || IS_LITERAL(*tok, "__byte__")
885     ) {
886         ct->type = INT8_TYPE;
887         ct->base_size = sizeof(uint8_t);
888         ct->align_mask = 0;
889
890     } else if (IS_LITERAL(*tok, "HI")
891         || IS_LITERAL(*tok, "__HI__")) {
892         ct->type = INT16_TYPE;
893         ct->base_size = sizeof(uint16_t);
894         ct->align_mask = ALIGNOF(a16);
895
896     } else if (IS_LITERAL(*tok, "SI")
897         || IS_LITERAL(*tok, "__SI__")
898     #if defined ARCH_X86 || defined ARCH_ARM
899         || IS_LITERAL(*tok, "word")
900         || IS_LITERAL(*tok, "__word__")
901         || IS_LITERAL(*tok, "pointer")
902         || IS_LITERAL(*tok, "__pointer__")
903     #endif
904     ) {

```

```

905         ct->type = INT32_TYPE;
906         ct->base_size = sizeof(uint32_t);
907         ct->align_mask = ALIGNOF(a32);
908
909     } else if (IS_LITERAL(*tok, "DI")
910             || IS_LITERAL(*tok, "__DI__")
911 #if defined ARCH_X64
912             || IS_LITERAL(*tok, "word")
913             || IS_LITERAL(*tok, "__word__")
914             || IS_LITERAL(*tok, "pointer")
915             || IS_LITERAL(*tok, "__pointer__")
916 #endif
917         ) {
918         ct->type = INT64_TYPE;
919         ct->base_size = sizeof(uint64_t);
920         ct->align_mask = ALIGNOF(a64);
921
922     } else {
923         cp_error("unexpected mode on line %d",
924                 P->line);
925     }
926
927     check_token(P, TOK_CLOSE_PAREN, NULL,
928                 "expected mode(MODE) on line %d", P->line);
929
930 } else if (IS_LITERAL(*tok, "cdecl")
931         || IS_LITERAL(*tok, "__cdecl__")) {
932     ct->calling_convention = C_CALL;
933
934 } else if (IS_LITERAL(*tok, "fastcall")
935         || IS_LITERAL(*tok, "__fastcall__")) {
936     ct->calling_convention = FAST_CALL;
937
938 } else if (IS_LITERAL(*tok, "stdcall")
939         || IS_LITERAL(*tok, "__stdcall__")) {
940     ct->calling_convention = STD_CALL;
941 }
942 /* ignore unknown tokens within parentheses */
943 }
944 return 1;
945
946 } else if (IS_LITERAL(*tok, "__cdecl")) {
947     ct->calling_convention = C_CALL;
948     return 1;
949
950 } else if (IS_LITERAL(*tok, "__fastcall")) {
951     ct->calling_convention = FAST_CALL;
952     return 1;
953
954 } else if (IS_LITERAL(*tok, "__stdcall")) {
955     ct->calling_convention = STD_CALL;
956     return 1;
957
958 } else if (IS_LITERAL(*tok, "__extension__")
959         || IS_LITERAL(*tok, "extern")) {
960     /* ignore */
961     return 1;
962 } else {
963     return 0;
964 }
965 }
966
967 /* parses from after the opening paranthesis to after the closing paranthesis */
968 static void parse_function_arguments(struct parser* P, struct cp_ctype* ct)
969 {
970     struct token tok;
971     int args = 0;
972
973     for (;;) {
974         require_token(P, &tok);
975
976         if (tok.type == TOK_CLOSE_PAREN)
977             break;
978
979         if (args) {
980             if (tok.type != TOK_COMMA) {

```

```

981         cp_error("unexpected token in function "
982                 "argument %d on line %d",
983                 args, P->line);
984     }
985     require_token(P, &tok);
986 }
987
988 if (tok.type == TOK_VA_ARG) {
989     ct->has_var_arg = true;
990     check_token(P, TOK_CLOSE_PAREN, "",
991                 "unexpected token after ... in "
992                 "function on line %d",
993                 P->line);
994     break;
995 } else if (tok.type == TOK_TOKEN) {
996     struct cp_ctype at;
997
998     put_back(P);
999     parse_type(P, &at);
1000    parse_argument(P, &at, NULL, NULL);
1001
1002    /* array arguments are just treated as their
1003       * base pointer type */
1004    at.is_array = 0;
1005
1006    /* check for the c style int func(void) and error
1007       * on other uses of arguments of type void */
1008    if (at.type == VOID_TYPE && at.pointers == 0) {
1009        if (args) {
1010            cp_error("can't have argument of type "
1011                    "void on line %d",
1012                    P->line);
1013        }
1014
1015        check_token(P, TOK_CLOSE_PAREN, "",
1016                    "unexpected void in function on line %d",
1017                    P->line);
1018        break;
1019    }
1020    cp_push_ctype(&at);
1021    args++;
1022 } else {
1023     cp_error("unexpected token in function argument %d "
1024             "on line %d", args+1, P->line);
1025 }
1026 }
1027 }
1028
1029 static int max_bitfield_size(int type)
1030 {
1031     switch (type) {
1032     case BOOL_TYPE:
1033         return 1;
1034     case INT8_TYPE:
1035         return 8;
1036     case INT16_TYPE:
1037         return 16;
1038     case INT32_TYPE:
1039         return 32;
1040     case ENUM_TYPE:
1041         return 32;
1042     case INT64_TYPE:
1043         return 64;
1044     default:
1045         return -1;
1046     }
1047 }
1048
1049 static struct cp_ctype *parse_argument2(struct parser *P, struct cp_ctype *ct,
1050     struct token *name, struct parser *asmname);
1051
1052 /* parses from after the first ( in a function declaration or function pointer
1053    * can be one of:
1054    * void foo(...) before ...
1055    * void (foo)(...) before foo
1056    * void (* <>)(...) before <> which is the inner type */
1057 static struct cp_ctype *parse_function(struct parser *P, struct cp_ctype *ct,

```

```

1057         struct token *name, struct parser *asmname)
1058     {
1059         /* We have a function pointer or a function. The usr table will
1060          * get replaced by the canonical one (if there is one) in
1061          * find_canonical_usr after all the arguments and returns have
1062          * been parsed. */
1063         struct token tok;
1064         struct cp_ctype *ret = ct;
1065
1066         cp_push_ctype(ct);
1067
1068         memset(ct, 0, sizeof(*ct));
1069         ct->base_size = sizeof(void (*)());
1070         ct->align_mask = min(FUNCTION_ALIGN_MASK, P->align_mask);
1071         ct->type = FUNCTION_TYPE;
1072         ct->is_defined = 1;
1073
1074         if (name->type == TOK_NIL) {
1075             for (;;) {
1076                 require_token(P, &tok);
1077
1078                 if (tok.type == TOK_STAR) {
1079                     if (ct->type == FUNCTION_TYPE) {
1080                         ct->type = FUNCTION_PTR_TYPE;
1081                     } else {
1082                         increase_ptr_deref_level(P, ct);
1083                     }
1084                 } else if (parse_attribute(P, &tok, ct, asmname)) {
1085                     /* parse_attribute sets the appropriate fields */
1086                 } else {
1087                     /* call parse_argument to handle the inner
1088                      * contents e.g. the <> in "void (* <>)"
1089                      * (...)". Note that the inner contents can
1090                      * itself be a function, a function ptr,
1091                      * array, etc (e.g. "void (*signal(int sig,
1092                      * void (*func)(int)))(int) "). */
1093                     cp_error("TODO: inner function not supported for now.");
1094                     put_back(P);
1095                     ct = parse_argument2(P, ct, name, asmname);
1096                     break;
1097                 }
1098             }
1099
1100             check_token(P, TOK_CLOSE_PAREN, NULL,
1101                 "unexpected token in function on line %d", P->line);
1102             check_token(P, TOK_OPEN_PAREN, NULL,
1103                 "unexpected token in function on line %d", P->line);
1104         }
1105
1106         parse_function_arguments(P, ct);
1107
1108         /*@TODO support for inner function 24.11 2013 (houqp)*/
1109         /* if we have an inner function then set the outer function ptr as its
1110          * return type and return the inner function
1111          * e.g. for void (* <signal(int, void (*)(int))> )(int) inner is
1112          * surrounded by <>, return type is void (*)(int) */
1113
1114         return ret;
1115     }
1116
1117     static struct cp_ctype *parse_argument2(struct parser *P, struct cp_ctype *ct,
1118         struct token *name, struct parser *asmname)
1119     {
1120         struct token tok;
1121
1122         for (;;) {
1123             if (!next_token(P, &tok)) {
1124                 /* we've reached the end of the string */
1125                 break;
1126             } else if (tok.type == TOK_STAR) {
1127                 increase_ptr_deref_level(P, ct);
1128
1129                 /* __declspec(align(#)) may come before the type in a
1130                  * member */
1131                 if (!ct->is_packed) {
1132                     ct->align_mask = max(min(PTR_ALIGN_MASK, P->align_mask),

```

```

1133         ct->align_mask);
1134     }
1135 } else if (tok.type == TOK_REFERENCE) {
1136     cp_error("NYI: c++ reference types");
1137     return 0;
1138 } else if (parse_attribute(P, &tok, ct, asmname)) {
1139     /* parse_attribute has filled out appropriate fields in type */
1140
1141 } else if (tok.type == TOK_OPEN_PAREN) {
1142     ct = parse_function(P, ct, name, asmname);
1143 } else if (tok.type == TOK_OPEN_SQUARE) {
1144     /* array */
1145     if (ct->pointers == POINTER_MAX) {
1146         cp_error("maximum number of pointer derefs "
1147                 "reached - use a struct to break up "
1148                 "the pointers");
1149     }
1150     ct->is_array = 1;
1151     ct->pointers++;
1152     ct->const_mask <<= 1;
1153     require_token(P, &tok);
1154
1155     if (ct->pointers == 1 && !ct->is_defined) {
1156         cp_error("array of undefined type on line %d",
1157                 P->line);
1158     }
1159
1160     if (ct->is_variable_struct || ct->is_variable_array) {
1161         cp_error("can't have an array of a variably "
1162                 "sized type on line %d", P->line);
1163     }
1164
1165     if (tok.type == TOK_QUESTION) {
1166         ct->is_variable_array = 1;
1167         ct->variable_increment = (ct->pointers > 1) ?
1168             sizeof(void*) : ct->base_size;
1169         check_token(P, TOK_CLOSE_SQUARE, "",
1170                 "invalid character in array on line %d",
1171                 P->line);
1172     }
1173     } else if (tok.type == TOK_CLOSE_SQUARE) {
1174         ct->array_size = 0;
1175
1176     } else if (tok.type == TOK_TOKEN && IS_RESTRICT(tok)) {
1177         /* odd gcc extension foo[__restrict] for arguments */
1178         ct->array_size = 0;
1179         check_token(P, TOK_CLOSE_SQUARE, "",
1180                 "invalid character in array on line %d",
1181                 P->line);
1182     } else {
1183         int64_t asize;
1184         put_back(P);
1185         asize = calculate_constant(P);
1186         if (asize < 0) {
1187             cp_error("array size can not be "
1188                     "negative on line %d", P->line);
1189             return 0;
1190         }
1191         ct->array_size = (size_t) asize;
1192         check_token(P, TOK_CLOSE_SQUARE, "",
1193                 "invalid character in array on line %d",
1194                 P->line);
1195     }
1196
1197 } else if (tok.type == TOK_COLON) {
1198     int64_t bsize = calculate_constant(P);
1199
1200     if (ct->pointers || bsize < 0
1201         || bsize > max_bitfield_size(ct->type)) {
1202         cp_error("invalid bitfield on line %d", P->line);
1203     }
1204
1205     ct->is_bitfield = 1;
1206     ct->bit_size = (unsigned) bsize;
1207
1208 } else if (tok.type != TOK_TOKEN) {

```

```

1209         /* we've reached the end of the declaration */
1210         put_back(P);
1211         break;
1212
1213     } else if (IS_CONST(tok)) {
1214         ct->const_mask |= 1;
1215
1216     } else if (IS_VOLATILE(tok) || IS_RESTRICT(tok)) {
1217         /* ignored for now */
1218
1219     } else {
1220         *name = tok;
1221     }
1222 }
1223
1224 return ct;
1225 }
1226
1227
1228
1229 /* parses after the main base type of a typedef, function argument or
1230 * struct/union member
1231 * eg for const void* bar[3] the base type is void with the subtype so far of
1232 * const, this parses the "" bar[3]" and updates the type argument
1233 *
1234 * type must be as filled out by parse_type
1235 *
1236 * pushes the updated user value on the top of the stack
1237 */
1238 void parse_argument(struct parser *P, struct cp_ctype *ct, struct token *pname,
1239                     struct parser *asmname)
1240 {
1241     struct token tok, name;
1242
1243     memset(&name, 0, sizeof(name));
1244     parse_argument2(P, ct, &name, asmname);
1245
1246     for (;;) {
1247         if (!next_token(P, &tok)) {
1248             break;
1249         } else if (parse_attribute(P, &tok, ct, asmname)) {
1250             /* parse_attribute sets the appropriate fields */
1251         } else {
1252             put_back(P);
1253             break;
1254         }
1255     }
1256
1257     if (pname) {
1258         *pname = name;
1259     }
1260 }
1261
1262 static void parse_typedef(struct parser *P)
1263 {
1264     struct token tok;
1265     struct cp_ctype base_type;
1266     char typedef_name[MAX_TYPE_NAME_LEN];
1267
1268     parse_type(P, &base_type);
1269
1270     for (;;) {
1271         struct cp_ctype arg_type = base_type;
1272         struct token name;
1273
1274         memset(&name, 0, sizeof(name));
1275
1276         parse_argument(P, &arg_type, &name, NULL);
1277
1278         if (!name.size) {
1279             cp_error("Can't have a typedef without a name on line %d",
1280                     P->line);
1281         } else if (arg_type.is_variable_array) {
1282             cp_error("Can't typedef a variable length array on line %d",
1283                     P->line);
1284         }

```

```

1285     memset(typedef_name, 0, sizeof(typedef_name));
1286     strncpy(typedef_name, name.str, name.size);
1287     /* link typedef name with ctype for parser */
1288     cp_ctype_reg_type(typedef_name, &arg_type);
1289
1290
1291     require_token(P, &tok);
1292
1293     if (tok.type == TOK_SEMICOLON) {
1294         break;
1295     } else if (tok.type != TOK_COMMA) {
1296         cp_error("Unexpected character in typedef on line %d",
1297             P->line);
1298     }
1299 }
1300 }
1301
1302 #define END 0
1303 #define PRAGMA_POP 1
1304
1305 static int parse_root(struct parser *P)
1306 {
1307     struct token tok;
1308
1309     while (next_token(P, &tok)) {
1310         /* we can have:
1311          * struct definition
1312          * enum definition
1313          * union definition
1314          * struct/enum/union declaration
1315          * typedef
1316          * function declaration
1317          * pragma pack
1318          */
1319
1320         if (tok.type == TOK_SEMICOLON) {
1321             /* empty semicolon in root continue on */
1322
1323         } else if (tok.type == TOK_POUND) {
1324
1325             check_token(P, TOK_TOKEN, "pragma",
1326                 "unexpected pre processor directive on line %d",
1327                 P->line);
1328             check_token(P, TOK_TOKEN, "pack",
1329                 "unexpected pre processor directive on line %d",
1330                 P->line);
1331             check_token(P, TOK_OPEN_PAREN, "",
1332                 "invalid pack directive on line %d",
1333                 P->line);
1334             require_token(P, &tok);
1335
1336             if (tok.type == TOK_NUMBER) {
1337                 if (tok.integer != 1 && tok.integer != 2
1338                     && tok.integer != 4
1339                     && tok.integer != 8
1340                     && tok.integer != 16) {
1341                     cp_error("pack directive with invalid "
1342                         "pack size on line %d",
1343                         P->line);
1344                     return 0;
1345                 }
1346
1347                 P->align_mask = (unsigned) (tok.integer - 1);
1348                 check_token(P, TOK_CLOSE_PAREN, "",
1349                     "invalid pack directive on line %d",
1350                     P->line);
1351             } else if (tok.type == TOK_TOKEN && IS_LITERAL(tok, "push")) {
1352                 /*int line = P->line;*/
1353                 unsigned previous_alignment = P->align_mask;
1354
1355                 check_token(P, TOK_CLOSE_PAREN, "",
1356                     "invalid pack directive on line %d",
1357                     P->line);
1358
1359                 if (parse_root(P) != PRAGMA_POP) {
1360                     cp_error("reached end of string "

```

```

1361         "without a pragma pop to "
1362         "match the push on line %d",
1363         P->line);
1364     return 0;
1365 }
1366
1367     P->align_mask = previous_alignment;
1368
1369 } else if (tok.type == TOK_TOKEN && IS_LITERAL(tok, "pop")) {
1370     check_token(P, TOK_CLOSE_PAREN, "",
1371         "invalid pack directive on line %d",
1372         P->line);
1373     return PRAGMA_POP;
1374 } else {
1375     cp_error("invalid pack directive on line %d",
1376         P->line);
1377     return 0;
1378 }
1379 } else if (tok.type != TOK_TOKEN) {
1380     cp_error("unexpected character on line %d", P->line);
1381     return 0;
1382 } else if (IS_LITERAL(tok, "__extension__")) {
1383     /* ignore */
1384     continue;
1385 } else if (IS_LITERAL(tok, "extern")) {
1386     /* ignore extern as data and functions can only be
1387     * extern */
1388     continue;
1389 } else if (IS_LITERAL(tok, "typedef")) {
1390     parse_typedef(P);
1391 } else if (IS_LITERAL(tok, "static")) {
1392     /*@TODO we haven't tested static so far */
1393     cp_error("TODO: support static keyword.\n");
1394 } else {
1395     /* type declaration, type definition, or function
1396     * declaration */
1397     struct cp_ctype type;
1398     struct token name;
1399     struct parser asmname;
1400
1401     memset(&name, 0, sizeof(name));
1402     memset(&asmname, 0, sizeof(asmname));
1403
1404     put_back(P);
1405     parse_type(P, &type);
1406
1407     for (;;) {
1408         parse_argument(P, &type, &name, &asmname);
1409
1410         if (name.size) {
1411             /* global/function declaration */
1412             cp_symbol_build_func(&type, name.str, name.size);
1413             /* @TODO asmname is not used for now
1414             * since we are not supporting __asm__
1415             * as this point.
1416             * might need to bind it with function
1417             * name later. */
1418         } else {
1419             /* type declaration/definition -
1420             * already been processed */
1421         }
1422         require_token(P, &tok);
1423
1424         if (tok.type == TOK_SEMICOLON) {
1425             break;
1426         } else if (tok.type != TOK_COMMA) {
1427             cp_error("missing semicolon on line %d",
1428                 P->line);
1429         }
1430     }
1431 }
1432 }
1433
1434 return END;
1435 }
1436

```

```

1437 static int64_t calculate_constant2(struct parser *P, struct token *tok);
1438
1439 /* () */
1440 static int64_t calculate_constant1(struct parser *P, struct token *tok)
1441 {
1442     int64_t ret;
1443
1444     if (tok->type == TOK_NUMBER) {
1445         ret = tok->integer;
1446         next_token(P, tok);
1447         return ret;
1448
1449     } else if (tok->type == TOK_TOKEN) {
1450         /* look up name in constants table */
1451         cp_error("TODO: support name lookup in constant table\n");
1452         next_token(P, tok);
1453         return ret;
1454
1455     } else if (tok->type == TOK_OPEN_PAREN) {
1456         struct parser before_cast = *P;
1457         cp_error("TODO: handle open parent token in constant1\n");
1458         *P = before_cast;
1459         ret = calculate_constant(P);
1460
1461         require_token(P, tok);
1462         if (tok->type != TOK_CLOSE_PAREN) {
1463             cp_error("error whilst parsing constant at line %d",
1464                     P->line);
1465         }
1466
1467         next_token(P, tok);
1468         return ret;
1469     } else {
1470         cp_error("unexpected token whilst parsing constant at line %d",
1471                 P->line);
1472         return 0;
1473     }
1474 }
1475
1476 /* ! and ~, unary + and -, and sizeof */
1477 static int64_t calculate_constant2(struct parser *P, struct token *tok)
1478 {
1479     if (tok->type == TOK_LOGICAL_NOT) {
1480         require_token(P, tok);
1481         return !calculate_constant2(P, tok);
1482
1483     } else if (tok->type == TOK_BITWISE_NOT) {
1484         require_token(P, tok);
1485         return ~calculate_constant2(P, tok);
1486
1487     } else if (tok->type == TOK_PLUS) {
1488         require_token(P, tok);
1489         return calculate_constant2(P, tok);
1490
1491     } else if (tok->type == TOK_MINUS) {
1492         require_token(P, tok);
1493         return -calculate_constant2(P, tok);
1494
1495     } else if (tok->type == TOK_TOKEN &&
1496               (IS_LITERAL(*tok, "sizeof")
1497                || IS_LITERAL(*tok, "alignof")
1498                || IS_LITERAL(*tok, "__alignof__")
1499                || IS_LITERAL(*tok, "__alignof"))) {
1500         cp_error("TODO: support sizeof\n");
1501         bool issize = IS_LITERAL(*tok, "sizeof");
1502         struct cp_ctype type;
1503
1504         require_token(P, tok);
1505         if (tok->type != TOK_OPEN_PAREN) {
1506             cp_error("invalid sizeof at line %d", P->line);
1507         }
1508
1509         parse_type(P, &type);
1510         parse_argument(P, &type, NULL, NULL);
1511
1512         require_token(P, tok);

```

```

1513     if (tok->type != TOK_CLOSE_PAREN) {
1514         cp\_error("invalid sizeof at line %d", P->line);
1515     }
1516
1517     next\_token(P, tok);
1518
1519     return issize ? ctype\_size(&type) : type.align_mask + 1;
1520
1521 } else {
1522     return calculate\_constant1(P, tok);
1523 }
1524 }
1525
1526 /* binary * / and % (left associative) */
1527 static int64_t calculate\_constant3(struct parser *P, struct token *tok)
1528 {
1529     int64_t left = calculate\_constant2(P, tok);
1530
1531     for (;;) {
1532         if (tok->type == TOK_MULTIPLY) {
1533             require\_token(P, tok);
1534             left *= calculate\_constant2(P, tok);
1535
1536         } else if (tok->type == TOK_DIVIDE) {
1537             require\_token(P, tok);
1538             left /= calculate\_constant2(P, tok);
1539
1540         } else if (tok->type == TOK_MODULUS) {
1541             require\_token(P, tok);
1542             left %= calculate\_constant2(P, tok);
1543
1544         } else {
1545             return left;
1546         }
1547     }
1548 }
1549
1550 /* binary + and - (left associative) */
1551 static int64_t calculate\_constant4(struct parser *P, struct token *tok)
1552 {
1553     int64_t left = calculate\_constant3(P, tok);
1554
1555     for (;;) {
1556         if (tok->type == TOK_PLUS) {
1557             require\_token(P, tok);
1558             left += calculate\_constant3(P, tok);
1559
1560         } else if (tok->type == TOK_MINUS) {
1561             require\_token(P, tok);
1562             left -= calculate\_constant3(P, tok);
1563
1564         } else {
1565             return left;
1566         }
1567     }
1568 }
1569
1570 /* binary << and >> (left associative) */
1571 static int64_t calculate\_constant5(struct parser *P, struct token *tok)
1572 {
1573     int64_t left = calculate\_constant4(P, tok);
1574
1575     for (;;) {
1576         if (tok->type == TOK_LEFT_SHIFT) {
1577             require\_token(P, tok);
1578             left <<= calculate\_constant4(P, tok);
1579
1580         } else if (tok->type == TOK_RIGHT_SHIFT) {
1581             require\_token(P, tok);
1582             left >>= calculate\_constant4(P, tok);
1583
1584         } else {
1585             return left;
1586         }
1587     }
1588 }

```

```

1589
1590 /* binary <, <=, >, and >= (left associative) */
1591 static int64_t calculate_constant6(struct parser *P, struct token *tok)
1592 {
1593     int64_t left = calculate_constant5(P, tok);
1594
1595     for (;;) {
1596         if (tok->type == TOK_LESS) {
1597             require_token(P, tok);
1598             left = (left < calculate_constant5(P, tok));
1599
1600         } else if (tok->type == TOK_LESS_EQUAL) {
1601             require_token(P, tok);
1602             left = (left <= calculate_constant5(P, tok));
1603
1604         } else if (tok->type == TOK_GREATER) {
1605             require_token(P, tok);
1606             left = (left > calculate_constant5(P, tok));
1607
1608         } else if (tok->type == TOK_GREATER_EQUAL) {
1609             require_token(P, tok);
1610             left = (left >= calculate_constant5(P, tok));
1611
1612         } else {
1613             return left;
1614         }
1615     }
1616 }
1617
1618 /* binary ==, != (left associative) */
1619 static int64_t calculate_constant7(struct parser *P, struct token *tok)
1620 {
1621     int64_t left = calculate_constant6(P, tok);
1622
1623     for (;;) {
1624         if (tok->type == TOK_EQUAL) {
1625             require_token(P, tok);
1626             left = (left == calculate_constant6(P, tok));
1627
1628         } else if (tok->type == TOK_NOT_EQUAL) {
1629             require_token(P, tok);
1630             left = (left != calculate_constant6(P, tok));
1631
1632         } else {
1633             return left;
1634         }
1635     }
1636 }
1637
1638 /* binary & (left associative) */
1639 static int64_t calculate_constant8(struct parser *P, struct token *tok)
1640 {
1641     int64_t left = calculate_constant7(P, tok);
1642
1643     for (;;) {
1644         if (tok->type == TOK_BITWISE_AND) {
1645             require_token(P, tok);
1646             left = (left & calculate_constant7(P, tok));
1647
1648         } else {
1649             return left;
1650         }
1651     }
1652 }
1653
1654 /* binary ^ (left associative) */
1655 static int64_t calculate_constant9(struct parser *P, struct token *tok)
1656 {
1657     int64_t left = calculate_constant8(P, tok);
1658
1659     for (;;) {
1660         if (tok->type == TOK_BITWISE_XOR) {
1661             require_token(P, tok);
1662             left = (left ^ calculate_constant8(P, tok));
1663
1664         } else {

```

```

1665         return left;
1666     }
1667 }
1668 }
1669
1670 /* binary | (left associative) */
1671 static int64_t calculate_constant10(struct parser *P, struct token *tok)
1672 {
1673     int64_t left = calculate_constant9(P, tok);
1674
1675     for (;;) {
1676         if (tok->type == TOK_BITWISE_OR) {
1677             require_token(P, tok);
1678             left = (left | calculate_constant9(P, tok));
1679
1680         } else {
1681             return left;
1682         }
1683     }
1684 }
1685
1686 /* binary && (left associative) */
1687 static int64_t calculate_constant11(struct parser *P, struct token *tok)
1688 {
1689     int64_t left = calculate_constant10(P, tok);
1690
1691     for (;;) {
1692         if (tok->type == TOK_LOGICAL_AND) {
1693             require_token(P, tok);
1694             left = (left && calculate_constant10(P, tok));
1695
1696         } else {
1697             return left;
1698         }
1699     }
1700 }
1701
1702 /* binary || (left associative) */
1703 static int64_t calculate_constant12(struct parser *P, struct token *tok)
1704 {
1705     int64_t left = calculate_constant11(P, tok);
1706
1707     for (;;) {
1708         if (tok->type == TOK_LOGICAL_OR) {
1709             require_token(P, tok);
1710             left = (left || calculate_constant11(P, tok));
1711
1712         } else {
1713             return left;
1714         }
1715     }
1716 }
1717
1718 /* ternary ?: (right associative) */
1719 static int64_t calculate_constant13(struct parser *P, struct token *tok)
1720 {
1721     int64_t left = calculate_constant12(P, tok);
1722
1723     if (tok->type == TOK_QUESTION) {
1724         int64_t middle, right;
1725         require_token(P, tok);
1726         middle = calculate_constant13(P, tok);
1727         if (tok->type != TOK_COLON) {
1728             cp_error("invalid ternery (? :) in constant on line %d",
1729                 P->line);
1730         }
1731         require_token(P, tok);
1732         right = calculate_constant13(P, tok);
1733         return left ? middle : right;
1734
1735     } else {
1736         return left;
1737     }
1738 }
1739
1740 int64_t calculate_constant(struct parser* P)

```

```

1741 {
1742     struct token tok;
1743     int64_t ret;
1744     require\_token(P, &tok);
1745     ret = calculate\_constant13(P, &tok);
1746
1747     if (tok.type != TOK_NIL) {
1748         put\_back(P);
1749     }
1750
1751     return ret;
1752 }
1753
1754 void cp\_init\_parser(struct parser *P, const char *s)
1755 {
1756     memset(P, 0, sizeof(struct parser));
1757     P->line = 1;
1758     P->prev = P->next = s;
1759     P->align_mask = DEFAULT\_ALIGN\_MASK;
1760 }
1761
1762 /* used for ffi.cdef */
1763 int ffi\_parse\_cdef(const char *s)
1764 {
1765     struct parser P;
1766
1767     cp\_init\_parser(&P, s);
1768     if (parse\_root(&P) == PRAGMA\_POP) {
1769         cp\_error("pragma pop without an associated push on line %d",
1770                 P.line);
1771     }
1772
1773     return 0;
1774 }
1775
1776 /* used for ffi.new */
1777 void ffi\_parse\_new(const char *s, struct cp\_ctype *ct)
1778 {
1779     struct parser P;
1780
1781     cp\_init\_parser(&P, s);
1782     parse\_type(&P, ct);
1783     parse\_argument(&P, ct, NULL, NULL);
1784     cp\_update\_csym\_in\_ctype(ct);
1785 }
1786
1787 int ffi\_lookup\_csymbol\_id\_by\_name(const char *s)
1788 {
1789     struct parser P;
1790     struct cp\_ctype ct;
1791
1792     cp\_init\_parser(&P, s);
1793     parse\_type(&P, &ct);
1794     cp\_update\_csym\_in\_ctype(&ct);
1795     return ct.ffi_cs_id;
1796 }
1797
1798 void ffi\_cparser\_init(void)
1799 {
1800     cp\_ctype\_init();
1801 }
1802
1803 void ffi\_cparser\_free(void)
1804 {
1805     cp\_ctype\_free();
1806 }

```

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

userspace/ffi/ - ktap

- [cparser.c](#)
- [ctype.c](#)

[One Level Up](#)

[Top Level](#)

userspace/ffi/ctype.c - ktap

Global variables defined

- [csym_state](#)
- [cte_arr](#)
- [cte_arr_size](#)
- [cte_nr](#)
- [cts](#)

Data types defined

- [cp_ctype_entry](#)
- [cp_ctype_entry](#)
- [cp_ctype_stack](#)
- [ctype_stack](#)

Functions defined

- [__cp_symbol_dump_func](#)
- [__cp_symbol_dump_struct](#)
- [cp_ctype_dump_stack](#)
- [cp_ctype_free](#)
- [cp_ctype_init](#)
- [cp_ctype_reg_csymbols](#)
- [cp_ctype_reg_type](#)
- [cp_id_to_csymbols](#)
- [cp_push_ctype](#)
- [cp_push_ctype_with_name](#)
- [cp_set_defined](#)
- [cp_symbol_build_fake_record](#)
- [cp_symbol_build_func](#)
- [cp_symbol_build_pointer](#)
- [cp_symbol_build_record](#)
- [cp_symbol_dump_func](#)
- [cp_symbol_dump_struct](#)
- [cp_update_csymbols_in_ctype](#)

- [ct_set_type](#)
- [ctype_get_csym_state](#)
- [ctype_lookup_builtin_type](#)
- [ctype_lookup_csymbol_id](#)
- [ctype_lookup_type](#)
- [ctype_reg_table_grow](#)
- [ctype_size](#)
- [ctype_stack_free_space](#)
- [ctype_stack_grow](#)
- [ctype_stack_reset](#)
- [ctype_stack_top](#)
- [init_builtin_type](#)

Macros defined

- [DEFAULT_CTYPE_ARR_SIZE](#)
- [DEFAULT_STACK_SIZE](#)
- [DEFAULT_SYM_ARR_SIZE](#)
- [MAX_STACK_SIZE](#)
- [cs_arr](#)
- [cs_arr_size](#)
- [cs_nr](#)
- [ct_stack](#)
- [ct_stack_ct](#)

Source code

```

1 #include "../include/ktap_types.h"
2 #include "../include/ktap_opcodes.h"
3 #include "../ktapc.h"
4 #include "../cparser.h"
5
6
7 /* for ktap vm */
8 cp_csymbol_state csym_state;
9
10 #define cs_nr (csym_state.cs_nr)
11 #define cs_arr_size (csym_state.cs_arr_size)
12 #define cs_arr (csym_state.cs_arr)
13
14 csymbol *cp_id_to_csym(int id)
15 {
16     return &cs_arr[id];
17 }
18
19
20 typedef struct cp_ctype_entry {
21     char name[MAX_TYPE_NAME_LEN];

```

```

22     struct cp\_ctype ct;
23 } cp\_ctype\_entry;
24
25 #define DEFAULT\_CTYPE\_ARR\_SIZE 100
26 static int cte\_nr;
27 static int cte\_arr\_size;
28 static cp\_ctype\_entry *cte\_arr;
29
30
31 /* stack to help maintain state during parsing */
32 typedef struct cp\_ctype\_stack {
33     int size;
34     int top;
35     cp\_ctype\_entry *stack;
36 } ctype\_stack;
37
38
39 static ctype\_stack cts;
40
41 #define ct\_stack\(id\) (&(cts.stack[id]))
42 #define ct\_stack\_ct\(id\) (&(cts.stack[id].ct))
43
44
45
46 csymbol\_id cp\_ctype\_reg\_csymbol(csymbol *cs);
47 csymbol\_id ctype\_lookup\_csymbol\_id(const char *name);
48
49
50 size_t ctype\_size(const struct cp\_ctype *ct)
51 {
52     if (ct->pointers - ct->is_array) {
53         return sizeof(void*) * (ct->is_array ? ct->array_size : 1);
54     }
55     else if (!ct->is_defined || ct->type == VOID_TYPE) {
56         cp\_error("can't calculate size of an undefined type");
57         return 0;
58     }
59     else if (ct->variable_size_known) {
60         assert(ct->is_variable_struct && !ct->is_array);
61         return ct->base_size + ct->variable_increment;
62     }
63     else if (ct->is_variable_array || ct->is_variable_struct) {
64         cp\_error("internal error: calc size of variable type with "
65             "unknown size");
66         return 0;
67     }
68     else {
69         return ct->base_size * (ct->is_array ? ct->array_size : 1);
70     }
71 }
72
73 #define MAX\_STACK\_SIZE 100
74 int ctype\_stack\_grow(int size)
75 {
76     struct cp\_ctype\_entry *new_st;
77
78     assert(cts.size + size < MAX\_STACK\_SIZE);
79
80     new_st = realloc(cts.stack, (cts.size+size)*sizeof(cp\_ctype\_entry));
81     if (new_st)
82         cts.stack = new_st;
83     else
84         return -1;
85
86     cts.size += size;
87
88     return size;
89 }
90
91 int ctype\_stack\_free\_space()
92 {
93     return cts.size - cts.top;
94 }
95
96 int ctype\_stack\_top()
97 {
98     return cts.top;
99 }

```

```

98 void ctype_stack_reset(int top)
99 {
100     cts.top = top;
101 }
102
103 /* This function should be called before you would fetch
104 * ffi_cs_id from ctype */
105 void cp_update_csym_in_ctype(struct cp_ctype *ct)
106 {
107     int i;
108     struct cp_ctype *nct;
109
110     assert(ct->ffi_cs_id >= 0);
111     /* we have to check pointer here because cparser does type lookup by name
112     * before parsing '*', and for pointers, ct will always be the
113     * original type */
114     if (ct->pointers) {
115         for (i = 0; i < cte_nr; i++) {
116             nct = &(cte_arr[i].ct);
117             if (nct->type == ct->type &&
118                 ct->ffi_base_cs_id == nct->ffi_base_cs_id &&
119                 nct->pointers == ct->pointers) {
120                 break;
121             }
122         }
123
124         if (i == cte_nr) {
125             /* pointer type not found
126             * create a new pointer symbol for this type */
127             /* associate ctype with new csymbol */
128             ct->ffi_cs_id = cp_symbol_build_pointer(ct);
129             /* register wit new pointer name */
130             cp_ctype_reg_type(csym_name(ct_ffi_cs(ct)), ct);
131         } else {
132             /* pointer type already registered, reinstantiate ct */
133             ct->ffi_cs_id = cte_arr[i].ct.ffi_cs_id;
134         }
135     }
136 }
137
138 /* push ctype to stack, create new csymbol if needed */
139 void cp_push_ctype_with_name(struct cp_ctype *ct, const char *name, int nlen)
140 {
141     if (ctype_stack_free_space() < 1)
142         ctype_stack_grow(4);
143
144     cp_update_csym_in_ctype(ct);
145     memset(ct_stack(cts.top), 0, sizeof(cp_ctype_entry));
146     ct_stack(cts.top)->ct = *ct;
147     if (name)
148         strncpy(ct_stack(cts.top)->name, name, nlen);
149     cts.top++;
150 }
151
152 void cp_push_ctype(struct cp_ctype *ct)
153 {
154     cp_push_ctype_with_name(ct, NULL, 0);
155 }
156
157 void cp_set_defined(struct cp_ctype *ct)
158 {
159     ct->is_defined = 1;
160
161     /* @TODO: update ctypes and cdatas that were created before the
162     * definition came in */
163 }
164
165 void cp_ctype_dump_stack()
166 {
167     int i;
168     struct cp_ctype *ct;
169
170     printf("-----\n");
171     printf("start of ctype stack (%d) dump: \n", cts.top);
172     for (i = 0; i < cts.top; i++) {
173         ct = ct_stack_ct(i);

```

```

174     printf("[%d] -> cp_ctype: %d, sym_type: %d, pointer: %d "
175           "symbol_id: %d, name: %s\n",
176           i, ct->type,
177           csym_type(ct ffi_cs(ct)), ct->pointers, ct->ffi_cs_id,
178           ct_stack(i)->name);
179 }
180 }
181
182 int ctype_reg_table_grow()
183 {
184     cp_ctype_entry *new_arr;
185
186     new_arr = realloc(cte_arr, sizeof(cp_ctype_entry)*cte_arr_size*2);
187     if (!new_arr)
188         cp_error("failed to allocate memory for ctype array\n");
189
190     cte_arr_size = cte_arr_size * 2;
191     return 0;
192 }
193
194 /* return index in csymbol array */
195 csymbol_id cp_ctype_reg_csymbols(csymbol *cs)
196 {
197     if (cs_nr >= cs_arr_size) {
198         cs_arr_size *= 2;
199         cs_arr = realloc(cs_arr, cs_arr_size*sizeof(csymbol));
200         if (!cs_arr)
201             cp_error("failed to extend csymbol array!\n");
202     }
203
204     cs_arr[cs_nr] = *cs;
205     cs_nr++;
206
207     return cs_nr-1;
208 }
209
210 /* start csymbol reg table */
211 csymbol_id ctype_lookup_csymbols(const char *name)
212 {
213     int i;
214     csymbol *ct;
215
216     for (i = 0; i < cs_nr; i++) {
217         ct = cp_id_to_csym(i);
218         if (!strcmp(name, ct->name))
219             return i;
220     }
221
222     return -1;
223 }
224
225 void __cp_symbol_dump_struct(csymbol *cs)
226 {
227     int i;
228     csymbol *ncs;
229     csymbol_struct *stcs = csym_struct(cs);
230
231     printf("=== [%s] definition =====\n", csym_name(cs));
232     for (i = 0; i < stcs->memb_nr; i++) {
233         printf("\t(%d) ", i);
234         printf("csym_id: %d, ", stcs->members[i].id);
235         ncs = cp_id_to_csym(stcs->members[i].id);
236         printf("name: %s, ffi_ctype: %d, %s\n",
237               stcs->members[i].name, ncs->type, csym_name(ncs));
238     }
239 }
240
241 void cp_symbol_dump_struct(int id)
242 {
243     __cp_symbol_dump_struct(cp_id_to_csym(id));
244 }
245
246 int cp_symbol_build_record(const char *stname, int type, int start_top)
247 {
248     int i, id, memb_size;
249     cp_ctype_entry *cte;

```

```

250     csymbol nst;
251     struct\_member *st_membs;
252     csymbol\_struct *stcs;
253     struct\_cp\_ctype *ct;
254
255     if (cts.top <= start_top || !stname ||
256         (type != STRUCT_TYPE && type != UNION_TYPE)) {
257         cp\_error("invalid struct/union definition.\n");
258     }
259
260     id = ctype\_lookup\_csymbol\_id(stname);
261     if (id >= 0) {
262         assert(cp\_id\_to\_csym(id)->type == FFI_STRUCT ||
263             cp\_id\_to\_csym(id)->type == FFI_UNION);
264         assert(csym\_struct(cp\_id\_to\_csym(id))->memb_nr == -1);
265     }
266
267     memb_size = cts.top - start_top;
268     st_membs = malloc(memb_size*sizeof(struct\_member));
269     if (!st_membs)
270         cp\_error("failed to allocate memory for struct members.\n");
271     memset(st_membs, 0, memb_size*sizeof(struct\_member));
272
273     if (type == STRUCT_TYPE)
274         nst.type = FFI_STRUCT;
275     else
276         nst.type = FFI_UNION;
277     strcpy(nst.name, stname);
278
279     stcs = csym\_struct(&nst);
280     stcs->align = 0;
281     stcs->memb_nr = memb_size;
282     stcs->members = st_membs;
283
284     for (i = 0; i < memb_size; i++) {
285         cte = ct\_stack(i + start_top);
286         if (cte->name)
287             strcpy(st_membs[i].name, cte->name);
288         ct = ct\_stack\_ct(i + start_top);
289         st_membs[i].id = ct->ffi_cs_id;
290         if (!ct->is_array)
291             st_membs[i].len = -1;
292         else
293             st_membs[i].len = ct->array_size;
294     }
295
296     if (id < 0)
297         id = cp\_ctype\_reg\_csymbol(&nst);
298     else
299         cs\_arr[id] = nst;
300
301     ctype\_stack\_reset(start_top);
302
303     return id;
304 }
305
306 int cp\_symbol\_build\_fake\_record(const char *stname, int type)
307 {
308     int id;
309     csymbol nst;
310     csymbol\_struct *stcs;
311
312     if (!stname || (type != STRUCT_TYPE && type != UNION_TYPE)) {
313         cp\_error("invalid fake struct/union definition.\n");
314     }
315
316     id = ctype\_lookup\_csymbol\_id(stname);
317     if (id >= 0)
318         return id;
319
320     if (type == STRUCT_TYPE)
321         nst.type = FFI_STRUCT;
322     else
323         nst.type = FFI_UNION;
324     strcpy(nst.name, stname);
325

```

```

326     stcs = csym\_struct(&nst);
327     stcs->align = 0;
328     stcs->memb_nr = -1;
329     stcs->members = NULL;
330
331     id = cp\_ctype\_reg\_csymbol(&nst);
332
333     return id;
334 }
335
336
337 /* build pointer symbol from given csymbol */
338 int cp\_symbol\_build\_pointer(struct cp\_ctype *ct)
339 {
340     int id, ret;
341     csymbol ncspt;
342     csymbol *ref_cs = ct\_ffi\_cs(ct);
343
344     /* TODO: Check correctness of multi-level pointer 24.11.2013(unihorn) */
345     memset(&ncspt, 0, sizeof(csymbol));
346     ncspt.type = FFI_PTR;
347     ret = sprintf(ncspt.name, "%s *", csym\_name(ref_cs));
348     assert(ret < MAX\_TYPE\_NAME\_LEN);
349
350     csym\_set\_ptr\_deref\_id(&ncspt, ct->ffi_cs_id);
351     id = cp\_ctype\_reg\_csymbol(&ncspt);
352
353     return id;
354 }
355
356 void \_\_cp\_symbol\_dump\_func(csymbol *cs)
357 {
358     int i;
359     csymbol *ncs;
360     csymbol\_func *fcs = csym\_func(cs);
361
362     printf("=== [%s] function definition =====\n", csym\_name(cs));
363     ncs = cp\_csymf\_ret(fcs);
364     printf("address: %p\n", fcs->addr);
365     printf("return type: \n");
366     printf("\tcsym_id: %d, ffi_ctype: %d, %s\n",
367           fcs->ret_id, ncs->type, csym\_name(ncs));
368     printf("args type (%d): \n", fcs->arg_nr);
369     for (i = 0; i < csymf\_arg\_nr(fcs); i++) {
370         printf("\t (%d) ", i);
371         printf("csym_id: %d, ", fcs->arg_ids[i]);
372         ncs = cp\_csymf\_arg(fcs, i);
373         printf("ffi_ctype: %d, %s\n", ncs->type, csym\_name(ncs));
374     }
375 }
376
377 void cp\_symbol\_dump\_func(int id)
378 {
379     \_\_cp\_symbol\_dump\_func(cp\_id\_to\_csym(id));
380 }
381
382 int cp\_symbol\_build\_func(struct cp\_ctype *type, const char *fname, int fn_size)
383 {
384     int i = 1, arg_nr, id;
385     int *argsym_id_arr;
386     csymbol nfcs;
387     csymbol\_func *fcs;
388
389     if (cts.top == 0 || fn_size < 0 || !fname) {
390         cp\_error("invalid function definition.\n");
391     }
392
393     argsym_id_arr = NULL;
394     memset(&nfcs, 0, sizeof(csymbol));
395     csym\_type(&nfcs) = FFI_FUNC;
396
397     strncpy(csym\_name(&nfcs), fname, fn_size);
398
399     fcs = csym\_func(&nfcs);
400     fcs->has_var_arg = type->has_var_arg;
401     /* Type needed for handling variable args handle */

```

```

402 if (fcs->has_var_arg && !ctype_lookup_type("void *"))
403     cp_symbol_build_pointer(ctype_lookup_type("void"));
404
405 /* Fetch start address of function */
406 fcs->addr = (void *)find_kernel_symbol(csym_name(&nfcs));
407 if (!fcs->addr)
408     cp_error("wrong function address for %s\n", csym_name(&nfcs));
409
410 /* bottom of the stack is return type */
411 fcs->ret_id = ct_stack_ct(0)->ffi_cs_id;
412
413 /* the rest is argument type */
414 if (cts.top == 1) {
415     /* function takes no argument */
416     arg_nr = 0;
417 } else {
418     arg_nr = cts.top - 1;
419     argsym_id_arr = malloc(arg_nr * sizeof(int));
420     if (!argsym_id_arr)
421         cp_error("failed to allocate memory for function args.\n");
422     for (i = 0; i < arg_nr; i++) {
423         argsym_id_arr[i] = ct_stack_ct(i+1)->ffi_cs_id;
424     }
425 }
426 fcs->arg_nr = arg_nr;
427 fcs->arg_ids = argsym_id_arr;
428
429 id = cp_ctype_reg_csymbol(&nfcs);
430
431 /* clear stack since we have consumed all the ctypes */
432 ctype_stack_reset(0);
433
434 return id;
435 }
436
437 struct cp_ctype *cp_ctype_reg_type(char *name, struct cp_ctype *ct)
438 {
439     if (cte_nr >= cte_arr_size)
440         ctype_reg_table_grow();
441
442     memset(cte_arr[cte_nr].name, 0, MAX_TYPE_NAME_LEN);
443     strcpy(cte_arr[cte_nr].name, name);
444
445     cte_arr[cte_nr].ct = *ct;
446     cte_nr++;
447
448     return &(cte_arr[cte_nr-1].ct);
449 }
450
451 #if 0
452 /* TODO: used for size calculation */
453 static ffi_type ffi_int_type(ktap_state *ks, int size, bool sign)
454 {
455     switch(size) {
456     case 1:
457         if (!sign)
458             return FFI_UINT8;
459         else
460             return FFI_INT8;
461     case 2:
462         if (!sign)
463             return FFI_UINT16;
464         else
465             return FFI_INT16;
466     case 4:
467         if (!sign)
468             return FFI_UINT32;
469         else
470             return FFI_INT32;
471     case 8:
472         if (!sign)
473             return FFI_UINT64;
474         else
475             return FFI_INT64;
476     default:
477         kp_error(ks, "Error: Have not support int type of size %d\n", size);

```

```

478     return FFI_UNKNOWN;
479 }
480
481 /* NEVER reach here, silence compiler */
482 return -1;
483 }
484 #endif
485
486
487 static inline void ct_set_type(struct cp_ctype *ct, int type, int is_unsigned)
488 {
489     ct->type = type;
490     ct->is_unsigned = is_unsigned;
491 }
492
493 static void init_builtin_type(struct cp_ctype *ct, ffi_type ftype)
494 {
495     csymbol cs;
496     int cs_id;
497
498     csym_type(&cs) = ftype;
499     strncpy(csym_name(&cs), ffi_type_name(ftype), CSYM_NAME_MAX_LEN);
500     cs_id = cp_ctype_reg_csymbol(&cs);
501
502     memset(ct, 0, sizeof(*ct));
503     ct->ffi_base_cs_id = ct->ffi_cs_id = cs_id;
504     switch (ftype) {
505     case FFI_VOID:         ct_set_type(ct, VOID_TYPE, 0); break;
506     case FFI_UINT8:       ct_set_type(ct, INT8_TYPE, 1); break;
507     case FFI_INT8:        ct_set_type(ct, INT8_TYPE, 0); break;
508     case FFI_UINT16:      ct_set_type(ct, INT16_TYPE, 1); break;
509     case FFI_INT16:       ct_set_type(ct, INT16_TYPE, 0); break;
510     case FFI_UINT32:      ct_set_type(ct, INT32_TYPE, 1); break;
511     case FFI_INT32:       ct_set_type(ct, INT32_TYPE, 0); break;
512     case FFI_UINT64:     ct_set_type(ct, INT64_TYPE, 1); break;
513     case FFI_INT64:      ct_set_type(ct, INT64_TYPE, 0); break;
514     default:              break;
515     }
516     ct->base_size = ffi_type_size(ftype);
517     ct->align_mask = ffi_type_align(ftype) - 1;
518     ct->is_defined = 1;
519 }
520
521 /*
522  * lookup and register builtin C type on demand
523  * You should ensure that the type with name doesn't appear in
524  * csymbol table before calling.
525  */
526 struct cp_ctype *ctype_lookup_builtin_type(char *name)
527 {
528     struct cp_ctype ct;
529
530     if (!strcmp(name, "void", sizeof("void"))) {
531         init_builtin_type(&ct, FFI_VOID);
532         return cp_ctype_reg_type("void", &ct);
533     } else if (!strcmp(name, "int8_t", sizeof("int8_t"))) {
534         init_builtin_type(&ct, FFI_INT8);
535         return cp_ctype_reg_type("int8_t", &ct);
536     } else if (!strcmp(name, "uint8_t", sizeof("uint8_t"))) {
537         init_builtin_type(&ct, FFI_UINT8);
538         return cp_ctype_reg_type("uint8_t", &ct);
539     } else if (!strcmp(name, "int16_t", sizeof("int16_t"))) {
540         init_builtin_type(&ct, FFI_INT16);
541         return cp_ctype_reg_type("int16_t", &ct);
542     } else if (!strcmp(name, "uint16_t", sizeof("uint16_t"))) {
543         init_builtin_type(&ct, FFI_UINT16);
544         return cp_ctype_reg_type("uint16_t", &ct);
545     } else if (!strcmp(name, "int32_t", sizeof("int32_t"))) {
546         init_builtin_type(&ct, FFI_INT32);
547         return cp_ctype_reg_type("int32_t", &ct);
548     } else if (!strcmp(name, "uint32_t", sizeof("uint32_t"))) {
549         init_builtin_type(&ct, FFI_UINT32);
550         return cp_ctype_reg_type("uint32_t", &ct);
551     } else if (!strcmp(name, "int64_t", sizeof("int64_t"))) {
552         init_builtin_type(&ct, FFI_INT64);
553         return cp_ctype_reg_type("int64_t", &ct);

```

```

554 } else if (!strncmp(name, "uint64_t", sizeof("uint64_t"))) {
555     init\_builtin\_type(&ct, FFI_UINT64);
556     return cp\_ctype\_reg\_type("uint64_t", &ct);
557 } else {
558     /* no builtin type matched */
559     return NULL;
560 }
561 }
562
563 /* start ctype reg table */
564 struct cp\_ctype *ctype\_lookup\_type(char *name)
565 {
566     int i;
567     struct cp\_ctype *ct;
568
569     for (i = 0; i < cte\_nr; i++) {
570         ct = &cte\_arr[i].ct;
571         if (!strcmp(name, cte\_arr[i].name))
572             return ct;
573     }
574
575     /* see if it's a builtin C type
576     * return NULL if still no match */
577     return ctype\_lookup\_builtin\_type(name);
578 }
579
580 cp\_csymbol\_state *ctype\_get\_csym\_state(void)
581 {
582     return &csym\_state;
583 }
584
585 #define DEFAULT_STACK_SIZE 20
586 #define DEFAULT_SYM_ARR_SIZE 20
587 int cp\_ctype\_init()
588 {
589     cts.size = DEFAULT\_STACK\_SIZE;
590     cts.top = 0;
591     cts.stack = malloc(sizeof(cp\_ctype\_entry)*DEFAULT\_STACK\_SIZE);
592
593     cs\_nr = 0;
594     cs\_arr\_size = DEFAULT\_SYM\_ARR\_SIZE;
595     cs\_arr = malloc(sizeof(csymbol)*DEFAULT\_SYM\_ARR\_SIZE);
596     memset(cs\_arr, 0, sizeof(csymbol)*DEFAULT\_SYM\_ARR\_SIZE);
597
598     cte\_nr = 0;
599     cte\_arr\_size = DEFAULT\_CTYPE\_ARR\_SIZE;
600     cte\_arr = malloc(sizeof(cp\_ctype\_entry)*DEFAULT\_CTYPE\_ARR\_SIZE);
601
602     return 0;
603 }
604
605 int cp\_ctype\_free()
606 {
607     int i;
608     csymbol *cs;
609
610     if (cts.stack)
611         free(cts.stack);
612
613     if (cs\_arr) {
614         for (i = 0; i < cs\_nr; i++) {
615             cs = cp\_id\_to\_csym(i);
616             if (csym\_type(cs) == FFI_FUNC) {
617                 if (csym\_func(cs)->arg_ids)
618                     free(csym\_func(cs)->arg_ids);
619             } else if (csym\_type(cs) == FFI_STRUCT) {
620                 if (csym\_struct(cs)->members)
621                     free(csym\_struct(cs)->members);
622             }
623         }
624         free(cs\_arr);
625     }
626
627     if (cte\_arr) {
628         free(cte\_arr);
629     }

```

```
630  
631  
632 }
```

```
return 0;
```

[One Level Up](#)

[Top Level](#)

userspace/kp_util.h - ktap

Data types defined

- [SBuf](#)
- [SBuf](#)
- [bool](#)
- [cp_csymbol_state](#)
- [cp_csymbol_state](#)
- [ktap_writer](#)

Macros defined

- [KP_CHAR_ALNUM](#)
- [KP_CHAR_ALPHA](#)
- [KP_CHAR_CNTRL](#)
- [KP_CHAR_DIGIT](#)
- [KP_CHAR_GRAPH](#)
- [KP_CHAR_IDENT](#)
- [KP_CHAR_LOWER](#)
- [KP_CHAR_PUNCT](#)
- [KP_CHAR_SPACE](#)
- [KP_CHAR_UPPER](#)
- [KP_CHAR_XDIGIT](#)
- [__KTAP_UTIL_H](#)
- [err2msg](#)
- [false](#)
- [kp_char_isa](#)
- [kp_char_isalnum](#)
- [kp_char_isalpha](#)
- [kp_char_iscntrl](#)
- [kp_char_isdigit](#)
- [kp_char_isgraph](#)
- [kp_char_isident](#)
- [kp_char_islower](#)

- [kp_char_ispunct](#)
- [kp_char_isspace](#)
- [kp_char_isupper](#)
- [kp_char_isxdigit](#)
- [kp_char_tolower](#)
- [kp_char_toupper](#)
- [sbufB](#)
- [sbufE](#)
- [sbufP](#)
- [sbufleft](#)
- [sbuflen](#)
- [sbufsz](#)
- [setsbufP](#)
- [true](#)
- [verbose_printf](#)

Source code

```

1  #ifndef __KTAP_UTIL_H
2  #define __KTAP_UTIL_H
3
4  #include <stdarg.h>
5
6  #include "../include/ktap_bc.h"
7  #include "../include/ktap_err.h"
8
9  typedef int bool;
10 #define false 0
11 #define true 1
12
13 /* Resizable string buffer. */
14 typedef struct SBuf {
15     char *p; /* String buffer pointer. */
16     char *e; /* String buffer end pointer. */
17     char *b; /* String buffer base. */
18 } SBuf;
19
20 /* Resizable string buffers. Struct definition in kp_obj.h. */
21 #define sbufB(sb) ((char *)(sb)->b)
22 #define sbufP(sb) ((char *)(sb)->p)
23 #define sbufE(sb) ((char *)(sb)->e)
24 #define sbufsz(sb) ((int)(sbufE((sb)) - sbufB((sb))))
25 #define sbuflen(sb) ((int)(sbufP((sb)) - sbufB((sb))))
26 #define sbufleft(sb) ((int)(sbufE((sb)) - sbufP((sb))))
27 #define setsbufP(sb, q) ((sb)->p = (q))
28
29 void kp_buf_init(SBuf *sb);
30 void kp_buf_reset(SBuf *sb);
31 void kp_buf_free(SBuf *sb);
32 char *kp_buf_more(SBuf *sb, int sz);
33 char *kp_buf_need(SBuf *sb, int sz);
34 char *kp_buf_wmem(char *p, const void *q, int len);
35 void kp_buf_putb(SBuf *sb, int c);
36 ktap_str_t *kp_buf_str(SBuf *sb);
37
38

```

```

39 #define KP_CHAR_CNTRL      0x01
40 #define KP_CHAR_SPACE     0x02
41 #define KP_CHAR_PUNCT     0x04
42 #define KP_CHAR_DIGIT     0x08
43 #define KP_CHAR_XDIGIT    0x10
44 #define KP_CHAR_UPPER     0x20
45 #define KP_CHAR_LOWER     0x40
46 #define KP_CHAR_IDENT     0x80
47 #define KP_CHAR_ALPHA     (KP_CHAR_LOWER|KP_CHAR_UPPER)
48 #define KP_CHAR_ALNUM     (KP_CHAR_ALPHA|KP_CHAR_DIGIT)
49 #define KP_CHAR_GRAPH     (KP_CHAR_ALNUM|KP_CHAR_PUNCT)
50
51 /* Only pass -1 or 0..255 to these macros. Never pass a signed char! */
52 #define kp_char_isa(c, t)  ((kp_char_bits+1)[(c)] & t)
53 #define kp_char_iscntrl(c)  kp_char_isa((c), KP_CHAR_CNTRL)
54 #define kp_char_isspace(c)  kp_char_isa((c), KP_CHAR_SPACE)
55 #define kp_char_ispunct(c)  kp_char_isa((c), KP_CHAR_PUNCT)
56 #define kp_char_isdigit(c)  kp_char_isa((c), KP_CHAR_DIGIT)
57 #define kp_char_isxdigit(c)  kp_char_isa((c), KP_CHAR_XDIGIT)
58 #define kp_char_isupper(c)  kp_char_isa((c), KP_CHAR_UPPER)
59 #define kp_char_islower(c)  kp_char_isa((c), KP_CHAR_LOWER)
60 #define kp_char_isident(c)  kp_char_isa((c), KP_CHAR_IDENT)
61 #define kp_char_isalpha(c)  kp_char_isa((c), KP_CHAR_ALPHA)
62 #define kp_char_isalnum(c)  kp_char_isa((c), KP_CHAR_ALNUM)
63 #define kp_char_isgraph(c)  kp_char_isa((c), KP_CHAR_GRAPH)
64
65 #define kp_char_toupper(c)  ((c) - (kp_char_islower(c) >> 1))
66 #define kp_char_tolower(c)  ((c) + kp_char_isupper(c))
67
68 extern const char *kp_err_allmsg;
69 #define err2msg(em)         (kp_err_allmsg+(int)(em))
70
71 extern const uint8_t kp_char_bits[257];
72
73
74 char *strfmt_wuleb128(char *p, uint32_t v);
75 void kp_err_lex(ktap_str_t *src, const char *tok, BCLine line,
76               ErrMsg em, va_list argp);
77 char *kp_sprintf(const char *fmt, ...);
78 const char *kp_sprintfv(const char *fmt, va_list argp);
79
80 void *kp_reallocv(void *block, size_t osize, size_t nsize);
81
82 void kp_str_resize(void);
83 ktap_str_t *kp_str_newz(const char *str);
84 ktap_str_t *kp_str_new(const char *str, size_t l);
85
86 ktap_tab_t *kp_tab_new();
87 const ktap_val_t *kp_tab_get(ktap_tab_t *t, const ktap_val_t *key);
88 const ktap_val_t *kp_tab_getstr(ktap_tab_t *t, const ktap_str_t *ts);
89 void kp_tab_setvalue(ktap_tab_t *t, const ktap_val_t *key, ktap_val_t *val);
90 ktap_val_t *kp_tab_set(ktap_tab_t *t, const ktap_val_t *key);
91
92 int kp_obj_equal(const ktap_val_t *t1, const ktap_val_t *t2);
93
94 bool strglobmatch(const char *str, const char *pat);
95 int kallsyms_parse(void *arg,
96                  int(*process_symbol)(void *arg, const char *name,
97                                       char type, unsigned long start));
98
99 unsigned long find_kernel_symbol(const char *symbol);
100 void list_available_events(const char *match);
101 void process_available_tracepoints(const char *sys, const char *event,
102                                  int(*process)(const char *sys,
103                                                const char *event));
104 int kallsyms_parse(void *arg,
105                  int(*process_symbol)(void *arg, const char *name,
106                                       char type, unsigned long start));
107
108
109 #ifndef CONFIG_KTAP_FFI
110 #include "../include/ktap_ffi.h"
111
112 typedef struct cp_csymbol_state {
113     int cs_nr; /* number of c symbols */
114     int cs_arr_size; /* size of current symbol arrays */

```

```
115     csymbol *cs\_arr;  
116 } cp\_csymbol\_state;  
117  
118 cp\_csymbol\_state *ctype\_get\_csym\_state(void);  
119 void kp\_dump\_csymbols();  
120 #endif  
121  
122 ktap\_eventdesc\_t *kp\_parse\_events(const char *eventdef);  
123 void cleanup\_event\_resources(void);  
124  
125 extern int verbose;  
126 #define verbose\_printf(...) \  
127     if (verbose) \  
128         printf("[verbose] " __VA_ARGS__);  
129  
130  
131 void kp\_dump\_proto(ktap\_proto\_t *pt);  
132 typedef int (*ktap\_writer)(const void* p, size_t sz, void* ud);  
133 int kp\_bcwrite(ktap\_proto\_t *pt, ktap\_writer writer, void *data, int strip);  
134  
135 int kp\_create\_reader(const char *output);  
136 #endif
```

[One Level Up](#)

[Top Level](#)

userspace/kp_util.c - ktap

Global variables defined

- [kp_char_bits](#)
- [kp_err_allmsg](#)
- [kp_niltv](#)
- [strtab](#)
- [strtab_nr](#)
- [strtab_size](#)

Data types defined

- [ksym_addr_t](#)

Functions defined

- [__match_charclass](#)
- [__match_glob](#)
- [createstrobj](#)
- [find_kernel_symbol](#)
- [kallsyms_parse](#)
- [kp_buf_free](#)
- [kp_buf_init](#)
- [kp_buf_more](#)
- [kp_buf_need](#)
- [kp_buf_putb](#)
- [kp_buf_reset](#)
- [kp_buf_str](#)
- [kp_buf_wmem](#)
- [kp_err_lex](#)
- [kp_obj_equal](#)
- [kp_reallocv](#)
- [kp_sprintf](#)
- [kp_sprintfv](#)
- [kp_str_new](#)
- [kp_str_newz](#)

- [kp_str_resize](#)
- [kp_tab_get](#)
- [kp_tab_getstr](#)
- [kp_tab_new](#)
- [kp_tab_set](#)
- [kp_tab_setvalue](#)
- [list_available_events](#)
- [process_available_tracepoints](#)
- [strfmt_wuleb128](#)
- [strglobmatch](#)
- [stringtable_insert](#)
- [stringtable_search](#)
- [symbol_cmp](#)

Macros defined

- [AVAILABLE_EVENTS_PATH](#)
- [ERRDEF](#)
- [KALLSYMS_PATH](#)
- [gkey](#)
- [gnode](#)
- [gval](#)
- [handle_error](#)
- [niltv](#)

Source code

```

1 /*
2  * util.c
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Adapted from luajit and lua interpreter.
9  * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with

```

```

22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include <stdarg.h>
27 #include <stdio.h>
28 #include <stdlib.h>
29 #include <string.h>
30 #include <math.h>
31 #include <ctype.h>
32 #include "../include/ktap_types.h"
33 #include "../include/ktap_bc.h"
34 #include "kp_util.h"
35
36 /* Error message strings. */
37 const char *kp_err_allmsg =
38 #define ERRDEF(name, msg) msg "\0"
39 #include "../include/ktap_errmsg.h"
40 ;
41
42 const uint8_t kp_char_bits[257] = {
43     0,
44     1, 1, 1, 1, 1, 1, 1, 1, 3, 3, 3, 3, 3, 1, 1,
45     1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
46     2, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4,
47     152,152,152,152,152,152,152,152,152,152, 4, 4, 4, 4, 4, 4,
48     4,176,176,176,176,176,176,160,160,160,160,160,160,160,160,160,160,
49     160,160,160,160,160,160,160,160,160,160,160, 4, 4, 4, 4,132,
50     4,208,208,208,208,208,208,192,192,192,192,192,192,192,192,192,
51     192,192,192,192,192,192,192,192,192,192, 4, 4, 4, 4, 1,
52     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
53     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
54     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
55     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
56     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
57     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
58     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,
59     128,128,128,128,128,128,128,128,128,128,128,128,128,128,128,128
60 };
61
62 void kp_buf_init(SBuf *sb)
63 {
64     sb->b = (char *)malloc(200);
65     sb->p = NULL;
66     sb->e = sb->b + 200;
67 }
68
69 void kp_buf_reset(SBuf *sb)
70 {
71     sb->p = sb->b;
72 }
73
74 void kp_buf_free(SBuf *sb)
75 {
76     free(sbufB(sb));
77 }
78
79 char *kp_buf_more(SBuf *sb, int sz)
80 {
81     char *b;
82     int old_len = sbufLen(sb);
83
84     if (sz > sbufLeft(sb)) {
85         b = realloc(sbufB(sb), sbufLen(sb) * 2);
86         sb->b = b;
87         sb->p = b + old_len;
88         sb->e = b + old_len * 2;
89     }
90
91     return sbufP(sb);
92 }
93
94 char *kp_buf_need(SBuf *sb, int sz)
95 {
96     char *b;
97     int old_len = sbufLen(sb);

```

```

98
99     if (sz > sbufsz(sb)) {
100         b = realloc(sbufB(sb), sz);
101         sb->b = b;
102         sb->p = b + old_len;
103         sb->e = b + sz;
104     }
105
106     return sbufB(sb);
107 }
108
109 char *kp\_buf\_wmem(char *p, const void *q, int len)
110 {
111     return (char *)memcpy(p, q, len) + len;
112 }
113
114 void kp\_buf\_putb(SBuf *sb, int c)
115 {
116     char *p = kp\_buf\_more(sb, 1);
117     *p++ = (char)c;
118     setsbufP(sb, p);
119 }
120
121 ktap\_str\_t *kp\_buf\_str(SBuf *sb)
122 {
123     return kp\_str\_new(sbufB(sb), sbuflen(sb));
124 }
125
126 /* Write ULEB128 to buffer. */
127 char *strfmt\_wuleb128(char *p, uint32\_t v)
128 {
129     for (; v >= 0x80; v >>= 7)
130         *p++ = (char)((v & 0x7f) | 0x80);
131     *p++ = (char)v;
132     return p;
133 }
134
135 void kp\_err\_lex(ktap\_str\_t *src, const char *tok, BCLine line,
136               ErrMsg em, va\_list argp)
137 {
138     const char *msg;
139
140     msg = kp\_sprintfv(err2msg(em), argp);
141     msg = kp\_sprintf("%s:%d: %s", getstr(src), line, msg);
142     if (tok)
143         msg = kp\_sprintf(err2msg(KP_ERR_XNEAR), msg, tok);
144     fprintf(stderr, "%s: %s\n", err2msg(KP_ERR_XSYNTAX), msg);
145     exit(-1);
146 }
147
148 void *kp\_reallocv(void *block, size\_t osize, size\_t nsize)
149 {
150     return realloc(block, nsize);
151 }
152
153 static const ktap\_val\_t kp\_niltv = { {NULL}, {KTAP\_TNIL} } ;
154 #define niltv (&kp\_niltv)
155
156 #define gnode(t,i) (&(t)->node[i])
157 #define gkey(n) (&(n)->key)
158 #define gval(n) (&(n)->val)
159
160 const ktap\_val\_t *kp\_tab\_get(ktap\_tab\_t *t, const ktap\_val\_t *key)
161 {
162     int i;
163
164     switch (itype(key)) {
165     case KTAP\_TNIL:
166         return niltv;
167     case KTAP\_TNUM:
168         for (i = 0; i <= t->hmask; i++) {
169             ktap\_val\_t *v = gkey(gnode(t, i));
170             if (is\_number(v) && nvalue(key) == nvalue(v))
171                 return gval(gnode(t, i));
172         }
173         break;

```

```

174 case KTAP_TSTR:
175     for (i = 0; i <= t->hmask; i++) {
176         ktap_val_t *v = gkey(gnode(t, i));
177         if (is_string(v) && (rawtsvalue(key) == rawtsvalue(v)))
178             return gval(gnode(t, i));
179     }
180     break;
181 default:
182     for (i = 0; i <= t->hmask; i++) {
183         if (kp_obj_equal(key, gkey(gnode(t, i))))
184             return gval(gnode(t, i));
185     }
186     break;
187 }
188
189 return niltv;
190 }
191
192 const ktap_val_t *kp_tab_getstr(ktap_tab_t *t, const ktap_str_t *ts)
193 {
194     int i;
195
196     for (i = 0; i <= t->hmask; i++) {
197         ktap_val_t *v = gkey(gnode(t, i));
198         if (is_string(v) && (ts == rawtsvalue(v)))
199             return gval(gnode(t, i));
200     }
201
202     return niltv;
203 }
204
205 void kp_tab_setvalue(ktap_tab_t *t, const ktap_val_t *key, ktap_val_t *val)
206 {
207     const ktap_val_t *v = kp_tab_get(t, key);
208
209     if (v != niltv) {
210         set_obj((ktap_val_t *)v, val);
211     } else {
212         if (t->freetop == t->node) {
213             int size = (t->hmask + 1) * sizeof(ktap_node_t);
214             t->node = realloc(t->node, size * 2);
215             memset(t->node + t->hmask + 1, 0, size);
216             t->freetop = t->node + (t->hmask + 1) * 2;
217             t->hmask = (t->hmask + 1) * 2 - 1;
218         }
219
220         ktap_node_t *n = --t->freetop;
221         set_obj(gkey(n), key);
222         set_obj(gval(n), val);
223     }
224 }
225
226 ktap_val_t *kp_tab_set(ktap_tab_t *t, const ktap_val_t *key)
227 {
228     const ktap_val_t *v = kp_tab_get(t, key);
229
230     if (v != niltv) {
231         return (ktap_val_t *)v;
232     } else {
233         if (t->freetop == t->node) {
234             int size = (t->hmask + 1) * sizeof(ktap_node_t);
235             t->node = realloc(t->node, size * 2);
236             memset(t->node + t->hmask + 1, 0, size);
237             t->freetop = t->node + (t->hmask + 1) * 2;
238             t->hmask = (t->hmask + 1) * 2 - 1;
239         }
240
241         ktap_node_t *n = --t->freetop;
242         set_obj(gkey(n), key);
243         set_nil(gval(n));
244         return gval(n);
245     }
246 }
247
248
249 ktap_tab_t *kp_tab_new(void)

```

```

250 {
251     int hsize, i;
252
253     ktap_tab_t *t = malloc(sizeof(ktap_tab_t));
254     t->gct = ~KTAP_TTAB;
255     hsize = 1024;
256     t->hmask = hsize - 1;
257     t->node = (ktap_node_t *)malloc(hsize * sizeof(ktap_node_t));
258     t->freetop = &t->node[hsize];
259     t->asize = 0;
260
261     for (i = 0; i <= t->hmask; i++) {
262         set_nil(&t->node[i].val);
263         set_nil(&t->node[i].key);
264     }
265     return t;
266 }
267
268 /* simple interned string array, use hash table in future */
269 static ktap_str_t **strtab;
270 static int strtab_size = 1000; /* initial size */
271 static int strtab_nr;
272
273 void kp_str_resize(void)
274 {
275     int size = strtab_size * sizeof(ktap_str_t *);
276
277     strtab = malloc(size);
278     if (!strtab) {
279         fprintf(stderr, "cannot allocate stringtable\n");
280         exit(-1);
281     }
282
283     memset(strtab, 0, size);
284     strtab_nr = 0;
285 }
286
287 static ktap_str_t *stringtable_search(const char *str, int len)
288 {
289     int i;
290
291     for (i = 0; i < strtab_nr; i++) {
292         ktap_str_t *s = strtab[i];
293         if ((len == s->len) && !memcmp(str, getstr(s), len))
294             return s;
295     }
296
297     return NULL;
298 }
299
300 static void stringtable_insert(ktap_str_t *ts)
301 {
302     strtab[strtab_nr++] = ts;
303
304     if (strtab_nr == strtab_size) {
305         int size = strtab_size * sizeof(ktap_str_t *);
306         strtab = realloc(strtab, size * 2);
307         memset(strtab + strtab_size, 0, size);
308         strtab_size *= 2;
309     }
310 }
311
312 static ktap_str_t *createstrobj(const char *str, size_t l)
313 {
314     ktap_str_t *ts;
315     size_t totalsize; /* total size of TString object */
316
317     totalsize = sizeof(ktap_str_t) + ((l + 1) * sizeof(char));
318     ts = (ktap_str_t *)malloc(totalsize);
319     ts->gct = ~KTAP_TSTR;
320     ts->len = l;
321     ts->reserved = 0;
322     ts->extra = 0;
323     memcpy(ts + 1, str, l * sizeof(char));
324     ((char *) (ts + 1))[l] = '\0'; /* ending 0 */
325     return ts;

```

```

326 }
327
328 ktap_str_t *kp_str_new(const char *str, size_t l)
329 {
330     ktap_str_t *ts = stringtable_search(str, l);
331
332     if (ts)
333         return ts;
334
335     ts = createstrobj(str, l);
336     stringtable_insert(ts);
337     return ts;
338 }
339
340 ktap_str_t *kp_str_newz(const char *str)
341 {
342     return kp_str_new(str, strlen(str));
343 }
344
345 /*
346  * todo: memory leak here
347  */
348 char *kp_sprintf(const char *fmt, ...)
349 {
350     char *msg = malloc(128);
351
352     va_list argp;
353     va_start(argp, fmt);
354     vsprintf(msg, fmt, argp);
355     va_end(argp);
356     return msg;
357 }
358
359 const char *kp_sprintfv(const char *fmt, va_list argp)
360 {
361     char *msg = malloc(128);
362
363     vsprintf(msg, fmt, argp);
364     return msg;
365 }
366
367 int kp_obj_equal(const ktap_val_t *t1, const ktap_val_t *t2)
368 {
369     switch (itype(t1)) {
370     case KTAP_TNIL:
371         return 1;
372     case KTAP_TNUM:
373         return nvalue(t1) == nvalue(t2);
374     case KTAP_TTRUE:
375     case KTAP_TFALSE:
376         return itype(t1) == itype(t2);
377     case KTAP_TLIGHTUD:
378         return pvalue(t1) == pvalue(t2);
379     case KTAP_TFUNC:
380         return fvalue(t1) == fvalue(t2);
381     case KTAP_TSTR:
382         return rawtsvalue(t1) == rawtsvalue(t2);
383     default:
384         return gcvalue(t1) == gcvalue(t2);
385     }
386
387     return 0;
388 }
389
390 /*
391  * strglobmatch is copied from perf(linux/tools/perf/util/string.c)
392  */
393
394 /* Character class matching */
395 static bool __match_charclass(const char *pat, char c, const char **npat)
396 {
397     bool complement = false, ret = true;
398
399     if (*pat == '!') {
400         complement = true;
401         pat++;

```

```

402 }
403 if (*pat++ == c) /* First character is special */
404     goto end;
405
406 while (*pat && *pat != ']') { /* Matching */
407     if (*pat == '-' && *(pat + 1) != ']') { /* Range */
408         if (*(pat - 1) <= c && c <= *(pat + 1))
409             goto end;
410         if (*(pat - 1) > *(pat + 1))
411             goto error;
412         pat += 2;
413     } else if (*pat++ == c)
414         goto end;
415 }
416 if (!*pat)
417     goto error;
418 ret = false;
419
420 end:
421 while (*pat && *pat != ']') /* Searching closing */
422     pat++;
423 if (!*pat)
424     goto error;
425 *npat = pat + 1;
426 return complement ? !ret : ret;
427
428 error:
429     return false;
430 }
431
432 /* Glob/lazy pattern matching */
433 static bool __match_glob(const char *str, const char *pat, bool ignore_space)
434 {
435     while (*str && *pat && *pat != '*') {
436         if (ignore_space) {
437             /* Ignore spaces for lazy matching */
438             if (isspace(*str)) {
439                 str++;
440                 continue;
441             }
442             if (isspace(*pat)) {
443                 pat++;
444                 continue;
445             }
446         }
447         if (*pat == '?') { /* Matches any single character */
448             str++;
449             pat++;
450             continue;
451         } else if (*pat == '[') /* Character classes/Ranges */
452             if (__match_charclass(pat + 1, *str, &pat)) {
453                 str++;
454                 continue;
455             } else
456                 return false;
457         else if (*pat == '\\') /* Escaped char match as normal char */
458             pat++;
459         if (*str++ != *pat++)
460             return false;
461     }
462     /* Check wild card */
463     if (*pat == '*') {
464         while (*pat == '*')
465             pat++;
466         if (!*pat) /* Tail wild card matches all */
467             return true;
468         while (*str)
469             if (__match_glob(str++, pat, ignore_space))
470                 return true;
471     }
472     return !*str && !*pat;
473 }
474
475 /**
476 * strglobmatch - glob expression pattern matching
477 * @str: the target string to match

```

```

478 * @pat: the pattern string to match
479 *
480 * This returns true if the @str matches @pat. @pat can includes wildcards
481 * ('*', '?') and character classes ([CHARS], complementation and ranges are
482 * also supported). Also, this supports escape character ('\') to use special
483 * characters as normal character.
484 *
485 * Note: if @pat syntax is broken, this always returns false.
486 */
487 bool strglobmatch(const char *str, const char *pat)
488 {
489     return __match_glob(str, pat, false);
490 }
491
492 #define handle_error(str) do { perror(str); exit(-1); } while(0)
493
494 #define KALLSYMS_PATH "/proc/kallsyms"
495 /*
496 * read kernel symbol from /proc/kallsyms
497 */
498 int kallsyms_parse(void *arg,
499                   int(*process_symbol)(void *arg, const char *name,
500                   char type, unsigned long start))
501 {
502     FILE *file;
503     char *line = NULL;
504     int ret = 0;
505     int found = 0;
506
507     file = fopen(KALLSYMS_PATH, "r");
508     if (file == NULL)
509         handle_error("open " KALLSYMS_PATH " failed");
510
511     while (!feof(file)) {
512         char *symbol_addr, *symbol_name;
513         char symbol_type;
514         unsigned long start;
515         int line_len;
516         size_t n;
517
518         line_len = getline(&line, &n, file);
519         if (line_len < 0 || !line)
520             break;
521
522         line[--line_len] = '\0'; /* \n */
523
524         symbol_addr = strtok(line, " \t");
525         start = strtoul(symbol_addr, NULL, 16);
526
527         symbol_type = *strtok(NULL, " \t");
528         symbol_name = strtok(NULL, " \t");
529
530         ret = process_symbol(arg, symbol_name, symbol_type, start);
531         if (!ret)
532             found = 1;
533     }
534
535     free(line);
536     fclose(file);
537
538     return found;
539 }
540
541 struct ksym_addr_t {
542     const char *name;
543     unsigned long addr;
544 };
545
546 static int symbol_cmp(void *arg, const char *name, char type,
547                     unsigned long start)
548 {
549     struct ksym_addr_t *base = arg;
550
551     if (strcmp(base->name, name) == 0) {
552         base->addr = start;
553         return 1;

```

```

554     }
555
556     return 0;
557 }
558
559 unsigned long find_kernel_symbol(const char *symbol)
560 {
561     int ret;
562     struct ksym_addr_t arg = {
563         .name = symbol,
564         .addr = 0
565     };
566
567     ret = kallsyms_parse(&arg, symbol_cmp);
568     if (ret < 0 || arg.addr == 0) {
569         fprintf(stderr, "cannot read kernel symbol \"%s\" in %s\n",
570             symbol, KALLSYMS_PATH);
571         exit(EXIT_FAILURE);
572     }
573
574     return arg.addr;
575 }
576
577
578 #define AVAILABLE_EVENTS_PATH "/sys/kernel/debug/tracing/available_events"
579
580 void list_available_events(const char *match)
581 {
582     FILE *file;
583     char *line = NULL;
584
585     file = fopen(AVAILABLE_EVENTS_PATH, "r");
586     if (file == NULL)
587         handle_error("open " AVAILABLE_EVENTS_PATH " failed");
588
589     while (!feof(file)) {
590         int line_len;
591         size_t n;
592
593         line_len = getline(&line, &n, file);
594         if (line_len < 0 || !line)
595             break;
596
597         if (!match || strglobmatch(line, match))
598             printf("%s", line);
599     }
600
601     free(line);
602     fclose(file);
603 }
604
605 void process_available_tracepoints(const char *sys, const char *event,
606     int (*process)(const char *sys,
607     const char *event))
608 {
609     char *line = NULL;
610     FILE *file;
611     char str[128] = {0};
612
613     /* add '\n' into tail */
614     snprintf(str, 64, "%s:%s\n", sys, event);
615
616     file = fopen(AVAILABLE_EVENTS_PATH, "r");
617     if (file == NULL)
618         handle_error("open " AVAILABLE_EVENTS_PATH " failed");
619
620     while (!feof(file)) {
621         int line_len;
622         size_t n;
623
624         line_len = getline(&line, &n, file);
625         if (line_len < 0 || !line)
626             break;
627
628         if (strglobmatch(line, str)) {
629             char match_sys[64] = {0};

```

```
630         char match_event[64] = {0};
631         char *sep;
632
633         sep = strchr(line, ':');
634         memcpy(match_sys, line, sep - line);
635         memcpy(match_event, sep + 1,
636                line_len - (sep - line) - 2);
637
638         if (process(match_sys, match_event))
639             break;
640     }
641 }
642
643 free(line);
644 fclose(file);
645 }
646
```

[One Level Up](#)

[Top Level](#)

runtime/kp_str.c - ktap

Functions defined

- [addlenmod](#)
- [arg_error](#)
- [kp_str_cmp](#)
- [kp_str_fmt](#)
- [kp_str_freall](#)
- [kp_str_hash](#)
- [kp_str_new](#)
- [kp_str_resize](#)
- [scanformat](#)
- [str_fastcmp](#)

Macros defined

- [FLAGS](#)
- [INTERMLEN](#)
- [INTERM T](#)
- [L_ESC](#)
- [MAX_FORMAT](#)
- [STRING_HASHLIMIT](#)
- [uchar](#)

Source code

```
1 /*
2  * kp_str.c - ktap string data struction manipulation
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Adapted from luajit and lua interpreter.
9  * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
```

```

24 */
25
26 #include "../include/ktap_types.h"
27 #include "kp_obj.h"
28 #include "kp_str.h"
29 #include "kp_mempool.h"
30
31 #include <linux/ctype.h>
32 #include <linux/module.h>
33 #include <linux/kallsyms.h>
34 #include "ktap.h"
35 #include "kp_transport.h"
36 #include "kp_vm.h"
37 #include "kp_events.h"
38
39 int kp_str_cmp(const ktap_str_t *ls, const ktap_str_t *rs)
40 {
41     const char *l = getstr(ls);
42     size_t ll = ls->len;
43     const char *r = getstr(rs);
44     size_t lr = rs->len;
45
46     for (;;) {
47         int temp = strcmp(l, r);
48         if (temp != 0)
49             return temp;
50         else {
51             /* strings are equal up to a '\0' */
52
53             /* index of first '\0' in both strings */
54             size_t len = strlen(l);
55
56             /* r is finished? */
57             if (len == lr)
58                 return (len == ll) ? 0 : 1;
59             else if (len == ll) /* l is finished? */
60                 return -1;
61
62             /*
63              * both strings longer than `len';
64              * go on comparing (after the '\0')
65              */
66             len++;
67             l += len; ll -= len; r += len; lr -= len;
68         }
69     }
70 }
71
72 /* Fast string data comparison. Caveat: unaligned access to 1st string! */
73 static __always_inline int str_fastcmp(const char *a, const char *b, int len)
74 {
75     int i = 0;
76
77     kp_assert(len > 0);
78     kp_assert((((uintptr_t)a + len - 1) & (PAGE_SIZE - 1)) <= PAGE_SIZE - 4);
79
80     do { /* Note: innocuous access up to end of string + 3. */
81         uint32_t v = *(uint32_t *) (a + i) ^ *(const uint32_t *) (b + i);
82         if (v) {
83             i -= len;
84 #if KP_LE
85             return (int32_t) i >= -3 ? (v << (32 + (i << 3))) : 1;
86 #else
87             return (int32_t) i >= -3 ? (v >> (32 + (i << 3))) : 1;
88 #endif
89         }
90         i += 4;
91     } while (i < len);
92     return 0;
93 }
94
95
96 //TODO: change hash algo
97
98 #define STRING_HASHLIMIT 5
99 static __always_inline unsigned int kp_str_hash(const char *str, size_t len)

```

```

100 {
101     unsigned int h = 201236 ^ len;
102     size_t step = (len >> STRING_HASHLIMIT) + 1;
103     size_t l1;
104
105     for (l1 = len; l1 >= step; l1 -= step)
106         h = h ^ ((h<<5) + (h>>2) + (u8)(str[l1 - 1]));
107
108     return h;
109 }
110
111
112 /*
113 * resizes the string table
114 */
115 int kp_str_resize(ktap_state_t *ks, int newmask)
116 {
117     ktap_global_state_t *g = G(ks);
118     ktap_str_t **newhash;
119
120     newhash = kp_zalloc(ks, (newmask + 1) * sizeof(ktap_str_t *));
121     if (!newhash)
122         return -ENOMEM;
123
124     g->strmask = newmask;
125     g->strhash = newhash;
126     return 0;
127 }
128
129 /*
130 * Intern a string and return string object.
131 */
132 ktap_str_t * kp_str_new(ktap_state_t *ks, const char *str, size_t len)
133 {
134     ktap_global_state_t *g = G(ks);
135     ktap_str_t *s;
136     ktap_obj_t *o;
137     unsigned int h = kp_str_hash(str, len);
138     unsigned long flags;
139
140     if (len >= KP_MAX_STR)
141         return NULL;
142
143     local_irq_save(flags);
144     arch_spin_lock(&g->str_lock);
145
146     o = (ktap_obj_t *)g->strhash[h & g->strmask];
147     if (likely((((uintptr_t)str+len-1) & (PAGE_SIZE-1)) <= PAGE_SIZE-4)) {
148         while (o != NULL) {
149             ktap_str_t *sx = (ktap_str_t *)o;
150             if (sx->len == len &&
151                 !str_fastcmp(str, getstr(sx), len)) {
152                 arch_spin_unlock(&g->str_lock);
153                 local_irq_restore(flags);
154                 return sx; /* Return existing string. */
155             }
156             o = gch(o)->nextgc;
157         }
158     } else { /* Slow path: end of string is too close to a page boundary */
159         while (o != NULL) {
160             ktap_str_t *sx = (ktap_str_t *)o;
161             if (sx->len == len &&
162                 !memcmp(str, getstr(sx), len)) {
163                 arch_spin_unlock(&g->str_lock);
164                 local_irq_restore(flags);
165                 return sx; /* Return existing string. */
166             }
167             o = gch(o)->nextgc;
168         }
169     }
170
171     /* create a new string, allocate it from mempool, not use kmalloc. */
172     s = kp_mempool_alloc(ks, sizeof(ktap_str_t) + len + 1);
173     if (unlikely(!s))
174         goto out;
175     s->gct = ~KTAP_TSTR;

```

```

176     s->len = len;
177     s->hash = h;
178     s->reserved = 0;
179     memcpy(s + 1, str, len);
180     ((char *) (s + 1))[len] = '\0'; /* ending 0 */
181
182     /* Add it to string hash table */
183     h &= g->strmask;
184     s->nextgc = (ktap_obj_t *)g->strhash[h];
185     g->strhash[h] = s;
186     if (g->strnum++ > KP_MAX_STRNUM) {
187         kp_error(ks, "exceed max string number %d\n", KP_MAX_STRNUM);
188         s = NULL;
189     }
190
191 out:
192     arch_spin_unlock(&g->str_lock);
193     local_irq_restore(flags);
194     return s; /* Return newly interned string. */
195 }
196
197 void kp_str_freeall(ktap_state_t *ks)
198 {
199     /* don't need to free string in here, it will handled by mempool */
200     kp_free(ks, G(ks)->strhash);
201 }
202
203 /* kp_str_fmt - printf implementation */
204
205 /* macro to `unsign' a character */
206 #define uchar(c)    ((unsigned char)(c))
207
208 #define L_ESC      '%'
209
210 /* valid flags in a format specification */
211 #define FLAGS      "-+ #0"
212
213 #define INTFRMLEN  "ll"
214 #define INTFRM_T   long long
215
216 /*
217  * maximum size of each format specification (such as '%-099.99d')
218  * (+10 accounts for %99.99x plus margin of error)
219  */
220 #define MAX_FORMAT    (sizeof(FLAGS) + sizeof(INTFRMLEN) + 10)
221
222 static const char *scanformat(ktap_state_t *ks, const char *strfmt, char *form)
223 {
224     const char *p = strfmt;
225     while (*p != '\0' && strchr(FLAGS, *p) != NULL)
226         p++; /* skip flags */
227
228     if ((size_t)(p - strfmt) >= sizeof(FLAGS)/sizeof(char)) {
229         kp_error(ks, "invalid format (repeated flags)\n");
230         return NULL;
231     }
232
233     if (isdigit(uchar(*p)))
234         p++; /* skip width */
235
236     if (isdigit(uchar(*p)))
237         p++; /* (2 digits at most) */
238
239     if (*p == '.') {
240         p++;
241         if (isdigit(uchar(*p)))
242             p++; /* skip precision */
243         if (isdigit(uchar(*p)))
244             p++; /* (2 digits at most) */
245     }
246
247     if (isdigit(uchar(*p))) {
248         kp_error(ks, "invalid format (width or precision too long)\n");
249         return NULL;
250     }
251 }

```

```

252     *(form++) = '%';
253     memcpy(form, strfrmt, (p - strfrmt + 1) * sizeof(char));
254     form += p - strfrmt + 1;
255     *form = '\\0';
256     return p;
257 }
258
259
260 /*
261  * add length modifier into formats
262  */
263 static void addlenmod(char *form, const char *lenmod)
264 {
265     size_t l = strlen(form);
266     size_t lm = strlen(lenmod);
267     char spec = form[l - 1];
268
269     strcpy(form + l - 1, lenmod);
270     form[l + lm - 1] = spec;
271     form[l + lm] = '\\0';
272 }
273
274
275 static void arg_error(ktap_state_t *ks, int nargs, const char *extrams)
276 {
277     kp_error(ks, "bad argument #%d: (%s)\\n", nargs, extrams);
278 }
279
280 int kp_str_fmt(ktap_state_t *ks, struct trace_seq *seq)
281 {
282     int arg = 1;
283     size_t sfl;
284     ktap_val_t *arg_fmt = kp_arg(ks, 1);
285     int argnum = kp_arg_nr(ks);
286     const char *strfrmt, *strfrmt_end;
287
288     strfrmt = svalue(arg_fmt);
289     sfl = rawtsvalue(arg_fmt)->len;
290     strfrmt_end = strfrmt + sfl;
291
292     while (strfrmt < strfrmt_end) {
293         if (*strfrmt != L_ESC)
294             trace_seq_putc(seq, *strfrmt++);
295         else if (*++strfrmt == L_ESC)
296             trace_seq_putc(seq, *strfrmt++);
297         else { /* format item */
298             char form[MAX_FORMAT];
299
300             if (++arg > argnum) {
301                 arg_error(ks, arg, "no value");
302                 return -1;
303             }
304
305             strfrmt = scanformat(ks, strfrmt, form);
306             switch (*strfrmt++) {
307                 case 'c':
308                     kp_arg_checknumber(ks, arg);
309
310                     trace_seq_printf(seq, form,
311                                     nvalue(kp_arg(ks, arg)));
312                     break;
313                 case 'd': case 'i': {
314                     ktap_number n;
315                     INTFRM_T ni;
316
317                     kp_arg_checknumber(ks, arg);
318
319                     n = nvalue(kp_arg(ks, arg));
320                     ni = (INTFRM_T)n;
321                     addlenmod(form, INTFRMLEN);
322                     trace_seq_printf(seq, form, ni);
323                     break;
324                 }
325                 case 'p': {
326                     char str[KSYM_SYMBOL_LEN];
327

```

```

328     kp\_arg\_checknumber(ks, arg);
329
330     SPRINT\_SYMBOL(str, nvalue(kp\_arg(ks, arg)));
331     \_trace\_seq\_puts(seq, str);
332     break;
333 }
334 case 'o': case 'u': case 'x': case 'X': {
335     ktap\_number n;
336     unsigned INTFRM\_I ni;
337
338     kp\_arg\_checknumber(ks, arg);
339
340     n = nvalue(kp\_arg(ks, arg));
341     ni = (unsigned INTFRM\_I)n;
342     addlenmod(form, INTFRMLEN);
343     trace\_seq\_printf(seq, form, ni);
344     break;
345 }
346 case 's': {
347     ktap\_val\_t *v = kp\_arg(ks, arg);
348     const char *s;
349     size_t l;
350
351     if (is\_nil(v)) {
352         \_trace\_seq\_puts(seq, "nil");
353         return 0;
354     }
355
356     if (is\_eventstr(v)) {
357         const char *str = kp\_event\_tostr(ks);
358         if (!str)
359             return -1;
360         \_trace\_seq\_puts(seq,
361             kp\_event\_tostr(ks));
362         return 0;
363     }
364
365     kp\_arg\_checkstring(ks, arg);
366
367     s = svalue(v);
368     l = rawtsvalue(v)->len;
369     if (!strchr(form, '.') && l >= 100) {
370         /*
371          * no precision and string is too long
372          * to be formatted;
373          * keep original string
374          */
375         \_trace\_seq\_puts(seq, s);
376         break;
377     } else {
378         trace\_seq\_printf(seq, form, s);
379         break;
380     }
381 }
382 default: /* also treat cases `pnLlh' */
383     kp\_error(ks, "invalid option " KTAP\_QL("%%c")
384         " to " KTAP\_QL("format"),
385         *(strfmt - 1));
386     return -1;
387 }
388 }
389 }
390
391 return 0;
392 }
393

```

[One Level Up](#)

[Top Level](#)

runtime/kp_obj.c - ktap

Global variables defined

- [kp_err_allmsg](#)

Functions defined

- [kp_free](#)
- [kp_malloc](#)
- [kp_obj_dump](#)
- [kp_obj_free_gclist](#)
- [kp_obj_freeall](#)
- [kp_obj_kstack2str](#)
- [kp_obj_len](#)
- [kp_obj_new](#)
- [kp_obj_rawequal](#)
- [kp_obj_show](#)
- [kp_zalloc](#)

Macros defined

- [ERRDEF](#)
- [KTAP_ALLOC_FLAGS](#)

Source code

```
1 /*
2  * kp_obj.c - ktap object generic operation
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Adapted from luajit and lua interpreter.
9  * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include <linux/stacktrace.h>
```

```

27 #include <linux/module.h>
28 #include <linux/kallsyms.h>
29 #include <linux/slab.h>
30 #include "../include/ktap_types.h"
31 #include "../include/ktap_ffi.h"
32 #include "kp_obj.h"
33 #include "kp_str.h"
34 #include "kp_tab.h"
35 #include "ktap.h"
36 #include "kp_vm.h"
37 #include "kp_transport.h"
38
39 /* Error message strings. */
40 const char *kp_err_allmsg =
41 #define ERRDEF(name, msg) msg "\0"
42 #include "../include/ktap_errmsg.h"
43 ;
44
45 /* memory allocation flag */
46 #define KTAP_ALLOC_FLAGS ((GFP_KERNEL | __GFP_NORETRY | __GFP_NOWARN) \
47 & ~__GFP_WAIT)
48
49 /*
50 * TODO: It's not safe to call into facilities in the kernel at-large,
51 * so we may need to use ktap own memory pool, not kmalloc.
52 */
53
54
55 void *kp_malloc(ktap_state_t *ks, int size)
56 {
57     void *addr;
58
59     addr = kmalloc(size, KTAP_ALLOC_FLAGS);
60     if (unlikely(!addr)) {
61         kp_error(ks, "kmalloc failed\n");
62     }
63     return addr;
64 }
65
66 void *kp_zalloc(ktap_state_t *ks, int size)
67 {
68     void *addr;
69
70     addr = kzalloc(size, KTAP_ALLOC_FLAGS);
71     if (unlikely(!addr))
72         kp_error(ks, "kzalloc failed\n");
73     return addr;
74 }
75
76 void kp_free(ktap_state_t *ks, void *addr)
77 {
78     kfree(addr);
79 }
80
81
82 void kp_obj_dump(ktap_state_t *ks, const ktap_val_t *v)
83 {
84     switch (itype(v)) {
85     case KTAP_TNIL:
86         kp_puts(ks, "NIL");
87         break;
88     case KTAP_TTRUE:
89         kp_printf(ks, "true");
90         break;
91     case KTAP_TFALSE:
92         kp_printf(ks, "false");
93         break;
94     case KTAP_TNUM:
95         kp_printf(ks, "NUM %ld", nvalue(v));
96         break;
97     case KTAP_TLIGHTUD:
98         kp_printf(ks, "LIGHTUD 0x%lx", (unsigned long)pvalue(v));
99         break;
100    case KTAP_TFUNC:
101        kp_printf(ks, "FUNCTION 0x%lx", (unsigned long)fvalue(v));
102        break;

```

```

103     case KTAP_TSTR:
104         kp_printf(ks, "STR #%"s", svalue(v));
105         break;
106     case KTAP_TTAB:
107         kp_printf(ks, "TABLE 0x%lx", (unsigned long)hvalue(v));
108         break;
109     default:
110         kp_printf(ks, "GCVALUE 0x%lx", (unsigned long)gcvalue(v));
111         break;
112 }
113 }
114
115 void kp_obj_show(ktap_state_t *ks, const ktap_val_t *v)
116 {
117     switch (itype(v)) {
118     case KTAP_TNIL:
119         kp_puts(ks, "nil");
120         break;
121     case KTAP_TTRUE:
122         kp_puts(ks, "true");
123         break;
124     case KTAP_TFALSE:
125         kp_puts(ks, "false");
126         break;
127     case KTAP_TNUM:
128         kp_printf(ks, "%ld", nvalue(v));
129         break;
130     case KTAP_TLIGHTUD:
131         kp_printf(ks, "lightud 0x%lx", (unsigned long)pvalue(v));
132         break;
133     case KTAP_TCFUNC:
134         kp_printf(ks, "cfunction 0x%lx", (unsigned long)fvalue(v));
135         break;
136     case KTAP_TFUNC:
137         kp_printf(ks, "function 0x%lx", (unsigned long)gcvalue(v));
138         break;
139     case KTAP_TSTR:
140         kp_puts(ks, svalue(v));
141         break;
142     case KTAP_TTAB:
143         kp_printf(ks, "table 0x%lx", (unsigned long)hvalue(v));
144         break;
145 #ifdef CONFIG_KTAP_FFI
146     case KTAP_TCDATA:
147         kp_cdata_dump(ks, cdvalue(v));
148         break;
149 #endif
150     case KTAP_TEVENTSTR:
151         /* check event context */
152         if (!ks->current_event) {
153             kp_error(ks,
154                 "cannot stringify event str in invalid context\n");
155             return;
156         }
157
158         kp_transport_event_write(ks, ks->current_event);
159         break;
160     case KTAP_TKSTACK:
161         kp_transport_print_kstack(ks, v->val.stack.depth,
162             v->val.stack.skip);
163         break;
164     default:
165         kp_error(ks, "print unknown value type: %d\n", itype(v));
166         break;
167     }
168 }
169
170
171 /*
172 * equality of ktap values.
173 */
174 int kp_obj_rawequal(const ktap_val_t *t1, const ktap_val_t *t2)
175 {
176     switch (itype(t1)) {
177     case KTAP_TNIL:
178     case KTAP_TTRUE:

```

```

179     case KTAP_TFALSE:
180         return 1;
181     case KTAP_TNUM:
182         return nvalue(t1) == nvalue(t2);
183     case KTAP_TLIGHTUD:
184         return pvalue(t1) == pvalue(t2);
185     case KTAP_TFUNC:
186         return fvalue(t1) == fvalue(t2);
187     case KTAP_TSTR:
188         return rawtsvalue(t1) == rawtsvalue(t2);
189     case KTAP_TTAB:
190         return hvalue(t1) == hvalue(t2);
191     default:
192         return gcvalue(t1) == gcvalue(t2);
193 }
194
195 return 0;
196 }
197
198 /*
199 * ktap will not use lua's length operator for table,
200 * also # is not for length operator any more in ktap.
201 */
202 int kp_obj_len(ktap_state_t *ks, const ktap_val_t *v)
203 {
204     switch(itype(v)) {
205     case KTAP_TTAB:
206         return kp_tab_len(ks, hvalue(v));
207     case KTAP_TSTR:
208         return rawtsvalue(v)->len;
209     default:
210         kp_printf(ks, "cannot get length of type %d\n", v->type);
211         return -1;
212     }
213     return 0;
214 }
215
216 /* need to protect allgc field? */
217 ktap_obj_t *kp_obj_new(ktap_state_t *ks, size_t size)
218 {
219     ktap_obj_t *o, **list;
220
221     if (ks != G(ks)->mainthread) {
222         kp_error(ks, "kp_obj_new only can be called in mainthread\n");
223         return NULL;
224     }
225
226     o = kp_malloc(ks, size);
227     if (unlikely(!o))
228         return NULL;
229
230     list = &G(ks)->allgc;
231     gch(o)->nextgc = *list;
232     *list = o;
233
234     return o;
235 }
236
237
238 /* this function may be time consuming, move out from table set/get? */
239 ktap_str_t *kp_obj_kstack2str(ktap_state_t *ks, uint16_t depth, uint16_t skip)
240 {
241     struct stack_trace trace;
242     unsigned long *bt;
243     char *btstr, *p;
244     int i;
245
246     bt = kp_this_cpu_print_buffer(ks); /* use print percpu buffer */
247     trace.nr_entries = 0;
248     trace.skip = skip;
249     trace.max_entries = depth;
250     trace.entries = (unsigned long *) (bt + 1);
251     save_stack_trace(&trace);
252
253     /* convert backtrace to string */
254     p = btstr = kp_this_cpu_temp_buffer(ks);

```

```

255     for (i = 0; i < trace.nr_entries; i++) {
256         unsigned long addr = trace.entries[i];
257
258         if (addr == ULONG_MAX)
259             break;
260
261         p += sprint_symbol(p, addr);
262         *p++ = '\n';
263     }
264
265     return kp\_str\_new(ks, btstr, p - btstr);
266 }
267
268 void kp\_obj\_free\_gclist(ktap\_state\_t *ks, ktap\_obj\_t *o)
269 {
270     while (o) {
271         ktap\_obj\_t *next;
272
273         next = gch(o)->nextgc;
274         switch (gch(o)->gct) {
275             case ~KTAP\_TTAB:
276                 kp\_tab\_free(ks, (ktap\_tab\_t *)o);
277                 break;
278             case ~KTAP\_TUPVAL:
279                 kp\_freeupval(ks, (ktap\_upval\_t *)o);
280                 break;
281             default:
282                 kp\_free(ks, o);
283         }
284         o = next;
285     }
286 }
287
288 void kp\_obj\_freeall(ktap\_state\_t *ks)
289 {
290     kp\_obj\_free\_gclist(ks, G(ks)->allgc);
291     G(ks)->allgc = NULL;
292 }
293

```

[One Level Up](#)

[Top Level](#)

runtime/kp_transport.c - ktap

Global variables defined

- [ftrace find event](#)
- [tracing pipe fops](#)

Data types defined

- [ktap ftrace entry](#)
- [ktap trace iterator](#)
- [ktap trace type](#)

Functions defined

- [__find_next_entry](#)
- [__kp_bputs](#)
- [__kp_puts](#)
- [_trace_seq_puts](#)
- [_trace_seq_to_user](#)
- [kp_printf](#)
- [kp_transport_event_write](#)
- [kp_transport_exit](#)
- [kp_transport_init](#)
- [kp_transport_print_kstack](#)
- [kp_transport_write](#)
- [ns2usecs](#)
- [peek_next_entry](#)
- [poll_wait_pipe](#)
- [print_trace_bputs](#)
- [print_trace_fmt](#)
- [print_trace_fn](#)
- [print_trace_line](#)
- [print_trace_stack](#)
- [trace_consume](#)
- [trace_empty](#)
- [trace_find_next_entry_inc](#)

- [trace_print_timestamp](#)
- [tracing_open_pipe](#)
- [tracing_read_pipe](#)
- [tracing_release_pipe](#)
- [tracing_wait_pipe](#)

Macros defined

- [KTAP_TRACE_ITER](#)
- [TRACE_BUF_SIZE_DEFAULT](#)

Source code

```

1  /*
2  * kp_transport.c - ktap transport functionality
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <linux/debugfs.h>
23 #include <linux/ftrace_event.h>
24 #include <linux/stacktrace.h>
25 #include <linux/clocksource.h>
26 #include <asm/uaccess.h>
27 #include <linux/slab.h>
28 #include <linux/module.h>
29 #include <linux/kallsyms.h>
30 #include "../include/ktap_types.h"
31 #include "ktap.h"
32 #include "kp_events.h"
33 #include "kp_transport.h"
34
35 struct ktap_trace_iterator {
36     struct ring_buffer *buffer;
37     int print_timestamp;
38     void *private;
39
40     struct trace_iterator iter;
41 };
42
43 enum ktap_trace_type {
44     __TRACE_FIRST_TYPE = 0,
45
46     TRACE_FN = 1, /* must be same as ftrace definition in kernel */
47     TRACE_PRINT,
48     TRACE_BPUTS,
49     TRACE_STACK,
50     TRACE_USER_STACK,
51
52     __TRACE_LAST_TYPE,
53 };

```

```

54
55 #define KTAP_TRACE_ITER(iter) \
56     container_of(iter, struct ktap_trace_iterator, iter)
57
58 static
59 ssize_t _trace_seq_to_user(struct trace_seq *s, char __user *ubuf, size_t cnt)
60 {
61     int len;
62     int ret;
63
64     if (!cnt)
65         return 0;
66
67     if (s->len <= s->readpos)
68         return -EBUSY;
69
70     len = s->len - s->readpos;
71     if (cnt > len)
72         cnt = len;
73     ret = copy_to_user(ubuf, s->buffer + s->readpos, cnt);
74     if (ret == cnt)
75         return -EFAULT;
76
77     cnt -= ret;
78
79     s->readpos += cnt;
80     return cnt;
81 }
82
83 int _trace_seq_puts(struct trace_seq *s, const char *str)
84 {
85     int len = strlen(str);
86
87     if (s->full)
88         return 0;
89
90     if (len > ((PAGE_SIZE - 1) - s->len)) {
91         s->full = 1;
92         return 0;
93     }
94
95     memcpy(s->buffer + s->len, str, len);
96     s->len += len;
97
98     return len;
99 }
100
101 static int trace_empty(struct trace_iterator *iter)
102 {
103     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
104     int cpu;
105
106     for_each_online_cpu(cpu) {
107         if (!ring_buffer_empty_cpu(ktap_iter->buffer, cpu))
108             return 0;
109     }
110
111     return 1;
112 }
113
114 static void trace_consume(struct trace_iterator *iter)
115 {
116     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
117
118     ring_buffer_consume(ktap_iter->buffer, iter->cpu, &iter->ts,
119                        &iter->lost_events);
120 }
121
122 unsigned long long ns2usecs(cycle_t nsec)
123 {
124     nsec += 500;
125     do_div(nsec, 1000);
126     return nsec;
127 }
128
129 static int trace_print_timestamp(struct trace_iterator *iter)

```

```

130 {
131     struct trace_seq *s = &iter->seq;
132     unsigned long long t;
133     unsigned long secs, usec_rem;
134
135     t = ns2usecs(iter->ts);
136     usec_rem = do_div(t, USEC_PER_SEC);
137     secs = (unsigned long)t;
138
139     return trace_seq_printf(s, "%5lu.%06lu: ", secs, usec_rem);
140 }
141
142 /* todo: export kernel function ftrace_find_event in future, and make faster */
143 static struct trace_event *(*ftrace_find_event)(int type);
144
145 static enum print_line_t print_trace_fmt(struct trace_iterator *iter)
146 {
147     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
148     struct trace_entry *entry = iter->ent;
149     struct trace_event *ev;
150
151     ev = ftrace_find_event(entry->type);
152
153     if (ktap_iter->print_timestamp && !trace_print_timestamp(iter))
154         return TRACE_TYPE_PARTIAL_LINE;
155
156     if (ev) {
157         int ret = ev->funcs->trace(iter, 0, ev);
158
159         /* overwrite '\n' at the ending */
160         iter->seq.buffer[iter->seq.len - 1] = '\0';
161         iter->seq.len--;
162         return ret;
163     }
164
165     return TRACE_TYPE_PARTIAL_LINE;
166 }
167
168 static enum print_line_t print_trace_stack(struct trace_iterator *iter)
169 {
170     struct trace_entry *entry = iter->ent;
171     struct stack_trace trace;
172     char str[KSYM_SYMBOL_LEN];
173     int i;
174
175     trace.entries = (unsigned long *) (entry + 1);
176     trace.nr_entries = (iter->ent_size - sizeof(*entry)) /
177         sizeof(unsigned long);
178
179     if (!trace_seq_puts(&iter->seq, "<stack trace>\n"))
180         return TRACE_TYPE_PARTIAL_LINE;
181
182     for (i = 0; i < trace.nr_entries; i++) {
183         unsigned long p = trace.entries[i];
184
185         if (p == ULONG_MAX)
186             break;
187
188         sprint_symbol(str, p);
189         if (!trace_seq_printf(&iter->seq, " => %s\n", str))
190             return TRACE_TYPE_PARTIAL_LINE;
191     }
192
193     return TRACE_TYPE_HANDLED;
194 }
195
196 struct ktap_ftrace_entry {
197     struct trace_entry entry;
198     unsigned long ip;
199     unsigned long parent_ip;
200 };
201
202 static enum print_line_t print_trace_fn(struct trace_iterator *iter)
203 {
204     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
205     struct ktap_ftrace_entry *field = (struct ktap_ftrace_entry *)iter->ent;

```

```

206 char str[KSYM_SYMBOL_LEN];
207
208 if (ktap_iter->print_timestamp && !trace_print_timestamp(iter))
209     return TRACE_TYPE_PARTIAL_LINE;
210
211 sprint_symbol(str, field->ip);
212 if (!trace_seq_puts(&iter->seq, str))
213     return TRACE_TYPE_PARTIAL_LINE;
214
215 if (!trace_seq_puts(&iter->seq, " <- "))
216     return TRACE_TYPE_PARTIAL_LINE;
217
218 sprint_symbol(str, field->parent_ip);
219 if (!trace_seq_puts(&iter->seq, str))
220     return TRACE_TYPE_PARTIAL_LINE;
221
222 return TRACE_TYPE_HANDLED;
223 }
224
225 static enum print_line_t print_trace_bputs(struct trace_iterator *iter)
226 {
227     if (!trace_seq_puts(&iter->seq,
228         (const char *)*(unsigned long *)(&iter->ent + 1)))
229         return TRACE_TYPE_PARTIAL_LINE;
230
231     return TRACE_TYPE_HANDLED;
232 }
233
234 static enum print_line_t print_trace_line(struct trace_iterator *iter)
235 {
236     struct trace_entry *entry = iter->ent;
237     char *str = (char *)(&entry + 1);
238
239     if (entry->type == TRACE_PRINT) {
240         if (!trace_seq_printf(&iter->seq, "%s", str))
241             return TRACE_TYPE_PARTIAL_LINE;
242
243         return TRACE_TYPE_HANDLED;
244     }
245
246     if (entry->type == TRACE_BPUTS)
247         return print_trace_bputs(iter);
248
249     if (entry->type == TRACE_STACK)
250         return print_trace_stack(iter);
251
252     if (entry->type == TRACE_FN)
253         return print_trace_fn(iter);
254
255     return print_trace_fmt(iter);
256 }
257
258 static struct trace_entry *
259 peek_next_entry(struct trace_iterator *iter, int cpu, u64 *ts,
260     unsigned long *lost_events)
261 {
262     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
263     struct ring_buffer_event *event;
264
265     event = ring_buffer_peek(ktap_iter->buffer, cpu, ts, lost_events);
266     if (event) {
267         iter->ent_size = ring_buffer_event_length(event);
268         return ring_buffer_event_data(event);
269     }
270
271     return NULL;
272 }
273
274 static struct trace_entry *
275 __find_next_entry(struct trace_iterator *iter, int *ent_cpu,
276     unsigned long *missing_events, u64 *ent_ts)
277 {
278     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
279     struct ring_buffer *buffer = ktap_iter->buffer;
280     struct trace_entry *ent, *next = NULL;
281     unsigned long lost_events = 0, next_lost = 0;

```

```

282     u64 next_ts = 0, ts;
283     int next_cpu = -1;
284     int next_size = 0;
285     int cpu;
286
287     for_each_online_cpu(cpu) {
288         if (ring_buffer_empty_cpu(buffer, cpu))
289             continue;
290
291         ent = peek_next_entry(iter, cpu, &ts, &lost_events);
292         /*
293          * Pick the entry with the smallest timestamp:
294          */
295         if (ent && (!next || ts < next_ts)) {
296             next = ent;
297             next_cpu = cpu;
298             next_ts = ts;
299             next_lost = lost_events;
300             next_size = iter->ent_size;
301         }
302     }
303
304     iter->ent_size = next_size;
305
306     if (ent_cpu)
307         *ent_cpu = next_cpu;
308
309     if (ent_ts)
310         *ent_ts = next_ts;
311
312     if (missing_events)
313         *missing_events = next_lost;
314
315     return next;
316 }
317
318 /* Find the next real entry, and increment the iterator to the next entry */
319 static void *trace_find_next_entry_inc(struct trace_iterator *iter)
320 {
321     iter->ent = __find_next_entry(iter, &iter->cpu,
322                                 &iter->lost_events, &iter->ts);
323     if (iter->ent)
324         iter->idx++;
325
326     return iter->ent ? iter : NULL;
327 }
328
329 static void poll_wait_pipe(void)
330 {
331     set_current_state(TASK_INTERRUPTIBLE);
332     /* sleep for 100 msecs, and try again. */
333     schedule_timeout(HZ / 10);
334 }
335
336 static int tracing_wait_pipe(struct file *filp)
337 {
338     struct trace_iterator *iter = filp->private_data;
339     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
340     ktap_state_t *ks = ktap_iter->private;
341
342     while (trace_empty(iter)) {
343
344         if ((filp->f_flags & O_NONBLOCK)) {
345             return -EAGAIN;
346         }
347
348         mutex_unlock(&iter->mutex);
349
350         poll_wait_pipe();
351
352         mutex_lock(&iter->mutex);
353
354         if (G(ks)->wait_user && trace_empty(iter))
355             return -EINTR;
356     }
357

```

```

358     return 1;
359 }
360
361 static ssize_t
362 tracing_read_pipe(struct file *filp, char __user *ubuf, size_t cnt,
363                 loff_t *ppos)
364 {
365     struct trace_iterator *iter = filp->private_data;
366     ssize_t sret;
367
368     /* return any leftover data */
369     sret = trace_seq_to_user(&iter->seq, ubuf, cnt);
370     if (sret != -EBUSY)
371         return sret;
372
373     /*
374      * Avoid more than one consumer on a single file descriptor
375      * This is just a matter of traces coherency, the ring buffer itself
376      * is protected.
377     */
378     mutex_lock(&iter->mutex);
379
380 waitagain:
381     sret = tracing_wait_pipe(filp);
382     if (sret <= 0)
383         goto out;
384
385     /* stop when tracing is finished */
386     if (trace_empty(iter)) {
387         sret = 0;
388         goto out;
389     }
390
391     if (cnt >= PAGE_SIZE)
392         cnt = PAGE_SIZE - 1;
393
394     /* reset all but tr, trace, and overruns */
395     memset(&iter->seq, 0,
396           sizeof(struct trace_iterator) -
397           offsetof(struct trace_iterator, seq));
398     iter->pos = -1;
399
400 while (trace_find_next_entry_inc(iter) != NULL) {
401     enum print_line_t ret;
402     int len = iter->seq.len;
403
404     ret = print_trace_line(iter);
405     if (ret == TRACE_TYPE_PARTIAL_LINE) {
406         /* don't print partial lines */
407         iter->seq.len = len;
408         break;
409     }
410     if (ret != TRACE_TYPE_NO_CONSUME)
411         trace_consume(iter);
412
413     if (iter->seq.len >= cnt)
414         break;
415
416     /*
417      * Setting the full flag means we reached the trace_seq buffer
418      * size and we should leave by partial output condition above.
419      * One of the trace_seq_* functions is not used properly.
420     */
421     WARN_ONCE(iter->seq.full, "full flag set for trace type %d",
422              iter->ent->type);
423 }
424
425 /* Now copy what we have to the user */
426 sret = trace_seq_to_user(&iter->seq, ubuf, cnt);
427 if (iter->seq.readpos >= iter->seq.len)
428     trace_seq_init(&iter->seq);
429
430 /*
431  * If there was nothing to send to user, in spite of consuming trace
432  * entries, go back to wait for more entries.
433 */
434 if (sret == -EBUSY)

```

```

434     goto waitagain;
435
436 out:
437     mutex_unlock(&iter->mutex);
438
439     return sret;
440 }
441
442 static int tracing_open_pipe(struct inode *inode, struct file *filp)
443 {
444     struct ktap_trace_iterator *ktap_iter;
445     ktap_state_t *ks = inode->i_private;
446
447     /* create a buffer to store the information to pass to userspace */
448     ktap_iter = kzalloc(sizeof(*ktap_iter), GFP_KERNEL);
449     if (!ktap_iter)
450         return -ENOMEM;
451
452     ktap_iter->private = ks;
453     ktap_iter->buffer = G(ks)->buffer;
454     ktap_iter->print_timestamp = G(ks)->parm->print_timestamp;
455     mutex_init(&ktap_iter->iter.mutex);
456     filp->private_data = &ktap_iter->iter;
457
458     nonseekable_open(inode, filp);
459
460     return 0;
461 }
462
463 static int tracing_release_pipe(struct inode *inode, struct file *file)
464 {
465     struct trace_iterator *iter = file->private_data;
466     struct ktap_trace_iterator *ktap_iter = KTAP_TRACE_ITER(iter);
467
468     mutex_destroy(&iter->mutex);
469     kfree(ktap_iter);
470     return 0;
471 }
472
473 static const struct file_operations tracing_pipe_fops = {
474     .open      = tracing_open_pipe,
475     .read      = tracing_read_pipe,
476     .splice_read = NULL,
477     .release   = tracing_release_pipe,
478     .llseek   = no_llseek,
479 };
480
481 /*
482  * preempt disabled in ring_buffer_lock_reserve
483  *
484  * The implementation is similar with funtion __ftrace_trace_stack.
485  */
486 void kp_transport_print_kstack(ktap_state_t *ks, uint16_t depth, uint16_t skip)
487 {
488     struct ring_buffer *buffer = G(ks)->buffer;
489     struct ring_buffer_event *event;
490     struct trace_entry *entry;
491     int size;
492
493     size = depth * sizeof(unsigned long);
494     event = ring_buffer_lock_reserve(buffer, sizeof(*entry) + size);
495     if (!event) {
496         KTAP_STATS(ks)->events_missed += 1;
497         return;
498     } else {
499         struct stack_trace trace;
500
501         entry = ring_buffer_event_data(event);
502         tracing_generic_entry_update(entry, 0, 0);
503         entry->type = TRACE_STACK;
504
505         trace.nr_entries = 0;
506         trace.skip = skip;
507         trace.max_entries = depth;
508         trace.entries = (unsigned long *)(entry + 1);
509         save_stack_trace(&trace);

```

```

510         ring_buffer_unlock_commit(buffer, event);
511     }
512 }
513 }
514
515 void kp_transport_event_write(ktap_state_t *ks, struct ktap_event_data *e)
516 {
517     struct ring_buffer *buffer = G(ks)->buffer;
518     struct ring_buffer_event *event;
519     struct trace_entry *ev_entry = e->data->raw->data;
520     struct trace_entry *entry;
521     int entry_size = e->data->raw->size;
522
523     event = ring_buffer_lock_reserve(buffer, entry_size +
524                                     sizeof(struct ftrace_event_call *));
525     if (!event) {
526         KTAP_STATS(ks)->events_missed += 1;
527         return;
528     } else {
529         entry = ring_buffer_event_data(event);
530
531         memcpy(entry, ev_entry, entry_size);
532
533         ring_buffer_unlock_commit(buffer, event);
534     }
535 }
536
537 void kp_transport_write(ktap_state_t *ks, const void *data, size_t length)
538 {
539     struct ring_buffer *buffer = G(ks)->buffer;
540     struct ring_buffer_event *event;
541     struct trace_entry *entry;
542     int size;
543
544     size = sizeof(struct trace_entry) + length;
545
546     event = ring_buffer_lock_reserve(buffer, size);
547     if (!event) {
548         KTAP_STATS(ks)->events_missed += 1;
549         return;
550     } else {
551         entry = ring_buffer_event_data(event);
552
553         tracing_generic_entry_update(entry, 0, 0);
554         entry->type = TRACE_PRINT;
555         memcpy(entry + 1, data, length);
556
557         ring_buffer_unlock_commit(buffer, event);
558     }
559 }
560
561 /* general print function */
562 void kp_printf(ktap_state_t *ks, const char *fmt, ...)
563 {
564     char buff[1024];
565     va_list args;
566     int len;
567
568     va_start(args, fmt);
569     len = vsnprintf(buff, 1024, fmt, args);
570     va_end(args);
571
572     buff[len] = '\0';
573     kp_transport_write(ks, buff, len + 1);
574 }
575
576 void __kp_puts(ktap_state_t *ks, const char *str)
577 {
578     kp_transport_write(ks, str, strlen(str) + 1);
579 }
580
581 void __kp_bputs(ktap_state_t *ks, const char *str)
582 {
583     struct ring_buffer *buffer = G(ks)->buffer;
584     struct ring_buffer_event *event;
585     struct trace_entry *entry;

```

```

586     int size;
587
588     size = sizeof(struct trace_entry) + sizeof(unsigned long *);
589
590     event = ring_buffer_lock_reserve(buffer, size);
591     if (!event) {
592         KTAP_STATS(ks)->events_missed += 1;
593         return;
594     } else {
595         entry = ring_buffer_event_data(event);
596
597         tracing_generic_entry_update(entry, 0, 0);
598         entry->type = TRACE_BPUTS;
599         *(unsigned long *)(entry + 1) = (unsigned long)str;
600
601         ring_buffer_unlock_commit(buffer, event);
602     }
603 }
604
605 void kp_transport_exit(ktap_state_t *ks)
606 {
607     if (G(ks)->buffer)
608         ring_buffer_free(G(ks)->buffer);
609     debugfs_remove(G(ks)->trace_pipe_dentry);
610 }
611
612 #define TRACE_BUF_SIZE_DEFAULT    1441792UL /* 16384 * 88 (sizeof(entry)) */
613
614 int kp_transport_init(ktap_state_t *ks, struct dentry *dir)
615 {
616     struct ring_buffer *buffer;
617     struct dentry *dentry;
618     char filename[32] = {0};
619
620 #ifndef CONFIG_PPC64
621     ftrace_find_event = (void *)kallsyms_lookup_name(".ftrace_find_event");
622 #else
623     ftrace_find_event = (void *)kallsyms_lookup_name("ftrace_find_event");
624 #endif
625     if (!ftrace_find_event) {
626         printk("ktap: cannot lookup ftrace_find_event in kallsyms\n");
627         return -EINVAL;
628     }
629
630     buffer = ring_buffer_alloc(TRACE_BUF_SIZE_DEFAULT, RB_FL_OVERWRITE);
631     if (!buffer)
632         return -ENOMEM;
633
634     sprintf(filename, "trace_pipe_%d", (int)task_tgid_vnr(current));
635
636     dentry = debugfs_create_file(filename, 0444, dir,
637                                 ks, &tracing_pipe_fops);
638     if (!dentry) {
639         pr_err("ktapvm: cannot create trace_pipe file in debugfs\n");
640         ring_buffer_free(buffer);
641         return -1;
642     }
643
644     G(ks)->buffer = buffer;
645     G(ks)->trace_pipe_dentry = dentry;
646
647     return 0;
648 }
649

```

[One Level Up](#)

[Top Level](#)

runtime/kp_events.h - ktap

Data types defined

- [KTAP_EVENT_FIELD_TYPE](#)
- [KTAP_EVENT_TYPE](#)
- [ktap_event](#)
- [ktap_event_data](#)
- [ktap_event_field](#)

Macros defined

- [__KTAP_EVENTS_H__](#)

Source code

```
1 #ifndef \_\_KTAP\_EVENTS\_H\_\_
2 #define \_\_KTAP\_EVENTS\_H\_\_
3
4 #include <linux/ftrace_event.h>
5 #include <trace/syscall.h>
6 #include <trace/events/syscalls.h>
7 #include <linux/syscalls.h>
8 #include <linux/kprobes.h>
9
10 enum KTAP\_EVENT\_FIELD\_TYPE {
11     KTAP\_EVENT\_FIELD\_TYPE\_INVALID = 0, /* arg type not support yet */
12
13     KTAP\_EVENT\_FIELD\_TYPE\_INT,
14     KTAP\_EVENT\_FIELD\_TYPE\_LONG,
15     KTAP\_EVENT\_FIELD\_TYPE\_STRING,
16
17     KTAP\_EVENT\_FIELD\_TYPE\_REGISTER,
18     KTAP\_EVENT\_FIELD\_TYPE\_CONST,
19     KTAP\_EVENT\_FIELD\_TYPE\_NIL /* arg not exist */
20 };
21
22 struct ktap\_event\_field {
23     enum KTAP\_EVENT\_FIELD\_TYPE type;
24     int offset;
25 };
26
27 enum KTAP\_EVENT\_TYPE {
28     KTAP\_EVENT\_TYPE\_PERF,
29     KTAP\_EVENT\_TYPE\_TRACEPOINT,
30     KTAP\_EVENT\_TYPE\_SYSCALL\_ENTER,
31     KTAP\_EVENT\_TYPE\_SYSCALL\_EXIT,
32     KTAP\_EVENT\_TYPE\_KPROBE,
33 };
34
35 struct ktap\_event {
36     struct list_head list;
37     int type;
38     ktap\_state\_t *ks;
39     ktap\_func\_t *fn;
40     struct perf_event *perf;
41     int syscall_nr; /* for syscall event */
42     struct ktap\_event\_field fields[9]; /* arg1..arg9 */
43     ktap\_str\_t *name; /* intern probename string */
44
45     struct kprobe kp; /* kprobe event */
46 };
47
48 /* this structure allocate on stack */
```

```
49 struct ktap_event_data {
50     struct ktap_event *event;
51     struct perf_sample_data *data;
52     struct pt_regs *regs;
53     ktap_str_t *argstr; /* for cache argstr intern string */
54 };
55
56 int kp_events_init(ktap_state_t *ks);
57 void kp_events_exit(ktap_state_t *ks);
58
59 int kp_event_create(ktap_state_t *ks, struct perf_event_attr *attr,
60     struct task_struct *task, const char *filter,
61     ktap_func_t *fn);
62 int kp_event_create_tracepoint(ktap_state_t *ks, const char *event_name,
63     ktap_func_t *fn);
64
65 int kp_event_create_kprobe(ktap_state_t *ks, const char *event_name,
66     ktap_func_t *fn);
67 void kp_event_getarg(ktap_state_t *ks, ktap_val_t *ra, int idx);
68 const char *kp_event_tostr(ktap_state_t *ks);
69 const ktap_str_t *kp_event_stringify(ktap_state_t *ks);
70
71 #endif /* __KTAP_EVENTS_H */
```

[One Level Up](#)

[Top Level](#)

runtime/kp_events.c - ktap

Data types defined

- [ftrace_event_field](#)

Functions defined

- [call_probe_closure](#)
- [events_destroy](#)
- [get_fields](#)
- [init_event_fields](#)
- [kp_event_create](#)
- [kp_event_create_kprobe](#)
- [kp_event_getarg](#)
- [kp_event_stringify](#)
- [kp_event_tostr](#)
- [kp_events_exit](#)
- [kp_events_init](#)
- [perf_callback](#)
- [pre_handler_kprobe](#)

Source code

```
1 /*
2  * kp_events.c - ktap events management (registry, destroy, event callback)
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <linux/module.h>
23 #include <linux/ctype.h>
24 #include <linux/slab.h>
25 #include <linux/version.h>
26 #include <asm/syscall.h>
27 #include "../include/ktap_types.h"
28 #include "ktap.h"
29 #include "kp_obj.h"
```

```

30 #include "kp_str.h"
31 #include "kp_transport.h"
32 #include "kp_vm.h"
33 #include "kp_events.h"
34
35 const char *kp_event_tostr(ktap_state_t *ks)
36 {
37     struct ktap_event_data *e = ks->current_event;
38     struct ftrace_event_call *call;
39     struct trace_iterator *iter;
40     struct trace_event *ev;
41     enum print_line_t ret = TRACE_TYPE_NO_CONSUME;
42     static const char *dummy_msg = "argstr_not_available";
43
44     /* need to check current context is valid tracing context */
45     if (!ks->current_event) {
46         kp_error(ks, "cannot stringify event str in invalid context\n");
47         return NULL;
48     }
49
50     /*check if stringified before */
51     if (ks->current_event->argstr)
52         return getstr(ks->current_event->argstr);
53
54     /* timer event and raw tracepoint don't have associated argstr */
55     if (e->event->type == KTAP_EVENT_TYPE_PERF && e->event->perf->tp_event)
56         call = e->event->perf->tp_event;
57     else
58         return dummy_msg;
59
60     /* Simulate the iterator */
61
62     /*
63      * use temp percpu buffer as trace_iterator
64      * we cannot use same print_buffer because we may called from printf.
65      */
66     iter = kp_this_cpu_temp_buffer(ks);
67
68     trace_seq_init(&iter->seq);
69     iter->ent = e->data->raw->data;
70
71     ev = &(call->event);
72     if (ev)
73         ret = ev->funcs->trace(iter, 0, ev);
74
75     if (ret != TRACE_TYPE_NO_CONSUME) {
76         struct trace_seq *s = &iter->seq;
77         int len = s->len >= PAGE_SIZE ? PAGE_SIZE - 1 : s->len;
78
79         s->buffer[len] = '\0';
80         return &s->buffer[0];
81     }
82
83     return dummy_msg;
84 }
85
86 /* return string repr of 'argstr' */
87 const ktap_str_t *kp_event_stringify(ktap_state_t *ks)
88 {
89     const char *str;
90     ktap_str_t *ts;
91
92     /*check if stringified before */
93     if (ks->current_event->argstr)
94         return ks->current_event->argstr;
95
96     str = kp_event_tostr(ks);
97     if (!str)
98         return NULL;
99
100     ts = kp_str_newz(ks, str);
101     ks->current_event->argstr = ts;
102     return ts;
103 }
104
105 /*

```

```

106 * This definition should keep update with kernel/trace/trace.h
107 * TODD: export this struct in kernel
108 */
109 struct ftrace_event_field {
110     struct list_head    link;
111     const char          *name;
112     const char          *type;
113     int                 filter_type;
114     int                 offset;
115     int                 size;
116     int                 is_signed;
117 };
118
119 static struct list_head *get_fields(struct ftrace_event_call *event_call)
120 {
121     if (!event_call->class->get_fields)
122         return &event_call->class->fields;
123     return event_call->class->get_fields(event_call);
124 }
125
126 void kp_event_getarg(ktap_state_t *ks, ktap_val_t *ra, int idx)
127 {
128     struct ktap_event_data *e = ks->current_event;
129     struct ktap_event *event = e->event;
130     struct ktap_event_field *event_fields = &event->fields[idx];
131
132     switch (event_fields->type) {
133     case KTAP_EVENT_FIELD_TYPE_INT: {
134         struct trace_entry *entry = e->data->raw->data;
135         void *value = (unsigned char *)entry + event_fields->offset;
136         int n = *(int *)value;
137         set_number(ra, n);
138         return;
139     }
140     case KTAP_EVENT_FIELD_TYPE_LONG: {
141         struct trace_entry *entry = e->data->raw->data;
142         void *value = (unsigned char *)entry + event_fields->offset;
143         long n = *(long *)value;
144         set_number(ra, n);
145         return;
146     }
147     case KTAP_EVENT_FIELD_TYPE_STRING: {
148         struct trace_entry *entry = e->data->raw->data;
149         ktap_str_t *ts;
150         void *value = (unsigned char *)entry + event_fields->offset;
151         ts = kp_str_newz(ks, (char *)value);
152         if (ts)
153             set_string(ra, ts);
154         else
155             set_nil(ra);
156         return;
157     }
158     case KTAP_EVENT_FIELD_TYPE_CONST: {
159         set_number(ra, (ktap_number)event_fields->offset);
160         return;
161     }
162     case KTAP_EVENT_FIELD_TYPE_REGISTER: {
163         unsigned long *reg = (unsigned long *)((u8 *)e->regs +
164             event_fields->offset);
165         set_number(ra, *reg);
166         return;
167     }
168     case KTAP_EVENT_FIELD_TYPE_NIL:
169         set_nil(ra);
170         return;
171     case KTAP_EVENT_FIELD_TYPE_INVALID:
172         kp_error(ks, "the field type is not supported yet\n");
173         set_nil(ra);
174         return;
175     }
176 }
177
178 /* init all fields of event, for quick arg1..arg9 access */
179 static int init_event_fields(ktap_state_t *ks, struct ktap_event *event)
180 {
181     struct ftrace_event_call *event_call = event->perf->tp_event;

```

```

182 struct ktap\_event\_field *event_fields = &event->fields[0];
183 struct ftrace\_event\_field *field;
184 struct list_head *head;
185 int idx = 0, n = 0;
186
187 /* only init fields for tracepoint, not timer event */
188 if (!event_call)
189     return 0;
190
191 /* intern probename */
192 event->name = kp\_str\_newz(ks, event_call->name);
193 if (unlikely(!event->name))
194     return -ENOMEM;
195
196 head = get\_fields(event_call);
197 list_for_each_entry_reverse(field, head, link) {
198     if (n++ == 9) {
199         /*
200             * For some events have fields more than 9, just ignore
201             * those rest fields at present.
202             *
203             * TODO: support access all fields in tracepoint event
204             *
205             * Examples: mce:mce_record, ext4:ext4_writepages, ...
206             */
207         return 0;
208     }
209
210     event_fields[idx].offset = field->offset;
211
212     if (field->size == 4) {
213         event_fields[idx].type = KTAP_EVENT_FIELD_TYPE_INT;
214         idx++;
215         continue;
216     } else if (field->size == 8) {
217         event_fields[idx].type = KTAP_EVENT_FIELD_TYPE_LONG;
218         idx++;
219         continue;
220     }
221     if (!strncmp(field->type, "char", 4)) {
222         event_fields[idx].type = KTAP_EVENT_FIELD_TYPE_STRING;
223         idx++;
224         continue;
225     }
226
227     /* TODO: add more type check */
228     event_fields[idx++].type = KTAP_EVENT_FIELD_TYPE_INVALID;
229 }
230
231 /* init all rest fields as NIL */
232 while (idx < 9)
233     event_fields[idx++].type = KTAP_EVENT_FIELD_TYPE_NIL;
234
235 return 0;
236 }
237
238 static inline void call\_probe\_closure(ktap\_state\_t *mainthread,
239                                     ktap\_func\_t *fn,
240                                     struct ktap\_event\_data *e, int rctx)
241 {
242     ktap\_state\_t *ks;
243     ktap\_val\_t *func;
244
245     ks = kp\_vm\_new\_thread(mainthread, rctx);
246     set\_func(ks->top, fn);
247     func = ks->top;
248     incr\_top(ks);
249
250     ks->current_event = e;
251
252     kp\_vm\_call(ks, func, 0);
253
254     ks->current_event = NULL;
255     kp\_vm\_exit\_thread(ks);
256 }
257

```

```

258 /*
259  * Callback tracing function for perf event subsystem.
260  *
261  * make ktap reentrant, don't disable irq in callback function,
262  * same as perf and ftrace. to make reentrant, we need some
263  * percpu data to be context isolation(irq/sirq/nmi/process)
264  *
265  * The recursion checking in here is mainly purpose for avoiding
266  * corrupt ktap_state_t with timer closure callback. For tracepoint
267  * recursion, perf core already handle it.
268  *
269  * Note tracepoint handler is calling with rcu_read_lock.
270  */
271 static void perf_callback(struct perf_event *perf_event,
272                          struct perf_sample_data *data,
273                          struct pt_regs *regs)
274 {
275     struct ktap_event *event;
276     struct ktap_event_data e;
277     ktap_state_t *ks;
278     int rctx;
279
280     event = perf_event->overflow_handler_context;
281     ks = event->ks;
282
283     if (unlikely(ks->stop))
284         return;
285
286     rctx = get_recursion_context(ks);
287     if (unlikely(rctx < 0))
288         return;
289
290     e.event = event;
291     e.data = data;
292     e.regs = regs;
293     e.argstr = NULL;
294
295     call_probe_closure(ks, event->fn, &e, rctx);
296
297     put_recursion_context(ks, rctx);
298 }
299
300 /*
301  * Generic ktap event creation function (based on perf callback)
302  * purpose for tracepoints/kprobe/uprobe/profile-timer/hw_breakpoint/pmu.
303  */
304 int kp_event_create(ktap_state_t *ks, struct perf_event_attr *attr,
305                   struct task_struct *task, const char *filter,
306                   ktap_func_t *fn)
307 {
308     struct ktap_event *event;
309     struct perf_event *perf_event;
310     void *callback = perf_callback;
311     int cpu, ret;
312
313     if (G(ks)->parm->dry_run)
314         callback = NULL;
315
316     /*
317      * don't tracing until ktap_wait, the reason is:
318      * 1). some event may hit before apply filter
319      * 2). more simple to manage tracing thread
320      * 3). avoid race with mainthread.
321      *
322      * Another way to do this is make attr.disabled as 1, then use
323      * perf_event_enable after filter apply, however, perf_event_enable
324      * was not exported in kernel older than 3.3, so we drop this method.
325      */
326     ks->stop = 1;
327
328     for_each_cpu(cpu, G(ks)->cpumask) {
329         event = kzalloc(sizeof(struct ktap_event), GFP_KERNEL);
330         if (!event)
331             return -ENOMEM;
332
333         event->type = KTAP_EVENT_TYPE_PERF;

```

```

334     event->ks = ks;
335     event->fn = fn;
336     perf_event = perf_event_create_kernel_counter(attr, cpu, task,
337                                                 callback, event);
338     if (IS_ERR(perf_event)) {
339         int err = PTR_ERR(perf_event);
340         kp_error(ks, "unable register perf event: "
341                 "[cpu: %d; id: %d; err: %d]\n",
342                 cpu, attr->config, err);
343         kfree(event);
344         return err;
345     }
346
347     if (attr->type == PERF_TYPE_TRACEPOINT) {
348         const char *name = perf_event->tp_event->name;
349         kp_verbose_printf(ks, "enable perf event: "
350                             "[cpu: %d; id: %d; name: %s; "
351                             "filter: %s; pid: %d]\n",
352                             cpu, attr->config, name, filter,
353                             task ? task_tgid_vnr(task) : -1);
354     } else if (attr->type == PERF_TYPE_SOFTWARE &&
355               attr->config == PERF_COUNT_SW_CPU_CLOCK) {
356         kp_verbose_printf(ks, "enable profile event: "
357                             "[cpu: %d; sample_period: %d]\n",
358                             cpu, attr->sample_period);
359     } else {
360         kp_verbose_printf(ks, "unknown perf event type\n");
361     }
362
363     event->perf = perf_event;
364     INIT_LIST_HEAD(&event->list);
365     list_add_tail(&event->list, &G(ks)->events_head);
366
367     if (init_event_fields(ks, event)) {
368         kp_error(ks, "unable init event fields id %d\n",
369                 attr->config);
370         perf_event_release_kernel(event->perf);
371         list_del(&event->list);
372         kfree(event);
373         return ret;
374     }
375
376     if (!filter)
377         continue;
378
379     ret = kp_ftrace_profile_set_filter(perf_event, attr->config,
380                                       filter);
381     if (ret) {
382         kp_error(ks, "unable set event filter: "
383                 "[id: %d; filter: %s; ret: %d]\n",
384                 attr->config, filter, ret);
385         perf_event_release_kernel(event->perf);
386         list_del(&event->list);
387         kfree(event);
388         return ret;
389     }
390 }
391
392 return 0;
393 }
394
395 /*
396  * tracepoint_probe_register functions changed prototype by introduce
397  * 'struct tracepoint', this cause hard to refer tracepoint by name.
398  * And these ktap raw tracepoint interface is not courage to use, so disable
399  * it now.
400  */
401 #if 0
402 /*
403  * Ignore function proto in here, just use first argument.
404  */
405 static void probe_callback(void *__data)
406 {
407     struct ktap_event *event = __data;
408     ktap_state_t *ks = event->ks;
409     struct ktap_event_data e;

```

```

410     struct pt_regs regs; /* pt_regs maybe is large for stack */
411     int rctx;
412
413     if (unlikely(ks->stop))
414         return;
415
416     rctx = get\_recursion\_context(ks);
417     if (unlikely(rctx < 0))
418         return;
419
420     perf_fetch_caller_regs(&regs);
421
422     e.event = event;
423     e.regs = &regs;
424     e.argstr = NULL;
425
426     call\_probe\_closure(ks, event->fn, &e, rctx);
427
428     put\_recursion\_context(ks, rctx);
429 }
430
431 /*
432  * syscall events handling
433  */
434
435 static DEFINE_MUTEX(syscall_trace_lock);
436 static DECLARE_BITMAP(enabled_enter_syscalls, NR_syscalls);
437 static DECLARE_BITMAP(enabled_exit_syscalls, NR_syscalls);
438 static int sys_refcount_enter;
439 static int sys_refcount_exit;
440
441 static int get_syscall_num(const char *name)
442 {
443     int i;
444
445     for (i = 0; i < NR_syscalls; i++) {
446         if (syscalls\_metadata[i] &&
447             !strcmp(name, syscalls\_metadata[i]->name + 4))
448             return i;
449     }
450     return -1;
451 }
452
453 static void trace_syscall_enter(void *data, struct pt_regs *regs, long id)
454 {
455     struct ktap\_event *event = data;
456     ktap\_state\_t *ks = event->ks;
457     struct ktap\_event\_data e;
458     int syscall_nr;
459     int rctx;
460
461     if (unlikely(ks->stop))
462         return;
463
464     syscall_nr = syscall_get_nr(current, regs);
465     if (unlikely(syscall_nr < 0))
466         return;
467     if (!test_bit(syscall_nr, enabled_enter_syscalls))
468         return;
469
470     rctx = get\_recursion\_context(ks);
471     if (unlikely(rctx < 0))
472         return;
473
474     e.event = event;
475     e.regs = regs;
476     e.argstr = NULL;
477
478     call\_probe\_closure(ks, event->fn, &e, rctx);
479
480     put\_recursion\_context(ks, rctx);
481 }
482
483 static void trace_syscall_exit(void *data, struct pt_regs *regs, long id)
484 {
485     struct ktap\_event *event = data;

```

```

486 ktap\_state\_t *ks = event->ks;
487 struct ktap\_event\_data e;
488 int syscall_nr;
489 int rctx;
490
491 syscall_nr = syscall_get_nr(current, regs);
492 if (unlikely(syscall_nr < 0))
493     return;
494 if (!test\_bit(syscall_nr, enabled\_exit\_syscalls))
495     return;
496
497 if (unlikely(ks->stop))
498     return;
499
500 rctx = get\_recursion\_context(ks);
501 if (unlikely(rctx < 0))
502     return;
503
504 e.event = event;
505 e.regs = regs;
506 e.argstr = NULL;
507
508 call\_probe\_closure(ks, event->fn, &e, rctx);
509
510 put\_recursion\_context(ks, rctx);
511 }
512
513 /* called in dry-run mode, purpose for compare overhead with normal vm call */
514 static void dry_run_callback(void *data, struct pt\_regs *regs, long id)
515 {
516
517 }
518
519 static void init_syscall_event_fields(struct ktap\_event *event, int is_enter)
520 {
521     struct ftrace\_event\_call *event_call;
522     struct ktap\_event\_field *event_fields = &event->fields[0];
523     struct syscall\_metadata *meta = syscalls\_metadata[event->syscall_nr];
524     int idx = 0;
525
526     event_call = is_enter ? meta->enter_event : meta->exit_event;
527
528     event_fields[0].type = KTAP\_EVENT\_FIELD\_TYPE\_CONST;
529     event_fields[0].offset = event->syscall_nr;
530
531     if (!is_enter) {
532 #ifdef CONFIG\_X86\_64
533         event_fields[1].type = KTAP\_EVENT\_FIELD\_TYPE\_REGISTER;
534         event_fields[1].offset = offsetof(struct pt\_regs, ax);
535 #endif
536         return;
537     }
538
539     while (idx++ < meta->nb_args) {
540         event_fields[idx].type = KTAP\_EVENT\_FIELD\_TYPE\_REGISTER;
541 #ifdef CONFIG\_X86\_64
542         switch (idx) {
543             case 1:
544                 event_fields[idx].offset = offsetof(struct pt\_regs, di);
545                 break;
546             case 2:
547                 event_fields[idx].offset = offsetof(struct pt\_regs, si);
548                 break;
549             case 3:
550                 event_fields[idx].offset = offsetof(struct pt\_regs, dx);
551                 break;
552             case 4:
553                 event_fields[idx].offset =
554                     offsetof(struct pt\_regs, r10);
555                 break;
556             case 5:
557                 event_fields[idx].offset = offsetof(struct pt\_regs, r8);
558                 break;
559             case 6:
560                 event_fields[idx].offset = offsetof(struct pt\_regs, r9);
561                 break;

```

```

562     }
563 #else
564 #warning "don't support syscall tracepoint event register access in this arch, use 'trace syscalls:* {}'
instead"
565     break;
566 #endif
567 }
568
569 /* init all rest fields as NIL */
570 while (idx < 9)
571     event_fields[idx++].type = KTAP_EVENT_FIELD_TYPE_NIL;
572 }
573
574 static int syscall_event_register(ktap_state_t *ks, const char *event_name,
575                                 struct ktap_event *event)
576 {
577     int syscall_nr = 0, is_enter = 0;
578     void *callback = NULL;
579     int ret = 0;
580
581     if (!strncmp(event_name, "sys_enter_", 10)) {
582         is_enter = 1;
583         event->type = KTAP_EVENT_TYPE_SYSCALL_ENTER;
584         syscall_nr = get_syscall_num(event_name + 10);
585         callback = trace_syscall_enter;
586     } else if (!strncmp(event_name, "sys_exit_", 9)) {
587         is_enter = 0;
588         event->type = KTAP_EVENT_TYPE_SYSCALL_EXIT;
589         syscall_nr = get_syscall_num(event_name + 9);
590         callback = trace_syscall_exit;
591     }
592
593     if (G(ks)->parm->dry_run)
594         callback = dry_run_callback;
595
596     if (syscall_nr < 0)
597         return -1;
598
599     event->syscall_nr = syscall_nr;
600
601     init_syscall_event_fields(event, is_enter);
602
603     mutex_lock(&syscall_trace_lock);
604     if (is_enter) {
605         if (!sys_refcount_enter)
606             ret = register_trace_sys_enter(callback, event);
607         if (!ret) {
608             set_bit(syscall_nr, enabled_enter_syscalls);
609             sys_refcount_enter++;
610         }
611     } else {
612         if (!sys_refcount_exit)
613             ret = register_trace_sys_exit(callback, event);
614         if (!ret) {
615             set_bit(syscall_nr, enabled_exit_syscalls);
616             sys_refcount_exit++;
617         }
618     }
619     mutex_unlock(&syscall_trace_lock);
620
621     return ret;
622 }
623
624 static int syscall_event_unregister(ktap_state_t *ks, struct ktap_event *event)
625 {
626     int ret = 0;
627     void *callback;
628
629     if (event->type == KTAP_EVENT_TYPE_SYSCALL_ENTER)
630         callback = trace_syscall_enter;
631     else
632         callback = trace_syscall_exit;
633
634     if (G(ks)->parm->dry_run)
635         callback = dry_run_callback;
636 }

```

```

637 mutex_lock(&syscall_trace_lock);
638 if (event->type == KTAP_EVENT_TYPE_SYSCALL_ENTER){
639     sys_refcount_enter--;
640     clear_bit(event->syscall_nr, enabled_enter_syscalls);
641     if (!sys_refcount_enter)
642         unregister_trace_sys_enter(callback, event);
643 } else {
644     sys_refcount_exit--;
645     clear_bit(event->syscall_nr, enabled_exit_syscalls);
646     if (!sys_refcount_exit)
647         unregister_trace_sys_exit(callback, event);
648 }
649 mutex_unlock(&syscall_trace_lock);
650
651 return ret;
652 }
653
654 /*
655  * Register tracepoint event directly, not based on perf callback
656  *
657  * This tracing method would be more faster than perf callback,
658  * because it won't need to write trace data into any temp buffer,
659  * and code path is much shorter than perf callback.
660  */
661 int kprobe_create_tracepoint(ktap_state_t *ks, const char *event_name,
662                             ktap_func_t *fn)
663 {
664     struct ktap_event *event;
665     void *callback = probe_callback;
666     int is_syscall = 0;
667     int ret;
668
669     if (G(ks)->parm->dry_run)
670         callback = NULL;
671
672     if (!strncmp(event_name, "sys_enter_", 10) ||
673         !strncmp(event_name, "sys_exit_", 9))
674         is_syscall = 1;
675
676     event = kzalloc(sizeof(struct ktap_event), GFP_KERNEL);
677     if (!event)
678         return -ENOMEM;
679
680     event->ks = ks;
681     event->fn = fn;
682     event->name = kp_str_newz(ks, event_name);
683     if (unlikely(!event->name)) {
684         kfree(event);
685         return -ENOMEM;
686     }
687
688     INIT_LIST_HEAD(&event->list);
689     list_add_tail(&event->list, &G(ks)->events_head);
690
691     if (is_syscall) {
692         ret = syscall_event_register(ks, event_name, event);
693     } else {
694         event->type = KTAP_EVENT_TYPE_TRACEPOINT;
695         ret = tracepoint_probe_register(event_name, callback, event);
696     }
697
698     if (ret) {
699         kp_error(ks, "register tracepoint %s failed, ret: %d\n",
700                 event_name, ret);
701         list_del(&event->list);
702         kfree(event);
703         return ret;
704     }
705     return 0;
706 }
707
708 #endif
709
710 /* kprobe handler */
711 static int __kprobes pre_handler_kprobe(struct kprobe *p, struct pt_regs *regs)
712 {

```

```

713     struct ktap_event *event = container_of(p, struct ktap_event, kp);
714     ktap_state_t *ks = event->ks;
715     struct ktap_event_data e;
716     int rctx;
717
718     if (unlikely(ks->stop))
719         return 0;
720
721     rctx = get_recursion_context(ks);
722     if (unlikely(rctx < 0))
723         return 0;
724
725     e.event = event;
726     e.regs = regs;
727     e.argstr = NULL;
728
729     call_probe_closure(ks, event->fn, &e, rctx);
730
731     put_recursion_context(ks, rctx);
732     return 0;
733 }
734
735 /*
736  * Register kprobe event directly, not based on perf callback
737  *
738  * This tracing method would be more faster than perf callback,
739  * because it won't need to write trace data into any temp buffer,
740  * and code path is much shorter than perf callback.
741  */
742 int kp_event_create_kprobe(ktap_state_t *ks, const char *event_name,
743                          ktap_func_t *fn)
744 {
745     struct ktap_event *event;
746     void *callback = pre_handler_kprobe;
747     int ret;
748
749     if (G(ks)->parm->dry_run)
750         callback = NULL;
751
752     event = kzalloc(sizeof(struct ktap_event), GFP_KERNEL);
753     if (!event)
754         return -ENOMEM;
755
756     event->ks = ks;
757     event->fn = fn;
758     event->name = kp_str_newz(ks, event_name);
759     if (unlikely(!event->name)) {
760         kfree(event);
761         return -ENOMEM;
762     }
763
764     INIT_LIST_HEAD(&event->list);
765     list_add_tail(&event->list, &G(ks)->events_head);
766
767     event->type = KTAP_EVENT_TYPE_KPROBE;
768
769     event->kp.symbol_name = event_name;
770     event->kp.pre_handler = callback;
771     ret = register_kprobe(&event->kp);
772     if (ret) {
773         kp_error(ks, "register kprobe event %s failed, ret: %d\n",
774                event_name, ret);
775         list_del(&event->list);
776         kfree(event);
777         return ret;
778     }
779     return 0;
780 }
781
782
783 static void events_destroy(ktap_state_t *ks)
784 {
785     struct ktap_event *event;
786     struct list_head *tmp, *pos;
787     struct list_head *head = &G(ks)->events_head;
788

```

```

789 list_for_each(pos, head) {
790     event = container_of(pos, struct ktap_event,
791                          list);
792     if (event->type == KTAP_EVENT_TYPE_PERF)
793         perf_event_release_kernel(event->perf);
794 #if 0
795     else if (event->type == KTAP_EVENT_TYPE_TRACEPOINT)
796         tracepoint_probe_unregister(getstr(event->name),
797                                    probe_callback, event);
798     else if (event->type == KTAP_EVENT_TYPE_SYSCALL_ENTER ||
799            event->type == KTAP_EVENT_TYPE_SYSCALL_EXIT )
800         syscall_event_unregister(ks, event);
801 #endif
802     else if (event->type == KTAP_EVENT_TYPE_KPROBE)
803         unregister_kprobe(&event->kp);
804     }
805     /*
806     * Ensure our callback won't be called anymore. The buffers
807     * will be freed after that.
808     */
809     tracepoint_synchronize_unregister();
810
811     list_for_each_safe(pos, tmp, head) {
812         event = container_of(pos, struct ktap_event,
813                             list);
814         list_del(&event->list);
815         kfree(event);
816     }
817 }
818
819 void kp_events_exit(ktap_state_t *ks)
820 {
821     if (!G(ks)->trace_enabled)
822         return;
823
824     events_destroy(ks);
825
826     /* call trace_end_closure after all event unregistered */
827     if ((G(ks)->state != KTAP_ERROR) && G(ks)->trace_end_closure) {
828         G(ks)->state = KTAP_TRACE_END;
829         set_func(ks->top, G(ks)->trace_end_closure);
830         incr_top(ks);
831         kp_vm_call(ks, ks->top - 1, 0);
832         G(ks)->trace_end_closure = NULL;
833     }
834
835     G(ks)->trace_enabled = 0;
836 }
837
838 int kp_events_init(ktap_state_t *ks)
839 {
840     G(ks)->trace_enabled = 1;
841     return 0;
842 }
843

```

[One Level Up](#)

[Top Level](#)

runtime/kp_str.h - ktap

Macros defined

- [__KTAP_STR_H__](#)
- [kp_str_newz](#)

Source code

```
1 #ifndef __KTAP_STR_H__
2 #define __KTAP_STR_H__
3
4 int kp_str_cmp(const ktap_str_t *ls, const ktap_str_t *rs);
5 int kp_str_resize(ktap_state_t *ks, int newmask);
6 void kp_str_freeall(ktap_state_t *ks);
7 ktap_str_t * kp_str_new(ktap_state_t *ks, const char *str, size_t len);
8
9 #define kp_str_newz(ks, s)    (kp_str_new(ks, s, strlen(s)))
10
11 #include <linux/trace_seq.h>
12 int kp_str_fmt(ktap_state_t *ks, struct trace_seq *seq);
13
14 #endif /* __KTAP_STR_H__ */
```

runtime/kp_vm.h - ktap

Functions defined

- [kp_vm_exit_thread](#)
- [kp_vm_new_thread](#)
- [kp_vm_try_to_exit](#)

Macros defined

- [__KTAP_VM_H__](#)

Source code

```
1 #ifndef __KTAP_VM_H__
2 #define __KTAP_VM_H__
3
4 #include "kp_obj.h"
5
6 void kp_vm_call_proto(ktap_state_t *ks, ktap_proto_t *pt);
7 void kp_vm_call(ktap_state_t *ks, StkId func, int nresults);
8 int kp_vm_validate_code(ktap_state_t *ks, ktap_proto_t *pt, ktap_val_t *base);
9 void kp_vm_exit(ktap_state_t *ks);
10 ktap_state_t *kp_vm_new_state(ktap_option_t *parm, struct dentry *dir);
11 void kp_optimize_code(ktap_state_t *ks, int level, ktap_proto_t *f);
12 int kp_vm_register_lib(ktap_state_t *ks, const char *libname,
13                       const ktap_libfunc_t *funcs);
14
15
16 static __always_inline
17 ktap_state_t *kp_vm_new_thread(ktap_state_t *mainthread, int rctx)
18 {
19     ktap_state_t *ks;
20
21     ks = kp_this_cpu_state(mainthread, rctx);
22     ks->top = ks->stack;
23     return ks;
24 }
25
26 static __always_inline
27 void kp_vm_exit_thread(ktap_state_t *ks)
28 {
29 }
30
31 /*
32  * This function only tell ktapvm this thread want to exit,
33  * let mainthread handle real exit work later.
34  */
35 static __always_inline
36 void kp_vm_try_to_exit(ktap_state_t *ks)
37 {
38     G(ks)->mainthread->stop = 1;
39     G(ks)->state = KTAP_EXIT;
40 }
41
42
43 #endif /* __KTAP_VM_H__ */
```

runtime/kp_vm.c - ktap

Global variables defined

- [bc_names](#)

Functions defined

- [adjust_varargs](#)
- [check_hot_loop](#)
- [checkstack](#)
- [findupval](#)
- [free_preserved_data](#)
- [func_closeuv](#)
- [func_new](#)
- [func_new_empty](#)
- [gfunc_add](#)
- [gfunc_get](#)
- [gfunc_getidx](#)
- [init_arguments](#)
- [init_preserved_data](#)
- [init_registry](#)
- [kp_freeupval](#)
- [kp_vm_call](#)
- [kp_vm_call_proto](#)
- [kp_vm_exit](#)
- [kp_vm_new_state](#)
- [kp_vm_register_lib](#)
- [kp_vm_validate_code](#)
- [poscall](#)
- [sl_wait_task_exit_actor](#)
- [sl_wait_task_pause_actor](#)
- [sleep_loop](#)
- [str_concat](#)
- [unlinkupval](#)

- [wait_user_completion](#)
- [wait_user_interrupt](#)

Macros defined

- [ALLOC_PERCPU](#)
- [BCNAME](#)
- [BCNAME](#)
- [BCNAME](#)
- [BCNAME](#)
- [DISPATCH](#)
- [KTAP_MIN_RESERVED_STACK_SIZE](#)
- [KTAP_PERCPU_BUFFER_SIZE](#)
- [KTAP_STACK_SIZE](#)
- [KTAP_STACK_SIZE_BYTES](#)
- [NUMADD](#)
- [NUMDIV](#)
- [NUMEQ](#)
- [NUMLE](#)
- [NUMLT](#)
- [NUMMOD](#)
- [NUMMUL](#)
- [NUMSUB](#)
- [NUMUNM](#)
- [RA](#)
- [RA](#)
- [RB](#)
- [RB](#)
- [RC](#)
- [RC](#)
- [RD](#)
- [RD](#)
- [RKD](#)
- [SIZE_KTAP_FUNC](#)
- [arith_NV](#)

- [arith_VN](#)
- [arith_VV](#)
- [dojump](#)
- [donextjump](#)

Source code

```

1  /*
2  * kp_vm.c - ktap script virtual machine in Linux kernel
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Adapted from luajit and lua interpreter.
9  * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include <linux/slab.h>
27 #include <linux/ftrace_event.h>
28 #include <linux/signal.h>
29 #include <linux/sched.h>
30 #include <linux/uaccess.h>
31 #include "../include/ktap_types.h"
32 #include "../include/ktap_bc.h"
33 #include "../include/ktap_ffi.h"
34 #include "ktap.h"
35 #include "kp_obj.h"
36 #include "kp_str.h"
37 #include "kp_mempool.h"
38 #include "kp_tab.h"
39 #include "kp_transport.h"
40 #include "kp_vm.h"
41 #include "kp_events.h"
42
43 #define KTAP_MIN_RESERVED_STACK_SIZE 20
44 #define KTAP_STACK_SIZE 120 /* enlarge this value for big stack */
45 #define KTAP_STACK_SIZE_BYTES (KTAP_STACK_SIZE * sizeof(ktap_val_t))
46
47 #define KTAP_PERCPU_BUFFER_SIZE (3 * PAGE_SIZE)
48
49 static ktap_cfunction gfunc_get(ktap_state_t *ks, int idx);
50 static int gfunc_getidx(ktap_global_state_t *g, ktap_cfunction cfunc);
51
52 static ktap_str_t *str_concat(ktap_state_t *ks, StkId top, int start, int end)
53 {
54     int i, len = 0;
55     ktap_str_t *ts;
56     char *ptr, *buffer;
57
58     for (i = start; i <= end; i++) {
59         if (!is_string(top + i)) {
60             kp_error(ks, "cannot concat non-string\n");
61             return NULL;
62         }
63     }

```

```

64     len += rawtsvalue(top + i)->len;
65 }
66
67 if (len >= KTAP_PERCPU_BUFFER_SIZE) {
68     kp_error(ks, "Error: too long string concatenation\n");
69     return NULL;
70 }
71
72 preempt_disable_notrace();
73
74 buffer = kp_this_cpu_print_buffer(ks);
75 ptr = buffer;
76
77 for (i = start; i <= end; i++) {
78     int len = rawtsvalue(top + i)->len;
79     strncpy(ptr, svalue(top + i), len);
80     ptr += len;
81 }
82 ts = kp_str_new(ks, buffer, len);
83
84 preempt_enable_notrace();
85
86 return ts;
87 }
88
89 static ktap_upval_t *findupval(ktap_state_t *ks, StkId slot)
90 {
91     ktap_global_state_t *g = G(ks);
92     ktap_upval_t **pp = &ks->openupval;
93     ktap_upval_t *p;
94     ktap_upval_t *uv;
95
96     while (*pp != NULL && (p = *pp)->v >= slot) {
97         if (p->v == slot) { /* found a corresponding upvalue? */
98             return p;
99         }
100        pp = (ktap_upval_t **)&p->nextgc;
101    }
102
103    /* not found: create a new one */
104    uv = (ktap_upval_t *)kp_malloc(ks, sizeof(ktap_upval_t));
105    if (!uv)
106        return NULL;
107    uv->gct = ~KTAP_TUPVAL;
108    uv->closed = 0; /* still open */
109    uv->v = slot; /* current value lives in the stack */
110    /* Insert into sorted list of open upvalues. */
111    uv->nextgc = (ktap_obj_t *)*pp;
112    *pp = uv;
113    uv->prev = &g->uvhead; /* double link it in `uvhead' list */
114    uv->next = g->uvhead.next;
115    uv->next->prev = uv;
116    g->uvhead.next = uv;
117    return uv;
118 }
119
120 static void unlinkupval(ktap_upval_t *uv)
121 {
122     uv->next->prev = uv->prev; /* remove from `uvhead' list */
123     uv->prev->next = uv->next;
124 }
125
126 void kp_freeupval(ktap_state_t *ks, ktap_upval_t *uv)
127 {
128     if (!uv->closed) /* is it open? */
129         unlinkupval(uv); /* remove from open list */
130     kp_free(ks, uv); /* free upvalue */
131 }
132
133 /* close upvals */
134 static void func_closeuv(ktap_state_t *ks, StkId level)
135 {
136     ktap_upval_t *uv;
137     ktap_global_state_t *g = G(ks);
138     while (ks->openupval != NULL &&
139         (uv = ks->openupval)->v >= level) {

```

```

140     ktap_obj_t *o = obj2gco(uv);
141     /* remove from 'open' list */
142     ks->openupval = (ktap_upval_t *)uv->nextgc;
143     unlinkupval(uv); /* remove upvalue from 'uvhead' list */
144     set_obj(&uv->tv, uv->v); /* move value to upvalue slot */
145     uv->v = &uv->tv; /* now current value lives here */
146     uv->closed = 1;
147     gch(o)->nextgc = g->allgc; /* link upvalue into 'allgc' list */
148     g->allgc = o;
149 }
150 }
151
152 #define SIZE_KTAP_FUNC(n) (sizeof(ktap_func_t) - sizeof(ktap_obj_t *) + \
153     sizeof(ktap_obj_t *) * (n))
154 static ktap_func_t *func_new_empty(ktap_state_t *ks, ktap_proto_t *pt)
155 {
156     ktap_func_t *fn;
157
158     /* only mainthread can create new function */
159     if (ks != G(ks)->mainthread) {
160         kp_error(ks, "only mainthread can create function\n");
161         return NULL;
162     }
163
164     fn = (ktap_func_t *)kp_obj_new(ks, SIZE_KTAP_FUNC(pt->sizeuv));
165     if (!fn)
166         return NULL;
167     fn->gct = ~KTAP_TFUNC;
168     fn->nupvalues = 0; /* Set to zero until upvalues are initialized. */
169     fn->pc = proto_bc(pt);
170     fn->p = pt;
171
172     return fn;
173 }
174
175 static ktap_func_t *func_new(ktap_state_t *ks, ktap_proto_t *pt,
176     ktap_func_t *parent, ktap_val_t *base)
177 {
178     ktap_func_t *fn;
179     int nuv = pt->sizeuv, i;
180
181     fn = func_new_empty(ks, pt);
182     if (!fn)
183         return NULL;
184
185     fn->nupvalues = nuv;
186     for (i = 0; i < nuv; i++) {
187         uint32_t v = proto_uv(pt)[i];
188         ktap_upval_t *uv;
189
190         if (v & PROTO_UV_LOCAL) {
191             uv = findupval(ks, base + (v & 0xff));
192             if (!uv)
193                 return NULL;
194             uv->immutable = ((v / PROTO_UV_IMMUTABLE) & 1);
195         } else {
196             uv = parent->upvals[v];
197         }
198         fn->upvals[i] = uv;
199     }
200     return fn;
201 }
202
203 static inline int checkstack(ktap_state_t *ks, int n)
204 {
205     if (unlikely(ks->stack_last - ks->top <= n)) {
206         kp_error(ks, "stack overflow, please enlarge stack size\n");
207         return -1;
208     }
209     return 0;
210 }
211
212 static StkId adjust_varargs(ktap_state_t *ks, ktap_proto_t *p, int actual)
213 {
214     int i;
215     int nfixargs = p->numparams;

```

```

216     StkId base, fixed;
217
218     /* move fixed parameters to final position */
219     fixed = ks->top - actual; /* first fixed argument */
220     base = ks->top; /* final position of first argument */
221
222     for (i=0; i < nfixargs; i++) {
223         set_obj(ks->top++, fixed + i);
224         set_nil(fixed + i);
225     }
226
227     return base;
228 }
229
230 static void poscall(ktap_state_t *ks, StkId func, StkId first_result,
231                    int wanted)
232 {
233     int i;
234
235     for (i = wanted; i != 0 && first_result < ks->top; i--)
236         set_obj(func++, first_result++);
237
238     while(i-- > 0)
239         set_nil(func++);
240 }
241
242 void kp_vm_call_proto(ktap_state_t *ks, ktap_proto_t *pt)
243 {
244     ktap_func_t *fn;
245
246     fn = func_new_empty(ks, pt);
247     if (!fn)
248         return;
249     set_func(ks->top++, fn);
250     kp_vm_call(ks, ks->top - 1, 0);
251 }
252
253 /*
254 * Hot loop detaction
255 *
256 * Check hot loop detaction in three cases:
257 * 1. jmp -x: this happens in 'while (expr) { ... }'
258 * 2. FORPREP-FORLOOP
259 * 3. TFORCALL-TFORLOOP
260 */
261 static __always_inline int check_hot_loop(ktap_state_t *ks, int loop_count)
262 {
263     if (unlikely(loop_count == kp_max_loop_count)) {
264         kp_error(ks, "loop execute count exceed max limit(%d)\n",
265                 kp_max_loop_count);
266         return -1;
267     }
268
269     return 0;
270 }
271
272 #define dojump(i, e) { pc += (int)bc_d(i) - BCBIAS_J + e; }
273 #define donextjump { instr = *pc; dojump(instr, 1); }
274
275 #define NUMADD(a, b) ((a) + (b))
276 #define NUMSUB(a, b) ((a) - (b))
277 #define NUMMUL(a, b) ((a) * (b))
278 #define NUMDIV(a, b) ((a) / (b))
279 #define NUMUNM(a) (-a)
280 #define NUMEQ(a, b) ((a) == (b))
281 #define NUMLT(a, b) ((a) < (b))
282 #define NUMLE(a, b) ((a) <= (b))
283 #define NUMMOD(a, b) ((a) % (b))
284
285 #define arith_vv(ks, op) { \
286     ktap_val_t *rb = RB; \
287     ktap_val_t *rc = RC; \
288     if (is_number(rb) && is_number(rc)) { \
289         ktap_number nb = nvalue(rb), nc = nvalue(rc); \
290         set_number(RA, op(nb, nc)); \
291     } else { \

```

```

292     kp_puts(ks, "Error: Cannot make arith operation\n"); \
293     return; \
294 } }
295
296 #define arith_VN(ks, op) { \
297     ktap_val_t *rb = RB; \
298     if (is_number(rb)) { \
299         ktap_number nb = nvalue(rb);\
300         ktap_number nc = nvalue((ktap_val_t *)kbase + bc_c(instr));\
301         set_number(RA, op(nb, nc)); \
302     } else { \
303         kp_puts(ks, "Error: Cannot make arith operation\n"); \
304         return; \
305     } }
306
307 #define arith_NV(ks, op) { \
308     ktap_val_t *rb = RB; \
309     if (is_number(rb)) { \
310         ktap_number nb = nvalue(rb);\
311         ktap_number nc = nvalue((ktap_val_t *)kbase + bc_c(instr));\
312         set_number(RA, op(nc, nb)); \
313     } else { \
314         kp_puts(ks, "Error: Cannot make arith operation\n"); \
315         return; \
316     } }
317
318
319 static const char * const bc_names[] = {
320 #define BCNAME(name, ma, mb, mc, mt)      #name,
321     BCDEF(BCNAME)
322 #undef BCNAME
323     NULL
324 };
325
326
327 /*
328  * ktap bytecode interpreter routine
329  *
330  *
331  * kp_vm_call only can be used for:
332  * 1). call ktap function, not light C function
333  * 2). accept fixed argument function
334  */
335 void kp_vm_call(ktap_state_t *ks, StkId func, int nresults)
336 {
337     int loop_count = 0;
338     ktap_func_t *fn;
339     ktap_proto_t *pt;
340     ktap_obj_t **kbase;
341     unsigned int instr, op;
342     const unsigned int *pc;
343     StkId base; /* stack pointer */
344     int multres = 0; /* temp variable */
345     ktap_tab_t *gtab = G(ks)->gtab;
346
347     /* use computed goto for opcode dispatch */
348
349     static void *dispatch_table[] = {
350 #define BCNAME(name, ma, mb, mc, mt)      &&DO_BC_##name,
351     BCDEF(BCNAME)
352 #undef BCNAME
353     };
354
355 #define DISPATCH() \
356     do { \
357         instr = *(pc++); \
358         op = bc_op(instr); \
359         goto *dispatch_table[op]; \
360     } while (0)
361
362 #define RA    (base + bc_a(instr))
363 #define RB    (base + bc_b(instr))
364 #define RC    (base + bc_c(instr))
365 #define RD    (base + bc_d(instr))
366 #define RKD   ((ktap_val_t *)kbase + bc_d(instr))
367

```

```

368  /*TODO: fix argument number mismatch, example: sort cmp closure */
369
370  fn = clvalue(func);
371  pt = fn->p;
372  kbase = fn->p->k;
373  base = func + 1;
374  pc = proto_bc(pt) + 1;
375  ks->top = base + pt->framesize;
376  func->pcr = 0; /* no previous frame */
377
378  /* main loop of interpreter */
379  DISPATCH();
380
381  while (1) {
382  DO_BC_ISLT: /* Jump if A < D */
383      if (!is_number(RA) || !is_number(RD)) {
384          kp_error(ks, "compare with non-number\n");
385          return;
386      }
387
388      if (nvalue(RA) >= nvalue(RD))
389          pc++;
390      else
391          donextjump;
392      DISPATCH();
393  DO_BC_ISGE: /* Jump if A >= D */
394      if (!is_number(RA) || !is_number(RD)) {
395          kp_error(ks, "compare with non-number\n");
396          return;
397      }
398
399      if (nvalue(RA) < nvalue(RD))
400          pc++;
401      else
402          donextjump;
403      DISPATCH();
404  DO_BC_ISLE: /* Jump if A <= D */
405      if (!is_number(RA) || !is_number(RD)) {
406          kp_error(ks, "compare with non-number\n");
407          return;
408      }
409
410      if (nvalue(RA) > nvalue(RD))
411          pc++;
412      else
413          donextjump;
414      DISPATCH();
415  DO_BC_ISGT: /* Jump if A > D */
416      if (!is_number(RA) || !is_number(RD)) {
417          kp_error(ks, "compare with non-number\n");
418          return;
419      }
420
421      if (nvalue(RA) <= nvalue(RD))
422          pc++;
423      else
424          donextjump;
425      DISPATCH();
426  DO_BC_ISEQV: /* Jump if A = D */
427      if (!kp_obj_equal(RA, RD))
428          pc++;
429      else
430          donextjump;
431      DISPATCH();
432  DO_BC_ISNEV: /* Jump if A != D */
433      if (kp_obj_equal(RA, RD))
434          pc++;
435      else
436          donextjump;
437      DISPATCH();
438  DO_BC_ISEQS: { /* Jump if A = D */
439      int idx = ~bc_d(instr);
440
441      if (!is_string(RA) ||
442          rawtsvalue(RA) != (ktap_str_t *)kbase[idx])
443          pc++;

```

```

444     else
445         donextjump;
446     DISPATCH();
447 }
448 DO_BC_ISNES: { /* Jump if A != D */
449     int idx = ~bc_d(instr);
450
451     if (is_string(RA) &&
452         rawtsvalue(RA) == (ktap_str_t *)kbase[idx])
453         pc++;
454     else
455         donextjump;
456     DISPATCH();
457 }
458 DO_BC_ISEQN: /* Jump if A = D */
459     if (!is_number(RA) || nvalue(RA) != nvalue(RKD))
460         pc++;
461     else
462         donextjump;
463     DISPATCH();
464 DO_BC_ISNEN: /* Jump if A != D */
465     if (is_number(RA) && nvalue(RA) == nvalue(RKD))
466         pc++;
467     else
468         donextjump;
469     DISPATCH();
470 DO_BC_ISEQP: /* Jump if A = D */
471     if (itype(RA) != ~bc_d(instr))
472         pc++;
473     else
474         donextjump;
475     DISPATCH();
476 DO_BC_ISNEP: /* Jump if A != D */
477     if (itype(RA) == ~bc_d(instr))
478         pc++;
479     else
480         donextjump;
481     DISPATCH();
482 DO_BC_ISTC: /* Copy D to A and jump, if D is true */
483     if (itype(RD) == KTAP_TNIL || itype(RD) == KTAP_TFALSE)
484         pc++;
485     else {
486         set_obj(RA, RD);
487         donextjump;
488     }
489     DISPATCH();
490 DO_BC_ISFC: /* Copy D to A and jump, if D is false */
491     if (itype(RD) != KTAP_TNIL && itype(RD) != KTAP_TFALSE)
492         pc++;
493     else {
494         set_obj(RA, RD);
495         donextjump;
496     }
497     DISPATCH();
498 DO_BC_IST: /* Jump if D is true */
499     if (itype(RD) == KTAP_TNIL || itype(RD) == KTAP_TFALSE)
500         pc++;
501     else
502         donextjump;
503     DISPATCH();
504 DO_BC_ISF: /* Jump if D is false */
505     /* only nil and false are considered false,
506      * all other values are true */
507     if (itype(RD) != KTAP_TNIL && itype(RD) != KTAP_TFALSE)
508         pc++;
509     else
510         donextjump;
511     DISPATCH();
512 DO_BC_ISTYPE: /* generated by genlibbc, not compiler; not used now */
513 DO_BC_ISNUM:
514     return;
515 DO_BC_MOV: /* Copy D to A */
516     set_obj(RA, RD);
517     DISPATCH();
518 DO_BC_NOT: /* Set A to boolean not of D */
519     if (itype(RD) == KTAP_TNIL || itype(RD) == KTAP_TFALSE)

```

```

520     setitype(RA, KTAP_TTRUE);
521     else
522         setitype(RA, KTAP_TFALSE);
523
524     DISPATCH();
525 DO_BC_UNM: /* Set A to -D (unary minus) */
526     if (!is_number(RD)) {
527         kp_error(ks, "use '-' operator on non-number\n");
528         return;
529     }
530
531     set_number(RA, -nvalue(RD));
532     DISPATCH();
533 DO_BC_ADDVN: /* A = B + C */
534     arith_VN(ks, NUMADD);
535     DISPATCH();
536 DO_BC_SUBVN: /* A = B - C */
537     arith_VN(ks, NUMSUB);
538     DISPATCH();
539 DO_BC_MULVN: /* A = B * C */
540     arith_VN(ks, NUMMUL);
541     DISPATCH();
542 DO_BC_DIVVN: /* A = B / C */
543     /* divide 0 checking */
544     if (!nvalue((ktap_val_t *)kbase + bc_c(instr))) {
545         kp_error(ks, "divide 0 arith operation\n");
546         return;
547     }
548     arith_VN(ks, NUMDIV);
549     DISPATCH();
550 DO_BC_MODVN: /* A = B % C */
551     /* divide 0 checking */
552     if (!nvalue((ktap_val_t *)kbase + bc_c(instr))) {
553         kp_error(ks, "mod 0 arith operation\n");
554         return;
555     }
556     arith_VN(ks, NUMMOD);
557     DISPATCH();
558 DO_BC_ADDNV: /* A = C + B */
559     arith_NV(ks, NUMADD);
560     DISPATCH();
561 DO_BC_SUBNV: /* A = C - B */
562     arith_NV(ks, NUMSUB);
563     DISPATCH();
564 DO_BC_MULNV: /* A = C * B */
565     arith_NV(ks, NUMMUL);
566     DISPATCH();
567 DO_BC_DIVNV: /* A = C / B */
568     /* divide 0 checking */
569     if (!nvalue(RB)){
570         kp_error(ks, "divide 0 arith operation\n");
571         return;
572     }
573     arith_NV(ks, NUMDIV);
574     DISPATCH();
575 DO_BC_MODNV: /* A = C % B */
576     /* divide 0 checking */
577     if (!nvalue(RB)){
578         kp_error(ks, "mod 0 arith operation\n");
579         return;
580     }
581     arith_NV(ks, NUMMOD);
582     DISPATCH();
583 DO_BC_ADDVV: /* A = B + C */
584     arith_VV(ks, NUMADD);
585     DISPATCH();
586 DO_BC_SUBVV: /* A = B - C */
587     arith_VV(ks, NUMSUB);
588     DISPATCH();
589 DO_BC_MULVV: /* A = B * C */
590     arith_VV(ks, NUMMUL);
591     DISPATCH();
592 DO_BC_DIVVV: /* A = B / C */
593     arith_VV(ks, NUMDIV);
594     DISPATCH();
595 DO_BC_MODVV: /* A = B % C */

```

```

596     arith_VV(ks, NUMMOD);
597     DISPATCH();
598 DO_BC_POW: /* A = B ^ C, rejected */
599     return;
600 DO_BC_CAT: { /* A = B .. ~ .. C */
601     /* The CAT instruction concatenates all values in
602     * variable slots B to C inclusive. */
603     ktap_str_t *ts = str_concat(ks, base, bc_b(instr),
604                               bc_c(instr));
605     if (!ts)
606         return;
607
608     set_string(RA, ts);
609     DISPATCH();
610 }
611 DO_BC_KSTR: { /* Set A to string constant D */
612     int idx = ~bc_d(instr);
613     set_string(RA, (ktap_str_t *)kbase[idx]);
614     DISPATCH();
615 }
616 DO_BC_KCDATA: /* not used now */
617     DISPATCH();
618 DO_BC_KSHORT: /* Set A to 16 bit signed integer D */
619     set_number(RA, bc_d(instr));
620     DISPATCH();
621 DO_BC_KNUM: /* Set A to number constant D */
622     set_number(RA, nvalue(RKD));
623     DISPATCH();
624 DO_BC_KPRI: /* Set A to primitive D */
625     setitype(RA, ~bc_d(instr));
626     DISPATCH();
627 DO_BC_KNIL: { /* Set slots A to D to nil */
628     int i;
629     for (i = 0; i <= bc_d(instr) - bc_a(instr); i++) {
630         set_nil(RA + i);
631     }
632     DISPATCH();
633 }
634 DO_BC_UGET: /* Set A to upvalue D */
635     set_obj(RA, fn->upvals[bc_d(instr)]->v);
636     DISPATCH();
637 DO_BC_USETV: /* Set upvalue A to D */
638     set_obj(fn->upvals[bc_a(instr)]->v, RD);
639     DISPATCH();
640 DO_BC_UINCV: { /* upvalus[A] += D */
641     ktap_val_t *v = fn->upvals[bc_a(instr)]->v;
642     if (unlikely(!is_number(RD) || !is_number(v))) {
643         kp_error(ks, "use '+=' on non-number\n");
644         return;
645     }
646     set_number(v, nvalue(v) + nvalue(RD));
647     DISPATCH();
648 }
649 DO_BC_USETS: { /* Set upvalue A to string constant D */
650     int idx = ~bc_d(instr);
651     set_string(fn->upvals[bc_a(instr)]->v,
652              (ktap_str_t *)kbase[idx]);
653     DISPATCH();
654 }
655 DO_BC_USETN: /* Set upvalue A to number constant D */
656     set_number(fn->upvals[bc_a(instr)]->v, nvalue(RKD));
657     DISPATCH();
658 DO_BC_UINCN: { /* upvalus[A] += D */
659     ktap_val_t *v = fn->upvals[bc_a(instr)]->v;
660     if (unlikely(!is_number(v))) {
661         kp_error(ks, "use '+=' on non-number\n");
662         return;
663     }
664     set_number(v, nvalue(v) + nvalue(RKD));
665     DISPATCH();
666 }
667 DO_BC_USETP: /* Set upvalue A to primitive D */
668     setitype(fn->upvals[bc_a(instr)]->v, ~bc_d(instr));
669     DISPATCH();
670 DO_BC_UCLO: /* Close upvalues for slots . rbase and jump to target D */
671     if (ks->openupval != NULL)

```

```

672     func_closeuv(ks, RA);
673     dojump(instr, 0);
674     DISPATCH();
675 DO_BC_FNEW: {
676     /* Create new closure from prototype D and store it in A */
677     int idx = ~bc_d(instr);
678     ktap_func_t *subfn = func_new(ks, (ktap_proto_t *)kbase[idx],
679     fn, base);
680     if (unlikely(!subfn))
681         return;
682     set_func(RA, subfn);
683     DISPATCH();
684 }
685 DO_BC_TNEW: { /* Set A to new table with size D */
686     /*
687     * preallocate default narr and nrec,
688     * op_b and op_c is not used
689     * This would allocate more memory for some static table.
690     */
691     ktap_tab_t *t = kp_tab_new_ah(ks, 0, 0);
692     if (unlikely(!t))
693         return;
694     set_table(RA, t);
695     DISPATCH();
696 }
697 DO_BC_TDUP: { /* Set A to duplicated template table D */
698     int idx = ~bc_d(instr);
699     ktap_tab_t *t = kp_tab_dup(ks, (ktap_tab_t *)kbase[idx]);
700     if (!t)
701         return;
702     set_table(RA, t);
703     DISPATCH();
704 }
705 DO_BC_GGET: { /* A = _G[D] */
706     int idx = ~bc_d(instr);
707     kp_tab_getstr(gtab, (ktap_str_t *)kbase[idx], RA);
708     DISPATCH();
709 }
710 DO_BC_GSET: /* _G[D] = A, rejected. */
711 DO_BC_GINC: /* _G[D] += A, rejected. */
712     return;
713 DO_BC_TGETV: /* A = B[C] */
714     if (unlikely(!is_table(RB))) {
715         kp_error(ks, "get key from non-table\n");
716         return;
717     }
718
719     kp_tab_get(ks, hvalue(RB), RC, RA);
720     DISPATCH();
721 DO_BC_TGETS: { /* A = B[C] */
722     int idx = ~bc_c(instr);
723
724     if (unlikely(!is_table(RB))) {
725         kp_error(ks, "get key from non-table\n");
726         return;
727     }
728     kp_tab_getstr(hvalue(RB), (ktap_str_t *)kbase[idx], RA);
729     DISPATCH();
730 }
731 DO_BC_TGETB: { /* A = B[C] */
732     /* 8 bit literal C operand as an unsigned integer
733     * index (0..255) */
734     uint8_t idx = bc_c(instr);
735
736     if (unlikely(!is_table(RB))) {
737         kp_error(ks, "set key to non-table\n");
738         return;
739     }
740     kp_tab_getint(hvalue(RB), idx, RA);
741     DISPATCH();
742 }
743 DO_BC_TGETR: /* generated by genlibbc, not compiler, not used */
744     return;
745 DO_BC_TSETV: /* B[C] = A */
746     if (unlikely(!is_table(RB))) {
747         kp_error(ks, "set key to non-table\n");

```

```

748     return;
749 }
750 kp_tab_set(ks, hvalue(RB), RC, RA);
751 DISPATCH();
752 DO_BC_TINCV: /* B[C] += A */
753 if (unlikely(!is_table(RB))) {
754     kp_error(ks, "set key to non-table\n");
755     return;
756 }
757 if (unlikely(!is_number(RA))) {
758     kp_error(ks, "use '+=' on non-number\n");
759     return;
760 }
761 kp_tab_incr(ks, hvalue(RB), RC, nvalue(RA));
762 DISPATCH();
763 DO_BC_TSETS: { /* B[C] = A */
764     int idx = ~bc_c(instr);
765
766     if (unlikely(!is_table(RB))) {
767         kp_error(ks, "set key to non-table\n");
768         return;
769     }
770     kp_tab_setstr(ks, hvalue(RB), (ktap_str_t *)kbase[idx], RA);
771     DISPATCH();
772 }
773 DO_BC_TINCS: { /* B[C] += A */
774     int idx = ~bc_c(instr);
775
776     if (unlikely(!is_table(RB))) {
777         kp_error(ks, "set key to non-table\n");
778         return;
779     }
780     if (unlikely(!is_number(RA))) {
781         kp_error(ks, "use '+=' on non-number\n");
782         return;
783     }
784     kp_tab_incrstr(ks, hvalue(RB), (ktap_str_t *)kbase[idx],
785         nvalue(RA));
786     DISPATCH();
787 }
788 DO_BC_TSETB: { /* B[C] = A */
789     /* 8 bit literal C operand as an unsigned integer
790      * index (0..255) */
791     uint8_t idx = bc_c(instr);
792
793     if (unlikely(!is_table(RB))) {
794         kp_error(ks, "set key to non-table\n");
795         return;
796     }
797     kp_tab_setint(ks, hvalue(RB), idx, RA);
798     DISPATCH();
799 }
800 DO_BC_TINCB: { /* B[C] = A */
801     uint8_t idx = bc_c(instr);
802
803     if (unlikely(!is_table(RB))) {
804         kp_error(ks, "set key to non-table\n");
805         return;
806     }
807     if (unlikely(!is_number(RA))) {
808         kp_error(ks, "use '+=' on non-number\n");
809         return;
810     }
811     kp_tab_incrint(ks, hvalue(RB), idx, nvalue(RA));
812     DISPATCH();
813 }
814 DO_BC_TSETM: /* don't support */
815     return;
816 DO_BC_TSETR: /* generated by genlibbc, not compiler, not used */
817     return;
818 DO_BC_CALLM:
819 DO_BC_CALL: { /* b: return_number + 1; c: argument + 1 */
820     int c = bc_c(instr);
821     int nresults = bc_b(instr) - 1;
822     StkId oldtop = ks->top;
823     StkId newfunc = RA;

```

```

824     if (op == BC_CALL && c != 0)
825         ks->top = RA + c;
826     else if (op == BC_CALLM)
827         ks->top = RA + c + multres;
828
829
830     if (itype(newfunc) == KTAP_TCFUNC) { /* light C function */
831         ktap_cfunction f = fvalue(newfunc);
832         int n;
833
834         if (unlikely(checkstack(ks,
835             KTAP_MIN_RESERVED_STACK_SIZE)))
836             return;
837
838         ks->func = newfunc;
839         n = (*f)(ks);
840         if (unlikely(n < 0)) /* error occured */
841             return;
842         poscall(ks, newfunc, ks->top - n, nresults);
843
844         ks->top = oldtop;
845         multres = n + 1; /* set to multres */
846         DISPATCH();
847     } else if (itype(newfunc) == KTAP_TFUNC) { /* ktap function */
848         int n;
849
850         func = newfunc;
851         pt = clvalue(func)->p;
852
853         if (unlikely(checkstack(ks, pt->framesize)))
854             return;
855
856         /* get number of real arguments */
857         n = (int)(ks->top - func) - 1;
858
859         /* complete missing arguments */
860         for (; n < pt->numparams; n++)
861             set_nil(ks->top++);
862
863         base = (!(pt->flags & PROTO_VARARG)) ? func + 1 :
864             adjust_varargs(ks, pt, n);
865
866         fn = clvalue(func);
867         pt = fn->p;
868         kbase = pt->k;
869         func->pcr = pc - 1; /* save pc */
870         ks->top = base + pt->framesize;
871         pc = proto_bc(pt) + 1; /* starting point */
872         DISPATCH();
873     } else {
874         kp_error(ks, "attempt to call nil function\n");
875         return;
876     }
877 }
878 DO_BC_CALLMT: /* don't support */
879     return;
880 DO_BC_CALLT: { /* Tailcall: return A(A+1, ..., A+D-1) */
881     StkId nfunc = RA;
882
883     if (itype(nfunc) == KTAP_TCFUNC) { /* light C function */
884         kp_error(ks, "don't support callt for C function");
885         return;
886     } else if (itype(nfunc) == KTAP_TFUNC) { /* ktap function */
887         int aux;
888
889         /*
890          * tail call: put called frame (n) in place of
891          * caller one (o)
892          */
893         StkId ofunc = func; /* caller function */
894         /* last stack slot filled by 'precall' */
895         StkId lim = nfunc + 1 + clvalue(nfunc)->p->numparams;
896
897         fn = clvalue(nfunc);
898         ofunc->val = nfunc->val;
899

```

```

900     /* move new frame into old one */
901     for (aux = 1; nfunc + aux < lim; aux++)
902         set_obj(ofunc + aux, nfunc + aux);
903
904     pt = fn->p;
905     kbase = pt->k;
906     ks->top = base + pt->framesize;
907     pc = proto_bc(pt) + 1; /* starting point */
908     DISPATCH();
909 } else {
910     kp_error(ks, "attempt to call nil function\n");
911     return;
912 }
913 }
914 DO_BC_ITERC: /* don't support it now */
915     return;
916 DO_BC_ITERN: /* Specialized ITERC, if iterator function A-3 is next()*/
917     /* detect hot loop */
918     if (unlikely(check_hot_loop(ks, loop_count++) < 0))
919         return;
920
921     if (kp_tab_next(ks, hvalue(RA - 2), RA)) {
922         donextjump; /* Get jump target from ITERL */
923     } else {
924         pc++; /* jump to ITERL + 1 */
925     }
926     DISPATCH();
927 DO_BC_VARG: /* don't support */
928     return;
929 DO_BC_ISNEXT: /* Verify ITERN specialization and jump */
930     if (!is_cfunc(RA - 3) || !is_table(RA - 2) || !is_nil(RA - 1)
931         || fvalue(RA - 3) != (ktap_cfunction)kp_tab_next) {
932         /* Despecialize bytecode if any of the checks fail. */
933         setbc_op(pc - 1, BC_JMP);
934         dojump(instr, 0);
935         setbc_op(pc, BC_ITERC);
936     } else {
937         dojump(instr, 0);
938         set_nil(RA); /* init control variable */
939     }
940     DISPATCH();
941 DO_BC_RETM: /* don't support return multiple values */
942 DO_BC_RET:
943     return;
944 DO_BC_RET0:
945     /* if it's called from external invocation, just return */
946     if (!func->pcr)
947         return;
948
949     pc = func->pcr; /* restore PC */
950
951     multres = bc_d(instr);
952     set_nil(func);
953
954     base = func - bc_a(*pc);
955     func = base - 1;
956     fn = clvalue(func);
957     kbase = fn->p->k;
958     ks->top = base + pt->framesize;
959     pc++;
960
961     DISPATCH();
962 DO_BC_RET1:
963     /* if it's called from external invocation, just return */
964     if (!func->pcr)
965         return;
966
967     pc = func->pcr; /* restore PC */
968
969     multres = bc_d(instr);
970     set_obj(base - 1, RA); /* move result */
971
972     base = func - bc_a(*pc);
973     func = base - 1;
974     fn = clvalue(func);
975     kbase = fn->p->k;

```

```

976     ks->top = base + pt->framesize;
977     pc++;
978
979     DISPATCH();
980 DO_BC_FORI: { /* Numeric 'for' loop init */
981     ktap_number idx;
982     ktap_number limit;
983     ktap_number step;
984
985     if (unlikely(!is_number(RA) || !is_number(RA + 1) ||
986                !is_number(RA + 2))) {
987         kp_error(ks, KTAP_QL("for")
988                " init/limit/step value must be a number\n");
989         return;
990     }
991
992     idx = nvalue(RA);
993     limit = nvalue(RA + 1);
994     step = nvalue(RA + 2);
995
996     if (NUMLT(0, step) ? NUMLE(idx, limit) : NUMLE(limit, idx)) {
997         set_number(RA + 3, nvalue(RA));
998     } else {
999         dojump(instr, 0);
1000     }
1001     DISPATCH();
1002 }
1003 DO_BC_JFORI: /* not used */
1004     return;
1005 DO_BC_FORL: { /* Numeric 'for' loop */
1006     ktap_number step = nvalue(RA + 2);
1007     /* increment index */
1008     ktap_number idx = NUMADD(nvalue(RA), step);
1009     ktap_number limit = nvalue(RA + 1);
1010     if (NUMLT(0, step) ? NUMLE(idx, limit) : NUMLE(limit, idx)) {
1011         dojump(instr, 0); /* jump back */
1012         set_number(RA, idx); /* update internal index... */
1013         set_number(RA + 3, idx); /* ...and external index */
1014     }
1015
1016     if (unlikely(check_hot_loop(ks, loop_count++) < 0))
1017         return;
1018
1019     DISPATCH();
1020 }
1021 DO_BC_IFORL: /* not used */
1022 DO_BC_JFORL:
1023 DO_BC_ITERL:
1024 DO_BC_IITERL:
1025 DO_BC_JITERL:
1026     return;
1027 DO_BC_LOOP: /* Generic loop */
1028     /* ktap use this bc to detect hot loop */
1029     if (unlikely(check_hot_loop(ks, loop_count++) < 0))
1030         return;
1031     DISPATCH();
1032 DO_BC_ILOOP: /* not used */
1033 DO_BC_JLOOP:
1034     return;
1035 DO_BC_JMP: /* Jump */
1036     dojump(instr, 0);
1037     DISPATCH();
1038 DO_BC_FUNCF: /* function header, not used */
1039 DO_BC_IFUNCF:
1040 DO_BC_JFUNCF:
1041 DO_BC_FUNCV:
1042 DO_BC_IFUNCV:
1043 DO_BC_JFUNCV:
1044 DO_BC_FUNCCL:
1045 DO_BC_FUNCCLW:
1046     return;
1047 DO_BC_VARGN: /* arg0 .. arg9*/
1048     if (unlikely(!ks->current_event)) {
1049         kp_error(ks, "invalid event context\n");
1050         return;
1051     }

```

```

1052     kp\_event\_getarg(ks, RA, bc\_d(instr));
1053     DISPATCH();
1054 DO_BC_VARGSTR: { /* argstr */
1055     /*
1056      * If you pass argstr to print/printf function directly,
1057      * then no extra string generated, so don't worry string
1058      * poll size for below case:
1059      *   print(argstr)
1060      *
1061      * If you use argstr as table key like below, then it may
1062      * overflow your string pool size, so be care of on it.
1063      *   table[argstr] = v
1064      *
1065      * If you assign argstr to upval or table value like below,
1066      * it don't really write string, just write type KTAP\_TEVENTSTR,
1067      * the value will be interpreted when value print out in valid
1068      * event context, if context mismatch, error will report.
1069      *   table[V] = argstr
1070      *   upval = argstr
1071      *
1072      * If you want to save real string of argstr, then use it like
1073      * below, again, be care of string pool size in this case.
1074      *   table[V] = stringof(argstr)
1075      *   upval = stringof(argstr)
1076      */
1077     struct ktap\_event\_data *e = ks->current_event;
1078
1079     if (unlikely(!e)) {
1080         kp\_error(ks, "invalid event context\n");
1081         return;
1082     }
1083
1084     if (e->argstr) /* argstr been stringified */
1085         set\_string(RA, e->argstr);
1086     else
1087         set\_eventstr(RA);
1088     DISPATCH();
1089 }
1090 DO_BC_VPROBENAME: { /* probename */
1091     struct ktap\_event\_data *e = ks->current_event;
1092
1093     if (unlikely(!e)) {
1094         kp\_error(ks, "invalid event context\n");
1095         return;
1096     }
1097     set\_string(RA, e->event->name);
1098     DISPATCH();
1099 }
1100 DO_BC_VPID: /* pid */
1101     set\_number(RA, (int)current->pid);
1102     DISPATCH();
1103 DO_BC_VTID: /* tid */
1104     set\_number(RA, (int)task_pid_vnr(current));
1105     DISPATCH();
1106 DO_BC_VUID: { /* uid */
1107 #if LINUX_VERSION_CODE >= KERNEL_VERSION(3, 5, 0)
1108     uid_t uid = from_kuid_munged(current_user_ns(), current_uid());
1109 #else
1110     uid_t uid = current_uid();
1111 #endif
1112     set\_number(RA, (int)uid);
1113     DISPATCH();
1114 }
1115 DO_BC_VCPU: /* cpu */
1116     set\_number(RA, smp_processor_id());
1117     DISPATCH();
1118 DO_BC_VEEXECNAME: { /* execname */
1119     ktap\_str\_t *ts = kp\_str\_newz(ks, current->comm);
1120     if (unlikely(!ts))
1121         return;
1122     set\_string(RA, ts);
1123     DISPATCH();
1124 }
1125 DO_BC_GFUNC: { /* Call built-in C function, patched by BC_GGET */
1126     ktap\_cfunction cfunc = gfunc\_get(ks, bc\_d(instr));
1127

```

```

1128     set_cfunc(RA, cfunc);
1129     DISPATCH();
1130 }
1131 }
1132 }
1133
1134 /*
1135 * Validate byte code and static analysis.
1136 *
1137 * TODD: more type checking before real running.
1138 */
1139 int kp_vm_validate_code(ktap_state_t *ks, ktap_proto_t *pt, ktap_val_t *base)
1140 {
1141     const unsigned int *pc = proto_bc(pt) + 1;
1142     unsigned int instr, op;
1143     ktap_obj_t **kbase = pt->k;
1144     ktap_tab_t *gtab = G(ks)->gtab;
1145     int i;
1146
1147     #define RA    (base + bc_a(instr))
1148     #define RB    (base + bc_b(instr))
1149     #define RC    (base + bc_c(instr))
1150     #define RD    (base + bc_d(instr))
1151
1152     if (pt->framesize > KP_MAX_SLOTS) {
1153         kp_error(ks, "exceed max frame size %d\n", pt->framesize);
1154         return -1;
1155     }
1156
1157     if (base + pt->framesize > ks->stack_last) {
1158         kp_error(ks, "stack overflow\n");
1159         return -1;
1160     }
1161
1162     for (i = 0; i < pt->sizebc - 1; i++) {
1163         instr = *pc++;
1164         op = bc_op(instr);
1165
1166         if (op >= BC_MAX) {
1167             kp_error(ks, "unknown byte code %d\n", op);
1168             return -1;
1169         }
1170
1171         switch (op) {
1172         case BC_FNEW: {
1173             int idx = -bc_d(instr);
1174             ktap_proto_t *newpt = (ktap_proto_t *)kbase[idx];
1175             if (kp_vm_validate_code(ks, newpt, RA + 1))
1176                 return -1;
1177
1178             break;
1179         }
1180         case BC_RETM: case BC_RET:
1181             kp_error(ks, "don't support return multiple values\n");
1182             return -1;
1183         case BC_GSET: case BC_GINC: { /* _G[D] = A, _G[D] += A */
1184             int idx = -bc_d(instr);
1185             ktap_str_t *ts = (ktap_str_t *)kbase[idx];
1186             kp_error(ks, "cannot set global variable '%s'\n",
1187                 getstr(ts));
1188             return -1;
1189         }
1190         case BC_GGET: {
1191             int idx = -bc_d(instr);
1192             ktap_str_t *ts = (ktap_str_t *)kbase[idx];
1193             ktap_val_t val;
1194             kp_tab_getstr(gtab, ts, &val);
1195             if (is_nil(&val)) {
1196                 kp_error(ks, "undefined global variable"
1197                     " '%s'\n", getstr(ts));
1198                 return -1;
1199             } else if (is_cfunc(&val)) {
1200                 int idx = gfunc_getidx(G(ks), fvalue(&val));
1201                 if (idx >= 0) {
1202                     /* patch BC_GGET bytecode to BC_GFUNC */

```

```

1204         setbc\_op(pc - 1, BC_GFUNC);
1205         setbc\_d(pc - 1, idx);
1206     }
1207 }
1208 break;
1209 }
1210 case BC_ITERC:
1211     kp\_error(ks, "ktap only support pairs iteraor\n");
1212     return -1;
1213 case BC_POW:
1214     kp\_error(ks, "ktap don't support pow arith\n");
1215     return -1;
1216 }
1217 }
1218
1219 return 0;
1220 }
1221
1222 /* return cfunction by idx */
1223 static ktap\_cfunction gfunc\_get(ktap\_state\_t *ks, int idx)
1224 {
1225     return G(ks)->gfunc_tbl[idx];
1226 }
1227
1228 /* get cfunction index, the index is for fast get cfunction in runtime */
1229 static int gfunc\_getidx(ktap\_global\_state\_t *g, ktap\_cfunction cfunc)
1230 {
1231     int nr = g->nr_builtin_cfunction;
1232     ktap\_cfunction *gfunc_tbl = g->gfunc_tbl;
1233     int i;
1234
1235     for (i = 0; i < nr; i++) {
1236         if (gfunc_tbl[i] == cfunc)
1237             return i;
1238     }
1239
1240     return -1;
1241 }
1242
1243 static void gfunc\_add(ktap\_state\_t *ks, ktap\_cfunction cfunc)
1244 {
1245     int nr = G(ks)->nr_builtin_cfunction;
1246
1247     if (nr == KP\_MAX\_CACHED\_CFUNCTION) {
1248         kp\_error(ks, "please enlarge KP\_MAX\_CACHED\_CFUNCTION %d\n",
1249                 KP\_MAX\_CACHED\_CFUNCTION);
1250         return;
1251     }
1252     G(ks)->gfunc_tbl[nr] = cfunc;
1253     G(ks)->nr_builtin_cfunction++;
1254 }
1255
1256 /* function for register library */
1257 int kp\_vm\_register\_lib(ktap\_state\_t *ks, const char *libname,
1258                        const ktap\_libfunc\_t *funcs)
1259 {
1260     ktap\_tab\_t *gtab = G(ks)->gtab;
1261     ktap\_tab\_t *target_tbl;
1262     int i;
1263
1264     /* lib is null when register baselib function */
1265     if (libname == NULL)
1266         target_tbl = gtab;
1267     else {
1268         ktap\_val\_t key, val;
1269         ktap\_str\_t *ts = kp\_str\_newz(ks, libname);
1270         if (!ts)
1271             return -ENOMEM;
1272
1273         /* calculate the function number contained by this library */
1274         for (i = 0; funcs[i].name != NULL; i++) {
1275         }
1276
1277         target_tbl = kp\_tab\_new\_ah(ks, 0, i + 1);
1278         if (!target_tbl)
1279             return -ENOMEM;

```

```

1280     set\_string(&key, ts);
1281     set\_table(&val, target_tbl);
1282     kp\_tab\_set(ks, gtab, &key, &val);
1283 }
1284
1285
1286 /* TODO: be care of same function name issue, foo() and tbl.foo() */
1287 for (i = 0; funcs[i].name != NULL; i++) {
1288     ktap\_str\_t *func_name = kp\_str\_newz(ks, funcs[i].name);
1289     ktap\_val\_t fn;
1290
1291     if (unlikely(!func_name))
1292         return -ENOMEM;
1293
1294     set\_cfunc(&fn, funcs[i].func);
1295     kp\_tab\_setstr(ks, target_tbl, func_name, &fn);
1296
1297     gfunc\_add(ks, funcs[i].func);
1298 }
1299
1300 return 0;
1301 }
1302
1303 static int init\_registry(ktap\_state\_t *ks)
1304 {
1305     ktap\_tab\_t *registry = kp\_tab\_new\_ah(ks, 2, 0);
1306     ktap\_val\_t gtbl;
1307     ktap\_tab\_t *t;
1308
1309     if (!registry)
1310         return -1;
1311
1312     set\_table(&G(ks)->registry, registry);
1313
1314     /* assume there will have max 1024 global variables */
1315     t = kp\_tab\_new\_ah(ks, 0, 1024);
1316     if (!t)
1317         return -1;
1318
1319     set\_table(&gtbl, t);
1320     kp\_tab\_setint(ks, registry, KTAP\_RIDX\_GLOBALS, &gtbl);
1321     G(ks)->gtab = t;
1322
1323     return 0;
1324 }
1325
1326 static int init\_arguments(ktap\_state\_t *ks, int argc, char __user **user_argv)
1327 {
1328     ktap\_tab\_t *gtbl = G(ks)->gtab;
1329     ktap\_tab\_t *arg_tbl = kp\_tab\_new\_ah(ks, argc, 1);
1330     ktap\_val\_t arg_tblval;
1331     ktap\_val\_t arg_tsval;
1332     ktap\_str\_t *argts = kp\_str\_newz(ks, "arg");
1333     char **argv;
1334     int i, ret;
1335
1336     if (!arg_tbl)
1337         return -1;
1338
1339     if (unlikely(!argts))
1340         return -ENOMEM;
1341
1342     set\_string(&arg_tsval, argts);
1343     set\_table(&arg_tblval, arg_tbl);
1344     kp\_tab\_set(ks, gtbl, &arg_tsval, &arg_tblval);
1345
1346     if (!argc)
1347         return 0;
1348
1349     if (argc > 1024)
1350         return -EINVAL;
1351
1352     argv = kzalloc(argc * sizeof\(char \*\), GFP_KERNEL);
1353     if (!argv)
1354         return -ENOMEM;
1355

```

```

1356     ret = copy_from_user(argv, user_argv, argc * sizeof(char *));
1357     if (ret < 0) {
1358         kfree(argv);
1359         return -EFAULT;
1360     }
1361
1362     ret = 0;
1363     for (i = 0; i < argc; i++) {
1364         ktap_val_t val;
1365         char __user *ustr = argv[i];
1366         char *kstr;
1367         int len;
1368         int res;
1369
1370         len = strlen_user(ustr);
1371         if (len > 0x1000) {
1372             ret = -EINVAL;
1373             break;
1374         }
1375
1376         kstr = kmalloc(len + 1, GFP_KERNEL);
1377         if (!kstr) {
1378             ret = -ENOMEM;
1379             break;
1380         }
1381
1382         if (strncpy_from_user(kstr, ustr, len) < 0) {
1383             kfree(kstr);
1384             ret = -EFAULT;
1385             break;
1386         }
1387
1388         kstr[len] = '\0';
1389
1390         if (!kstrtoint(kstr, 10, &res)) {
1391             set_number(&val, res);
1392         } else {
1393             ktap_str_t *ts = kp_str_newz(ks, kstr);
1394             if (unlikely(!ts)) {
1395                 kfree(kstr);
1396                 ret = -ENOMEM;
1397                 break;
1398             }
1399
1400             set_string(&val, ts);
1401         }
1402
1403         kp_tab_setint(ks, arg_tbl, i, &val);
1404
1405         kfree(kstr);
1406     }
1407
1408     kfree(argv);
1409     return ret;
1410 }
1411
1412 static void free_preserved_data(ktap_state_t *ks)
1413 {
1414     int cpu, i, j;
1415
1416     /* free stack for each allocated ktap_state */
1417     for_each_possible_cpu(cpu) {
1418         for (j = 0; j < PERF_NR_CONTEXTS; j++) {
1419             void *percpu_state = G(ks)->percpu_state[j];
1420             ktap_state_t *pks;
1421
1422             if (!percpu_state)
1423                 break;
1424             pks = per_cpu_ptr(percpu_state, cpu);
1425             if (!ks)
1426                 break;
1427             kfree(pks->stack);
1428         }
1429     }
1430
1431     /* free percpu ktap_state */

```

```

1432 for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1433     if (G(ks)->percpu_state[i])
1434         free_percpu(G(ks)->percpu_state[i]);
1435 }
1436
1437 /* free percpu ktap print buffer */
1438 for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1439     if (G(ks)->percpu_print_buffer[i])
1440         free_percpu(G(ks)->percpu_print_buffer[i]);
1441 }
1442
1443 /* free percpu ktap temp buffer */
1444 for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1445     if (G(ks)->percpu_temp_buffer[i])
1446         free_percpu(G(ks)->percpu_temp_buffer[i]);
1447 }
1448
1449 /* free percpu ktap recursion context flag */
1450 for (i = 0; i < PERF_NR_CONTEXTS; i++)
1451     if (G(ks)->recursion_context[i])
1452         free_percpu(G(ks)->recursion_context[i]);
1453 }
1454
1455 #define ALLOC_PERCPU(size) __alloc_percpu(size, __alignof__(char))
1456 static int init_preserved_data(ktap_state_t *ks)
1457 {
1458     void __percpu *data;
1459     int cpu, i, j;
1460
1461     /* init percpu ktap_state */
1462     for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1463         data = ALLOC_PERCPU(sizeof(ktap_state_t));
1464         if (!data)
1465             goto fail;
1466         G(ks)->percpu_state[i] = data;
1467     }
1468
1469     /* init stack for each allocated ktap_state */
1470     for_each_possible_cpu(cpu) {
1471         for (j = 0; j < PERF_NR_CONTEXTS; j++) {
1472             void *percpu_state = G(ks)->percpu_state[j];
1473             ktap_state_t *pks;
1474
1475             if (!percpu_state)
1476                 break;
1477             pks = per_cpu_ptr(percpu_state, cpu);
1478             if (!ks)
1479                 break;
1480             pks->stack = kzalloc(KTAP_STACK_SIZE_BYTES, GFP_KERNEL);
1481             if (!pks->stack)
1482                 goto fail;
1483
1484             pks->stack_last = pks->stack + KTAP_STACK_SIZE;
1485             G(pks) = G(ks);
1486         }
1487     }
1488
1489     /* init percpu ktap print buffer */
1490     for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1491         data = ALLOC_PERCPU(KTAP_PERCPU_BUFFER_SIZE);
1492         if (!data)
1493             goto fail;
1494         G(ks)->percpu_print_buffer[i] = data;
1495     }
1496
1497     /* init percpu ktap temp buffer */
1498     for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1499         data = ALLOC_PERCPU(KTAP_PERCPU_BUFFER_SIZE);
1500         if (!data)
1501             goto fail;
1502         G(ks)->percpu_temp_buffer[i] = data;
1503     }
1504
1505     /* init percpu ktap recursion context flag */
1506     for (i = 0; i < PERF_NR_CONTEXTS; i++) {
1507         data = alloc_percpu(int);

```

```

1508     if (!data)
1509         goto fail;
1510     G(ks)->recursion_context[i] = data;
1511 }
1512
1513     return 0;
1514
1515 fail:
1516     free_preserved_data(ks);
1517     return -ENOMEM;
1518 }
1519
1520 /*
1521  * wait ktapio thread read all content in ring buffer.
1522  *
1523  * Here we use stupid approach to sync with ktapio thread,
1524  * note that we cannot use semaphore/completion/other sync method,
1525  * because ktapio thread could be killed by SIG_KILL in anytime, there
1526  * have no safe way to up semaphore or wake waitqueue before thread exit.
1527  *
1528  * we also cannot use waitqueue of current->signal->wait_chldexit to sync
1529  * exit, because mainthread and ktapio thread are in same thread group.
1530  *
1531  * Also ktap mainthread must wait ktapio thread exit, otherwise ktapio
1532  * thread will oops when access ktap structure.
1533  */
1534 static void wait_user_completion(ktap_state_t *ks)
1535 {
1536     struct task_struct *tsk = G(ks)->task;
1537     G(ks)->wait_user = 1;
1538
1539     while (1) {
1540         set_current_state(TASK_INTERRUPTIBLE);
1541         /* sleep for 100 msecs, and try again. */
1542         schedule_timeout(HZ / 10);
1543
1544         if (get_nr_threads(tsk) == 1)
1545             break;
1546     }
1547 }
1548
1549 static void sleep_loop(ktap_state_t *ks,
1550                      int (*actor)(ktap_state_t *ks, void *arg), void *arg)
1551 {
1552     while (!ks->stop) {
1553         set_current_state(TASK_INTERRUPTIBLE);
1554         /* sleep for 100 msecs, and try again. */
1555         schedule_timeout(HZ / 10);
1556
1557         if (actor(ks, arg))
1558             return;
1559     }
1560 }
1561
1562 static int sl_wait_task_pause_actor(ktap_state_t *ks, void *arg)
1563 {
1564     struct task_struct *task = (struct task_struct *)arg;
1565
1566     if (task->state)
1567         return 1;
1568     else
1569         return 0;
1570 }
1571
1572 static int sl_wait_task_exit_actor(ktap_state_t *ks, void *arg)
1573 {
1574     struct task_struct *task = (struct task_struct *)arg;
1575
1576     if (signal_pending(current)) {
1577         flush_signals(current);
1578
1579         /* newline for handle CTRL+C display as ^C */
1580         kp_puts(ks, "\n");
1581         return 1;
1582     }
1583 }

```

```

1584     /* stop waiting if target pid is exited */
1585     if (task && task->state == TASK_DEAD)
1586         return 1;
1587
1588     return 0;
1589 }
1590
1591 /* wait user interrupt, signal killed */
1592 static void wait_user_interrupt(ktap_state_t *ks)
1593 {
1594     struct task_struct *task = G(ks)->trace_task;
1595
1596     if (G(ks)->state == KTAP_EXIT || G(ks)->state == KTAP_ERROR)
1597         return;
1598
1599     /* let tracing goes now. */
1600     ks->stop = 0;
1601
1602     if (G(ks)->parm->workload) {
1603         /* make sure workload is in pause state
1604          * so it won't miss the signal */
1605         sleep_loop(ks, sl_wait_task_pause_actor, task);
1606         /* tell workload process to start executing */
1607         send_sig(SIGINT, G(ks)->trace_task, 0);
1608     }
1609
1610     if (!G(ks)->parm->quiet)
1611         kp_printf(ks, "Tracing... Hit Ctrl-C to end.\n");
1612
1613     sleep_loop(ks, sl_wait_task_exit_actor, task);
1614 }
1615
1616 /*
1617 * ktap exit, free all resources.
1618 */
1619 void kp_vm_exit(ktap_state_t *ks)
1620 {
1621     if (!list_empty(&G(ks)->events_head) ||
1622         !list_empty(&G(ks)->timers))
1623         wait_user_interrupt(ks);
1624
1625     kp_exit_timers(ks);
1626     kp_events_exit(ks);
1627
1628     /* free all resources got by ktap */
1629 #ifdef CONFIG_KTAP_FFI
1630     ffi_free_symbols(ks);
1631 #endif
1632     kp_str_freeall(ks);
1633     kp_mempool_destroy(ks);
1634
1635     func_closeuv(ks, 0); /* close all open upvals, let below call free it */
1636     kp_obj_freeall(ks);
1637
1638     kp_vm_exit_thread(ks);
1639     kp_free(ks, ks->stack);
1640
1641     free_preserved_data(ks);
1642     free_cpumask_var(G(ks)->cpumask);
1643
1644     wait_user_completion(ks);
1645
1646     /* should invoke after wait_user_completion */
1647     if (G(ks)->trace_task)
1648         put_task_struct(G(ks)->trace_task);
1649
1650     kp_transport_exit(ks);
1651     kp_free(ks, ks); /* free self */
1652 }
1653
1654 /*
1655 * ktap mainthread initialization
1656 */
1657 ktap_state_t *kp_vm_new_state(ktap_option_t *parm, struct dentry *dir)
1658 {
1659     ktap_state_t *ks;

```

```

1660 ktap\_global\_state\_t *g;
1661 pid_t pid;
1662 int cpu;
1663
1664 ks = kzalloc(sizeof(ktap\_state\_t) + sizeof(ktap\_global\_state\_t),
1665             GFP_KERNEL);
1666 if (!ks)
1667     return NULL;
1668
1669 g(ks) = (ktap\_global\_state\_t *)(ks + 1);
1670 g = g(ks);
1671 g->mainthread = ks;
1672 g->task = current;
1673 g->parm = parm;
1674 g->str_lock = (arch_spinlock_t)__ARCH_SPIN_LOCK_UNLOCKED;
1675 g->strmask = ~(int)0;
1676 g->uvhead.prev = &g->uvhead;
1677 g->uvhead.next = &g->uvhead;
1678 g->state = KTAP\_RUNNING;
1679 INIT_LIST_HEAD(&(g->timers));
1680 INIT_LIST_HEAD(&(g->events_head));
1681
1682 if (kp\_transport\_init(ks, dir))
1683     goto out;
1684
1685 ks->stack = kp\_malloc(ks, KTAP\_STACK\_SIZE\_BYTES);
1686 if (!ks->stack)
1687     goto out;
1688
1689 ks->stack_last = ks->stack + KTAP\_STACK\_SIZE;
1690 ks->top = ks->stack;
1691
1692 pid = (pid_t)parm->trace_pid;
1693 if (pid != -1) {
1694     struct task_struct *task;
1695
1696     rcu_read_lock();
1697     task = pid_task(find_vpid(pid), PIDTYPE_PID);
1698     if (!task) {
1699         kp\_error(ks, "cannot find pid %d\n", pid);
1700         rcu_read_unlock();
1701         goto out;
1702     }
1703     g->trace_task = task;
1704     get_task_struct(task);
1705     rcu_read_unlock();
1706 }
1707
1708 if( !alloc_cpumask_var(&g->cpumask, GFP_KERNEL))
1709     goto out;
1710
1711 cpumask_copy(g->cpumask, cpu_online_mask);
1712
1713 cpu = parm->trace_cpu;
1714 if (cpu != -1) {
1715     if (!cpu_online(cpu)) {
1716         kp\_error(ks, "ktap: cpu %d is not online\n", cpu);
1717         goto out;
1718     }
1719
1720     cpumask_clear(g->cpumask);
1721     cpumask_set_cpu(cpu, g->cpumask);
1722 }
1723
1724 if (kp\_mempool\_init(ks, KP\_MAX\_MEMPOOL\_SIZE))
1725     goto out;
1726
1727 if (kp\_str\_resize(ks, 1024 - 1)) /* set string hashtable size */
1728     goto out;
1729
1730 if (init\_registry(ks))
1731     goto out;
1732 if (init\_arguments(ks, parm->argc, parm->argv))
1733     goto out;
1734
1735 /* init librarys */

```

```
1736     if (kp\_lib\_init\_base(ks))
1737         goto out;
1738     if (kp\_lib\_init\_kdebug(ks))
1739         goto out;
1740     if (kp\_lib\_init\_timer(ks))
1741         goto out;
1742     if (kp\_lib\_init\_ansi(ks))
1743         goto out;
1744     #ifdef CONFIG_KTAP_FFI
1745         if (kp\_lib\_init\_ffi(ks))
1746             goto out;
1747     #endif
1748     if (kp\_lib\_init\_table(ks))
1749         goto out;
1750
1751     if (kp\_lib\_init\_net(ks))
1752         goto out;
1753
1754     if (init\_preserved\_data(ks))
1755         goto out;
1756
1757     if (kp\_events\_init(ks))
1758         goto out;
1759
1760     return ks;
1761
1762 out:
1763     g->state = KTAP\_ERROR;
1764     kp\_vm\_exit(ks);
1765     return NULL;
1766 }
1767
```

[One Level Up](#)

[Top Level](#)

userspace/kp_parse.c - ktap

Global variables defined

- [priority](#)

Data types defined

- [BinOpr](#)
- [BinOpr](#)
- [ExpDesc](#)
- [ExpDesc](#)
- [ExpKind](#)
- [FuncScope](#)
- [FuncScope](#)
- [FuncState](#)
- [FuncState](#)
- [LHSVarList](#)
- [LHSVarList](#)
- [VarIndex](#)

Functions defined

- [assign_adjust](#)
- [assign_hazard](#)
- [bcemit_INS](#)
- [bcemit_arith](#)
- [bcemit_binop](#)
- [bcemit_binop_left](#)
- [bcemit_branch](#)
- [bcemit_branch_f](#)
- [bcemit_branch_t](#)
- [bcemit_comp](#)
- [bcemit_jmp](#)
- [bcemit_method](#)
- [bcemit_nil](#)
- [bcemit_store](#)

- [bcemit_store_incr](#)
- [bcemit_unop](#)
- [bcopisret](#)
- [bcreg_bump](#)
- [bcreg_free](#)
- [bcreg_reserve](#)
- [const_gc](#)
- [const_num](#)
- [const_str](#)
- [err_limit](#)
- [err_syntax](#)
- [err_token](#)
- [expr](#)
- [expr_binop](#)
- [expr_bracket](#)
- [expr_cond](#)
- [expr_discharge](#)
- [expr_field](#)
- [expr_free](#)
- [expr_index](#)
- [expr_init](#)
- [expr_kvalue](#)
- [expr_list](#)
- [expr_next](#)
- [expr_numiszero](#)
- [expr_primary](#)
- [expr_simple](#)
- [expr_str](#)
- [expr_table](#)
- [expr_toanyreg](#)
- [expr_tonextreg](#)
- [expr_toreg](#)
- [expr_toreg_nobbranch](#)

- [expr_toval](#)
- [expr_unop](#)
- [foldarith](#)
- [fs_finish](#)
- [fs_fixup_bc](#)
- [fs_fixup_k](#)
- [fs_fixup_line](#)
- [fs_fixup_ret](#)
- [fs_fixup_uv1](#)
- [fs_fixup_uv2](#)
- [fs_fixup_var](#)
- [fs_init](#)
- [fs_prep_line](#)
- [fs_prep_var](#)
- [fscope_begin](#)
- [fscope_end](#)
- [fscope_uvmark](#)
- [gola_close](#)
- [gola_findlabel](#)
- [gola_fixup](#)
- [gola_new](#)
- [gola_patch](#)
- [gola_resolve](#)
- [invertcond](#)
- [jmp_append](#)
- [jmp_dropval](#)
- [jmp_next](#)
- [jmp_novalue](#)
- [jmp_patch](#)
- [jmp_patchins](#)
- [jmp_patchtestreg](#)
- [jmp_patchval](#)
- [jmp_tohere](#)

- [kp_parse](#)
- [kp_parse_keepstr](#)
- [lex_check](#)
- [lex_match](#)
- [lex_opt](#)
- [lex_str](#)
- [number_foldarith](#)
- [parse_args](#)
- [parse_assignment](#)
- [parse_block](#)
- [parse_body](#)
- [parse_body_no_args](#)
- [parse_break](#)
- [parse_call_assign](#)
- [parse_chunk](#)
- [parse_for](#)
- [parse_for_iter](#)
- [parse_for_num](#)
- [parse_func](#)
- [parse_if](#)
- [parse_isend](#)
- [parse_label](#)
- [parse_local](#)
- [parse_params](#)
- [parse_repeat](#)
- [parse_return](#)
- [parse_stmt](#)
- [parse_then](#)
- [parse_timer](#)
- [parse_trace](#)
- [parse_while](#)
- [predict_next](#)
- [synlevel_begin](#)

- [token2binop](#)
- [var_add](#)
- [var_lookup](#)
- [var_lookup_local](#)
- [var_lookup_uv](#)
- [var_new](#)
- [var_remove](#)

Macros defined

- [FLS](#)
- [FSCOPE_BREAK](#)
- [FSCOPE_GOLA](#)
- [FSCOPE_LOOP](#)
- [FSCOPE_NOCLOSE](#)
- [FSCOPE_UPVAL](#)
- [KP_MAX_VSTACK](#)
- [NAME_BREAK](#)
- [UNARY_PRIORITY](#)
- [VARNAMEDEF](#)
- [VARNAMEENUM](#)
- [VARNAMEENUM](#)
- [VSTACK_GOTO](#)
- [VSTACK_LABEL](#)
- [VSTACK_VAR_RW](#)
- [bcemit_ABC](#)
- [bcemit_AD](#)
- [bcemit_AJ](#)
- [bcptr](#)
- [checkcond](#)
- [checklimit](#)
- [checklimitgt](#)
- [checku8](#)
- [const_pri](#)
- [expr_hasjump](#)

- [expr_isk](#)
- [expr_isk_nojump](#)
- [expr_isnumk](#)
- [expr_isnumk_nojump](#)
- [expr_isstrk](#)
- [expr_numberV](#)
- [expr_numtv](#)
- [fs_fixup_line](#)
- [fs_fixup_var](#)
- [fs_prep_line](#)
- [fs_prep_var](#)
- [gola_isgoto](#)
- [gola_isgotolabel](#)
- [gola_islabel](#)
- [hsize2hbits](#)
- [synlevel_end](#)
- [tvhaskslot](#)
- [tvkslot](#)
- [var_get](#)
- [var_lookup](#)
- [var_new_fixed](#)
- [var_new_lit](#)

Source code

```

1  /*
2  * ktap parser (source code -> bytecode).
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Adapted from luajit and lua interpreter.
9  * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,

```

```

23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include "../include/ktap_types.h"
27 #include "../include/ktap_err.h"
28 #include "kp_util.h"
29 #include "kp_lex.h"
30
31 /* Fixed internal variable names. */
32 #define VARNAMEDEF(_) \
33     _(FOR_IDX, "(for index)") \
34     _(FOR_STOP, "(for limit)") \
35     _(FOR_STEP, "(for step)") \
36     _(FOR_GEN, "(for generator)") \
37     _(FOR_STATE, "(for state)") \
38     _(FOR_CTL, "(for control)")
39
40 enum {
41     VARNAME_END,
42 #define VARNAMEENUM(name, str) VARNAME_##name,
43     VARNAMEDEF(VARNAMEENUM)
44 #undef VARNAMEENUM
45     VARNAME__MAX
46 };
47
48 /* -- Parser structures and definitions ----- */
49
50 /* Expression kinds. */
51 typedef enum {
52     /* Constant expressions must be first and in this order: */
53     VKNIL,
54     VKFALSE,
55     VKTRUE,
56     VKSTR, /* sval = string value */
57     VKNUM, /* nval = number value */
58     VKLAST = VKNUM,
59     VKCDATA, /* nval = cdata value, not treated as a constant expression */
60     /* Non-constant expressions follow: */
61     VLOCAL, /* info = local register, aux = vstack index */
62     VUPVAL, /* info = upvalue index, aux = vstack index */
63     VGLOBAL, /* sval = string value */
64     VINDEXED, /* info = table register, aux = index reg/byte/string const */
65     VJMP, /* info = instruction PC */
66     VRELOCABLE, /* info = instruction PC */
67     VNONRELOC, /* info = result register */
68     VCALL, /* info = instruction PC, aux = base */
69     VVOID,
70
71     VARGN,
72     VARGSTR,
73     VARGNAME,
74     VPID,
75     VTID,
76     VUID,
77     VCPU,
78     VEEXECNAME,
79     VMAX
80 } ExpKind;
81
82 /* Expression descriptor. */
83 typedef struct ExpDesc {
84     union {
85         struct {
86             uint32_t info; /* Primary info. */
87             uint32_t aux; /* Secondary info. */
88         } s;
89         ktap_val_t nval; /* Number value. */
90         ktap_str_t *sval; /* String value. */
91     } u;
92     ExpKind k;
93     BCPos t; /* True condition jump list. */
94     BCPos f; /* False condition jump list. */
95 } ExpDesc;
96
97 /* Macros for expressions. */
98 #define expr_hasjump(e) ((e)->t != (e)->f)

```

```

99
100 #define expr_isk(e) ((e)->k <= VKLAST)
101 #define expr_isk_nojump(e) (expr_isk(e) && !expr_hasjump(e))
102 #define expr_isnumk(e) ((e)->k == VKNUM)
103 #define expr_isnumk_nojump(e) (expr_isnumk(e) && !expr_hasjump(e))
104 #define expr_isstrk(e) ((e)->k == VKSTR)
105
106 #define expr_numtv(e) (&(e)->u.nval)
107 #define expr_numberV(e) nvalue(expr_numtv((e)))
108
109 /* Initialize expression. */
110 static inline void expr_init(ExpDesc *e, ExpKind k, uint32_t info)
111 {
112     e->k = k;
113     e->u.s.info = info;
114     e->f = e->t = NO_JMP;
115 }
116
117 /* Check number constant for +-0. */
118 static int expr_numiszero(ExpDesc *e)
119 {
120     ktap_val_t *o = expr_numtv(e);
121     return (nvalue(o) == 0);
122 }
123
124 /* Per-function linked list of scope blocks. */
125 typedef struct FuncScope {
126     struct FuncScope *prev; /* Link to outer scope. */
127     int vstart; /* Start of block-local variables. */
128     uint8_t nactvar; /* Number of active vars outside the scope. */
129     uint8_t flags; /* Scope flags. */
130 } FuncScope;
131
132 #define FSCOPE_LOOP 0x01 /* Scope is a (breakable) loop. */
133 #define FSCOPE_BREAK 0x02 /* Break used in scope. */
134 #define FSCOPE_GOLA 0x04 /* Goto or label used in scope. */
135 #define FSCOPE_UPVAL 0x08 /* Upvalue in scope. */
136 #define FSCOPE_NOCLOSE 0x10 /* Do not close upvalues. */
137
138 #define NAME_BREAK ((ktap_str_t*)(uintptr_t)1)
139
140 /* Index into variable stack. */
141 typedef uint16_t VarIndex;
142 #define KP_MAX_VSTACK (65536 - KP_MAX_UPVAL)
143
144 /* Variable/goto/label info. */
145 #define VSTACK_VAR_RW 0x01 /* R/W variable. */
146 #define VSTACK_GOTO 0x02 /* Pending goto. */
147 #define VSTACK_LABEL 0x04 /* Label. */
148
149 /* Per-function state. */
150 typedef struct FuncState {
151     ktap_tab_t *kt; /* Hash table for constants. */
152     LexState *ls; /* Lexer state. */
153     FuncScope *bl; /* Current scope. */
154     struct FuncState *prev; /* Enclosing function. */
155     BCPos pc; /* Next bytecode position. */
156     BCPos lasttarget; /* Bytecode position of last jump target. */
157     BCPos jmpc; /* Pending jump list to next bytecode. */
158     BCReg freereg; /* First free register. */
159     BCReg nactvar; /* Number of active local variables. */
160     BCReg nkn, nkgc; /* Number of ktap_number/ktap_obj_t constants */
161     BCLine linedefined; /* First line of the function definition. */
162     BCInsLine *bcbase; /* Base of bytecode stack. */
163     BCPos bclim; /* Limit of bytecode stack. */
164     int vbase; /* Base of variable stack for this function. */
165     uint8_t flags; /* Prototype flags. */
166     uint8_t numparams; /* Number of parameters. */
167     uint8_t framesize; /* Fixed frame size. */
168     uint8_t nuv; /* Number of upvalues */
169     VarIndex varmap[KP_MAX_LOCVAR]; /* Map from register to variable idx. */
170     VarIndex uvmap[KP_MAX_UPVAL]; /* Map from upvalue to variable idx. */
171     VarIndex uvtmp[KP_MAX_UPVAL]; /* Temporary upvalue map. */
172 } FuncState;
173
174 /* Binary and unary operators. ORDER OPR */

```

```

175 typedef enum BinOpr {
176     OPR_ADD, OPR_SUB, OPR_MUL, OPR_DIV, OPR_MOD, OPR_POW, /* ORDER ARITH */
177     OPR_CONCAT,
178     OPR_NE, OPR_EQ,
179     OPR_LT, OPR_GE, OPR_LE, OPR_GT,
180     OPR_AND, OPR_OR,
181     OPR_NOBINOPR
182 } BinOpr;
183
184 KP_STATIC_ASSERT((int)BC_ISGE-(int)BC_ISLT == (int)OPR_GE-(int)OPR_LT);
185 KP_STATIC_ASSERT((int)BC_ISLE-(int)BC_ISLT == (int)OPR_LE-(int)OPR_LT);
186 KP_STATIC_ASSERT((int)BC_ISGT-(int)BC_ISLT == (int)OPR_GT-(int)OPR_LT);
187 KP_STATIC_ASSERT((int)BC_SUBVV-(int)BC_ADDVV == (int)OPR_SUB-(int)OPR_ADD);
188 KP_STATIC_ASSERT((int)BC_MULVV-(int)BC_ADDVV == (int)OPR_MUL-(int)OPR_ADD);
189 KP_STATIC_ASSERT((int)BC_DIVVV-(int)BC_ADDVV == (int)OPR_DIV-(int)OPR_ADD);
190 KP_STATIC_ASSERT((int)BC_MODVV-(int)BC_ADDVV == (int)OPR_MOD-(int)OPR_ADD);
191
192 /* -- Error handling ----- */
193
194 static void err_syntax(LexState *ls, ErrMsg em)
195 {
196     kp_lex_error(ls, ls->tok, em);
197 }
198
199 static void err_token(LexState *ls, LexToken tok)
200 {
201     kp_lex_error(ls, ls->tok, KP_ERR_XTOKEN, kp_lex_token2str(ls, tok));
202 }
203
204 static void err_limit(FuncState *fs, uint32_t limit, const char *what)
205 {
206     if (fs->linedefined == 0)
207         kp_lex_error(fs->ls, 0, KP_ERR_XLIMM, limit, what);
208     else
209         kp_lex_error(fs->ls, 0, KP_ERR_XLIMF, fs->linedefined,
210                     limit, what);
211 }
212
213 #define checklimit(fs, v, l, m)    if ((v) >= (l)) err_limit(fs, l, m)
214 #define checklimitgt(fs, v, l, m) if ((v) > (l)) err_limit(fs, l, m)
215 #define checkcond(ls, c, em)      { if (!(c)) err_syntax(ls, em); }
216
217 /* -- Management of constants ----- */
218
219 /* Return bytecode encoding for primitive constant. */
220 #define const_pri(e)    ((e)->k)
221
222 #define tvhaskslot(o)  (is_number(o))
223 #define tvkslot(o)    (nvalue(o))
224
225 /* Add a number constant. */
226 static BCREg const_num(FuncState *fs, ExpDesc *e)
227 {
228     ktap_val_t *o;
229
230     kp_assert(expr_isnumk(e));
231     o = kp_tab_set(fs->kt, &e->u.nval);
232     if (tvhaskslot(o))
233         return tvkslot(o);
234     set_number(o, fs->nkn);
235     return fs->nkn++;
236 }
237
238 /* Add a GC object constant. */
239 static BCREg const_gc(FuncState *fs, ktap_obj_t *gc, uint32_t itype)
240 {
241     ktap_val_t key, *o;
242
243     setitype(&key, itype);
244     key.val.gc = gc;
245     o = kp_tab_set(fs->kt, &key);
246     if (tvhaskslot(o))
247         return tvkslot(o);
248     set_number(o, fs->nkgc);
249     return fs->nkgc++;
250 }

```

```

251
252 /* Add a string constant. */
253 static BCReg const_str(FuncState *fs, ExpDesc *e)
254 {
255     kp_assert(expr_isstrk(e) || e->k == VGLOBAL);
256     return const_gc(fs, obj2gc(e->u.sval), KTAP_TSTR);
257 }
258
259 /* Anchor string constant. */
260 ktap_str_t *kp_parse_keepstr(LexState *ls, const char *str, size_t len)
261 {
262     ktap_val_t v;
263     ktap_str_t *s = kp_str_new(str, len);
264
265     set_string(&v, s);
266     ktap_val_t *tv = kp_tab_set(ls->fs->kt, &v);
267     if (is_nil(tv))
268         set_bool(tv, 1);
269     return s;
270 }
271
272 /* -- Jump list handling ----- */
273
274 /* Get next element in jump list. */
275 static BCPos jmp_next(FuncState *fs, BCPos pc)
276 {
277     ptrdiff_t delta = bc_j(fs->bcbase[pc].ins);
278     if ((BCPos)delta == NO_JMP)
279         return NO_JMP;
280     else
281         return (BCPos)(((ptrdiff_t)pc+1)+delta);
282 }
283
284 /* Check if any of the instructions on the jump list produce no value. */
285 static int jmp_novalue(FuncState *fs, BCPos list)
286 {
287     for (; list != NO_JMP; list = jmp_next(fs, list)) {
288         BCIns p = fs->bcbase[list >= 1 ? list-1 : list].ins;
289         if (!(bc_op(p) == BC_ISTC || bc_op(p) == BC_ISFC ||
290             bc_a(p) == NO_REG))
291             return 1;
292     }
293     return 0;
294 }
295
296 /* Patch register of test instructions. */
297 static int jmp_patchtestreg(FuncState *fs, BCPos pc, BCReg reg)
298 {
299     BCInsLine *ilp = &fs->bcbase[pc >= 1 ? pc-1 : pc];
300     BCOp op = bc_op(ilp->ins);
301
302     if (op == BC_ISTC || op == BC_ISFC) {
303         if (reg != NO_REG && reg != bc_d(ilp->ins)) {
304             setbc_a(&ilp->ins, reg);
305         } else /* Nothing to store or already in the right register */
306             setbc_op(&ilp->ins, op+(BC_IST-BC_ISTC));
307             setbc_a(&ilp->ins, 0);
308         }
309     } else if (bc_a(ilp->ins) == NO_REG) {
310         if (reg == NO_REG) {
311             ilp->ins =
312                 BCINS_AJ(BC_JMP, bc_a(fs->bcbase[pc].ins), 0);
313         } else {
314             setbc_a(&ilp->ins, reg);
315             if (reg >= bc_a(ilp[1].ins))
316                 setbc_a(&ilp[1].ins, reg+1);
317         }
318     } else {
319         return 0; /* Cannot patch other instructions. */
320     }
321     return 1;
322 }
323
324 /* Drop values for all instructions on jump list. */
325 static void jmp_dropval(FuncState *fs, BCPos list)
326 {

```

```

327     for (; list != NO_JUMP; list = jmp_next(fs, list))
328         jmp_patchtestreg(fs, list, NO_REG);
329 }
330
331 /* Patch jump instruction to target. */
332 static void jmp_patchins(FuncState *fs, BCPos pc, BCPos dest)
333 {
334     BCIns *jmp = &fs->bcbase[pc].ins;
335     BCPos offset = dest-(pc+1)+BCBIAS_J;
336
337     kp_assert(dest != NO_JUMP);
338     if (offset > BCMAX_D)
339         err_syntax(fs->ls, KP_ERR_XJUMP);
340     setbc_d(jmp, offset);
341 }
342
343 /* Append to jump list. */
344 static void jmp_append(FuncState *fs, BCPos *l1, BCPos l2)
345 {
346     if (l2 == NO_JUMP) {
347         return;
348     } else if (*l1 == NO_JUMP) {
349         *l1 = l2;
350     } else {
351         BCPos list = *l1;
352         BCPos next;
353         /* Find last element. */
354         while ((next = jmp_next(fs, list)) != NO_JUMP)
355             list = next;
356         jmp_patchins(fs, list, l2);
357     }
358 }
359
360 /* Patch jump list and preserve produced values. */
361 static void jmp_patchval(FuncState *fs, BCPos list, BCPos vtarget,
362     BCReg reg, BCPos dtarget)
363 {
364     while (list != NO_JUMP) {
365         BCPos next = jmp_next(fs, list);
366         if (jmp_patchtestreg(fs, list, reg)) {
367             /* Jump to target with value. */
368             jmp_patchins(fs, list, vtarget);
369         } else {
370             /* Jump to default target. */
371             jmp_patchins(fs, list, dtarget);
372         }
373         list = next;
374     }
375 }
376
377 /* Jump to following instruction. Append to list of pending jumps. */
378 static void jmp_tohere(FuncState *fs, BCPos list)
379 {
380     fs->lasttarget = fs->pc;
381     jmp_append(fs, &fs->jpc, list);
382 }
383
384 /* Patch jump list to target. */
385 static void jmp_patch(FuncState *fs, BCPos list, BCPos target)
386 {
387     if (target == fs->pc) {
388         jmp_tohere(fs, list);
389     } else {
390         kp_assert(target < fs->pc);
391         jmp_patchval(fs, list, target, NO_REG, target);
392     }
393 }
394
395 /* -- Bytecode register allocator ----- */
396
397 /* Bump frame size. */
398 static void bcreg_bump(FuncState *fs, BCReg n)
399 {
400     BCReg sz = fs->freereg + n;
401
402     if (sz > fs->framesize) {

```

```

403     if (sz >= KP_MAX_SLOTS)
404         err_syntax(fs->ls, KP_ERR_XSLOTS);
405     fs->framesize = (uint8_t)sz;
406 }
407 }
408
409 /* Reserve registers. */
410 static void bcreg_reserve(FuncState *fs, BCReg n)
411 {
412     bcreg_bump(fs, n);
413     fs->freereg += n;
414 }
415
416 /* Free register. */
417 static void bcreg_free(FuncState *fs, BCReg reg)
418 {
419     if (reg >= fs->nactvar) {
420         fs->freereg--;
421         kp_assert(reg == fs->freereg);
422     }
423 }
424
425 /* Free register for expression. */
426 static void expr_free(FuncState *fs, ExpDesc *e)
427 {
428     if (e->k == VNONRELOC)
429         bcreg_free(fs, e->u.s.info);
430 }
431
432 /* -- Bytecode emitter ----- */
433
434 /* Emit bytecode instruction. */
435 static BCPos bcemit_INS(FuncState *fs, BCIns ins)
436 {
437     BCPos pc = fs->pc;
438     LexState *ls = fs->ls;
439
440     jmp_patchval(fs, fs->jpc, pc, NO_REG, pc);
441     fs->jpc = NO_JUMP;
442     if (pc >= fs->bclim) {
443         ptrdiff_t base = fs->bcbase - ls->bcstack;
444         checklimit(fs, ls->sizebcstack, KP_MAX_BCINS,
445                 "bytecode instructions");
446         if (!ls->bcstack) {
447             ls->bcstack = malloc(sizeof(BCInsLine) * 20);
448             ls->sizebcstack = 20;
449         } else {
450             ls->bcstack = realloc(ls->bcstack,
451                 ls->sizebcstack * sizeof(BCInsLine) * 2);
452             ls->sizebcstack = ls->sizebcstack * 2;
453         }
454         fs->bclim = (BCPos)(ls->sizebcstack - base);
455         fs->bcbase = ls->bcstack + base;
456     }
457     fs->bcbase[pc].ins = ins;
458     fs->bcbase[pc].line = ls->lastline;
459     fs->pc = pc+1;
460     return pc;
461 }
462
463 #define bcemit_ABC(fs, o, a, b, c)    bcemit_INS(fs, BCINS_ABC(o, a, b, c))
464 #define bcemit_AD(fs, o, a, d)      bcemit_INS(fs, BCINS_AD(o, a, d))
465 #define bcemit_AJ(fs, o, a, j)      bcemit_INS(fs, BCINS_AJ(o, a, j))
466
467 #define bcptr(fs, e)                (&(fs->bcbase[(e)->u.s.info].ins))
468
469 /* -- Bytecode emitter for expressions ----- */
470
471 /* Discharge non-constant expression to any register. */
472 static void expr_discharge(FuncState *fs, ExpDesc *e)
473 {
474     BCIns ins;
475
476     if (e->k == VUPVAL) {
477         ins = BCINS_AD(BC_UGET, 0, e->u.s.info);
478     } else if (e->k == VGLOBAL) {

```

```

479     ins = BCINS_AD(BC_GGET, 0, const_str(fs, e));
480 } else if (e->k == VINDEXED) {
481     BCReg rc = e->u.s.aux;
482     if ((int32_t)rc < 0) {
483         ins = BCINS_ABC(BC_TGETS, 0, e->u.s.info, ~rc);
484     } else if (rc > BCMAX_C) {
485         ins = BCINS_ABC(BC_TGETB, 0, e->u.s.info,
486             rc-(BCMAX_C+1));
487     } else {
488         bcreg_free(fs, rc);
489         ins = BCINS_ABC(BC_TGETV, 0, e->u.s.info, rc);
490     }
491     bcreg_free(fs, e->u.s.info);
492 } else if (e->k == VCALL) {
493     e->u.s.info = e->u.s.aux;
494     e->k = VNONRELOC;
495     return;
496 } else if (e->k == VLOCAL) {
497     e->k = VNONRELOC;
498     return;
499 } else {
500     return;
501 }
502
503 e->u.s.info = bcemit_INS(fs, ins);
504 e->k = VRELOCABLE;
505 }
506
507 /* Emit bytecode to set a range of registers to nil. */
508 static void bcemit_nil(FuncState *fs, BCReg from, BCReg n)
509 {
510     if (fs->pc > fs->lasttarget) { /* No jumps to current position? */
511         BCIns *ip = &fs->bcbase[fs->pc-1].ins;
512         BCReg pto, pfrom = bc_a(*ip);
513         /* Try to merge with the previous instruction. */
514         switch (bc_op(*ip)) {
515             case BC_KPRI:
516                 if (bc_d(*ip) != -KTAP_TNIL) break;
517                 if (from == pfrom) {
518                     if (n == 1)
519                         return;
520                 } else if (from == pfrom+1) {
521                     from = pfrom;
522                     n++;
523                 } else {
524                     break;
525                 }
526                 /* Replace KPRI. */
527                 *ip = BCINS_AD(BC_KNIL, from, from+n-1);
528                 return;
529             case BC_KNIL:
530                 pto = bc_d(*ip);
531                 /* Can we connect both ranges? */
532                 if (pfrom <= from && from <= pto+1) {
533                     if (from+n-1 > pto) {
534                         /* Patch previous instruction range. */
535                         setbc_d(ip, from+n-1);
536                     }
537                     return;
538                 }
539                 break;
540             default:
541                 break;
542         }
543     }
544
545     /* Emit new instruction or replace old instruction. */
546     bcemit_INS(fs, n == 1 ? BCINS_AD(BC_KPRI, from, VKNIL) :
547         BCINS_AD(BC_KNIL, from, from+n-1));
548 }
549
550 /* Discharge an expression to a specific register. Ignore branches. */
551 static void expr_toreg_nobbranch(FuncState *fs, ExpDesc *e, BCReg reg)
552 {
553     BCIns ins;
554

```

```

555     expr\_discharge(fs, e);
556     if (e->k == VKSTR) {
557         ins = BCINS\_AD(BC_KSTR, reg, const\_str(fs, e));
558     } else if (e->k == VKNUM) {
559         ktag\_number n = expr\_numberV(e);
560         if (n >= 0 && n <= 0xffff) {
561             ins = BCINS\_AD(BC_KSHORT, reg, (BCReg)(uint16_t)n);
562         } else
563             ins = BCINS\_AD(BC_KNUM, reg, const\_num(fs, e));
564     } else if (e->k == VRELOCABLE) {
565         setbc\_a(bcptr(fs, e), reg);
566         goto noins;
567     } else if (e->k == VNONRELOC) {
568         if (reg == e->u.s.info)
569             goto noins;
570         ins = BCINS\_AD(BC_MOV, reg, e->u.s.info);
571     } else if (e->k == VKNIL) {
572         bcemit\_nil(fs, reg, 1);
573         goto noins;
574     } else if (e->k <= VKTRUE) {
575         ins = BCINS\_AD(BC_KPRI, reg, const\_pri(e));
576     } else if (e->k == VARGN) {
577         ins = BCINS\_AD(BC_VARGN, reg, e->u.s.info);
578     } else if (e->k > VARGN && e->k < VMAX) {
579         ins = BCINS\_AD(e->k - VARGN + BC_VARGN, reg, 0);
580     } else {
581         kp\_assert(e->k == VVOID || e->k == VJMP);
582         return;
583     }
584     bcemit\_INS(fs, ins);
585 noins:
586     e->u.s.info = reg;
587     e->k = VNONRELOC;
588 }
589
590 /* Forward declaration. */
591 static BCPos bcemit\_jmp(FuncState *fs);
592
593 /* Discharge an expression to a specific register. */
594 static void expr\_toreg(FuncState *fs, ExpDesc *e, BCReg reg)
595 {
596     expr\_toreg\_nobran(fs, e, reg);
597     if (e->k == VJMP) {
598         /* Add it to the true jump list. */
599         jmp\_append(fs, &e->t, e->u.s.info);
600     }
601     if (expr\_hasjump(e)) { /* Discharge expression with branches. */
602         BCPos jend, jfalse = NO\_JMP, jtrue = NO\_JMP;
603         if (jmp\_novalue(fs, e->t) || jmp\_novalue(fs, e->f)) {
604             BCPos jval = (e->k == VJMP) ? NO\_JMP : bcemit\_jmp(fs);
605             jfalse = bcemit\_AD(fs, BC_KPRI, reg, VKFALSE);
606             bcemit\_AJ(fs, BC_JMP, fs->freereg, 1);
607             jtrue = bcemit\_AD(fs, BC_KPRI, reg, VKTRUE);
608             jmp\_tohere(fs, jval);
609         }
610         jend = fs->pc;
611         fs->lasttarget = jend;
612         jmp\_patchval(fs, e->f, jend, reg, jfalse);
613         jmp\_patchval(fs, e->t, jend, reg, jtrue);
614     }
615     e->f = e->t = NO\_JMP;
616     e->u.s.info = reg;
617     e->k = VNONRELOC;
618 }
619
620 /* Discharge an expression to the next free register. */
621 static void expr\_tonextreg(FuncState *fs, ExpDesc *e)
622 {
623     expr\_discharge(fs, e);
624     expr\_free(fs, e);
625     bcreg\_reserve(fs, 1);
626     expr\_toreg(fs, e, fs->freereg - 1);
627 }
628
629 /* Discharge an expression to any register. */
630 static BCReg expr\_toanyreg(FuncState *fs, ExpDesc *e)

```

```

631 {
632     expr_discharge(fs, e);
633     if (e->k == VNONRELOC) {
634         if (!expr_hasjump(e))
635             return e->u.s.info; /* Already in a register. */
636         if (e->u.s.info >= fs->nactvar) {
637             /* Discharge to temp. register. */
638             expr_toreg(fs, e, e->u.s.info);
639             return e->u.s.info;
640         }
641     }
642     expr_tonextreg(fs, e); /* Discharge to next register. */
643     return e->u.s.info;
644 }
645
646 /* Partially discharge expression to a value. */
647 static void expr_toval(FuncState *fs, ExpDesc *e)
648 {
649     if (expr_hasjump(e))
650         expr_toanyreg(fs, e);
651     else
652         expr_discharge(fs, e);
653 }
654
655 /* Emit store for LHS expression. */
656 static void bcemit_store(FuncState *fs, ExpDesc *var, ExpDesc *e)
657 {
658     BCIns ins;
659
660     if (var->k == VLOCAL) {
661         fs->ls->vstack[var->u.s.aux].info |= VSTACK_VAR_RW;
662         expr_free(fs, e);
663         expr_toreg(fs, e, var->u.s.info);
664         return;
665     } else if (var->k == VUPVAL) {
666         fs->ls->vstack[var->u.s.aux].info |= VSTACK_VAR_RW;
667         expr_toval(fs, e);
668         if (e->k <= VKTRUE)
669             ins = BCINS_AD(BC_USETP, var->u.s.info, const_pri(e));
670         else if (e->k == VKSTR)
671             ins = BCINS_AD(BC_USETS, var->u.s.info,
672                 const_str(fs, e));
673         else if (e->k == VKNUM)
674             ins = BCINS_AD(BC_USETN, var->u.s.info,
675                 const_num(fs, e));
676         else
677             ins = BCINS_AD(BC_USETV, var->u.s.info,
678                 expr_toanyreg(fs, e));
679     } else if (var->k == VGLOBAL) {
680         BCReg ra = expr_toanyreg(fs, e);
681         ins = BCINS_AD(BC_GSET, ra, const_str(fs, var));
682     } else {
683         BCReg ra, rc;
684         kp_assert(var->k == VINDEXED);
685         ra = expr_toanyreg(fs, e);
686         rc = var->u.s.aux;
687         if ((int32_t)rc < 0) {
688             ins = BCINS_ABC(BC_TSETS, ra, var->u.s.info, ~rc);
689         } else if (rc > BCMAX_C) {
690             ins = BCINS_ABC(BC_TSETB, ra, var->u.s.info,
691                 rc-(BCMAX_C+1));
692         } else {
693             /*
694              * Free late allocated key reg to avoid assert on
695              * free of value reg. This can only happen when
696              * called from expr_table().
697              */
698             kp_assert(e->k != VNONRELOC || ra < fs->nactvar ||
699                 rc < ra || (bcreg_free(fs, rc),1));
700             ins = BCINS_ABC(BC_TSETV, ra, var->u.s.info, rc);
701         }
702     }
703     bcemit_INS(fs, ins);
704     expr_free(fs, e);
705 }
706

```

```

707 /* Emit store for '+=' expression. */
708 static void bcemit_store_incr(FuncState *fs, ExpDesc *var, ExpDesc *e)
709 {
710     BCIns ins;
711
712     if (var->k == VLOCAL) {
713         /* don't need to do like "var a=0; a+=1", just use 'a=a+1' */
714         err_syntax(fs->ls, KP_ERR_XSYMBOL);
715         return;
716     } else if (var->k == VUPVAL) {
717         fs->ls->vstack[var->u.s.aux].info |= VSTACK_VAR_RW;
718         expr_toval(fs, e);
719         if (e->k == VKNUM) {
720             ins = BCINS_AD(BC_UINCN, var->u.s.info,
721                 const_num(fs, e));
722         } else if (e->k <= VKTRUE || e->k == VKSTR) {
723             err_syntax(fs->ls, KP_ERR_XSYMBOL);
724             return;
725         } else
726             ins = BCINS_AD(BC_UINCV, var->u.s.info,
727                 expr_toanyreg(fs, e));
728     } else if (var->k == VGLOBAL) {
729         BCReg ra = expr_toanyreg(fs, e);
730         ins = BCINS_AD(BC_GINC, ra, const_str(fs, var));
731     } else {
732         BCReg ra, rc;
733         kp_assert(var->k == VINDEXED);
734         ra = expr_toanyreg(fs, e);
735         rc = var->u.s.aux;
736         if ((int32_t)rc < 0) {
737             ins = BCINS_ABC(BC_TINCS, ra, var->u.s.info, ~rc);
738         } else if (rc > BCMAX_C) {
739             ins = BCINS_ABC(BC_TINCB, ra, var->u.s.info,
740                 rc-(BCMAX_C+1));
741         } else {
742             /*
743             * Free late allocated key reg to avoid assert on
744             * free of value reg. This can only happen when
745             * called from expr_table().
746             */
747             kp_assert(e->k != VNONRELOC || ra < fs->nactvar ||
748                 rc < ra || (bcreg_free(fs, rc),1));
749             ins = BCINS_ABC(BC_TINCV, ra, var->u.s.info, rc);
750         }
751     }
752     bcemit_INS(fs, ins);
753     expr_free(fs, e);
754 }
755
756 /* Emit method lookup expression. */
757 static void bcemit_method(FuncState *fs, ExpDesc *e, ExpDesc *key)
758 {
759     BCReg idx, func, obj = expr_toanyreg(fs, e);
760
761     expr_free(fs, e);
762     func = fs->freereg;
763     bcemit_AD(fs, BC_MOV, func+1, obj); /* Copy object to first argument. */
764     kp_assert(expr_isstrk(key));
765     idx = const_str(fs, key);
766     if (idx <= BCMAX_C) {
767         bcreg_reserve(fs, 2);
768         bcemit_ABC(fs, BC_TGETS, func, obj, idx);
769     } else {
770         bcreg_reserve(fs, 3);
771         bcemit_AD(fs, BC_KSTR, func+2, idx);
772         bcemit_ABC(fs, BC_TGETV, func, obj, func+2);
773         fs->freereg--;
774     }
775     e->u.s.info = func;
776     e->k = VNONRELOC;
777 }
778
779 /* -- Bytecode emitter for branches ----- */
780 /* Emit unconditional branch. */

```

```

783 static BCPos bcemit_jump(FuncState *fs)
784 {
785     BCPos jpc = fs->jpc;
786     BCPos j = fs->pc - 1;
787     BCIns *ip = &fs->bcbase[j].ins;
788
789     fs->jpc = NO_JUMP;
790     if ((int32_t)j >= (int32_t)fs->lasttarget && bc_op(*ip) == BC_UCLO)
791         setbc_j(ip, NO_JUMP);
792     else
793         j = bcemit_AJ(fs, BC_JMP, fs->freereg, NO_JUMP);
794     jmp_append(fs, &j, jpc);
795     return j;
796 }
797
798 /* Invert branch condition of bytecode instruction. */
799 static void invertcond(FuncState *fs, ExpDesc *e)
800 {
801     BCIns *ip = &fs->bcbase[e->u.s.info - 1].ins;
802     setbc_op(ip, bc_op(*ip)^1);
803 }
804
805 /* Emit conditional branch. */
806 static BCPos bcemit_branch(FuncState *fs, ExpDesc *e, int cond)
807 {
808     BCPos pc;
809
810     if (e->k == VRELOCABLE) {
811         BCIns *ip = bcptr(fs, e);
812         if (bc_op(*ip) == BC_NOT) {
813             *ip = BCINS_AD(cond ? BC_ISF : BC_IST, 0, bc_d(*ip));
814             return bcemit_jump(fs);
815         }
816     }
817     if (e->k != VNONRELOC) {
818         bcreg_reserve(fs, 1);
819         expr_toreg_nobbranch(fs, e, fs->freereg-1);
820     }
821     bcemit_AD(fs, cond ? BC_ISTC : BC_ISFC, NO_REG, e->u.s.info);
822     pc = bcemit_jump(fs);
823     expr_free(fs, e);
824     return pc;
825 }
826
827 /* Emit branch on true condition. */
828 static void bcemit_branch_t(FuncState *fs, ExpDesc *e)
829 {
830     BCPos pc;
831
832     expr_discharge(fs, e);
833     if (e->k == VKSTR || e->k == VKNUM || e->k == VKTRUE)
834         pc = NO_JUMP; /* Never jump. */
835     else if (e->k == VJMP)
836         invertcond(fs, e), pc = e->u.s.info;
837     else if (e->k == VKFALSE || e->k == VKNIL)
838         expr_toreg_nobbranch(fs, e, NO_REG), pc = bcemit_jump(fs);
839     else
840         pc = bcemit_branch(fs, e, 0);
841     jmp_append(fs, &e->f, pc);
842     jmp_tohere(fs, e->t);
843     e->t = NO_JUMP;
844 }
845
846 /* Emit branch on false condition. */
847 static void bcemit_branch_f(FuncState *fs, ExpDesc *e)
848 {
849     BCPos pc;
850
851     expr_discharge(fs, e);
852     if (e->k == VKNIL || e->k == VKFALSE)
853         pc = NO_JUMP; /* Never jump. */
854     else if (e->k == VJMP)
855         pc = e->u.s.info;
856     else if (e->k == VKSTR || e->k == VKNUM || e->k == VKTRUE)
857         expr_toreg_nobbranch(fs, e, NO_REG), pc = bcemit_jump(fs);
858     else

```

```

859     pc = bcemit_branch(fs, e, 1);
860     jmp_append(fs, &e->t, pc);
861     jmp_tohere(fs, e->f);
862     e->f = NO JMP;
863 }
864
865 /* -- Bytecode emitter for operators ----- */
866
867 static ktap_number number_foldarith(ktap_number x, ktap_number y, int op)
868 {
869     switch (op) {
870     case OPR_ADD - OPR_ADD: return x + y;
871     case OPR_SUB - OPR_ADD: return x - y;
872     case OPR_MUL - OPR_ADD: return x * y;
873     case OPR_DIV - OPR_ADD: return x / y;
874     default: return x;
875     }
876 }
877
878 /* Try constant-folding of arithmetic operators. */
879 static int foldarith(BinOpr opr, ExpDesc *e1, ExpDesc *e2)
880 {
881     ktap_val_t o;
882     ktap_number n;
883
884     if (!expr_isnumk_nojump(e1) || !expr_isnumk_nojump(e2))
885         return 0;
886
887     if (opr == OPR_DIV && expr_numberV(e2) == 0)
888         return 0; /* do not attempt to divide by 0 */
889
890     if (opr == OPR_MOD)
891         return 0; /* ktap current do not support pow arith */
892
893     n = number_foldarith(expr_numberV(e1), expr_numberV(e2),
894                          (int)opr-OPR_ADD);
895     set_number(&o, n);
896     set_number(&e1->u.nval, n);
897     return 1;
898 }
899
900 /* Emit arithmetic operator. */
901 static void bcemit_arith(FuncState *fs, BinOpr opr, ExpDesc *e1, ExpDesc *e2)
902 {
903     BCReg rb, rc, t;
904     uint32_t op;
905
906     if (foldarith(opr, e1, e2))
907         return;
908     if (opr == OPR_POW) {
909         op = BC_POW;
910         rc = expr_toanyreg(fs, e2);
911         rb = expr_toanyreg(fs, e1);
912     } else {
913         op = opr-OPR_ADD+BC_ADDVV;
914         /*
915          * Must discharge 2nd operand first since VINDEXXED
916          * might free regs.
917          */
918         expr_toval(fs, e2);
919         if (expr_isnumk(e2) && (rc = const_num(fs, e2)) <= BCMAX_C)
920             op -= BC_ADDVV-BC_ADDVN;
921         else
922             rc = expr_toanyreg(fs, e2);
923         /* 1st operand discharged by bcemit_binop_left,
924          * but need KNUM/KSHORT. */
925         kp_assert(expr_isnumk(e1) || e1->k == VNONRELOC);
926         expr_toval(fs, e1);
927         /* Avoid two consts to satisfy bytecode constraints. */
928         if (expr_isnumk(e1) && !expr_isnumk(e2) &&
929             (t = const_num(fs, e1)) <= BCMAX_B) {
930             rb = rc; rc = t; op -= BC_ADDVV-BC_ADDVN;
931         } else {
932             rb = expr_toanyreg(fs, e1);
933         }
934     }

```

```

935  /* Using expr free might cause asserts if the order is wrong. */
936  if (e1->k == VNONRELOC && e1->u.s.info >= fs->nactvar)
937      fs->freereg--;
938  if (e2->k == VNONRELOC && e2->u.s.info >= fs->nactvar)
939      fs->freereg--;
940  e1->u.s.info = bcemit_ABC(fs, op, 0, rb, rc);
941  e1->k = VRELOCABLE;
942  }
943
944  /* Emit comparison operator. */
945  static void bcemit_comp(FuncState *fs, BinOpr opr, ExpDesc *e1, ExpDesc *e2)
946  {
947      ExpDesc *eret = e1;
948      BCIns ins;
949
950      expr_toval(fs, e1);
951      if (opr == OPR_EQ || opr == OPR_NE) {
952          BCOp op = opr == OPR_EQ ? BC_ISEQV : BC_ISNEV;
953          BCReg ra;
954
955          if (expr_isk(e1)) { /* Need constant in 2nd arg. */
956              e1 = e2;
957              e2 = eret;
958          }
959          ra = expr_toanyreg(fs, e1); /* First arg must be in a reg. */
960          expr_toval(fs, e2);
961          switch (e2->k) {
962              case VKNIL: case VKFALSE: case VKTRUE:
963                  ins = BCINS_AD(op+(BC_ISEQP-BC_ISEQV), ra,
964                      const_pri(e2));
965                  break;
966              case VKSTR:
967                  ins = BCINS_AD(op+(BC_ISEQS-BC_ISEQV), ra,
968                      const_str(fs, e2));
969                  break;
970              case VKNUM:
971                  ins = BCINS_AD(op+(BC_ISEQN-BC_ISEQV), ra,
972                      const_num(fs, e2));
973                  break;
974              default:
975                  ins = BCINS_AD(op, ra, expr_toanyreg(fs, e2));
976                  break;
977          }
978      } else {
979          uint32_t op = opr-OPR_LT+BC_ISLT;
980          BCReg ra, rd;
981          if ((op-BC_ISLT) & 1) { /* GT -> LT, GE -> LE */
982              e1 = e2; e2 = eret; /* Swap operands. */
983              op = ((op-BC_ISLT)^3)+BC_ISLT;
984              expr_toval(fs, e1);
985          }
986          rd = expr_toanyreg(fs, e2);
987          ra = expr_toanyreg(fs, e1);
988          ins = BCINS_AD(op, ra, rd);
989      }
990      /* Using expr free might cause asserts if the order is wrong. */
991      if (e1->k == VNONRELOC && e1->u.s.info >= fs->nactvar)
992          fs->freereg--;
993      if (e2->k == VNONRELOC && e2->u.s.info >= fs->nactvar)
994          fs->freereg--;
995      bcemit_INS(fs, ins);
996      eret->u.s.info = bcemit_jmp(fs);
997      eret->k = VJMP;
998  }
999
1000  /* Fixup left side of binary operator. */
1001  static void bcemit_binop_left(FuncState *fs, BinOpr op, ExpDesc *e)
1002  {
1003      if (op == OPR_AND) {
1004          bcemit_branch_t(fs, e);
1005      } else if (op == OPR_OR) {
1006          bcemit_branch_f(fs, e);
1007      } else if (op == OPR_CONCAT) {
1008          expr_tonextreg(fs, e);
1009      } else if (op == OPR_EQ || op == OPR_NE) {
1010          if (!expr_isk_nojump(e))

```

```

1011     expr\_toanyreg(fs, e);
1012 } else {
1013     if (!expr\_isnumk\_nojump(e))
1014         expr\_toanyreg(fs, e);
1015 }
1016 }
1017
1018 /* Emit binary operator. */
1019 static void bcemit\_binop(FuncState *fs, BinOpr op, ExpDesc *e1, ExpDesc *e2)
1020 {
1021     if (op <= OPR_POW) {
1022         bcemit\_arith(fs, op, e1, e2);
1023     } else if (op == OPR_AND) {
1024         kp\_assert(e1->t == NO\_JMP); /* List must be closed. */
1025         expr\_discharge(fs, e2);
1026         jmp\_append(fs, &e2->f, e1->f);
1027         *e1 = *e2;
1028     } else if (op == OPR_OR) {
1029         kp\_assert(e1->f == NO\_JMP); /* List must be closed. */
1030         expr\_discharge(fs, e2);
1031         jmp\_append(fs, &e2->t, e1->t);
1032         *e1 = *e2;
1033     } else if (op == OPR_CONCAT) {
1034         expr\_toval(fs, e2);
1035         if (e2->k == VRELOCABLE && bc\_op(*bcptr(fs, e2)) == BC_CAT) {
1036             kp\_assert(e1->u.s.info == bc\_b(*bcptr(fs, e2))-1);
1037             expr\_free(fs, e1);
1038             setbc\_b(bcptr(fs, e2), e1->u.s.info);
1039             e1->u.s.info = e2->u.s.info;
1040         } else {
1041             expr\_tonextreg(fs, e2);
1042             expr\_free(fs, e2);
1043             expr\_free(fs, e1);
1044             e1->u.s.info = bcemit\_ABC(fs, BC_CAT, 0, e1->u.s.info,
1045                                     e2->u.s.info);
1046         }
1047         e1->k = VRELOCABLE;
1048     } else {
1049         kp\_assert(op == OPR_NE || op == OPR_EQ || op == OPR_LT ||
1050                op == OPR_GE || op == OPR_LE || op == OPR_GT);
1051         bcemit\_comp(fs, op, e1, e2);
1052     }
1053 }
1054
1055 /* Emit unary operator. */
1056 static void bcemit\_unop(FuncState *fs, BCOp op, ExpDesc *e)
1057 {
1058     if (op == BC_NOT) {
1059         /* Swap true and false lists. */
1060         { BCPos temp = e->f; e->f = e->t; e->t = temp; }
1061         jmp\_dropval(fs, e->f);
1062         jmp\_dropval(fs, e->t);
1063         expr\_discharge(fs, e);
1064         if (e->k == VKNIL || e->k == VKFALSE) {
1065             e->k = VKTRUE;
1066             return;
1067         } else if (expr\_isk(e)) {
1068             e->k = VKFALSE;
1069             return;
1070         } else if (e->k == VJMP) {
1071             invertcond(fs, e);
1072             return;
1073         } else if (e->k == VRELOCABLE) {
1074             bcreg\_reserve(fs, 1);
1075             setbc\_a(bcptr(fs, e), fs->freereg-1);
1076             e->u.s.info = fs->freereg-1;
1077             e->k = VNONRELOC;
1078         } else {
1079             kp\_assert(e->k == VNONRELOC);
1080         }
1081     } else {
1082         kp\_assert(op == BC_UNM || op == BC_LEN);
1083         /* Constant-fold negations. */
1084         if (op == BC_UNM && !expr\_hasjump(e)) {
1085             /* Avoid folding to -0. */
1086             if (expr\_isnumk(e) && !expr\_numiszero(e)) {

```

```

1087         ktap_val_t *o = expr_numtv(e);
1088         if (is_number(o))
1089             set_number(o, -nvalue(o));
1090         return;
1091     }
1092 }
1093 expr_toanyreg(fs, e);
1094 }
1095 expr_free(fs, e);
1096 e->u.s.info = bcemit_AD(fs, op, 0, e->u.s.info);
1097 e->k = VRELOCABLE;
1098 }
1099
1100 /* -- Lexer support ----- */
1101
1102 /* Check and consume optional token. */
1103 static int lex_opt(LexState *ls, LexToken tok)
1104 {
1105     if (ls->tok == tok) {
1106         kp_lex_next(ls);
1107         return 1;
1108     }
1109     return 0;
1110 }
1111
1112 /* Check and consume token. */
1113 static void lex_check(LexState *ls, LexToken tok)
1114 {
1115     if (ls->tok != tok)
1116         err_token(ls, tok);
1117     kp_lex_next(ls);
1118 }
1119
1120 /* Check for matching token. */
1121 static void lex_match(LexState *ls, LexToken what, LexToken who, BCLine line)
1122 {
1123     if (!lex_opt(ls, what)) {
1124         if (line == ls->linenumber) {
1125             err_token(ls, what);
1126         } else {
1127             const char *swhat = kp_lex_token2str(ls, what);
1128             const char *swho = kp_lex_token2str(ls, who);
1129             kp_lex_error(ls, ls->tok, KP_ERR_XMATCH, swhat, swho,
1130                          line);
1131         }
1132     }
1133 }
1134
1135 /* Check for string token. */
1136 static ktap_str_t *lex_str(LexState *ls)
1137 {
1138     ktap_str_t *s;
1139
1140     if (ls->tok != TK_name)
1141         err_token(ls, TK_name);
1142     s = rawtsvalue(&ls->tokval);
1143     kp_lex_next(ls);
1144     return s;
1145 }
1146
1147 /* -- Variable handling ----- */
1148
1149 #define var_get(ls, fs, i) ((ls)->vstack[(fs)->varmap[(i)]]
1150
1151 /* Define a new local variable. */
1152 static void var_new(LexState *ls, BCReg n, ktap_str_t *name)
1153 {
1154     FuncState *fs = ls->fs;
1155     int vtop = ls->vtop;
1156
1157     checklimit(fs, fs->nactvar+n, KP_MAX_LOCVAR, "local variables");
1158     if (vtop >= ls->sizevstack) {
1159         if (ls->sizevstack >= KP_MAX_VSTACK)
1160             kp_lex_error(ls, 0, KP_ERR_XLIMC, KP_MAX_VSTACK);
1161         if (!ls->vstack) {
1162             ls->vstack = malloc(sizeof(VarInfo) * 20);

```

```

1163     ls->sizevstack = 20;
1164 } else {
1165     ls->vstack = realloc(ls->vstack,
1166         ls->sizevstack * sizeof(VarInfo) * 2);
1167     ls->sizevstack = ls->sizevstack * 2;
1168 }
1169 }
1170 kp\_assert((uintptr_t)name < VARNAME__MAX ||
1171     kp\_tab\_getstr(fs->kt, name) != NULL);
1172 ls->vstack[vtop].name = name;
1173 fs->varmap[fs->nactvar+n] = (uint16_t)vtop;
1174 ls->vtop = vtop+1;
1175 }
1176
1177 #define var_new_lit(ls, n, v) \
1178     var\_new(ls, (n), kp\_parse\_keepstr(ls, "" v, sizeof(v)-1))
1179
1180 #define var_new_fixed(ls, n, vn) \
1181     var\_new(ls, (n), (ktap\_str\_t *)(uintptr_t)(vn))
1182
1183 /* Add local variables. */
1184 static void var\_add(LexState *ls, BCReg nvars)
1185 {
1186     FuncState *fs = ls->fs;
1187     BCReg nactvar = fs->nactvar;
1188
1189     while (nvars-->0) {
1190         VarInfo *v = &var\_get(ls, fs, nactvar);
1191         v->startpc = fs->pc;
1192         v->slot = nactvar++;
1193         v->info = 0;
1194     }
1195     fs->nactvar = nactvar;
1196 }
1197
1198 /* Remove local variables. */
1199 static void var\_remove(LexState *ls, BCReg tolevel)
1200 {
1201     FuncState *fs = ls->fs;
1202     while (fs->nactvar > tolevel)
1203         var\_get(ls, fs, --fs->nactvar).endpc = fs->pc;
1204 }
1205
1206 /* Lookup local variable name. */
1207 static BCReg var\_lookup\_local(FuncState *fs, ktap\_str\_t *n)
1208 {
1209     int i;
1210
1211     for (i = fs->nactvar-1; i >= 0; i--) {
1212         if (n == var\_get(fs->ls, fs, i).name)
1213             return (BCReg)i;
1214     }
1215     return (BCReg)-1; /* Not found. */
1216 }
1217
1218 /* Lookup or add upvalue index. */
1219 static int var\_lookup\_uv(FuncState *fs, int vidx, ExpDesc *e)
1220 {
1221     int i, n = fs->nuv;
1222
1223     for (i = 0; i < n; i++)
1224         if (fs->uvmmap[i] == vidx)
1225             return i; /* Already exists. */
1226
1227     /* Otherwise create a new one. */
1228     checklimit(fs, fs->nuv, KP\_MAX\_UPVAL, "upvalues");
1229     kp\_assert(e->k == VLOCAL || e->k == VUPVAL);
1230     fs->uvmmap[n] = (uint16_t)vidx;
1231     fs->uvtmp[n] = (uint16_t)(e->k == VLOCAL ? vidx :
1232         KP\_MAX\_VSTACK+e->u.s.info);
1233     fs->nuv = n+1;
1234     return n;
1235 }
1236
1237 /* Forward declaration. */
1238 static void fscope\_uvmark(FuncState *fs, BCReg level);

```

```

1239
1240 /* Recursively lookup variables in enclosing functions. */
1241 static int var_lookup(FuncState *fs, ktag_str_t *name, ExpDesc *e,
1242 int first)
1243 {
1244     if (fs) {
1245         BCReg reg = var_lookup_local(fs, name);
1246         if ((int32_t)reg >= 0) { /* Local in this function? */
1247             expr_init(e, VLOCAL, reg);
1248             if (!first) {
1249                 /* Scope now has an upvalue. */
1250                 fscope_uvmark(fs, reg);
1251             }
1252             return (int)(e->u.s.aux = (uint32_t)fs->varmap[reg]);
1253         } else {
1254             /* Var in outer func? */
1255             int vidx = var_lookup(fs->prev, name, e, 0);
1256             if ((int32_t)vidx >= 0) {
1257                 /* Yes, make it an upvalue here. */
1258                 e->u.s.info =
1259                     (uint8_t)var_lookup_uv(fs, vidx, e);
1260                 e->k = VUPVAL;
1261                 return vidx;
1262             }
1263         }
1264     } else { /* Not found in any function, must be a global. */
1265         expr_init(e, VGLOBAL, 0);
1266         e->u.sval = name;
1267     }
1268     return (int)-1; /* Global. */
1269 }
1270
1271 /* Lookup variable name. */
1272 #define var_lookup(ls, e) \
1273     var_lookup((ls)->fs, lex_str(ls), (e), 1)
1274
1275 /* -- Goto an label handling ----- */
1276
1277 /* Add a new goto or label. */
1278 static int gola_new(LexState *ls, ktag_str_t *name, uint8_t info, BCPos pc)
1279 {
1280     FuncState *fs = ls->fs;
1281     int vtop = ls->vtop;
1282
1283     if (vtop >= ls->sizevstack) {
1284         if (ls->sizevstack >= KP_MAX_VSTACK)
1285             kp_lex_error(ls, 0, KP_ERR_XLIMC, KP_MAX_VSTACK);
1286         if (!ls->vstack) {
1287             ls->vstack = malloc(sizeof(VarInfo) * 20);
1288             ls->sizevstack = 20;
1289         } else {
1290             ls->vstack = realloc(ls->vstack,
1291                 ls->sizevstack * sizeof(VarInfo) * 2);
1292             ls->sizevstack = ls->sizevstack * 2;
1293         }
1294     }
1295     kp_assert(name == NAME_BREAK ||
1296         kp_tab_getstr(fs->kt, name) != NULL);
1297     ls->vstack[vtop].name = name;
1298     ls->vstack[vtop].startpc = pc;
1299     ls->vstack[vtop].slot = (uint8_t)fs->nactvar;
1300     ls->vstack[vtop].info = info;
1301     ls->vtop = vtop+1;
1302     return vtop;
1303 }
1304
1305 #define gola_isgoto(v) ((v)->info & VSTACK_GOTO)
1306 #define gola_islabel(v) ((v)->info & VSTACK_LABEL)
1307 #define gola_isgotolabel(v) ((v)->info & (VSTACK_GOTO|VSTACK_LABEL))
1308
1309 /* Patch goto to jump to label. */
1310 static void gola_patch(LexState *ls, VarInfo *vg, VarInfo *vl)
1311 {
1312     FuncState *fs = ls->fs;
1313     BCPos pc = vg->startpc;
1314

```

```

1315     vg->name = NULL; /* Invalidate pending goto. */
1316     setbc_a(&fs->bcbase[pc].ins, vl->slot);
1317     jmp_patch(fs, pc, vl->startpc);
1318 }
1319
1320 /* Patch goto to close upvalues. */
1321 static void gola_close(LexState *ls, VarInfo *vg)
1322 {
1323     FuncState *fs = ls->fs;
1324     BCPos pc = vg->startpc;
1325     BCIns *ip = &fs->bcbase[pc].ins;
1326     kp_assert(gola_isgoto(vg));
1327     kp_assert(bc_op(*ip) == BC_JMP || bc_op(*ip) == BC_UCLO);
1328     setbc_a(ip, vg->slot);
1329     if (bc_op(*ip) == BC_JMP) {
1330         BCPos next = jmp_next(fs, pc);
1331         if (next != NO_JMP)
1332             jmp_patch(fs, next, pc); /* Jump to UCLO. */
1333         setbc_op(ip, BC_UCLO); /* Turn into UCLO. */
1334         setbc_j(ip, NO_JMP);
1335     }
1336 }
1337
1338 /* Resolve pending forward gotos for label. */
1339 static void gola_resolve(LexState *ls, FuncScope *bl, int idx)
1340 {
1341     VarInfo *vg = ls->vstack + bl->vstart;
1342     VarInfo *vl = ls->vstack + idx;
1343     for (; vg < vl; vg++)
1344         if (vg->name == vl->name && gola_isgoto(vg)) {
1345             if (vg->slot < vl->slot) {
1346                 ktap_str_t *name =
1347                     var_get(ls, ls->fs, vg->slot).name;
1348                 kp_assert((uintptr_t)name >= VARNAME__MAX);
1349                 ls->linenumber =
1350                     ls->fs->bcbase[vg->startpc].line;
1351                 kp_assert(vg->name != NAME_BREAK);
1352                 kp_lex_error(ls, 0, KP_ERR_XGSCOPE,
1353                     getstr(vg->name), getstr(name));
1354             }
1355             gola_patch(ls, vg, vl);
1356         }
1357 }
1358
1359 /* Fixup remaining gotos and labels for scope. */
1360 static void gola_fixup(LexState *ls, FuncScope *bl)
1361 {
1362     VarInfo *v = ls->vstack + bl->vstart;
1363     VarInfo *ve = ls->vstack + ls->vtop;
1364
1365     for (; v < ve; v++) {
1366         ktap_str_t *name = v->name;
1367         /* Only consider remaining valid gotos/labels. */
1368         if (name != NULL) {
1369             if (gola_islabel(v)) {
1370                 VarInfo *vg;
1371                 /* Invalidate label that goes out of scope. */
1372                 v->name = NULL;
1373                 /* Resolve pending backward gotos. */
1374                 for (vg = v+1; vg < ve; vg++)
1375                     if (vg->name == name &&
1376                         gola_isgoto(vg)) {
1377                         if ((bl->flags & FSCOPE_UPVAL) &&
1378                             vg->slot > v->slot)
1379                             gola_close(ls, vg);
1380                         gola_patch(ls, vg, v);
1381                     }
1382             } else if (gola_isgoto(v)) {
1383                 /* Propagate goto or break to outer scope. */
1384                 if (bl->prev) {
1385                     bl->prev->flags |= name == NAME_BREAK ?
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000

```

```

1390         ls->linenumber =
1391         ls->fs->bcbase[v->startpc].line;
1392         if (name == NAME_BREAK)
1393             kp_lex_error(ls, 0, KP_ERR_XBREAK);
1394         else
1395             kp_lex_error(ls, 0, KP_ERR_XLUNDEF, getstr(name));
1396     }
1397 }
1398 }
1399 }
1400 }
1401
1402 /* Find existing label. */
1403 static VarInfo *gola_findlabel(LexState *ls, ktap_str_t *name)
1404 {
1405     VarInfo *v = ls->vstack + ls->fs->bl->vstart;
1406     VarInfo *ve = ls->vstack + ls->vtop;
1407
1408     for (; v < ve; v++)
1409         if (v->name == name && gola_islabel(v))
1410             return v;
1411     return NULL;
1412 }
1413
1414 /* -- Scope handling ----- */
1415
1416 /* Begin a scope. */
1417 static void fscope_begin(FuncState *fs, FuncScope *bl, int flags)
1418 {
1419     bl->nactvar = (uint8_t)fs->nactvar;
1420     bl->flags = flags;
1421     bl->vstart = fs->ls->vtop;
1422     bl->prev = fs->bl;
1423     fs->bl = bl;
1424     kp_assert(fs->freereg == fs->nactvar);
1425 }
1426
1427 /* End a scope. */
1428 static void fscope_end(FuncState *fs)
1429 {
1430     FuncScope *bl = fs->bl;
1431     LexState *ls = fs->ls;
1432
1433     fs->bl = bl->prev;
1434     var_remove(ls, bl->nactvar);
1435     fs->freereg = fs->nactvar;
1436     kp_assert(bl->nactvar == fs->nactvar);
1437     if ((bl->flags & (FSCOPE_UPVAL|FSCOPE_NOCLOSE)) == FSCOPE_UPVAL)
1438         bcemit_AJ(fs, BC_UCLO, bl->nactvar, 0);
1439     if ((bl->flags & FSCOPE_BREAK) {
1440         if ((bl->flags & FSCOPE_LOOP) {
1441             int idx = gola_new(ls, NAME_BREAK, VSTACK_LABEL,
1442                 fs->pc);
1443             ls->vtop = idx; /* Drop break label immediately. */
1444             gola_resolve(ls, bl, idx);
1445             return;
1446         } /* else: need the fixup step to propagate the breaks. */
1447     } else if (!(bl->flags & FSCOPE_GOLA)) {
1448         return;
1449     }
1450     gola_fixup(ls, bl);
1451 }
1452
1453 /* Mark scope as having an upvalue. */
1454 static void fscope_uvmark(FuncState *fs, BCReg level)
1455 {
1456     FuncScope *bl;
1457
1458     for (bl = fs->bl; bl && bl->nactvar > level; bl = bl->prev);
1459     if (bl)
1460         bl->flags |= FSCOPE_UPVAL;
1461 }
1462
1463 /* -- Function state management ----- */
1464
1465 /* Fixup bytecode for prototype. */

```

```

1466 static void fs_fixup_bc(FuncState *fs, ktab_proto_t *pt, BCIns *bc, int n)
1467 {
1468     BCInsLine *base = fs->bcbase;
1469     int i;
1470
1471     pt->sizebc = n;
1472     bc[0] = BCINS_AD((fs->flags & PROTO_VARARG) ? BC_FUNCV : BC_FUNCF,
1473         fs->framesize, 0);
1474     for (i = 1; i < n; i++)
1475         bc[i] = base[i].ins;
1476 }
1477
1478 /* Fixup upvalues for child prototype, step #2. */
1479 static void fs_fixup_uv2(FuncState *fs, ktab_proto_t *pt)
1480 {
1481     VarInfo *vstack = fs->ls->vstack;
1482     uint16_t *uv = pt->uv;
1483     int i, n = pt->sizeuv;
1484
1485     for (i = 0; i < n; i++) {
1486         VarIndex vidx = uv[i];
1487         if (vidx >= KP_MAX_VSTACK)
1488             uv[i] = vidx - KP_MAX_VSTACK;
1489         else if ((vstack[vidx].info & VSTACK_VAR_RW))
1490             uv[i] = vstack[vidx].slot | PROTO_UV_LOCAL;
1491         else
1492             uv[i] = vstack[vidx].slot | PROTO_UV_LOCAL |
1493                 PROTO_UV_IMMUTABLE;
1494     }
1495 }
1496
1497 /* Fixup constants for prototype. */
1498 static void fs_fixup_k(FuncState *fs, ktab_proto_t *pt, void *kptr)
1499 {
1500     ktab_tab_t *kt;
1501     ktab_node_t *node;
1502     int i, hmask;
1503
1504     checklimitgt(fs, fs->nkn, BCMAX_D+1, "constants");
1505     checklimitgt(fs, fs->nkgc, BCMAX_D+1, "constants");
1506
1507     pt->k = kptr;
1508     pt->sizekn = fs->nkn;
1509     pt->sizekgc = fs->nkgc;
1510     kt = fs->kt;
1511     node = kt->node;
1512     hmask = kt->hmask;
1513     for (i = 0; i <= hmask; i++) {
1514         ktab_node_t *n = &node[i];
1515
1516         if (tvhaskslot(&n->val)) {
1517             ptrdiff_t kidx = (ptrdiff_t)tvkslot(&n->val);
1518             kp_assert(!is_number(&n->key));
1519             if (is_number(&n->key)) {
1520                 ktab_val_t *tv = &((ktab_val_t *)kptr)[kidx];
1521                 *tv = n->key;
1522             } else {
1523                 ktab_obj_t *o = n->key.val.gc;
1524                 ktab_obj_t **v = (ktab_obj_t **)kptr;
1525                 v[-kidx] = o;
1526                 if (is_proto(&n->key))
1527                     fs_fixup_uv2(fs, (ktab_proto_t *)o);
1528             }
1529         }
1530     }
1531 }
1532
1533 /* Fixup upvalues for prototype, step #1. */
1534 static void fs_fixup_uv1(FuncState *fs, ktab_proto_t *pt, uint16_t *uv)
1535 {
1536     pt->uv = uv;
1537     pt->sizeuv = fs->nuv;
1538     memcpy(uv, fs->uvtmp, fs->nuv*sizeof(VarIndex));
1539 }
1540
1541 #ifndef KTAP_DISABLE_LINEINFO

```

```

1542 /* Prepare lineinfo for prototype. */
1543 static size_t fs_prep_line(FuncState *fs, BCLine numline)
1544 {
1545     return (fs->pc-1) << (numline < 256 ? 0 : numline < 65536 ? 1 : 2);
1546 }
1547
1548 /* Fixup lineinfo for prototype. */
1549 static void fs_fixup_line(FuncState *fs, ktap_proto_t *pt,
1550     void *lineinfo, BCLine numline)
1551 {
1552     BCInsLine *base = fs->bcbase + 1;
1553     BCLine first = fs->linedefined;
1554     int i = 0, n = fs->pc-1;
1555
1556     pt->firstline = fs->linedefined;
1557     pt->numline = numline;
1558     pt->lineinfo = lineinfo;
1559     if (numline < 256) {
1560         uint8_t *li = (uint8_t *)lineinfo;
1561         do {
1562             BCLine delta = base[i].line - first;
1563             kp_assert(delta >= 0 && delta < 256);
1564             li[i] = (uint8_t)delta;
1565         } while (++i < n);
1566     } else if (numline < 65536) {
1567         uint16_t *li = (uint16_t *)lineinfo;
1568         do {
1569             BCLine delta = base[i].line - first;
1570             kp_assert(delta >= 0 && delta < 65536);
1571             li[i] = (uint16_t)delta;
1572         } while (++i < n);
1573     } else {
1574         uint32_t *li = (uint32_t *)lineinfo;
1575         do {
1576             BCLine delta = base[i].line - first;
1577             kp_assert(delta >= 0);
1578             li[i] = (uint32_t)delta;
1579         } while (++i < n);
1580     }
1581 }
1582
1583 /* Prepare variable info for prototype. */
1584 static size_t fs_prep_var(LexState *ls, FuncState *fs, size_t *ofsvar)
1585 {
1586     VarInfo *vs = ls->vstack, *ve;
1587     int i, n;
1588     BCPos lastpc;
1589
1590     kp_buf_reset(&ls->sb); /* Copy to temp. string buffer. */
1591     /* Store upvalue names. */
1592     for (i = 0, n = fs->nuv; i < n; i++) {
1593         ktap_str_t *s = vs[fs->uvmap[i]].name;
1594         int len = s->len+1;
1595         char *p = kp_buf_more(&ls->sb, len);
1596         p = kp_buf_wmem(p, getstr(s), len);
1597         setsbufP(&ls->sb, p);
1598     }
1599
1600     *ofsvar = sbuflen(&ls->sb);
1601     lastpc = 0;
1602     /* Store local variable names and compressed ranges. */
1603     for (ve = vs + ls->vtop, vs += fs->vbase; vs < ve; vs++) {
1604         if (!gola_isgotolabel(vs)) {
1605             ktap_str_t *s = vs->name;
1606             BCPos startpc;
1607             char *p;
1608             if ((uintptr_t)s < VARNAME__MAX) {
1609                 p = kp_buf_more(&ls->sb, 1 + 2*5);
1610                 *p++ = (char)(uintptr_t)s;
1611             } else {
1612                 int len = s->len+1;
1613                 p = kp_buf_more(&ls->sb, len + 2*5);
1614                 p = kp_buf_wmem(p, getstr(s), len);
1615             }
1616             startpc = vs->startpc;
1617             p = strfmt_wuleb128(p, startpc-lastpc);

```

```

1618         p = strfmt_wuleb128(p, vs->endpc-startpc);
1619         setsbufP(&ls->sb, p);
1620         lastpc = startpc;
1621     }
1622 }
1623
1624 kp_buf_putb(&ls->sb, '\0'); /* Terminator for varinfo. */
1625 return sbufLen(&ls->sb);
1626 }
1627
1628 /* Fixup variable info for prototype. */
1629 static void fs_fixup_var(LexState *ls, ktap_proto_t *pt, uint8_t *p,
1630                          size_t ofsvar)
1631 {
1632     pt->uvinfo = p;
1633     pt->varinfo = (char *)p + ofsvar;
1634     /* Copy from temp. buffer. */
1635     memcpy(p, sbufB(&ls->sb), sbufLen(&ls->sb));
1636 }
1637 #else
1638
1639 /* Initialize with empty debug info, if disabled. */
1640 #define fs_prep_line(fs, numline) (UNUSED(numline), 0)
1641 #define fs_fixup_line(fs, pt, li, numline) \
1642     pt->firstline = pt->numline = 0, (pt)->lineinfo = NULL
1643 #define fs_prep_var(ls, fs, ofsvar) (UNUSED(ofsvar), 0)
1644 #define fs_fixup_var(ls, pt, p, ofsvar) \
1645     (pt)->uvinfo = NULL, (pt)->varinfo = NULL
1646
1647 #endif
1648
1649 /* Check if bytecode op returns. */
1650 static int bcopisret(BCOp op)
1651 {
1652     switch (op) {
1653     case BC_CALLMT: case BC_CALLT:
1654     case BC_RETM: case BC_RET: case BC_RET0: case BC_RET1:
1655         return 1;
1656     default:
1657         return 0;
1658     }
1659 }
1660
1661 /* Fixup return instruction for prototype. */
1662 static void fs_fixup_ret(FuncState *fs)
1663 {
1664     BCPos lastpc = fs->pc;
1665
1666     if (lastpc <= fs->lasttarget ||
1667         !bcopisret(bc_op(fs->bcbase[lastpc-1].ins))) {
1668         if ((fs->bl->flags & FSCOPE_UPVAL))
1669             bcemit_AJ(fs, BC_UCLO, 0, 0);
1670         bcemit_AD(fs, BC_RET0, 0, 1); /* Need final return. */
1671     }
1672     fs->bl->flags |= FSCOPE_NOCLOSE; /* Handled above. */
1673     fscope_end(fs);
1674     kp_assert(fs->bl == NULL);
1675     /* May need to fixup returns encoded before first function
1676      * was created. */
1677     if (fs->flags & PROTO_FIXUP_RETURN) {
1678         BCPos pc;
1679         for (pc = 1; pc < lastpc; pc++) {
1680             BCIns ins = fs->bcbase[pc].ins;
1681             BCPos offset;
1682             switch (bc_op(ins)) {
1683             case BC_CALLMT: case BC_CALLT:
1684             case BC_RETM: case BC_RET: case BC_RET0: case BC_RET1:
1685                 /* Copy original instruction. */
1686                 offset = bcemit_INS(fs, ins);
1687                 fs->bcbase[offset].line = fs->bcbase[pc].line;
1688                 offset = offset-(pc+1)+BCBIAS_J;
1689                 if (offset > BCMAX_D)
1690                     err_syntax(fs->ls, KP_ERR_XFIXUP);
1691                 /* Replace with UCLO plus branch. */
1692                 fs->bcbase[pc].ins = BCINS_AD(BC_UCLO, 0,
1693                    offset);

```

```

1694         break;
1695     case BC_UCLLO:
1696         return; /* We're done. */
1697     default:
1698         break;
1699     }
1700 }
1701 }
1702 }
1703
1704 /* Finish a FuncState and return the new prototype. */
1705 static ktap_proto_t *fs_finish(LexState *ls, BCLine line)
1706 {
1707     FuncState *fs = ls->fs;
1708     BCLine numline = line - fs->linedefined;
1709     size_t sizept, ofsk, ofsuv, ofsli, ofsdbg, ofsvar;
1710     ktap_proto_t *pt;
1711
1712     /* Apply final fixups. */
1713     fs_fixup_ret(fs);
1714
1715     /* Calculate total size of prototype including all colocated arrays. */
1716     sizept = sizeof(ktap_proto_t) + fs->pc*sizeof(BCIns) +
1717             fs->nkgc*sizeof(ktap_obj_t *);
1718     sizept = (sizept + sizeof(ktap_val_t)-1) & ~(sizeof(ktap_val_t)-1);
1719     ofsk = sizept; sizept += fs->nkn*sizeof(ktap_val_t);
1720     ofsuv = sizept; sizept += ((fs->nuv+1)&-1)*2;
1721     ofsli = sizept; sizept += fs_prep_line(fs, numline);
1722     ofsdbg = sizept; sizept += fs_prep_var(ls, fs, &ofsvar);
1723
1724     /* Allocate prototype and initialize its fields. */
1725     pt = (ktap_proto_t *)malloc((int)sizept);
1726     pt->gct = ~KTAP_TPROTO;
1727     pt->sizept = (int)sizept;
1728     pt->flags =
1729         (uint8_t)(fs->flags & ~(PROTO_HAS_RETURN|PROTO_FIXUP_RETURN));
1730     pt->numparams = fs->numparams;
1731     pt->framesize = fs->framesize;
1732     pt->chunkname = ls->chunkname;
1733
1734     /* Close potentially uninitialized gap between bc and kgc. */
1735     *(uint32_t *)((char *)pt + ofsk - sizeof(ktap_obj_t *)*(fs->nkgc+1)) = 0;
1736     fs_fixup_bc(fs, pt, (BCIns *)((char *)pt + sizeof(ktap_proto_t)), fs->pc);
1737     fs_fixup_k(fs, pt, (void *)((char *)pt + ofsk));
1738     fs_fixup_uv1(fs, pt, (uint16_t *)((char *)pt + ofsuv));
1739     fs_fixup_line(fs, pt, (void *)((char *)pt + ofsli), numline);
1740     fs_fixup_var(ls, pt, (uint8_t *)((char *)pt + ofsdbg), ofsvar);
1741
1742     ls->vtop = fs->vbase; /* Reset variable stack. */
1743     ls->fs = fs->prev;
1744     kp_assert(ls->fs != NULL || ls->tok == TK_eof);
1745     return pt;
1746 }
1747
1748 /* Initialize a new FuncState. */
1749 static void fs_init(LexState *ls, FuncState *fs)
1750 {
1751     fs->prev = ls->fs; ls->fs = fs; /* Append to list. */
1752     fs->ls = ls;
1753     fs->vbase = ls->vtop;
1754     fs->pc = 0;
1755     fs->lasttarget = 0;
1756     fs->jpc = NO_JUMP;
1757     fs->freereg = 0;
1758     fs->nkgc = 0;
1759     fs->nkn = 0;
1760     fs->nactvar = 0;
1761     fs->nuv = 0;
1762     fs->bl = NULL;
1763     fs->flags = 0;
1764     fs->framesize = 1; /* Minimum frame size. */
1765     fs->kt = kp_tab_new();
1766 }
1767
1768 /* -- Expressions ----- */
1769

```

```

1770 /* Forward declaration. */
1771 static void expr(LexState *ls, ExpDesc *v);
1772
1773 /* Return string expression. */
1774 static void expr_str(LexState *ls, ExpDesc *e)
1775 {
1776     expr_init(e, VKSTR, 0);
1777     e->u.sval = lex_str(ls);
1778 }
1779
1780 #define checku8(x) ((x) == (int32_t)(uint8_t)(x))
1781
1782 /* Return index expression. */
1783 static void expr_index(FuncState *fs, ExpDesc *t, ExpDesc *e)
1784 {
1785     /* Already called: expr_toval(fs, e). */
1786     t->k = VININDEXED;
1787     if (expr_isnumk(e)) {
1788         ktap_number n = expr_numberV(e);
1789         int32_t k = (int)n;
1790         if (checku8(k) && n == (ktap_number)k) {
1791             /* 256..511: const byte key */
1792             t->u.s.aux = BCMAX_C+1+(uint32_t)k;
1793             return;
1794         }
1795     } else if (expr_isstrk(e)) {
1796         BCReg idx = const_str(fs, e);
1797         if (idx <= BCMAX_C) {
1798             /* -256..-1: const string key */
1799             t->u.s.aux = ~idx;
1800             return;
1801         }
1802     }
1803     t->u.s.aux = expr_toanyreg(fs, e); /* 0..255: register */
1804 }
1805
1806 /* Parse index expression with named field. */
1807 static void expr_field(LexState *ls, ExpDesc *v)
1808 {
1809     FuncState *fs = ls->fs;
1810     ExpDesc key;
1811
1812     expr_toanyreg(fs, v);
1813     kp_lex_next(ls); /* Skip dot or colon. */
1814     expr_str(ls, &key);
1815     expr_index(fs, v, &key);
1816 }
1817
1818 /* Parse index expression with brackets. */
1819 static void expr_bracket(LexState *ls, ExpDesc *v)
1820 {
1821     kp_lex_next(ls); /* Skip '['. */
1822     expr(ls, v);
1823     expr_toval(ls->fs, v);
1824     lex_check(ls, ']');
1825 }
1826
1827 /* Get value of constant expression. */
1828 static void expr_kvalue(ktap_val_t *v, ExpDesc *e)
1829 {
1830     if (e->k <= VKTRUE) {
1831         setitype(v, ~(uint32_t)e->k);
1832     } else if (e->k == VKSTR) {
1833         set_string(v, e->u.sval);
1834     } else {
1835         kp_assert(tvisnumber(expr_numtv(e)));
1836         *v = *expr_numtv(e);
1837     }
1838 }
1839
1840 #define FLS(x) ((uint32_t)(__builtin_clz(x)^31))
1841 #define hsize2hbits(s) ((s) ? ((s)==1 ? 1 : 1+FLS((uint32_t)((s)-1))) : 0)
1842
1843
1844 /* Parse table constructor expression. */
1845 static void expr_table(LexState *ls, ExpDesc *e)

```

```

1846 {
1847 FuncState *fs = ls->fs;
1848 BCLine line = ls->linenumber;
1849 ktap_tab_t *t = NULL;
1850 int vcall = 0, needarr = 0, fixt = 0;
1851 uint32_t narr = 1; /* First array index. */
1852 uint32_t nhash = 0; /* Number of hash entries. */
1853 BCReg freg = fs->freereg;
1854 BCPos pc = bcemit_AD(fs, BC_TNEW, freg, 0);
1855
1856 expr_init(e, VNONRELOC, freg);
1857 bcreg_reserve(fs, 1);
1858 freg++;
1859 lex_check(ls, '{');
1860 while (ls->tok != '}') {
1861     ExpDesc key, val;
1862     vcall = 0;
1863     if (ls->tok == '[') {
1864         expr_bracket(ls, &key); /* Already calls expr_toval. */
1865         if (!expr_isk(&key))
1866             expr_index(fs, e, &key);
1867         if (expr_isnumk(&key) && expr_numiszero(&key))
1868             needarr = 1;
1869         else
1870             nhash++;
1871         lex_check(ls, '=');
1872     } else if ((ls->tok == TK_name) &&
1873         kp_lex_lookahead(ls) == '=') {
1874         expr_str(ls, &key);
1875         lex_check(ls, '=');
1876         nhash++;
1877     } else {
1878         expr_init(&key, VKNUM, 0);
1879         set_number(&key.u.nval, (int)narr);
1880         narr++;
1881         needarr = vcall = 1;
1882     }
1883     expr(ls, &val);
1884     if (expr_isk(&key) && key.k != VKNIL &&
1885         (key.k == VKSTR || expr_isk_nojump(&val))) {
1886         ktap_val_t k, *v;
1887         if (!t) { /* Create template table on demand. */
1888             BCReg kidx;
1889             t = kp_tab_new();
1890             kidx = const_gc(fs, obj2gc(t), KTAP_TTAB);
1891             fs->bcbase[pc].ins = BCINS_AD(BC_TDUP, freg-1,
1892                 kidx);
1893         }
1894         vcall = 0;
1895         expr_kvalue(&k, &key);
1896         v = kp_tab_set(t, &k);
1897         /* Add const key/value to template table. */
1898         if (expr_isk_nojump(&val)) {
1899             expr_kvalue(v, &val);
1900         } else {
1901             /* Otherwise create dummy string key (avoids kp_tab_newkey). */
1902             set_table(v, t); /* Preserve key with table itself as value. */
1903             fixt = 1; /* Fix this later, after all resizes. */
1904             goto nonconst;
1905         }
1906     } else {
1907 nonconst:
1908         if (val.k != VCALL) {
1909             expr_toanyreg(fs, &val);
1910             vcall = 0;
1911         }
1912         if (expr_isk(&key))
1913             expr_index(fs, e, &key);
1914         bcemit_store(fs, e, &val);
1915     }
1916     fs->freereg = freg;
1917     if (!lex_opt(ls, ',', ')') && !lex_opt(ls, ';'))
1918         break;
1919 }
1920 lex_match(ls, '}', '{', line);
1921 if (vcall) {

```

```

1922     BCInsLine *ilp = &fs->bcbase[fs->pc-1];
1923     ExpDesc en;
1924     kp_assert(bc_a(ilp->ins) == freg &&
1925         bc_op(ilp->ins) == (narr > 256 ? BC_TSETV : BC_TSETB));
1926     expr_init(&en, VKNUM, 0);
1927     set_number(&en.u.nval, narr - 1);
1928     if (narr > 256) { fs->pc--; ilp--; }
1929     ilp->ins = BCINS_AD(BC_TSETM, freg, const_num(fs, &en));
1930     setbc_b(&ilp[-1].ins, 0);
1931 }
1932 if (pc == fs->pc-1) { /* Make expr relocable if possible. */
1933     e->u.s.info = pc;
1934     fs->freereg--;
1935     e->k = VRELOCABLE;
1936 } else {
1937     e->k = VNONRELOC; /* May have been changed by expr_index. */
1938 }
1939 if (!t) { /* Construct TNEW RD: hhhhhaaaaaaaaaa. */
1940     BCIns *ip = &fs->bcbase[pc].ins;
1941     if (!needarr) narr = 0;
1942     else if (narr < 3) narr = 3;
1943     else if (narr > 0x7ff) narr = 0x7ff;
1944     setbc_d(ip, narr|(hsize2hbits(nhash)<<11));
1945 } else {
1946     if (fixt) { /* Fix value for dummy keys in template table. */
1947         ktap_node_t *node = t->node;
1948         uint32_t i, hmask = t->hmask;
1949         for (i = 0; i <= hmask; i++) {
1950             ktap_node_t *n = &node[i];
1951             if (is_table(&n->val)) {
1952                 kp_assert(tabV(&n->val) == t);
1953                 /* Turn value into nil. */
1954                 set_nil(&n->val);
1955             }
1956         }
1957     }
1958 }
1959 }
1960
1961 /* Parse function parameters. */
1962 static BCReg parse_params(LexState *ls, int needself)
1963 {
1964     FuncState *fs = ls->fs;
1965     BCReg nparams = 0;
1966     lex_check(ls, '(');
1967     if (needself)
1968         var_new_lit(ls, nparams++, "self");
1969     if (ls->tok != ')') {
1970         do {
1971             if (ls->tok == TK_name) {
1972                 var_new(ls, nparams++, lex_str(ls));
1973             } else if (ls->tok == TK_dots) {
1974                 kp_lex_next(ls);
1975                 fs->flags |= PROTO_VARARG;
1976                 break;
1977             } else {
1978                 err_syntax(ls, KP_ERR_XPARAM);
1979             }
1980         } while (lex_opt(ls, ','));
1981     }
1982     var_add(ls, nparams);
1983     kp_assert(fs->nactvar == nparams);
1984     bcreg_reserve(fs, nparams);
1985     lex_check(ls, ')');
1986     return nparams;
1987 }
1988
1989 /* Forward declaration. */
1990 static void parse_chunk(LexState *ls);
1991
1992 /* Parse body of a function. */
1993 static void parse_body(LexState *ls, ExpDesc *e, int needself, BCLine line)
1994 {
1995     FuncState fs, *pfs = ls->fs;
1996     FuncScope bl;
1997     ktap_proto_t *pt;

```

```

1998 ptrdiff_t oldbase = pfs->bcbase - ls->bcstack;
1999
2000 fs_init(ls, &fs);
2001 fscope_begin(&fs, &bl, 0);
2002 fs.linedefined = line;
2003 fs.numparams = (uint8_t)parse_params(ls, needself);
2004 fs.bcbase = pfs->bcbase + pfs->pc;
2005 fs.bclim = pfs->bclim - pfs->pc;
2006 bcemit_AD(&fs, BC_FUNCF, 0, 0); /* Placeholder. */
2007 lex_check(ls, '{');
2008 parse_chunk(ls);
2009 lex_check(ls, '}');
2010 pt = fs_finish(ls, (ls->lastline = ls->linenumber));
2011 pfs->bcbase = ls->bcstack + oldbase; /* May have been reallocated. */
2012 pfs->bclim = (BCPos)(ls->sizebcstack - oldbase);
2013 /* Store new prototype in the constant array of the parent. */
2014 expr_init(e, VRELOCABLE,
2015          bcemit_AD(pfs, BC_FNEW, 0,
2016                  const_gc(pfs, (ktap_obj_t *)pt, KTAP_TPROTO)));
2017 if (!(pfs->flags & PROTO_CHILD)) {
2018     if (pfs->flags & PROTO_HAS_RETURN)
2019         pfs->flags |= PROTO_FIXUP_RETURN;
2020     pfs->flags |= PROTO_CHILD;
2021 }
2022 //kp_lex_next(ls);
2023 }
2024
2025 /* Parse body of a function, for 'trace/trace_end/profile/tick' closure */
2026 static void parse_body_no_args(LexState *ls, ExpDesc *e, int needself,
2027                               BCLine line)
2028 {
2029     FuncState fs, *pfs = ls->fs;
2030     FuncScope bl;
2031     ktap_proto_t *pt;
2032     ptrdiff_t oldbase = pfs->bcbase - ls->bcstack;
2033
2034     fs_init(ls, &fs);
2035     fscope_begin(&fs, &bl, 0);
2036     fs.linedefined = line;
2037     fs.numparams = 0;
2038     fs.bcbase = pfs->bcbase + pfs->pc;
2039     fs.bclim = pfs->bclim - pfs->pc;
2040     bcemit_AD(&fs, BC_FUNCF, 0, 0); /* Placeholder. */
2041     lex_check(ls, '{');
2042     parse_chunk(ls);
2043     lex_check(ls, '}');
2044     pt = fs_finish(ls, (ls->lastline = ls->linenumber));
2045     pfs->bcbase = ls->bcstack + oldbase; /* May have been reallocated. */
2046     pfs->bclim = (BCPos)(ls->sizebcstack - oldbase);
2047     /* Store new prototype in the constant array of the parent. */
2048     expr_init(e, VRELOCABLE,
2049             bcemit_AD(pfs, BC_FNEW, 0,
2050                     const_gc(pfs, (ktap_obj_t *)pt, KTAP_TPROTO)));
2051     if (!(pfs->flags & PROTO_CHILD)) {
2052         if (pfs->flags & PROTO_HAS_RETURN)
2053             pfs->flags |= PROTO_FIXUP_RETURN;
2054         pfs->flags |= PROTO_CHILD;
2055     }
2056     //kp_lex_next(ls);
2057 }
2058
2059
2060 /* Parse expression list. Last expression is left open. */
2061 static BCReg expr_list(LexState *ls, ExpDesc *v)
2062 {
2063     BCReg n = 1;
2064
2065     expr(ls, v);
2066     while (lex_opt(ls, ',')) {
2067         expr_tonextreg(ls->fs, v);
2068         expr(ls, v);
2069         n++;
2070     }
2071     return n;
2072 }
2073

```

```

2074 /* Parse function argument list. */
2075 static void parse_args(LexState *ls, ExpDesc *e)
2076 {
2077     FuncState *fs = ls->fs;
2078     ExpDesc args;
2079     BCIns ins;
2080     BCREg base;
2081     BCLine line = ls->linenumber;
2082
2083     if (ls->tok == '(') {
2084         if (line != ls->lastline)
2085             err_syntax(ls, KP_ERR_XAMBIG);
2086         kp_lex_next(ls);
2087         if (ls->tok == ')') { /* f(). */
2088             args.k = VVOID;
2089         } else {
2090             expr_list(ls, &args);
2091             /* f(a, b, g()) or f(a, b, ...). */
2092             if (args.k == VCALL) {
2093                 /* Pass on multiple results. */
2094                 setbc_b(bcptr(fs, &args), 0);
2095             }
2096             lex_match(ls, ')', '(', line);
2097         } else if (ls->tok == '{') {
2098             expr_table(ls, &args);
2099         } else if (ls->tok == TK_string) {
2100             expr_init(&args, VKSTR, 0);
2101             args.u.sval = rawtsvalue(&ls->tokval);
2102             kp_lex_next(ls);
2103         } else {
2104             err_syntax(ls, KP_ERR_XFUNARG);
2105             return; /* Silence compiler. */
2106         }
2107     }
2108
2109     kp_assert(e->k == VNONRELOC);
2110     base = e->u.s.info; /* Base register for call. */
2111     if (args.k == VCALL) {
2112         ins = BCINS_ABC(BC_CALLM, base, 2, args.u.s.aux - base - 1);
2113     } else {
2114         if (args.k != VVOID)
2115             expr_tonextreg(fs, &args);
2116         ins = BCINS_ABC(BC_CALL, base, 2, fs->freereg - base);
2117     }
2118     expr_init(e, VCALL, bcemit_INS(fs, ins));
2119     e->u.s.aux = base;
2120     fs->bcbase[fs->pc - 1].line = line;
2121     fs->freereg = base+1; /* Leave one result by default. */
2122 }
2123
2124 /* Parse primary expression. */
2125 static void expr_primary(LexState *ls, ExpDesc *v)
2126 {
2127     FuncState *fs = ls->fs;
2128
2129     /* Parse prefix expression. */
2130     if (ls->tok == '(') {
2131         BCLine line = ls->linenumber;
2132         kp_lex_next(ls);
2133         expr(ls, v);
2134         lex_match(ls, ')', '(', line);
2135         expr_discharge(ls->fs, v);
2136     } else if (ls->tok == TK_name) {
2137         var_lookup(ls, v);
2138     } else {
2139         err_syntax(ls, KP_ERR_XSYMBOL);
2140     }
2141
2142     for (;;) { /* Parse multiple expression suffixes. */
2143         if (ls->tok == '.') {
2144             expr_field(ls, v);
2145         } else if (ls->tok == '[') {
2146             ExpDesc key;
2147             expr_toanyreg(fs, v);
2148             expr_bracket(ls, &key);
2149             expr_index(fs, v, &key);

```

```

2150     } else if (ls->tok == ':') {
2151         ExpDesc key;
2152         kp\_lex\_next(ls);
2153         expr\_str(ls, &key);
2154         bcemit\_method(fs, v, &key);
2155         parse\_args(ls, v);
2156     } else if (ls->tok == '(' || ls->tok == TK_string ||
2157               ls->tok == '{') {
2158         expr\_tonextreg(fs, v);
2159         parse\_args(ls, v);
2160     } else {
2161         break;
2162     }
2163 }
2164 }
2165
2166 /* Parse simple expression. */
2167 static void expr\_simple(LexState *ls, ExpDesc *v)
2168 {
2169     switch (ls->tok) {
2170     case TK_number:
2171         expr\_init(v, VKNUM, 0);
2172         set\_obj(&v->u.nval, &ls->tokval);
2173         break;
2174     case TK_string:
2175         expr\_init(v, VKSTR, 0);
2176         v->u.sval = rawtsvalue(&ls->tokval);
2177         break;
2178     case TK_nil:
2179         expr\_init(v, VKNIL, 0);
2180         break;
2181     case TK_true:
2182         expr\_init(v, VKTRUE, 0);
2183         break;
2184     case TK_false:
2185         expr\_init(v, VKFALSE, 0);
2186         break;
2187     case TK_dots: { /* Vararg. */
2188         FuncState *fs = ls->fs;
2189         BCReg base;
2190         checkcond(ls, fs->flags & PROTO\_VARARG, KP_ERR_XDOTS);
2191         bcreg\_reserve(fs, 1);
2192         base = fs->freereg-1;
2193         expr\_init(v, VCALL, bcemit\_ABC(fs, BC_VARARG, base, 2,
2194         fs->numparams));
2195         v->u.s.aux = base;
2196         break;
2197     }
2198     case '{': /* Table constructor. */
2199         expr\_table(ls, v);
2200         return;
2201     case TK_function:
2202         kp\_lex\_next(ls);
2203         parse\_body(ls, v, 0, ls->linenumber);
2204         return;
2205     case TK_argstr:
2206         expr\_init(v, VARGSTR, 0);
2207         break;
2208     case TK_probename:
2209         expr\_init(v, VARGNAME, 0);
2210         break;
2211     case TK_arg0: case TK_arg1: case TK_arg2: case TK_arg3: case TK_arg4:
2212     case TK_arg5: case TK_arg6: case TK_arg7: case TK_arg8: case TK_arg9:
2213         expr\_init(v, VARGN, ls->tok - TK_arg0);
2214         break;
2215     case TK_pid:
2216         expr\_init(v, VPID, 0);
2217         break;
2218     case TK_tid:
2219         expr\_init(v, VTID, 0);
2220         break;
2221     case TK_uid:
2222         expr\_init(v, VUID, 0);
2223         break;
2224     case TK_cpu:
2225         expr\_init(v, VCPU, 0);

```

```

2226     break;
2227 case TK_execname:
2228     expr_init(v, VEXECNAME, 0);
2229     break;
2230 default:
2231     expr_primary(ls, v);
2232     return;
2233 }
2234 kp_lex_next(ls);
2235 }
2236
2237 /* Manage syntactic levels to avoid blowing up the stack. */
2238 static void synlevel_begin(LexState *ls)
2239 {
2240     if (++ls->level >= KP_MAX_XLEVEL)
2241         kp_lex_error(ls, 0, KP_ERR_XLEVELS);
2242 }
2243
2244 #define synlevel_end(ls)    ((ls)->level--)
2245
2246 /* Convert token to binary operator. */
2247 static BinOpr token2binop(LexToken tok)
2248 {
2249     switch (tok) {
2250     case '+':    return OPR_ADD;
2251     case '-':    return OPR_SUB;
2252     case '*':    return OPR_MUL;
2253     case '/':    return OPR_DIV;
2254     case '%':    return OPR_MOD;
2255     case '^':    return OPR_POW;
2256     case TK_concat: return OPR_CONCAT;
2257     case TK_ne:   return OPR_NE;
2258     case TK_eq:   return OPR_EQ;
2259     case '<':     return OPR_LT;
2260     case TK_le:   return OPR_LE;
2261     case '>':     return OPR_GT;
2262     case TK_ge:   return OPR_GE;
2263     case TK_and:  return OPR_AND;
2264     case TK_or:   return OPR_OR;
2265     default:     return OPR_NOBINOPR;
2266     }
2267 }
2268
2269 /* Priorities for each binary operator. ORDER OPR. */
2270 static const struct {
2271     uint8_t left;    /* Left priority. */
2272     uint8_t right;   /* Right priority. */
2273 } priority[] = {
2274     {6,6}, {6,6}, {7,7}, {7,7}, {7,7},    /* ADD SUB MUL DIV MOD */
2275     {10,9}, {5,4},    /* POW CONCAT (right associative) */
2276     {3,3}, {3,3},    /* EQ NE */
2277     {3,3}, {3,3}, {3,3}, {3,3},    /* LT GE GT LE */
2278     {2,2}, {1,1}    /* AND OR */
2279 };
2280
2281 #define UNARY_PRIORITY    8 /* Priority for unary operators. */
2282
2283 /* Forward declaration. */
2284 static BinOpr expr_binop(LexState *ls, ExpDesc *v, uint32_t limit);
2285
2286 /* Parse unary expression. */
2287 static void expr_unop(LexState *ls, ExpDesc *v)
2288 {
2289     BCOp op;
2290     if (ls->tok == TK_not) {
2291         op = BC_NOT;
2292     } else if (ls->tok == '-') {
2293         op = BC_UNM;
2294     } #if 0 /* ktap don't support lua length operator '#' */
2295     } else if (ls->tok == '#') {
2296         op = BC_LEN;
2297     } #endif
2298     } else {
2299         expr_simple(ls, v);
2300         return;
2301     }

```

```

2302     kp_lex_next(ls);
2303     expr_binop(ls, v, UNARY_PRIORITY);
2304     bcemit_unop(ls->fs, op, v);
2305 }
2306
2307 /* Parse binary expressions with priority higher than the limit. */
2308 static BinOpr expr_binop(LexState *ls, ExpDesc *v, uint32_t limit)
2309 {
2310     BinOpr op;
2311
2312     synlevel_begin(ls);
2313     expr_unop(ls, v);
2314     op = token2binop(ls->tok);
2315     while (op != OPR_NOBINOPR && priority[op].left > limit) {
2316         ExpDesc v2;
2317         BinOpr nextop;
2318         kp_lex_next(ls);
2319         bcemit_binop_left(ls->fs, op, v);
2320         /* Parse binary expression with higher priority. */
2321         nextop = expr_binop(ls, &v2, priority[op].right);
2322         bcemit_binop(ls->fs, op, v, &v2);
2323         op = nextop;
2324     }
2325     synlevel_end(ls);
2326     return op; /* Return unconsumed binary operator (if any). */
2327 }
2328
2329 /* Parse expression. */
2330 static void expr(LexState *ls, ExpDesc *v)
2331 {
2332     expr_binop(ls, v, 0); /* Priority 0: parse whole expression. */
2333 }
2334
2335 /* Assign expression to the next register. */
2336 static void expr_next(LexState *ls)
2337 {
2338     ExpDesc e;
2339     expr(ls, &e);
2340     expr_tonextreg(ls->fs, &e);
2341 }
2342
2343 /* Parse conditional expression. */
2344 static BCPos expr_cond(LexState *ls)
2345 {
2346     ExpDesc v;
2347
2348     lex_check(ls, '(');
2349     expr(ls, &v);
2350     if (v.k == VKNIL)
2351         v.k = VKFALSE;
2352     bcemit_branch_t(ls->fs, &v);
2353     lex_check(ls, ')');
2354     return v.f;
2355 }
2356
2357 /* -- Assignments ----- */
2358
2359 /* List of LHS variables. */
2360 typedef struct LHSVarList {
2361     ExpDesc v; /* LHS variable. */
2362     struct LHSVarList *prev; /* Link to previous LHS variable. */
2363 } LHSVarList;
2364
2365 /* Eliminate write-after-read hazards for local variable assignment. */
2366 static void assign_hazard(LexState *ls, LHSVarList *lh, const ExpDesc *v)
2367 {
2368     FuncState *fs = ls->fs;
2369     BCReg reg = v->u.s.info; /* Check against this variable. */
2370     BCReg tmp = fs->freereg; /* Rename to this temp. register(if needed) */
2371     int hazard = 0;
2372
2373     for (; lh; lh = lh->prev) {
2374         if (lh->v.k == VININDEXED) {
2375             if (lh->v.u.s.info == reg) { /* t[i], t = 1, 2 */
2376                 hazard = 1;
2377                 lh->v.u.s.info = tmp;

```

```

2378     }
2379     if (lh->v.u.s.aux == reg) { /* t[i], i = 1, 2 */
2380         hazard = 1;
2381         lh->v.u.s.aux = tmp;
2382     }
2383 }
2384 }
2385 if (hazard) {
2386     /* Rename conflicting variable. */
2387     bcemit_AD(fs, BC_MOV, tmp, reg);
2388     bcreg_reserve(fs, 1);
2389 }
2390 }
2391
2392 /* Adjust LHS/RHS of an assignment. */
2393 static void assign_adjust(LexState *ls, BCREg nvars, BCREg nexps, ExpDesc *e)
2394 {
2395     FuncState *fs = ls->fs;
2396     int32_t extra = (int32_t)nvars - (int32_t)nexps;
2397
2398     if (e->k == VCALL) {
2399         extra++; /* Compensate for the VCALL itself. */
2400         if (extra < 0)
2401             extra = 0;
2402         setbc_b(bcptr(fs, e), extra+1); /* Fixup call results. */
2403         if (extra > 1)
2404             bcreg_reserve(fs, (BCREg)extra-1);
2405     } else {
2406         if (e->k != VVOID)
2407             expr_tonextreg(fs, e); /* Close last expression. */
2408         if (extra > 0) { /* Leftover LHS are set to nil. */
2409             BCREg reg = fs->freereg;
2410             bcreg_reserve(fs, (BCREg)extra);
2411             bcemit_nil(fs, reg, (BCREg)extra);
2412         }
2413     }
2414 }
2415
2416 /* Recursively parse assignment statement. */
2417 static void parse_assignment(LexState *ls, LHSVarList *lh, BCREg nvars)
2418 {
2419     ExpDesc e;
2420
2421     checkcond(ls, VLOCAL <= lh->v.k && lh->v.k <= VINDEXED,
2422         KP_ERR_XSYNTAX);
2423     if (lex_opt(ls, ',', '')) { /* Collect LHS list and recurse upwards. */
2424         LHSVarList vl;
2425         vl.prev = lh;
2426         expr_primary(ls, &vl.v);
2427         if (vl.v.k == VLOCAL)
2428             assign_hazard(ls, lh, &vl.v);
2429         checklimit(ls->fs, ls->level + nvars, KP_MAX_XLEVEL,
2430             "variable names");
2431         parse_assignment(ls, &vl, nvars+1);
2432     } else { /* Parse RHS. */
2433         BCREg nexps;
2434         int assign_incr = 1;
2435
2436         if (lex_opt(ls, '='))
2437             assign_incr = 0;
2438         else if (lex_opt(ls, TK_incr))
2439             assign_incr = 1;
2440         else
2441             err_syntax(ls, KP_ERR_XSYMBOL);
2442
2443         nexps = expr_list(ls, &e);
2444         if (nexps == nvars) {
2445             if (e.k == VCALL) {
2446                 /* Vararg assignment. */
2447                 if (bc_op(*bcptr(ls->fs, &e)) == BC_VARG) {
2448                     ls->fs->freereg--;
2449                     e.k = VRELOCABLE;
2450                 } else { /* Multiple call results. */
2451                     /* Base of call is not relocatable. */
2452                     e.u.s.info = e.u.s.aux;
2453                     e.k = VNONRELOC;

```

```

2454     }
2455     }
2456     if (assign_incr == 0)
2457         bcemit_store(ls->fs, &lh->v, &e);
2458     else
2459         bcemit_store_incr(ls->fs, &lh->v, &e);
2460     return;
2461 }
2462 assign_adjust(ls, nvars, nexps, &e);
2463 if (nexps > nvars) {
2464     /* Drop leftover regs. */
2465     ls->fs->freereg -= nexps - nvars;
2466 }
2467 }
2468 /* Assign RHS to LHS and recurse downwards. */
2469 expr_init(&e, VNONRELOC, ls->fs->freereg-1);
2470 bcemit_store(ls->fs, &lh->v, &e);
2471 }
2472
2473 /* Parse call statement or assignment. */
2474 static void parse_call_assign(LexState *ls)
2475 {
2476     FuncState *fs = ls->fs;
2477     LHSVarList vl;
2478
2479     expr_primary(ls, &vl.v);
2480     if (vl.v.k == VCALL) { /* Function call statement. */
2481         setbc_b(bcptr(fs, &vl.v), 1); /* No results. */
2482     } else { /* Start of an assignment. */
2483         vl.prev = NULL;
2484         parse_assignment(ls, &vl, 1);
2485     }
2486 }
2487
2488 /* Parse 'var'(local in lua) statement. */
2489 static void parse_local(LexState *ls)
2490 {
2491     if (lex_opt(ls, TK_function)) { /* Local function declaration. */
2492         ExpDesc v, b;
2493         FuncState *fs = ls->fs;
2494         var_new(ls, 0, lex_str(ls));
2495         expr_init(&v, VLOCAL, fs->freereg);
2496         v.u.s.aux = fs->varmap[fs->freereg];
2497         bcreg_reserve(fs, 1);
2498         var_add(ls, 1);
2499         parse_body(ls, &b, 0, ls->linenumber);
2500         /* bcemit_store(fs, &v, &b) without setting VSTACK VAR_RW. */
2501         expr_free(fs, &b);
2502         expr_toreg(fs, &b, v.u.s.info);
2503         /* The upvalue is in scope, but the local is only valid
2504            * after the store. */
2505         var_get(ls, fs, fs->nactvar - 1).startpc = fs->pc;
2506     } else { /* Local variable declaration. */
2507         ExpDesc e;
2508         BCReg nexps, nvars = 0;
2509         do { /* Collect LHS. */
2510             var_new(ls, nvars++, lex_str(ls));
2511         } while (lex_opt(ls, ','));
2512         if (lex_opt(ls, '=')) { /* Optional RHS. */
2513             nexps = expr_list(ls, &e);
2514         } else { /* Or implicitly set to nil. */
2515             e.k = VVOID;
2516             nexps = 0;
2517         }
2518         assign_adjust(ls, nvars, nexps, &e);
2519         var_add(ls, nvars);
2520     }
2521 }
2522
2523 /* Parse 'function' statement. */
2524 static void parse_func(LexState *ls, BCLine line)
2525 {
2526     FuncState *fs = ls->fs;
2527     ExpDesc v, b;
2528
2529     kp_lex_next(ls); /* Skip 'function'. */

```

```

2530      /* function is declared as local */
2531  #if 1
2532      var_new(ls, 0, lex_str(ls));
2533      expr_init(&v, VLOCAL, fs->freereg);
2534      v.u.s.aux = fs->varmap[fs->freereg];
2535      bcreg_reserve(fs, 1);
2536      var_add(ls, 1);
2537      parse_body(ls, &b, 0, ls->linenumber);
2538      /* bcemit_store(fs, &v, &b) without setting VSTACK VAR RW. */
2539      expr_free(fs, &b);
2540      expr_toreg(fs, &b, v.u.s.info);
2541      /* The upvalue is in scope, but the local is only valid
2542      * after the store. */
2543      var_get(ls, fs, fs->nactvar - 1).startpc = fs->pc;
2544
2545  #else
2546      int needself = 0;
2547
2548      /* Parse function name. */
2549      var_lookup(ls, &v);
2550      while (ls->tok == '.') /* Multiple dot-separated fields. */
2551          expr_field(ls, &v);
2552      if (ls->tok == ':') { /* Optional colon to signify method call. */
2553          needself = 1;
2554          expr_field(ls, &v);
2555      }
2556      parse_body(ls, &b, needself, line);
2557      fs = ls->fs;
2558      bcemit_store(fs, &v, &b);
2559      fs->bcbase[fs->pc - 1].line = line; /* Set line for the store. */
2560  #endif
2561  }
2562
2563  /* -- Control transfer statements ----- */
2564
2565  /* Check for end of block. */
2566  static int parse_isend(LexToken tok)
2567  {
2568      switch (tok) {
2569          case TK_else: case TK_elseif: case TK_end: case TK_until: case TK_eof:
2570          case '}':
2571              return 1;
2572          default:
2573              return 0;
2574      }
2575  }
2576
2577  /* Parse 'return' statement. */
2578  static void parse_return(LexState *ls)
2579  {
2580      BCIns ins;
2581      FuncState *fs = ls->fs;
2582
2583      kp_lex_next(ls); /* Skip 'return'. */
2584      fs->flags |= PROTO_HAS_RETURN;
2585      if (parse_isend(ls->tok) || ls->tok == ';') { /* Bare return. */
2586          ins = BCINS_AD(BC_RET0, 0, 1);
2587      } else { /* Return with one or more values. */
2588          ExpDesc e; /* Receives the _last_ expression in the list. */
2589          BCReg nret = expr_list(ls, &e);
2590          if (nret == 1) { /* Return one result. */
2591              if (e.k == VCALL) { /* Check for tail call. */
2592                  BCIns *ip = bcptr(fs, &e);
2593                  /* It doesn't pay off to add BC_VARGT just
2594                  * for 'return ...'. */
2595                  if (bc_op(*ip) == BC_VARG)
2596                      goto notailcall;
2597                  fs->pc--;
2598                  ins = BCINS_AD(bc_op(*ip)-BC_CALL+BC_CALLT,
2599                              bc_a(*ip), bc_c(*ip));
2600              } else { /* Can return the result from any register. */
2601                  ins = BCINS_AD(BC_RET1,
2602                              expr_toanyreg(fs, &e), 2);
2603              }
2604          }
2605      } else {

```

```

2606         if (e.k == VCALL) {/* Append all results from a call */
2607 notailcall:
2608             setbc_b(bcptr(fs, &e), 0);
2609             ins = BCINS_AD(BC_RET, fs->nactvar,
2610                 e.u.s.aux - fs->nactvar);
2611         } else {
2612             /* Force contiguous registers. */
2613             expr_tonextreg(fs, &e);
2614             ins = BCINS_AD(BC_RET, fs->nactvar, nret+1);
2615         }
2616     }
2617 }
2618 if (fs->flags & PROTO_CHILD) {
2619     /* May need to close upvalues first. */
2620     bcemit_AJ(fs, BC_UCLO, 0, 0);
2621 }
2622 bcemit_INS(fs, ins);
2623 }
2624
2625 /* Parse 'break' statement. */
2626 static void parse_break(LexState *ls)
2627 {
2628     ls->fs->bl->flags |= FSCOPE_BREAK;
2629     gola_new(ls, NAME_BREAK, VSTACK_GOTO, bcemit_imp(ls->fs));
2630 }
2631
2632 /* Parse label. */
2633 static void parse_label(LexState *ls)
2634 {
2635     FuncState *fs = ls->fs;
2636     ktap_str_t *name;
2637     int idx;
2638
2639     fs->lasttarget = fs->pc;
2640     fs->bl->flags |= FSCOPE_GOLA;
2641     kp_lex_next(ls); /* Skip '::'. */
2642     name = lex_str(ls);
2643     if (gola_findlabel(ls, name))
2644         kp_lex_error(ls, 0, KP_ERR_XLDUP, getstr(name));
2645     idx = gola_new(ls, name, VSTACK_LABEL, fs->pc);
2646     lex_check(ls, TK_label);
2647     /* Recursively parse trailing statements: labels and ';'.
2648     for (;;) {
2649         if (ls->tok == TK_label) {
2650             synlevel_begin(ls);
2651             parse_label(ls);
2652             synlevel_end(ls);
2653         } else if (ls->tok == ';') {
2654             kp_lex_next(ls);
2655         } else {
2656             break;
2657         }
2658     }
2659     /* Trailing label is considered to be outside of scope. */
2660     if (parse_isend(ls->tok) && ls->tok != TK_until)
2661         ls->vstack[idx].slot = fs->bl->nactvar;
2662     gola_resolve(ls, fs->bl, idx);
2663 }
2664
2665 /* -- Blocks, loops and conditional statements ----- */
2666
2667 /* Parse a block. */
2668 static void parse_block(LexState *ls)
2669 {
2670     FuncState *fs = ls->fs;
2671     FuncScope bl;
2672
2673     fscope_begin(fs, &bl, 0);
2674     parse_chunk(ls);
2675     fscope_end(fs);
2676 }
2677
2678 /* Parse 'while' statement. */
2679 static void parse_while(LexState *ls, BCLine line)
2680 {
2681     FuncState *fs = ls->fs;

```

```

2682 BCPos start, loop, condexit;
2683 FuncScope bl;
2684
2685 kp_lex_next(ls); /* Skip 'while'. */
2686 start = fs->lasttarget = fs->pc;
2687 condexit = expr_cond(ls);
2688 fscope_begin(fs, &bl, FSCOPE_LOOP);
2689 //lex_check(ls, TK_do);
2690 lex_check(ls, '{');
2691 loop = bcemit_AD(fs, BC_LOOP, fs->nactvar, 0);
2692 parse_block(ls);
2693 jmp_patch(fs, bcemit_jmp(fs), start);
2694 //lex_match(ls, TK_end, TK_while, line);
2695 lex_check(ls, '}');
2696 fscope_end(fs);
2697 jmp_tohere(fs, condexit);
2698 jmp_patchins(fs, loop, fs->pc);
2699 }
2700
2701 /* Parse 'repeat' statement. */
2702 static void parse_repeat(LexState *ls, BCLine line)
2703 {
2704     FuncState *fs = ls->fs;
2705     BCPos loop = fs->lasttarget = fs->pc;
2706     BCPos condexit;
2707     FuncScope bl1, bl2;
2708
2709     fscope_begin(fs, &bl1, FSCOPE_LOOP); /* Breakable loop scope. */
2710     fscope_begin(fs, &bl2, 0); /* Inner scope. */
2711     kp_lex_next(ls); /* Skip 'repeat'. */
2712     bcemit_AD(fs, BC_LOOP, fs->nactvar, 0);
2713     parse_chunk(ls);
2714     lex_match(ls, TK_until, TK_repeat, line);
2715     /* Parse condition (still inside inner scope). */
2716     condexit = expr_cond(ls);
2717     /* No upvalues? Just end inner scope. */
2718     if (!(bl2.flags & FSCOPE_UPVAL)) {
2719         fscope_end(fs);
2720     } else {
2721         /* Otherwise generate: cond: UCLO+JMP out,
2722          * !cond: UCLO+JMP loop. */
2723         parse_break(ls); /* Break from loop and close upvalues. */
2724         jmp_tohere(fs, condexit);
2725         fscope_end(fs); /* End inner scope and close upvalues. */
2726         condexit = bcemit_jmp(fs);
2727     }
2728     jmp_patch(fs, condexit, loop); /* Jump backwards if !cond. */
2729     jmp_patchins(fs, loop, fs->pc);
2730     fscope_end(fs); /* End loop scope. */
2731 }
2732
2733 /* Parse numeric 'for'. */
2734 static void parse_for_num(LexState *ls, ktag_str_t *varname, BCLine line)
2735 {
2736     FuncState *fs = ls->fs;
2737     BCReg base = fs->freereg;
2738     FuncScope bl;
2739     BCPos loop, loopend;
2740
2741     /* Hidden control variables. */
2742     var_new_fixed(ls, FORL_IDX, VARNAME_FOR_IDX);
2743     var_new_fixed(ls, FORL_STOP, VARNAME_FOR_STOP);
2744     var_new_fixed(ls, FORL_STEP, VARNAME_FOR_STEP);
2745     /* Visible copy of index variable. */
2746     var_new(ls, FORL_EXT, varname);
2747     lex_check(ls, '=');
2748     expr_next(ls);
2749     lex_check(ls, ',');
2750     expr_next(ls);
2751     if (lex_opt(ls, ',')) {
2752         expr_next(ls);
2753     } else {
2754         /* Default step is 1. */
2755         bcemit_AD(fs, BC_KSHORT, fs->freereg, 1);
2756         bcreg_reserve(fs, 1);
2757     }

```

```

2758 var_add(ls, 3); /* Hidden control variables. */
2759 //lex_check(ls, TK_do);
2760 lex_check(ls, ' ');
2761 lex_check(ls, '{');
2762 loop = bcemit_AJ(fs, BC_FORI, base, NO_JMP);
2763 fscope_begin(fs, &bl, 0); /* Scope for visible variables. */
2764 var_add(ls, 1);
2765 bcreg_reserve(fs, 1);
2766 parse_block(ls);
2767 fscope_end(fs);
2768 /* Perform loop inversion. Loop control instructions are at the end. */
2769 loopend = bcemit_AJ(fs, BC_FORL, base, NO_JMP);
2770 fs->bcbase[loopend].line = line; /* Fix line for control ins. */
2771 jmp_patchins(fs, loopend, loop+1);
2772 jmp_patchins(fs, loop, fs->pc);
2773 }
2774
2775 /*
2776 * Try to predict whether the iterator is next() and specialize the bytecode.
2777 * Detecting next() and pairs() by name is simplistic, but quite effective.
2778 * The interpreter backs off if the check for the closure fails at runtime.
2779 */
2780 static int predict_next(LexState *ls, FuncState *fs, BCPos pc)
2781 {
2782     BCIns ins = fs->bcbase[pc].ins;
2783     ktap_str_t *name;
2784     const ktap_val_t *o;
2785
2786     switch (bc_op(ins)) {
2787     case BC_MOV:
2788         name = var_get(ls, fs, bc_d(ins)).name;
2789         break;
2790     case BC_UGET:
2791         name = ls->vstack[fs->uvmmap[bc_d(ins)]].name;
2792         break;
2793     case BC_GGET:
2794         /* There's no inverse index (yet), so lookup the strings. */
2795         o = kp_tab_getstr(fs->kt, kp_str_newz("pairs"));
2796         if (o && tvhaskslot(o) && tvkslot(o) == bc_d(ins))
2797             return 1;
2798         o = kp_tab_getstr(fs->kt, kp_str_newz("next"));
2799         if (o && tvhaskslot(o) && tvkslot(o) == bc_d(ins))
2800             return 1;
2801         return 0;
2802     default:
2803         return 0;
2804     }
2805
2806     return (name->len == 5 && !strcmp(getstr(name), "pairs")) ||
2807         (name->len == 4 && !strcmp(getstr(name), "next"));
2808 }
2809
2810 /* Parse 'for' iterator. */
2811 static void parse_for_iter(LexState *ls, ktap_str_t *indexname)
2812 {
2813     FuncState *fs = ls->fs;
2814     ExpDesc e;
2815     BCReg nvars = 0;
2816     BCLine line;
2817     BCReg base = fs->freereg + 3;
2818     BCPos loop, loopend, exprpc = fs->pc;
2819     FuncScope bl;
2820     int isnext;
2821
2822     /* Hidden control variables. */
2823     var_new_fixed(ls, nvars++, VARNAME_FOR_GEN);
2824     var_new_fixed(ls, nvars++, VARNAME_FOR_STATE);
2825     var_new_fixed(ls, nvars++, VARNAME_FOR_CTL);
2826
2827     /* Visible variables returned from iterator. */
2828     var_new(ls, nvars++, indexname);
2829     while (lex_opt(ls, ', '))
2830         var_new(ls, nvars++, lex_str(ls));
2831     lex_check(ls, TK_in);
2832     line = ls->linenumber;
2833     assign_adjust(ls, 3, expr_list(ls, &e), &e);

```

```

2834  /* The iterator needs another 3 slots (func + 2 args). */
2835  bcreg_bump(fs, 3);
2836  isnext = (nvars <= 5 && predict_next(ls, fs, exprpc));
2837  var_add(ls, 3); /* Hidden control variables. */
2838  //lex_check(ls, TK_do);
2839  lex_check(ls, '(');
2840  lex_check(ls, '{');
2841  loop = bcemit_AJ(fs, isnext ? BC_ISNEXT : BC_JMP, base, NO_JMP);
2842  fscope_begin(fs, &bl, 0); /* Scope for visible variables. */
2843  var_add(ls, nvars-3);
2844  bcreg_reserve(fs, nvars-3);
2845  parse_block(ls);
2846  fscope_end(fs);
2847  /* Perform loop inversion. Loop control instructions are at the end. */
2848  jmp_patchins(fs, loop, fs->pc);
2849  bcemit_ABC(fs, isnext ? BC_ITERN : BC_ITERC, base, nvars-3+1, 2+1);
2850  loopend = bcemit_AJ(fs, BC_ITERL, base, NO_JMP);
2851  fs->bcbase[loopend-1].line = line; /* Fix line for control ins. */
2852  fs->bcbase[loopend].line = line;
2853  jmp_patchins(fs, loopend, loop+1);
2854 }
2855
2856 /* Parse 'for' statement. */
2857 static void parse_for(LexState *ls, BCLine line)
2858 {
2859     FuncState *fs = ls->fs;
2860     ktab_str_t *varname;
2861     FuncScope bl;
2862
2863     fscope_begin(fs, &bl, FSCOPE_LOOP);
2864     kp_lex_next(ls); /* Skip 'for'. */
2865     lex_check(ls, '(');
2866     varname = lex_str(ls); /* Get first variable name. */
2867     if (ls->tok == '=')
2868         parse_for_num(ls, varname, line);
2869     else if (ls->tok == ',' || ls->tok == TK_in)
2870         parse_for_iter(ls, varname);
2871     else
2872         err_syntax(ls, KP_ERR_XFOR);
2873     //lex_check(ls, '}');
2874     //lex_match(ls, TK_end, TK_for, line);
2875     lex_match(ls, '}', TK_for, line);
2876     fscope_end(fs); /* Resolve break list. */
2877 }
2878
2879 /* Parse condition and 'then' block. */
2880 static BCPos parse_then(LexState *ls)
2881 {
2882     BCPos condexit;
2883     kp_lex_next(ls); /* Skip 'if' or 'elseif'. */
2884     condexit = expr_cond(ls);
2885     lex_check(ls, '{');
2886     parse_block(ls);
2887     lex_check(ls, '}');
2888     return condexit;
2889 }
2890
2891 /* Parse 'if' statement. */
2892 static void parse_if(LexState *ls, BCLine line)
2893 {
2894     FuncState *fs = ls->fs;
2895     BCPos flist;
2896     BCPos escapelist = NO_JMP;
2897     flist = parse_then(ls);
2898     while (ls->tok == TK_elseif) { /* Parse multiple 'elseif' blocks. */
2899         jmp_append(fs, &escapelist, bcemit_jmp(fs));
2900         jmp_tohere(fs, flist);
2901         flist = parse_then(ls);
2902     }
2903     if (ls->tok == TK_else) { /* Parse optional 'else' block. */
2904         jmp_append(fs, &escapelist, bcemit_jmp(fs));
2905         jmp_tohere(fs, flist);
2906         kp_lex_next(ls); /* Skip 'else'. */
2907         lex_check(ls, '{');
2908         parse_block(ls);
2909         lex_check(ls, '}');

```

```

2910 } else {
2911     jmp_append(fs, &escapelist, flist);
2912 }
2913 jmp_tohere(fs, escapelist);
2914 //lex_match(ls, TK_end, TK_if, line);
2915 }
2916
2917 /* Parse 'trace' and 'trace_end' statement. */
2918 static void parse_trace(LexState *ls)
2919 {
2920     ExpDesc v, key, args;
2921     ktap_str_t *kdebug_str = kp_str_newz("kdebug");
2922     ktap_str_t *probe_str = kp_str_newz("trace_by_id");
2923     ktap_str_t *probe_end_str = kp_str_newz("trace_end");
2924     FuncState *fs = ls->fs;
2925     int token = ls->tok;
2926     BCIns ins;
2927     BCReg base;
2928     BCLine line = ls->linenumber;
2929
2930     if (token == TK_trace)
2931         kp_lex_read_string_until(ls, '{}');
2932     else
2933         kp_lex_next(ls); /* skip "trace_end" keyword */
2934
2935     /* kdebug */
2936     expr_init(&v, VGLOBAL, 0);
2937     v.u.sval = kdebug_str;
2938     expr_toanyreg(fs, &v);
2939
2940     /* fieldsel: kdebug.probe */
2941     expr_init(&key, VKSTR, 0);
2942     key.u.sval = token == TK_trace ? probe_str : probe_end_str;
2943     expr_index(fs, &v, &key);
2944
2945     /* funcargs*/
2946     expr_tonextreg(fs, &v);
2947
2948     if (token == TK_trace) {
2949         ktap_eventdesc_t *evdef_info;
2950         const char *str;
2951
2952         /* argument: EVENTDEF string */
2953         lex_check(ls, TK_string);
2954         str = svalue(&ls->tokval);
2955         evdef_info = kp_parse_events(str);
2956         if (!evdef_info)
2957             kp_lex_error(ls, 0, KP_ERR_XEVENTDEF, str);
2958
2959         /* pass a userspace pointer to kernel */
2960         expr_init(&args, VKNUM, 0);
2961         set_number(&args.u.nval, (ktap_number)evdef_info);
2962
2963         expr_tonextreg(fs, &args);
2964     }
2965
2966     /* argument: callback function */
2967     parse_body_no_args(ls, &args, 0, ls->linenumber);
2968
2969     expr_tonextreg(fs, &args);
2970
2971     base = v.u.s.info; /* base register for call */
2972     ins = BCINS_ABC(BC_CALL, base, 2, fs->freereg - base);
2973
2974     expr_init(&v, VCALL, bcemit_INS(fs, ins));
2975     v.u.s.aux = base;
2976     fs->bcbase[fs->pc - 1].line = line;
2977     fs->freereg = base+1; /* Leave one result by default. */
2978
2979     setbc_b(bcptr(fs, &v), 1); /* No results. */
2980 }
2981
2982
2983
2984 /* Parse 'profile' and 'tick' statement. */
2985 static void parse_timer(LexState *ls)

```

```

2986 {
2987     FuncState *fs = ls->fs;
2988     ExpDesc v, key, args;
2989     ktap_str_t *token_str = rawtsvalue(&ls->tokval);
2990     ktap_str_t *interval_str;
2991     BCLine line = ls->linenumber;
2992     BCIns ins;
2993     BCReg base;
2994
2995     kp_lex_next(ls); /* skip '-' */
2996
2997     kp_lex_read_string_until(ls, '{');
2998     interval_str = rawtsvalue(&ls->tokval);
2999     lex_check(ls, TK_string);
3000
3001     /* timer */
3002     expr_init(&v, VGLOBAL, 0);
3003     v.u.sval = kp_str_newz("timer");
3004     expr_toanyreg(fs, &v);
3005
3006     /* fieldsel: timer.profile, timer.tick */
3007     expr_init(&key, VKSTR, 0);
3008     key.u.sval = token_str;
3009     expr_index(fs, &v, &key);
3010
3011     /* funcargs */
3012     expr_tonextreg(fs, &v);
3013
3014     /* argument: interval string */
3015     expr_init(&args, VKSTR, 0);
3016     args.u.sval = interval_str;
3017
3018     expr_tonextreg(fs, &args);
3019
3020     /* argument: callback function */
3021     parse_body_no_args(ls, &args, 0, ls->linenumber);
3022
3023     expr_tonextreg(fs, &args);
3024
3025     base = v.u.s.info; /* base register for call */
3026     ins = BCINS_ABC(BC_CALL, base, 2, fs->freereg - base);
3027
3028     expr_init(&v, VCALL, bcemit_INS(fs, ins));
3029     v.u.s.aux = base;
3030     fs->bcbase[fs->pc - 1].line = line;
3031     fs->freereg = base+1; /* Leave one result by default. */
3032
3033     setbc_b(bcptr(fs, &v), 1); /* No results. */
3034 }
3035
3036 /* -- Parse statements ----- */
3037
3038 /* Parse a statement. Returns 1 if it must be the last one in a chunk. */
3039 static int parse_stmt(LexState *ls)
3040 {
3041     BCLine line = ls->linenumber;
3042     switch (ls->tok) {
3043     case TK_if:
3044         parse_if(ls, line);
3045         break;
3046     case TK_while:
3047         parse_while(ls, line);
3048         break;
3049     case TK_do:
3050         kp_lex_next(ls);
3051         parse_block(ls);
3052         lex_match(ls, TK_end, TK_do, line);
3053         break;
3054     case TK_for:
3055         parse_for(ls, line);
3056         break;
3057     case TK_repeat:
3058         parse_repeat(ls, line);
3059         break;
3060     case TK_function:
3061         parse_func(ls, line);

```

```

3062     break;
3063 case TK_local:
3064     kp_lex_next(ls);
3065     parse_local(ls);
3066     break;
3067 case TK_return:
3068     parse_return(ls);
3069     return 1; /* Must be last. */
3070 case TK_break:
3071     kp_lex_next(ls);
3072     parse_break(ls);
3073     return 0; /* Must be last. */
3074 case ';':
3075     kp_lex_next(ls);
3076     break;
3077 case TK_label:
3078     parse_label(ls);
3079     break;
3080 case TK_trace:
3081 case TK_trace_end:
3082     parse_trace(ls);
3083     break;
3084 case TK_profile:
3085 case TK_tick:
3086     parse_timer(ls);
3087     break;
3088 default:
3089     parse_call_assign(ls);
3090     break;
3091 }
3092 return 0;
3093 }
3094
3095 /* A chunk is a list of statements optionally separated by semicolons. */
3096 static void parse_chunk(LexState *ls)
3097 {
3098     int islast = 0;
3099
3100     synlevel_begin(ls);
3101     while (!islast && !parse_isend(ls->tok)) {
3102         islast = parse_stmt(ls);
3103         lex_opt(ls, ';');
3104         kp_assert(ls->fs->framesize >= ls->fs->freereg &&
3105             ls->fs->freereg >= ls->fs->nactvar);
3106         /* Free registers after each stmt. */
3107         ls->fs->freereg = ls->fs->nactvar;
3108     }
3109     synlevel_end(ls);
3110 }
3111
3112 /* Entry point of bytecode parser. */
3113 ktap_proto_t *kp_parse(LexState *ls)
3114 {
3115     FuncState fs;
3116     FuncScope bl;
3117     ktap_proto_t *pt;
3118
3119     ls->chunkname = kp_str_newz(ls->chunkarg);
3120     ls->level = 0;
3121     fs_init(ls, &fs);
3122     fs.linedefined = 0;
3123     fs.numparams = 0;
3124     fs.bcbase = NULL;
3125     fs.bcylim = 0;
3126     fs.flags |= PROTO_VARARG; /* Main chunk is always a vararg func. */
3127     fscope_begin(&fs, &bl, 0);
3128     bcemit_AD(&fs, BC_FUNCV, 0, 0); /* Placeholder. */
3129     kp_lex_next(ls); /* Read-ahead first token. */
3130     parse_chunk(ls);
3131     if (ls->tok != TK_eof)
3132         err_token(ls, TK_eof);
3133     pt = fs_finish(ls, ls->linenumber);
3134     kp_assert(fs.prevp == NULL);
3135     kp_assert(ls->fs == NULL);
3136     kp_assert(pt->sizeuv == 0);
3137     return pt;

```

3138 }
3139

[One Level Up](#)

[Top Level](#)

userspace/kp_lex.h - ktap

Data types defined

- [BCInsLine](#)
- [BCInsLine](#)
- [LexChar](#)
- [LexState](#)
- [LexState](#)
- [LexToken](#)
- [VarInfo](#)
- [VarInfo](#)

Macros defined

- [TKDEF](#)
- [TKENUM1](#)
- [TKENUM1](#)
- [TKENUM2](#)
- [TKENUM2](#)
- [_KTAP_LEX_H](#)

Source code

```
1 /*
2  * Lexical analyzer.
3  *
4  * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
5  *
6  * Adapted from luajit and lua interpreter.
7  * Copyright (C) 2005-2014 Mike Pall.
8  * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
9  */
10
11 #ifndef _KTAP_LEX_H
12 #define _KTAP_LEX_H
13
14 #include <stdarg.h>
15 #include "../include/err.h"
16 #include "../include/ktap_bc.h"
17 #include "kp_util.h"
18
19 /* ktap lexer tokens. */
20 #define TKDEF(_, __) \
21     _(trace) _(trace_end) _(argstr) _(probename) _(ffi) \
22     _(arg0)_(arg1) _(arg2) _(arg3) _(arg4) _(arg5) _(arg6) _(arg7) \
23     _(arg8) _(arg9) _(profile) _(tick) \
24     _(pid) _(tid) _(uid) _(cpu) _(execname) __ (incr, +=) \
25     __ (and, &&) _(break) _(do) _(else) _(elseif) _(end) _(false) \
26     _(for) _(function) _(goto) _(if) _(in) __ (local, var) _(nil) \
27     __ (not, !) __ (or, ||) \
28     _(repeat) _(return) _(then) _(true) _(until) _(while) \
29     __ (concat, ...) __ (dots, ...) __ (eq, ==) __ (ge, >=) __ (le, <=) \
```

```

30  ____(ne, !=) __(label, ::) __(number, <number>) __(name, <name>) \
31  __(string, <string>) __(eof, <eof>)
32
33  enum {
34      TK_OFS = 256,
35      #define TKENUM1(name)      TK_##name,
36      #define TKENUM2(name, sym)  TK_##name,
37      TKDEF(TKENUM1, TKENUM2)
38      #undef TKENUM1
39      #undef TKENUM2
40      TK_RESERVED = TK_while - TK_OFS
41  };
42
43  typedef int LexChar;      /* Lexical character. Unsigned ext. from char. */
44  typedef int LexToken;    /* Lexical token. */
45
46  /* Combined bytecode ins/line. Only used during bytecode generation. */
47  typedef struct BCInsLine {
48      BCIns ins;           /* Bytecode instruction. */
49      BCLine line;        /* Line number for this bytecode. */
50  } BCInsLine;
51
52  /* Info for local variables. Only used during bytecode generation. */
53  typedef struct VarInfo {
54      ktap_str_t *name;    /* Local variable name or goto/label name. */
55      BCPos startpc;      /* First point where the local variable is active. */
56      BCPos endpc;        /* First point where the local variable is dead. */
57      uint8_t slot;       /* Variable slot. */
58      uint8_t info;       /* Variable/goto/label info. */
59  } VarInfo;
60
61  /* lexer state. */
62  typedef struct LexState {
63      struct FuncState *fs; /* Current FuncState. Defined in kp_parse.c. */
64      ktap_val_t tokval;    /* Current token value. */
65      ktap_val_t lookaheadval; /* Lookahead token value. */
66      const char *p;       /* Current position in input buffer. */
67      const char *pe;      /* End of input buffer. */
68      LexChar c;           /* Current character. */
69      LexToken tok;        /* Current token. */
70      LexToken lookahead; /* Lookahead token. */
71      SBuf sb;             /* String buffer for tokens. */
72      BCLine linenumber;   /* Input line counter. */
73      BCLine lastline;    /* Line of last token. */
74      ktap_str_t *chunkname; /* Current chunk name (interned string). */
75      const char *chunkarg; /* Chunk name argument. */
76      const char *mode;    /* Allow loading bytecode (b) and/or source text (t) */
77      VarInfo *vstack;     /* Stack for names and extents of local variables. */
78      int sizevstack;      /* Size of variable stack. */
79      int vtop;            /* Top of variable stack. */
80      BCInsLine *bcstack; /* Stack for bytecode instructions/line numbers. */
81      int sizebcstack;     /* Size of bytecode stack. */
82      uint32_t level;      /* Syntactical nesting level. */
83  } LexState;
84
85  int kp_lex_setup(LexState *ls, const char *str);
86  void kp_lex_cleanup(LexState *ls);
87  void kp_lex_next(LexState *ls);
88  void kp_lex_read_string_until(LexState *ls, int c);
89  LexToken kp_lex_lookahead(LexState *ls);
90  const char *kp_lex_token2str(LexState *ls, LexToken tok);
91  void kp_lex_error(LexState *ls, LexToken tok, ErrMsg em, ...);
92  void kp_lex_init(void);
93
94  #endif

```

[One Level Up](#)

[Top Level](#)

userspace/kp_lex.c - ktap

Global variables defined

- [tokenames](#)

Functions defined

- [kp_lex_cleanup](#)
- [kp_lex_error](#)
- [kp_lex_init](#)
- [kp_lex_lookahead](#)
- [kp_lex_next](#)
- [kp_lex_read_string_until](#)
- [kp_lex_setup](#)
- [kp_lex_token2str](#)
- [kp_str2d](#)
- [lex_longstring](#)
- [lex_newline](#)
- [lex_next](#)
- [lex_number](#)
- [lex_save](#)
- [lex_savenext](#)
- [lex_scan](#)
- [lex_skipeq](#)
- [lex_string](#)

Macros defined

- [LEX_EOF](#)
- [TKSTR1](#)
- [TKSTR1](#)
- [TKSTR2](#)
- [TKSTR2](#)
- [lex_iseol](#)

Source code

```

2 * Lexical analyzer.
3 *
4 * This file is part of ktap by Jovi Zhangwei.
5 *
6 * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7 *
8 * Adapted from luajit and lua interpreter.
9 * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include "../include/ktap_types.h"
27 #include "../include/ktap_err.h"
28 #include "kp_util.h"
29 #include "kp_lex.h"
30 #include "kp_parse.h"
31
32 /* lexer token names. */
33 static const char *const tokenames[] = {
34 #define TKSTR1(name) #name,
35 #define TKSTR2(name, sym) #sym,
36 TKDEF(TKSTR1, TKSTR2)
37 #undef TKSTR1
38 #undef TKSTR2
39 NULL
40 };
41
42 /* -- Buffer handling ----- */
43
44 #define LEX_EOF (-1)
45 #define lex_iseol(ls) (ls->c == '\n' || ls->c == '\r')
46
47 /* Get next character. */
48 static inline LexChar lex_next(LexState *ls)
49 {
50     return (ls->c = ls->p < ls->pe ? (LexChar)(uint8_t)*ls->p++ : LEX_EOF);
51 }
52
53 /* Save character. */
54 static inline void lex_save(LexState *ls, LexChar c)
55 {
56     kp_buf_putb(&ls->sb, c);
57 }
58
59 /* Save previous character and get next character. */
60 static inline LexChar lex_savenext(LexState *ls)
61 {
62     lex_save(ls, ls->c);
63     return lex_next(ls);
64 }
65
66 /* Skip line break. Handles "\n", "\r", "\r\n" or "\n\r". */
67 static void lex_newline(LexState *ls)
68 {
69     LexChar old = ls->c;
70
71     kp_assert(lex_iseol(ls));
72     lex_next(ls); /* Skip "\n" or "\r". */
73     if (lex_iseol(ls) && ls->c != old)
74         lex_next(ls); /* Skip "\n\r" or "\r\n". */
75     if (++ls->linenumber >= KP_MAX_LINE)
76         kp_lex_error(ls, ls->tok, KP_ERR_XLINES);
77 }

```

```

78
79 /* -- Scanner for terminals ----- */
80
81 static int kp_str2d(const char *s, size_t len, ktap_number *result)
82 {
83     char *endptr;
84
85     if (strpbrk(s, "nN")) /* reject 'inf' and 'nan' */
86         return 0;
87     else
88         *result = (long)strtoul(s, &endptr, 0);
89
90     if (endptr == s)
91         return 0; /* nothing recognized */
92     while (kp_char_isspace((unsigned char)(*endptr)))
93         endptr++;
94     return (endptr == s + len); /* OK if no trailing characters */
95 }
96
97
98 /* Parse a number literal. */
99 static void lex_number(LexState *ls, ktap_val_t *tv)
100 {
101     LexChar c, xp = 'e';
102     ktap_number n = 0;
103
104     kp_assert(kp_char_isdigit(ls->c));
105     if ((c = ls->c) == '0' && (lex_savenext(ls) | 0x20) == 'x')
106         xp = 'p';
107     while (kp_char_isident(ls->c) || ls->c == '.' ||
108           ((ls->c == '-' || ls->c == '+') && (c | 0x20) == xp)) {
109         c = ls->c;
110         lex_savenext(ls);
111     }
112     lex_save(ls, '\0');
113     if (!kp_str2d(sbufB(&ls->sb), sbufLen(&ls->sb) - 1, &n))
114         kp_lex_error(ls, ls->tok, KP_ERR_XNUMBER);
115     set_number(tv, n);
116 }
117
118 /* Skip equal signs for "[=...=]" and "[=...=]" and return their count. */
119 static int lex_skipeq(LexState *ls)
120 {
121     int count = 0;
122     LexChar s = ls->c;
123
124     kp_assert(s == '[' || s == ']');
125     while (lex_savenext(ls) == '=')
126         count++;
127     return (ls->c == s) ? count : (-count) - 1;
128 }
129
130 /* Parse a long string or long comment (tv set to NULL). */
131 static void lex_longstring(LexState *ls, ktap_val_t *tv, int sep)
132 {
133     lex_savenext(ls); /* Skip second '['. */
134     if (lex_iseol(ls)) /* Skip initial newline. */
135         lex_newline(ls);
136     for (;;) {
137         switch (ls->c) {
138             case LEX_EOF:
139                 kp_lex_error(ls, TK_eof,
140                             tv ? KP_ERR_XLSTR : KP_ERR_XLCOM);
141                 break;
142             case ']':
143                 if (lex_skipeq(ls) == sep) {
144                     lex_savenext(ls); /* Skip second ']'. */
145                     goto endloop;
146                 }
147                 break;
148             case '\n':
149             case '\r':
150                 lex_save(ls, '\n');
151                 lex_newline(ls);
152                 if (!tv) /* Don't waste space for comments. */
153                     kp_buf_reset(&ls->sb);

```



```

230         if (kp_char_isdigit(lex_next(ls))) {
231             c = c*10 + (ls->c - '0');
232             if (c > 255) {
233 err_xesc:
234                 kp_lex_error(ls,
235                     TK_string,
236                     KP_ERR_XESC);
237             }
238             lex_next(ls);
239         }
240     }
241     lex_save(ls, c);
242     continue;
243 }
244 lex_save(ls, c);
245 lex_next(ls);
246 continue;
247 }
248 default:
249     lex_savenext(ls);
250     break;
251 }
252 }
253 lex_savenext(ls); /* Skip trailing delimiter. */
254 set_string(tv,
255     kp_parse_keepstr(ls, sbufB(&ls->sb)+1, sbuflen(&ls->sb)-2));
256 }
257
258 /* lex helper for parse trace and parse timer */
259 void kp_lex_read_string_until(LexState *ls, int c)
260 {
261     ktap_str_t *ts;
262
263     kp_buf_reset(&ls->sb);
264
265     while (ls->c == ' ')
266         lex_next(ls);
267
268     do {
269         lex_savenext(ls);
270     } while (ls->c != c && ls->c != LEX_EOF);
271
272     if (ls->c != c)
273         kp_lex_error(ls, ls->tok, KP_ERR_XTOKEN, c);
274
275     ts = kp_parse_keepstr(ls, sbufB(&ls->sb), sbuflen(&ls->sb));
276     ls->tok = TK_string;
277     set_string(&ls->tokval, ts);
278 }
279
280
281 /* -- Main lexical scanner ----- */
282
283 /* Get next lexical token. */
284 static LexToken lex_scan(LexState *ls, ktap_val_t *tv)
285 {
286     kp_buf_reset(&ls->sb);
287     for (;;) {
288         if (kp_char_isident(ls->c)) {
289             ktap_str_t *s;
290             if (kp_char_isdigit(ls->c)) { /* Numeric literal. */
291                 lex_number(ls, tv);
292                 return TK_number;
293             }
294             /* Identifier or reserved word. */
295             do {
296                 lex_savenext(ls);
297             } while (kp_char_isident(ls->c));
298             s = kp_parse_keepstr(ls, sbufB(&ls->sb),
299                 sbuflen(&ls->sb));
300             set_string(tv, s);
301             if (s->reserved > 0) /* Reserved word? */
302                 return TK_OFs + s->reserved;
303             return TK_name;
304         }
305     }

```

```

306     switch (ls->c) {
307     case '\n':
308     case '\r':
309         lex_newline(ls);
310         continue;
311     case ' ':
312     case '\t':
313     case '\v':
314     case '\f':
315         lex_next(ls);
316         continue;
317
318     case '#':
319         while (!lex_iseol(ls) && ls->c != LEX_EOF)
320             lex_next(ls);
321         break;
322     case '-':
323         lex_next(ls);
324         if (ls->c != '-')
325             return '-';
326         lex_next(ls);
327         if (ls->c == '[') { /* Long comment "--[...]=*" */
328             int sep = lex_skipeg(ls);
329             /* 'lex_skipeg' may dirty the buffer */
330             kp_buf_reset(&ls->sb);
331             if (sep >= 0) {
332                 lex_longstring(ls, NULL, sep);
333                 kp_buf_reset(&ls->sb);
334                 continue;
335             }
336         }
337         /* Short comment "--.*\n" */
338         while (!lex_iseol(ls) && ls->c != LEX_EOF)
339             lex_next(ls);
340         continue;
341     case '[': {
342         int sep = lex_skipeg(ls);
343         if (sep >= 0) {
344             lex_longstring(ls, tv, sep);
345             return TK_string;
346         } else if (sep == -1) {
347             return '[';
348         } else {
349             kp_lex_error(ls, TK_string, KP_ERR_XLDELIM);
350             continue;
351         }
352     }
353     case '+': {
354         lex_next(ls);
355         if (ls->c != '=')
356             return '+';
357         else {
358             lex_next(ls);
359             return TK_incr;
360         }
361     }
362     case '=':
363         lex_next(ls);
364         if (ls->c != '=')
365             return '=';
366         else {
367             lex_next(ls);
368             return TK_eq;
369         }
370     case '<':
371         lex_next(ls);
372         if (ls->c != '=')
373             return '<';
374         else {
375             lex_next(ls);
376             return TK_le;
377         }
378     case '>':
379         lex_next(ls);
380         if (ls->c != '=')
381             return '>';

```

```

382     else {
383         lex_next(ls);
384         return TK_ge;
385     }
386 case '!':
387     lex_next(ls);
388     if (ls->c != '=')
389         return TK_not;
390     else {
391         lex_next(ls);
392         return TK_ne;
393     }
394 case ':':
395     lex_next(ls);
396     if (ls->c != ':')
397         return ':';
398     else {
399         lex_next(ls);
400         return TK_label;
401     }
402 case '"':
403 case '\':
404     lex_string(ls, tv);
405     return TK_string;
406 case '.':
407     if (lex_savenext(ls) == '.') {
408         lex_next(ls);
409         if (ls->c == '.') {
410             lex_next(ls);
411             return TK_dots; /* ... */
412         }
413         return TK_concat; /* .. */
414     } else if (!kp_char_isdigit(ls->c)) {
415         return '.';
416     } else {
417         lex_number(ls, tv);
418         return TK_number;
419     }
420 case LEX_EOF:
421     return TK_eof;
422 case '&':
423     lex_next(ls);
424     if (ls->c != '&')
425         return '&';
426     else {
427         lex_next(ls);
428         return TK_and;
429     }
430 case '|':
431     lex_next(ls);
432     if (ls->c != '|')
433         return '|';
434     else {
435         lex_next(ls);
436         return TK_or;
437     }
438 default: {
439     LexChar c = ls->c;
440     lex_next(ls);
441     return c; /* Single-char tokens (+ - / ...). */
442 }
443 }
444 }
445 }
446
447 /* -- Lexer API ----- */
448
449 /* Setup lexer state. */
450 int kp_lex_setup(LexState *ls, const char *str)
451 {
452     ls->fs = NULL;
453     ls->pe = ls->p = NULL;
454     ls->p = str;
455     ls->pe = str + strlen(str);
456     ls->vstack = NULL;
457     ls->sizevstack = 0;

```

```

458     ls->vtop = 0;
459     ls->bcstack = NULL;
460     ls->sizebcstack = 0;
461     ls->lookahead = TK_eof; /* No look-ahead token. */
462     ls->linenumber = 1;
463     ls->lastline = 1;
464     lex_next(ls); /* Read-ahead first char. */
465     if (ls->c == 0xef && ls->p + 2 <= ls->pe &&
466         (uint8_t)ls->p[0] == 0xbb &&
467         (uint8_t)ls->p[1] == 0xbf) { /* Skip UTF-8 BOM (if buffered). */
468         ls->p += 2;
469         lex_next(ls);
470     }
471     if (ls->c == '#') { /* Skip POSIX #! header line. */
472         do {
473             lex_next(ls);
474             if (ls->c == LEX_EOF)
475                 return 0;
476         } while (!lex_iseol(ls));
477         lex_newline(ls);
478     }
479     return 0;
480 }
481
482 /* Cleanup lexer state. */
483 void kp_lex_cleanup(LexState *ls)
484 {
485     free(ls->bcstack);
486     free(ls->vstack);
487     kp_buf_free(&ls->sb);
488 }
489
490 /* Return next lexical token. */
491 void kp_lex_next(LexState *ls)
492 {
493     ls->lastline = ls->linenumber;
494     if (ls->lookahead == TK_eof) { /* No lookahead token? */
495         ls->tok = lex_scan(ls, &ls->tokval); /* Get next token. */
496     } else { /* Otherwise return lookahead token. */
497         ls->tok = ls->lookahead;
498         ls->lookahead = TK_eof;
499         ls->tokval = ls->lookaheadval;
500     }
501 }
502
503 /* Look ahead for the next token. */
504 LexToken kp_lex_lookahead(LexState *ls)
505 {
506     kp_assert(ls->lookahead == TK_eof);
507     ls->lookahead = lex_scan(ls, &ls->lookaheadval);
508     return ls->lookahead;
509 }
510
511 /* Convert token to string. */
512 const char *kp_lex_token2str(LexState *ls, LexToken tok)
513 {
514     if (tok > TK_OFS)
515         return tokennames[tok-TK_OFS-1];
516     else if (!kp_char_iscntrl(tok))
517         return kp_sprintf("%c", tok);
518     else
519         return kp_sprintf("char(%d)", tok);
520 }
521
522 /* Lexer error. */
523 void kp_lex_error(LexState *ls, LexToken tok, ErrMsg em, ...)
524 {
525     const char *tokstr;
526     va_list argp;
527
528     if (tok == 0) {
529         tokstr = NULL;
530     } else if (tok == TK_name || tok == TK_string || tok == TK_number) {
531         lex_save(ls, '\0');
532         tokstr = sbufB(&ls->sb);
533     } else {

```

```
534     tokstr = kp\_lex\_token2str(ls, tok);
535 }
536
537 va_start(argp, em);
538 kp\_err\_lex(ls->chunkname, tokstr, ls->linenumber, em, argp);
539 va_end(argp);
540 }
541
542 /* Initialize strings for reserved words. */
543 void kp\_lex\_init()
544 {
545     uint32_t i;
546
547     for (i = 0; i < TK_RESERVED; i++) {
548         ktap\_str\_t *s = kp\_str\_newz(tokennames[i]);
549         s->reserved = (uint8_t)(i+1);
550     }
551 }
552
```

[One Level Up](#)

[Top Level](#)

runtime/kp_tab.c - ktap

Global variables defined

- [kp_niltv](#)

Data types defined

- [ktap_node2](#)
- [ktap_node2_t](#)

Functions defined

- [clearapart](#)
- [clearhpart](#)
- [hashkey](#)
- [hashmask](#)
- [hist_record_cmp](#)
- [keyindex](#)
- [kp_tab_clear](#)
- [kp_tab_dup](#)
- [kp_tab_free](#)
- [kp_tab_get](#)
- [kp_tab_getint](#)
- [kp_tab_getstr](#)
- [kp_tab_incr](#)
- [kp_tab_incrint](#)
- [kp_tab_incrstr](#)
- [kp_tab_len](#)
- [kp_tab_new](#)
- [kp_tab_new_ah](#)
- [kp_tab_newkey](#)
- [kp_tab_next](#)
- [kp_tab_print_hist](#)
- [kp_tab_set](#)
- [kp_tab_setint](#)
- [kp_tab_setstr](#)

- [newhpart](#)
- [newtab](#)
- [string_convert](#)
- [tab_get](#)
- [tab_getint](#)
- [tab_getinth](#)
- [tab_getstr](#)
- [tab_histdump](#)
- [tab_set](#)
- [tab_setint](#)
- [tab_setinth](#)
- [tab_setstr](#)

Macros defined

- [DISTRIBUTION_STR](#)
- [TABLE_NARR_ENTRIES](#)
- [TABLE_NREC_ENTRIES](#)
- [hashgcref](#)
- [hashlohi](#)
- [hashnum](#)
- [hashstr](#)
- [niltv](#)
- [tab_lock](#)
- [tab_lock_init](#)
- [tab_unlock](#)

Source code

```

1  /*
2  * kp_tab.c - Table handling.
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Adapted from luajit and lua interpreter.
9  * Copyright (C) 2005-2014 Mike Pall.
10 * Copyright (C) 1994-2008 Lua.org, PUC-Rio.
11 *
12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or

```

```

18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include <linux/spinlock.h>
27 #include <linux/module.h>
28 #include <linux/kallsyms.h>
29 #include <linux/slab.h>
30 #include <linux/sort.h>
31 #include "../include/ktap_types.h"
32 #include "ktap.h"
33 #include "kp_vm.h"
34 #include "kp_obj.h"
35 #include "kp_str.h"
36 #include "kp_events.h"
37 #include "kp_tab.h"
38
39 #define tab_lock_init(t) \
40     do { \
41         (t)->lock = (arch_spinlock_t)__ARCH_SPIN_LOCK_UNLOCKED; \
42     } while (0)
43 #define tab_lock(t) \
44     do { \
45         local_irq_save(flags); \
46         arch_spin_lock(&(t)->lock); \
47     } while (0)
48 #define tab_unlock(t) \
49     do { \
50         arch_spin_unlock(&(t)->lock); \
51         local_irq_restore(flags); \
52     } while (0)
53
54
55 const ktap_val_t kp_niltv = { {NULL}, {KTAP_TNIL} } ;
56 #define niltv (&kp_niltv)
57
58 /* -- Object hashing ----- */
59
60 /* Hash values are masked with the table hash mask and used as an index. */
61 static __always_inline
62 ktap_node_t *hashmask(const ktap_tab_t *t, uint32_t hash)
63 {
64     ktap_node_t *n = t->node;
65     return &n[hash & t->hmask];
66 }
67
68 /* String hashes are precomputed when they are interned. */
69 #define hashstr(t, s) hashmask(t, (s)->hash)
70
71 #define hashlohi(t, lo, hi) hashmask((t), hashrot((lo), (hi)))
72 #define hashnum(t, o) hashlohi((t), (o)->val.n & 0xffffffff, 0)
73 #define hashgcref(t, o) hashlohi((t), \
74     ((unsigned long)(o)->val.gc & 0xffffffff), \
75     ((unsigned long)(o)->val.gc & 0xffffffff) + HASH_BIAS)
76
77
78 /* Hash an arbitrary key and return its anchor position in the hash table. */
79 static ktap_node_t *hashkey(const ktap_tab_t *t, const ktap_val_t *key)
80 {
81     kp_assert(!tvisint(key));
82     if (is_string(key))
83         return hashstr(t, rawtsvalue(key));
84     else if (is_number(key))
85         return hashnum(t, key);
86     else if (is_bool(key))
87         return hashmask(t, boolvalue(key));
88     else
89         return hashgcref(t, key);
90 }
91
92 /* -- Table creation and destruction ----- */
93

```

```

94  /* Create new hash part for table. */
95  static __always_inline
96  int newhpart(ktap_state_t *ks, ktap_tab_t *t, uint32_t hbits)
97  {
98      uint32_t hsize;
99      ktap_node_t *node;
100     kp_assert(hbits != 0);
101
102     if (hbits > KP_MAX_HBITS) {
103         //kp_error(ks, LJ_ERR_TABOV);
104         kp_error(ks, "table overflow\n");
105         return -1;
106     }
107     hsize = 1u << hbits;
108     node = vmalloc(hsize * sizeof(ktap_node_t));
109     if (!node)
110         return -ENOMEM;
111     t->freetop = &node[hsize];
112     t->node = node;
113     t->hmask = hsize-1;
114
115     return 0;
116 }
117
118 /*
119 * Q: Why all of these copies of t->hmask, t->node etc. to local variables?
120 * A: Because alias analysis for C is _really_ tough.
121 * Even state-of-the-art C compilers won't produce good code without this.
122 */
123
124 /* Clear hash part of table. */
125 static __always_inline void clearhpart(ktap_tab_t *t)
126 {
127     uint32_t i, hmask = t->hmask;
128     ktap_node_t *node = t->node;
129     kp_assert(t->hmask != 0);
130
131     for (i = 0; i <= hmask; i++) {
132         ktap_node_t *n = &node[i];
133         n->next = NULL;
134         set_nil(&n->key);
135         set_nil(&n->val);
136     }
137
138     t->hnum = 0;
139 }
140
141 /* Clear array part of table. */
142 static __always_inline void clearapart(ktap_tab_t *t)
143 {
144     uint32_t i, asize = t->asize;
145     ktap_val_t *array = t->array;
146     for (i = 0; i < asize; i++)
147         set_nil(&array[i]);
148 }
149
150 /* Create a new table. Note: the slots are not initialized (yet). */
151 static ktap_tab_t *newtab(ktap_state_t *ks, uint32_t asize, uint32_t hbits)
152 {
153     ktap_tab_t *t;
154
155     t = (ktap_tab_t *)kp_obj_new(ks, sizeof(ktap_tab_t));
156     t->gct = ~KTAP_TTAB;
157     t->array = NULL;
158     t->asize = 0; /* In case the array allocation fails. */
159     t->hmask = 0;
160
161     tab_lock_init(t);
162
163     if (asize > 0) {
164         if (asize > KP_MAX_ASIZE) {
165             kp_error(ks, "table overflow\n");
166             return NULL;
167         }
168
169         t->array = vmalloc(asize * sizeof(ktap_val_t));

```

```

170     if (!t->array)
171         return NULL;
172     t->asize = asize;
173 }
174 if (hbits)
175     if (newhpart(ks, t, hbits)) {
176         vfree(t->array);
177         return NULL;
178     }
179     return t;
180 }
181
182 /* Create a new table.
183 *
184 * The array size is non-inclusive. E.g. asize=128 creates array slots
185 * for 0..127, but not for 128. If you need slots 1..128, pass asize=129
186 * (slot 0 is wasted in this case).
187 *
188 * The hash size is given in hash bits. hbits=0 means no hash part.
189 * hbits=1 creates 2 hash slots, hbits=2 creates 4 hash slots and so on.
190 */
191 ktap_tab_t *kp_tab_new(ktap_state_t *ks, uint32_t asize, uint32_t hbits)
192 {
193     ktap_tab_t *t = newtab(ks, asize, hbits);
194     if (!t)
195         return NULL;
196
197     clearapart(t);
198     if (t->hmask > 0)
199         clearhpart(t);
200     return t;
201 }
202
203 #define TABLE_NARR_ENTRIES    255 /* PAGE_SIZE / sizeof(ktap_value) - 1 */
204 #define TABLE_NREC_ENTRIES    2048 /* (PAGE_SIZE * 20) / sizeof(ktap_tnode)*/
205
206 ktap_tab_t *kp_tab_new_ah(ktap_state_t *ks, int32_t a, int32_t h)
207 {
208     if (a == 0 && h == 0) {
209         a = TABLE_NARR_ENTRIES;
210         h = TABLE_NREC_ENTRIES;
211     }
212
213     return kp_tab_new(ks, (uint32_t)(a > 0 ? a+1 : 0), hsize2hbits(h));
214 }
215
216 /* Duplicate a table. */
217 ktap_tab_t *kp_tab_dup(ktap_state_t *ks, const ktap_tab_t *kt)
218 {
219     ktap_tab_t *t;
220     uint32_t asize, hmask;
221     int i;
222
223     /* allocate default table size */
224     t = kp_tab_new_ah(ks, 0, 0);
225     if (!t)
226         return NULL;
227
228     asize = kt->asize;
229     if (asize > 0) {
230         ktap_val_t *array = t->array;
231         ktap_val_t *karray = kt->array;
232         if (asize < 64) {
233             /* An inlined loop beats memcpy for < 512 bytes. */
234             uint32_t i;
235             for (i = 0; i < asize; i++)
236                 set_obj(&array[i], &karray[i]);
237         } else {
238             memcpy(array, karray, asize*sizeof(ktap_val_t));
239         }
240     }
241
242     hmask = kt->hmask;
243     for (i = 0; i <= hmask; i++) {
244         ktap_node_t *knode = &kt->node[i];
245         if (is_nil(&knode->key))

```

```

246         continue;
247         kp_tab_set(ks, t, &knode->key, &knode->val);
248     }
249     return t;
250 }
251
252 /* Clear a table. */
253 void kp_tab_clear(ktap_tab_t *t)
254 {
255     clearapart(t);
256     if (t->hmask > 0) {
257         ktap_node_t *node = t->node;
258         t->freetop = &node[t->hmask+1];
259         clearhpart(t);
260     }
261 }
262
263 /* Free a table. */
264 void kp_tab_free(ktap_state_t *ks, ktap_tab_t *t)
265 {
266     if (t->hmask > 0)
267         vfree(t->node);
268     if (t->asize > 0)
269         vfree(t->array);
270     kp_free(ks, t);
271 }
272
273 /* -- Table getters ----- */
274
275 static const ktap_val_t *tab_getinth(ktap_tab_t *t, uint32_t key)
276 {
277     ktap_val_t k;
278     ktap_node_t *n;
279
280     set_number(&k, (ktap_number)key);
281     n = hashnum(t, &k);
282     do {
283         if (is_number(&n->key) && nvalue(&n->key) == key) {
284             return &n->val;
285         }
286     } while ((n = n->next));
287     return niltv;
288 }
289
290 static __always_inline
291 const ktap_val_t *tab_getint(ktap_tab_t *t, uint32_t key)
292 {
293     return ((key < t->asize) ? arrayslot(t, key) :
294            tab_getinth(t, key));
295 }
296
297 void kp_tab_getint(ktap_tab_t *t, uint32_t key, ktap_val_t *val)
298 {
299     unsigned long flags;
300
301     tab_lock(t);
302     set_obj(val, tab_getint(t, key));
303     tab_unlock(t);
304 }
305
306 static const ktap_val_t *tab_getstr(ktap_tab_t *t, ktap_str_t *key)
307 {
308     ktap_node_t *n = hashstr(t, key);
309     do {
310         if (is_string(&n->key) && rawtsvalue(&n->key) == key)
311             return &n->val;
312     } while ((n = n->next));
313     return niltv;
314 }
315
316 void kp_tab_getstr(ktap_tab_t *t, ktap_str_t *key, ktap_val_t *val)
317 {
318     unsigned long flags;
319
320     tab_lock(t);
321     set_obj(val, tab_getstr(t, key));

```

```

322     tab_unlock(t);
323 }
324
325 static const ktap_val_t *tab_get(ktap_state_t *ks, ktap_tab_t *t,
326     const ktap_val_t *key)
327 {
328     if (is_string(key)) {
329         return tab_getstr(t, rawtsvalue(key));
330     } else if (is_number(key)) {
331         ktap_number nk = nvalue(key);
332         uint32_t k = (uint32_t)nk;
333         if (nk == (ktap_number)k) {
334             return tab_getint(t, k);
335         } else {
336             goto genlookup;    /* Else use the generic lookup. */
337         }
338     } else if (is_eventstr(key)) {
339         const ktap_str_t *ts;
340
341         if (!ks->current_event) {
342             kp_error(ks,
343                 "cannot stringify event str in invalid context\n");
344             return niltv;
345         }
346
347         ts = kp_event_stringify(ks);
348         if (!ts)
349             return niltv;
350
351         return tab_getstr(t, rawtsvalue(key));
352     } else if (!is_nil(key)) {
353         ktap_node_t *n;
354     genlookup:
355         n = hashkey(t, key);
356         do {
357             if (kp_obj_equal(&n->key, key))
358                 return &n->val;
359         } while ((n = n->next));
360     }
361     return niltv;
362 }
363
364 void kp_tab_get(ktap_state_t *ks, ktap_tab_t *t, const ktap_val_t *key,
365     ktap_val_t *val)
366 {
367     unsigned long flags;
368
369     tab_lock(t);
370     set_obj(val, tab_get(ks, t, key));
371     tab_unlock(t);
372 }
373
374 /* -- Table setters ----- */
375
376 /* Insert new key. Use Brent's variation to optimize the chain length. */
377 static ktap_val_t *kp_tab_newkey(ktap_state_t *ks, ktap_tab_t *t,
378     const ktap_val_t *key)
379 {
380     ktap_node_t *n = hashkey(t, key);
381
382     if (!is_nil(&n->val) || t->hmask == 0) {
383         ktap_node_t *nodebase = t->node;
384         ktap_node_t *collide, *freenode = t->freetop;
385
386         kp_assert(freenode >= nodebase &&
387             freenode <= nodebase+t->hmask+1);
388         do {
389             if (freenode == nodebase) { /* No free node found? */
390                 // kp_error(ks, LJ_ERR_TABOV);
391                 kp_error(ks, "table overflow\n");
392                 return NULL;
393             }
394         } while (!is_nil(&(--freenode)->key));
395
396         t->freetop = freenode;
397         collide = hashkey(t, &n->key);

```

```

398     if (collide != n) { /* Colliding node not the main node? */
399         while (collide->next != n)
400             /* Find predecessor. */
401                 collide = collide->next;
402         collide->next = freenode; /* Relink chain. */
403         /* Copy colliding node into free node and
404          * free main node. */
405         freenode->val = n->val;
406         freenode->key = n->key;
407         freenode->next = n->next;
408         n->next = NULL;
409         set_nil(&n->val);
410         /* Rechain pseudo-resurrected string keys with
411          * colliding hashes. */
412         while (freenode->next) {
413             ktap_node_t *nn = freenode->next;
414             if (is_string(&nn->key) && !is_nil(&nn->val) &&
415                 hashstr(t, rawtsvalue(&nn->key)) == n) {
416                 freenode->next = nn->next;
417                 nn->next = n->next;
418                 n->next = nn;
419             } else {
420                 freenode = nn;
421             }
422         }
423     } else { /* Otherwise use free node. */
424         freenode->next = n->next; /* Insert into chain. */
425         n->next = freenode;
426         n = freenode;
427     }
428 }
429 set_obj(&n->key, key);
430 t->hnum++;
431 return &n->val;
432 }
433
434 static ktap_val_t *tab_setinth(ktap_state_t *ks, ktap_tab_t *t, uint32_t key)
435 {
436     ktap_val_t k;
437     ktap_node_t *n;
438
439     set_number(&k, (ktap_number)key);
440     n = hashnum(t, &k);
441     do {
442         if (is_number(&n->key) && nvalue(&n->key) == key)
443             return &n->val;
444     } while ((n = n->next));
445     return kp_tab_newkey(ks, t, &k);
446 }
447
448 static __always_inline
449 ktap_val_t *tab_setint(ktap_state_t *ks, ktap_tab_t *t, uint32_t key)
450 {
451     return ((key < t->asize) ? arrayslot(t, key) :
452             tab_setinth(ks, t, key));
453 }
454
455 void kp_tab_setint(ktap_state_t *ks, ktap_tab_t *t,
456                  uint32_t key, const ktap_val_t *val)
457 {
458     ktap_val_t *v;
459     unsigned long flags;
460
461     tab_lock(t);
462     v = tab_setint(ks, t, key);
463     if (likely(v))
464         set_obj(v, val);
465     tab_unlock(t);
466 }
467
468 void kp_tab_incrint(ktap_state_t *ks, ktap_tab_t *t, uint32_t key,
469                   ktap_number n)
470 {
471     ktap_val_t *v;
472     unsigned long flags;
473

```

```

474     tab_lock(t);
475     v = tab_setint(ks, t, key);
476     if (unlikely(!v))
477         goto out;
478
479     if (likely(is_number(v)))
480         set_number(v, nvalue(v) + n);
481     else if (is_nil(v))
482         set_number(v, n);
483     else
484         kp_error(ks, "use '+=' operator on non-number value\n");
485
486 out:
487     tab_unlock(t);
488 }
489
490 static ktap_val_t *tab_setstr(ktap_state_t *ks, ktap_tab_t *t,
491                             const ktap_str_t *key)
492 {
493     ktap_val_t k;
494     ktap_node_t *n = hashstr(t, key);
495     do {
496         if (is_string(&n->key) && rawtsvalue(&n->key) == key)
497             return &n->val;
498     } while ((n = n->next));
499     set_string(&k, key);
500     return kp_tab_newkey(ks, t, &k);
501 }
502
503 void kp_tab_setstr(ktap_state_t *ks, ktap_tab_t *t, const ktap_str_t *key,
504                  const ktap_val_t *val)
505 {
506     ktap_val_t *v;
507     unsigned long flags;
508
509     tab_lock(t);
510     v = tab_setstr(ks, t, key);
511     if (likely(v))
512         set_obj(v, val);
513     tab_unlock(t);
514 }
515
516 void kp_tab_incrstr(ktap_state_t *ks, ktap_tab_t *t, const ktap_str_t *key,
517                   ktap_number n)
518 {
519     ktap_val_t *v;
520     unsigned long flags;
521
522     tab_lock(t);
523     v = tab_setstr(ks, t, key);
524     if (unlikely(!v))
525         goto out;
526
527     if (likely(is_number(v)))
528         set_number(v, nvalue(v) + n);
529     else if (is_nil(v))
530         set_number(v, n);
531     else
532         kp_error(ks, "use '+=' operator on non-number value\n");
533 out:
534     tab_unlock(t);
535 }
536
537 static ktap_val_t *tab_set(ktap_state_t *ks, ktap_tab_t *t,
538                           const ktap_val_t *key)
539 {
540     ktap_node_t *n;
541
542     if (is_string(key)) {
543         return tab_setstr(ks, t, rawtsvalue(key));
544     } else if (is_number(key)) {
545         ktap_number nk = nvalue(key);
546         uint32_t k = (ktap_number)nk;
547         if (nk == (ktap_number)k)
548             return tab_setint(ks, t, k);
549     } else if (itype(key) == KTAP_TKSTACK) {

```

```

550     /* change stack into string */
551     ktap_str_t *bt = kp_obj_kstack2str(ks, key->val.stack.depth,
552                                     key->val.stack.skip);
553     if (!bt)
554         return NULL;
555     return tab_setstr(ks, t, bt);
556 } else if (is_eventstr(key)) {
557     const ktap_str_t *ts;
558
559     if (!ks->current_event) {
560         kp_error(ks,
561                 "cannot stringify event str in invalid context\n");
562         return NULL;
563     }
564
565     ts = kp_event_stringify(ks);
566     if (!ts)
567         return NULL;
568
569     return tab_setstr(ks, t, ts);
570     /* Else use the generic lookup. */
571 } else if (is_nil(key)) {
572     //kp_error(ks, LJ_ERR_NILIDX);
573     kp_error(ks, "table nil index\n");
574     return NULL;
575 }
576 n = hashkey(t, key);
577 do {
578     if (kp_obj_equal(&n->key, key))
579         return &n->val;
580 } while ((n = n->next));
581 return kp_tab_newkey(ks, t, key);
582 }
583
584 void kp_tab_set(ktap_state_t *ks, ktap_tab_t *t,
585               const ktap_val_t *key, const ktap_val_t *val)
586 {
587     ktap_val_t *v;
588     unsigned long flags;
589
590     tab_lock(t);
591     v = tab_set(ks, t, key);
592     if (likely(v))
593         set_obj(v, val);
594     tab_unlock(t);
595 }
596
597 void kp_tab_incr(ktap_state_t *ks, ktap_tab_t *t, ktap_val_t *key,
598                ktap_number n)
599 {
600     ktap_val_t *v;
601     unsigned long flags;
602
603     tab_lock(t);
604     v = tab_set(ks, t, key);
605     if (unlikely(!v))
606         goto out;
607
608     if (likely(is_number(v)))
609         set_number(v, nvalue(v) + n);
610     else if (is_nil(v))
611         set_number(v, n);
612     else
613         kp_error(ks, "use '++' operator on non-number value\n");
614 out:
615     tab_unlock(t);
616 }
617
618
619 /* -- Table traversal ----- */
620
621 /* Get the traversal index of a key. */
622 static uint32_t keyindex(ktap_state_t *ks, ktap_tab_t *t,
623                        const ktap_val_t *key)
624 {
625     if (is_number(key)) {

```

```

626     ktap_number nk = nvalue(key);
627     uint32_t k = (uint32_t)nk;
628     /* Array key indexes: [0..t->asize-1] */
629     if ((uint32_t)k < t->asize && nk == (ktap_number)k)
630         return (uint32_t)k;
631 }
632
633 if (!is_nil(key)) {
634     ktap_node_t *n = hashkey(t, key);
635     do {
636         if (kp_obj_equal(&n->key, key))
637             return t->asize + (uint32_t)(n - (t->node));
638         /* Hash key indexes: [t->asize..t->asize+t->nmask] */
639     } while ((n = n->next));
640     //kp_err_msg(ks, LJ_ERR_NEXTIDX);
641     kp_error(ks, "table next index\n");
642     return 0; /* unreachable */
643 }
644 return ~0u; /* A nil key starts the traversal. */
645 }
646
647 /* Advance to the next step in a table traversal. */
648 int kp_tab_next(ktap_state_t *ks, ktap_tab_t *t, ktap_val_t *key)
649 {
650     unsigned long flags;
651     uint32_t i;
652
653     tab_lock(t);
654     i = keyindex(ks, t, key); /* Find predecessor key index. */
655
656     /* First traverse the array keys. */
657     for (i++; i < t->asize; i++)
658         if (!is_nil(arrayslot(t, i))) {
659             set_number(key, i);
660             set_obj(key + 1, arrayslot(t, i));
661             tab_unlock(t);
662             return 1;
663         }
664     /* Then traverse the hash keys. */
665     for (i -= t->asize; i <= t->hmask; i++) {
666         ktap_node_t *n = &t->node[i];
667         if (!is_nil(&n->val)) {
668             set_obj(key, &n->key);
669             set_obj(key + 1, &n->val);
670             tab_unlock(t);
671             return 1;
672         }
673     }
674     tab_unlock(t);
675     return 0; /* End of traversal. */
676 }
677
678 /* -- Table length calculation ----- */
679
680 int kp_tab_len(ktap_state_t *ks, ktap_tab_t *t)
681 {
682     unsigned long flags;
683     int i, len = 0;
684
685     tab_lock(t);
686     for (i = 0; i < t->asize; i++) {
687         ktap_val_t *v = &t->array[i];
688
689         if (is_nil(v))
690             continue;
691         len++;
692     }
693
694     for (i = 0; i <= t->hmask; i++) {
695         ktap_node_t *n = &t->node[i];
696
697         if (is_nil(&n->key))
698             continue;
699
700         len++;
701     }

```

```

702     tab\_unlock(t);
703     return len;
704 }
705
706 static void string\_convert(char *output, const char *input)
707 {
708     if (strlen(input) > 32) {
709         strncpy(output, input, 32-4);
710         memset(output + 32-4, '.', 3);
711     } else
712         memcpy(output, input, strlen(input));
713 }
714
715 typedef struct ktap\_node2 {
716     ktap\_val\_t key;
717     ktap\_val\_t val;
718 } ktap\_node2\_t;
719
720 static int hist\_record\_cmp(const void *i, const void *j)
721 {
722     ktap\_number n1 = nvalue(&((const ktap\_node2\_t *)i)->val);
723     ktap\_number n2 = nvalue(&((const ktap\_node2\_t *)j)->val);
724
725     if (n1 == n2)
726         return 0;
727     else if (n1 < n2)
728         return 1;
729     else
730         return -1;
731 }
732
733 /* todo: make histdump to be faster, just need to sort n entries, not all */
734
735 /* print_hist: key should be number/string/ip, value must be number */
736 static void tab\_histdump(ktap\_state\_t *ks, ktap\_tab\_t *t, int shownums)
737 {
738     long start_time, delta_time;
739     uint32_t i, asize = t->asize;
740     ktap\_val\_t *array = t->array;
741     uint32_t hmask = t->hmask;
742     ktap\_node\_t *node = t->node;
743     ktap\_node2\_t *sort_mem;
744     char dist_str[39];
745     int total = 0, sum = 0;
746
747     start_time = gettimeofday\_ns();
748
749     sort_mem = kmalloc((t->asize + t->hnum) * sizeof(ktap\_node2\_t),
750                       GFP_KERNEL);
751     if (!sort_mem)
752         return;
753
754     /* copy all values in table into sort_mem. */
755     for (i = 0; i < asize; i++) {
756         ktap\_val\_t *val = &array[i];
757         if (is\_nil(val))
758             continue;
759
760         if (!is\_number(val)) {
761             kp\_error(ks, "print_hist only can print number\n");
762             goto out;
763         }
764
765         set\_number(&sort_mem[total].key, i);
766         set\_obj(&sort_mem[total].val, val);
767         sum += nvalue(val);
768         total++;
769     }
770
771     for (i = 0; i <= hmask; i++) {
772         ktap\_node\_t *n = &node[i];
773         ktap\_val\_t *val = &n->val;
774
775         if (is\_nil(val))
776             continue;
777

```

```

778     if (!is_number(val)) {
779         kp_error(ks, "print_hist only can print number\n");
780         goto out;
781     }
782
783     set_obj(&sort_mem[total].key, &n->key);
784     set_obj(&sort_mem[total].val, val);
785     sum += nvalue(val);
786     total++;
787 }
788
789 /* sort */
790 sort(sort_mem, total, sizeof(ktap_node2_t), hist_record_cmp, NULL);
791
792 dist_str[sizeof(dist_str) - 1] = '\0';
793
794 for (i = 0; i < total; i++) {
795     ktap_val_t *key = &sort_mem[i].key;
796     ktap_number num = nvalue(&sort_mem[i].val);
797     int ratio;
798
799     if (!--shownums)
800         break;
801
802     memset(dist_str, ' ', sizeof(dist_str) - 1);
803     ratio = (num * (sizeof(dist_str) - 1)) / sum;
804     memset(dist_str, '@', ratio);
805
806     if (is_string(key)) {
807         //char buf[32] = {0};
808
809         //string_convert(buf, svalue(key));
810         if (rawtsvalue(key)->len > 32) {
811             kp_puts(ks, svalue(key));
812             kp_printf(ks, "%s\n%d\n", dist_str, num);
813         } else {
814             kp_printf(ks, "%31s |%-7d\n", svalue(key),
815                     dist_str, num);
816         }
817     } else if (is_number(key)) {
818         kp_printf(ks, "%31d |%-7d\n", nvalue(key),
819                 dist_str, num);
820     } else if (is_kip(key)) {
821         char str[KSYM_SYMBOL_LEN];
822         char buf[32] = {0};
823
824         SPRINT_SYMBOL(str, nvalue(key));
825         string_convert(buf, str);
826         kp_printf(ks, "%31s |%-7d\n", buf, dist_str, num);
827     }
828 }
829
830 if (!shownums && total)
831     kp_printf(ks, "%31s |\n", "...");
832
833 out:
834 kfree(sort_mem);
835
836 delta_time = (gettimeofday_ns() - start_time) / NSEC_PER_USEC;
837 kp_verbose_printf(ks, "tab_histdump time: %d (us)\n", delta_time);
838 }
839
840 #define DISTRIBUTION_STR "----- Distribution -----"
841 void kp_tab_print_hist(ktap_state_t *ks, ktap_tab_t *t, int n)
842 {
843     kp_printf(ks, "%31s%s\n", "value ", DISTRIBUTION_STR, " count");
844     tab_histdump(ks, t, n);
845 }

```

[One Level Up](#)

[Top Level](#)

runtime/kp_tab.h - ktap

Functions defined

- [hashrot](#)

Macros defined

- [FLS](#)
- [HASH_BIAS](#)
- [HASH_ROT1](#)
- [HASH_ROT2](#)
- [HASH_ROT3](#)
- [__KTAP_TAB_H](#)
- [arrayslot](#)
- [hsize2hbits](#)
- [kp_rol](#)
- [kp_ror](#)

Source code

```
1 #ifndef __KTAP_TAB_H
2 #define __KTAP_TAB_H
3
4 /* Hash constants. Tuned using a brute force search. */
5 #define HASH_BIAS      (-0x04c11db7)
6 #define HASH_ROT1     14
7 #define HASH_ROT2     5
8 #define HASH_ROT3     13
9
10 /* Every half-decent C compiler transforms this into a rotate instruction. */
11 #define kp_rol(x, n)   (((x)<<(n)) | ((x)>>(-(int)(n)&(8*sizeof(x)-1))))
12 #define kp_ror(x, n)   (((x)<<(-(int)(n)&(8*sizeof(x)-1))) | ((x)>>(n)))
13
14 /* Scramble the bits of numbers and pointers. */
15 static __always_inline uint32_t hashrot(uint32_t lo, uint32_t hi)
16 {
17     /* Prefer variant that compiles well for a 2-operand CPU. */
18     lo ^= hi; hi = kp_rol(hi, HASH_ROT1);
19     lo -= hi; hi = kp_rol(hi, HASH_ROT2);
20     hi ^= lo; hi -= kp_rol(lo, HASH_ROT3);
21     return hi;
22 }
23
24
25 #define FLS(x)         ((uint32_t)(__builtin_clz(x)^31))
26 #define hsize2hbits(s) ((s) ? ((s)==1 ? 1 : 1+FLS((uint32_t)((s)-1))) : 0)
27
28 #define arrayslot(t, i)      (&(t)->array[(i)])
29
30 void kp_tab_set(ktap_state_t *ks, ktap_tab_t *t, const ktap_val_t *key,
31               const ktap_val_t *val);
32 void kp_tab_setstr(ktap_state_t *ks, ktap_tab_t *t,
33                  const ktap_str_t *key, const ktap_val_t *val);
34 void kp_tab_incrstr(ktap_state_t *ks, ktap_tab_t *t, const ktap_str_t *key,
35                   ktap_number n);
36 void kp_tab_get(ktap_state_t *ks, ktap_tab_t *t, const ktap_val_t *key,
```

```

37     ktap\_val\_t *val);
38 void kp\_tab\_getstr(ktap\_tab\_t *t, ktap\_str\_t *key, ktap\_val\_t *val);
39
40 void kp\_tab\_getint(ktap\_tab\_t *t, uint32_t key, ktap\_val\_t *val);
41 void kp\_tab\_setint(ktap\_state\_t *ks, ktap\_tab\_t *t,
42     uint32_t key, const ktap\_val\_t *val);
43 void kp\_tab\_incrint(ktap\_state\_t *ks, ktap\_tab\_t *t, uint32_t key,
44     ktap\_number n);
45 ktap\_tab\_t *kp\_tab\_new(ktap\_state\_t *ks, uint32_t asize, uint32_t hbits);
46 ktap\_tab\_t *kp\_tab\_new\_ah(ktap\_state\_t *ks, int32_t a, int32_t h);
47 ktap\_tab\_t *kp\_tab\_dup(ktap\_state\_t *ks, const ktap\_tab\_t *kt);
48
49 void kp\_tab\_free(ktap\_state\_t *ks, ktap\_tab\_t *t);
50 int kp\_tab\_len(ktap\_state\_t *ks, ktap\_tab\_t *t);
51 void kp\_tab\_dump(ktap\_state\_t *ks, ktap\_tab\_t *t);
52 void kp\_tab\_clear(ktap\_tab\_t *t);
53 void kp\_tab\_print\_hist(ktap\_state\_t *ks, ktap\_tab\_t *t, int n);
54 int kp\_tab\_next(ktap\_state\_t *ks, ktap\_tab\_t *t, StkId key);
55 int kp\_tab\_sort\_next(ktap\_state\_t *ks, ktap\_tab\_t *t, StkId key);
56 void kp\_tab\_sort(ktap\_state\_t *ks, ktap\_tab\_t *t, ktap\_func\_t *cmp_func);
57 void kp\_tab\_incr(ktap\_state\_t *ks, ktap\_tab\_t *t, ktap\_val\_t *key,
58     ktap\_number n);
59 #endif /* KTAP\_TAB\_H */

```

[One Level Up](#)

[Top Level](#)

runtime/kp_obj.h - ktap

Macros defined

- [__KTAP_OBJ_H](#)
- [kp_obj_equal](#)

Source code

```
1 #ifndef __KTAP_OBJ_H
2 #define __KTAP_OBJ_H
3
4 void *kp_malloc(ktap_state_t *ks, int size);
5 void *kp_zalloc(ktap_state_t *ks, int size);
6 void kp_free(ktap_state_t *ks, void *addr);
7
8 void kp_obj_dump(ktap_state_t *ks, const ktap_val_t *v);
9 void kp_obj_show(ktap_state_t *ks, const ktap_val_t *v);
10 int kp_obj_len(ktap_state_t *ks, const ktap_val_t *rb);
11 ktap_obj_t *kp_obj_new(ktap_state_t *ks, size_t size);
12 int kp_obj_rawequal(const ktap_val_t *t1, const ktap_val_t *t2);
13 ktap_str_t *kp_obj_kstack2str(ktap_state_t *ks, uint16_t depth, uint16_t skip);
14 void kp_obj_free_gclist(ktap_state_t *ks, ktap_obj_t *o);
15 void kp_obj_freeall(ktap_state_t *ks);
16
17 #define kp_obj_equal(o1, o2) \
18     (((o1)->type == (o2)->type) && kp_obj_rawequal(o1, o2))
19
20 #endif /* __KTAP_OBJ_H */
```

userspace/kp_parse_events.c - ktap

Global variables defined

- [id_nr](#)
- [idmap](#)
- [idmap_size](#)
- [kprobes_text_end](#)
- [kprobes_text_start](#)
- [probe_list_head](#)

Data types defined

- [probe_cb_base](#)
- [probe_list](#)

Functions defined

- [add_event](#)
- [add_tracepoint](#)
- [check_kprobe_addr_prohibited](#)
- [cleanup_event_resources](#)
- [format_symbol_name](#)
- [get_id_array](#)
- [get_next_eventdef](#)
- [get_sys_event_filter_str](#)
- [idmap_clear](#)
- [idmap_free](#)
- [idmap_get_max_id](#)
- [idmap_init](#)
- [idmap_is_set](#)
- [idmap_set](#)
- [init_kprobe_prohibited_area](#)
- [kp_parse_events](#)
- [kprobe_symbol_actor](#)
- [parse_events_add_kprobe](#)
- [parse_events_add_probe](#)

- [parse_events_add_sdt](#)
- [parse_events_add_tracepoint](#)
- [parse_events_add_uprobe](#)
- [parse_events_resolve_symbol](#)
- [parse_events_resolve_symbol](#)
- [strim](#)
- [uprobe_symbol_actor](#)
- [write_kprobe_event](#)
- [write_uprobe_event](#)

Macros defined

- [KPROBE_EVENTS_PATH](#)
- [TRACING_EVENTS_DIR](#)
- [UPROBE_EVENTS_PATH](#)

Source code

```

1 /*
2  * parse_events.c - ktap events parser
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <unistd.h>
23 #include <stdio.h>
24 #include <stdlib.h>
25 #include <string.h>
26 #include <dirent.h>
27 #include <fcntl.h>
28 #include <ctype.h>
29
30 #include "../include/ktap_types.h"
31 #include "../include/ktap_bc.h"
32 #include "kp_symbol.h"
33 #include "kp_util.h"
34
35 #define TRACING_EVENTS_DIR "/sys/kernel/debug/tracing/events"
36
37 static u8 *idmap;
38 static int idmap_size = 1024; /* set init size */
39 static int id_nr;
40
41 static int idmap_init(void)

```

```

42 {
43     idmap = malloc(idmap_size);
44     if (!idmap)
45         return -1;
46
47     memset(idmap, 0, idmap_size);
48     return 0;
49 }
50
51 static void idmap_free(void)
52 {
53     id_nr = 0;
54     free(idmap);
55 }
56
57 static inline int idmap_is_set(int id)
58 {
59     return idmap[id / 8] & (1 << (id % 8));
60 }
61
62 static void idmap_set(int id)
63 {
64     if (id >= idmap_size * 8) {
65         int newsize = id + 100; /* allocate extra 800 id */
66         idmap = realloc(idmap, newsize);
67         memset(idmap + idmap_size, 0, newsize - idmap_size);
68         idmap_size = newsize;
69     }
70
71     if (!idmap_is_set(id))
72         id_nr++;
73
74     idmap[id / 8] = idmap[id / 8] | (1 << (id % 8));
75 }
76
77 static void idmap_clear(int id)
78 {
79     if (!idmap_is_set(id))
80         return;
81
82     id_nr--;
83     idmap[id / 8] = idmap[id / 8] & ~ (1 << (id % 8));
84 }
85
86 static int idmap_get_max_id(void)
87 {
88     return idmap_size * 8;
89 }
90
91 static int *get_id_array()
92 {
93     int *id_array;
94     int i, j = 0;
95
96     id_array = malloc(sizeof(int) * id_nr);
97     if (!id_array)
98         return NULL;
99
100     for (i = 0; i < idmap_get_max_id(); i++) {
101         if (idmap_is_set(i))
102             id_array[j++] = i;
103     }
104
105     return id_array;
106 }
107
108 static int add_event(char *evtid_path)
109 {
110     char id_buf[24];
111     int id, fd;
112
113     fd = open(evtid_path, O_RDONLY);
114     if (fd < 0) {
115         /*
116          * some tracepoint doesn't have id file, like ftrace,
117          * return success in here, and don't print error.

```

```

118     */
119     verbose_printf("warning: cannot open file %s\n", evtid_path);
120     return 0;
121 }
122
123 if (read(fd, id_buf, sizeof(id_buf)) < 0) {
124     fprintf(stderr, "read file error %s\n", evtid_path);
125     close(fd);
126     return -1;
127 }
128
129 id = atoll(id_buf);
130
131 idmap_set(id);
132
133 close(fd);
134 return 0;
135 }
136
137 static int add_tracepoint(const char *sys_name, const char *evt_name)
138 {
139     char evtid_path[PATH_MAX] = {0};
140
141     snprintf(evtid_path, PATH_MAX, "%s/%s/%s/id", TRACING_EVENTS_DIR,
142             sys_name, evt_name);
143     return add_event(evtid_path);
144 }
145
146 static int parse_events_add_tracepoint(char *sys, char *event)
147 {
148     process_available_tracepoints(sys, event, add_tracepoint);
149     return 0;
150 }
151
152 enum {
153     KPROBE_EVENT,
154     UPROBE_EVENT,
155 };
156
157 struct probe_list {
158     struct probe_list *next;
159     int type;
160     char event[64];
161 };
162
163 static struct probe_list *probe_list_head; /* for cleanup resources */
164
165 /*
166  * Some symbol format cannot write to uprobe_events in debugfs, like:
167  * symbol "check_one_fd.part.0" in glibc.
168  * For those symbols, we change the format to:
169  * "check_one_fd.part.0" -> "check_one_fd_part_0"
170  */
171 static char *format_symbol_name(const char *old_symbol)
172 {
173     char *new_name = strdup(old_symbol);
174     char *name = new_name;
175     int changed = 0;
176
177     if (!isalpha(*name) && *name != '_') {
178         *name = '_';
179         changed = 1;
180     }
181
182     while (*++name != '\0') {
183         if (!isalpha(*name) && !isdigit(*name) && *name != '_') {
184             *name = '_';
185             changed = 1;
186             continue;
187         }
188     }
189
190     if (changed)
191         fprintf(stderr,
192             "Warning: symbol \"%s\" transformed to event \"%s\"\n",
193             old_symbol, new_name);

```

```

194
195     /* this is a good name */
196         return new_name;
197 }
198
199
200 #define KPROBE_EVENTS_PATH "/sys/kernel/debug/tracing/kprobe_events"
201
202 /**
203  * @return 0 on success, otherwise -1
204  */
205 static int
206 write_kprobe_event(int fd, int ret_probe, const char *symbol,
207                   unsigned long start, char *fetch_args)
208 {
209     char probe_event[128] = {0};
210     char event[64] = {0};
211     struct probe_list *pl;
212     char event_id_path[128] = {0};
213     char *symbol_name;
214     int id_fd, ret;
215
216     /* In case some symbols cannot write to uprobe_events debugfs file */
217     symbol_name = format_symbol_name(symbol);
218
219     if (!fetch_args)
220         fetch_args = " ";
221
222     if (ret_probe) {
223         snprintf(event, 64, "ktap_kprobes_%d/ret_%s",
224                 getpid(), symbol_name);
225         /* Return probe point must be a symbol */
226         snprintf(probe_event, 128, "r:%s %s %s",
227                 event, symbol, fetch_args);
228     } else {
229         snprintf(event, 64, "ktap_kprobes_%d/%s",
230                 getpid(), symbol_name);
231         snprintf(probe_event, 128, "p:%s 0x%lx %s",
232                 event, start, fetch_args);
233     }
234
235     sprintf(event_id_path, "/sys/kernel/debug/tracing/events/%s/id", event);
236     /* if event id already exist, then don't write to kprobes_event again */
237     id_fd = open(event_id_path, O_RDONLY);
238     if (id_fd > 0) {
239         close(id_fd);
240
241         /* remember add event id to ids_array */
242         ret = add_event(event_id_path);
243         if (ret)
244             goto error;
245
246         goto out;
247     }
248
249     verbose_printf("write kprobe event %s\n", probe_event);
250
251     if (write(fd, probe_event, strlen(probe_event)) <= 0) {
252         fprintf(stderr, "Cannot write %s to %s\n", probe_event,
253                 KPROBE_EVENTS_PATH);
254         goto error;
255     }
256
257     /* add to cleanup list */
258     pl = malloc(sizeof(struct probe_list));
259     if (!pl)
260         goto error;
261
262     pl->type = KPROBE_EVENT;
263     pl->next = probe_list_head;
264     memcpy(pl->event, event, 64);
265     probe_list_head = pl;
266
267     ret = add_event(event_id_path);
268     if (ret < 0)
269         goto error;

```

```

270
271 out:
272     free(symbol_name);
273     return 0;
274
275 error:
276     free(symbol_name);
277     return -1;
278 }
279
280 static unsigned long kprobes_text_start;
281 static unsigned long kprobes_text_end;
282
283 static void init_kprobe_prohibited_area(void)
284 {
285     static int once = 0;
286
287     if (once > 0)
288         return;
289
290     once = 1;
291     kprobes_text_start = find_kernel_symbol("__kprobes_text_start");
292     kprobes_text_end   = find_kernel_symbol("__kprobes_text_end");
293 }
294
295 static int check_kprobe_addr_prohibited(unsigned long addr)
296 {
297     if (addr >= kprobes_text_start && addr <= kprobes_text_end)
298         return -1;
299
300     return 0;
301 }
302
303 struct probe_cb_base {
304     int fd;
305     int ret_probe;
306     const char *event;
307     char *binary;
308     char *symbol;
309     char *fetch_args;
310 };
311
312 static int kprobe_symbol_actor(void *arg, const char *name, char type,
313                               unsigned long start)
314 {
315     struct probe_cb_base *base = (struct probe_cb_base *)arg;
316
317     /* only can probe text function */
318     if (type != 't' && type != 'T')
319         return -1;
320
321     if (!strglobmatch(name, base->symbol))
322         return -1;
323
324     if (check_kprobe_addr_prohibited(start))
325         return -1;
326
327     /* ignore return code of write debugfs */
328     write_kprobe_event(base->fd, base->ret_probe, name, start,
329                       base->fetch_args);
330
331     return 0; /* success */
332 }
333
334 static int parse_events_add_kprobe(char *event)
335 {
336     char *symbol, *end;
337     struct probe_cb_base base;
338     int fd, ret;
339
340     fd = open(KPROBE_EVENTS_PATH, O_WRONLY);
341     if (fd < 0) {
342         fprintf(stderr, "Cannot open %s\n", KPROBE_EVENTS_PATH);
343         return -1;
344     }
345

```

```

346     end = strpbrk(event, "% ");
347     if (end)
348         symbol = strndup(event, end - event);
349     else
350         symbol = strdup(event);
351
352     base.fd = fd;
353     base.ret_probe = !!strstr(event, "%return");
354     base.symbol = symbol;
355     base.fetch_args = strchr(event, ' ');
356
357     init_kprobe_prohibited_area();
358
359     ret = kallsyms_parse(&base, kprobe_symbol_actor);
360     if (ret <= 0) {
361         fprintf(stderr, "cannot parse symbol \"%s\"\n", symbol);
362         ret = -1;
363     } else {
364         ret = 0;
365     }
366
367     free(symbol);
368     close(fd);
369
370     return ret;
371 }
372
373 #define UPROBE_EVENTS_PATH "/sys/kernel/debug/tracing/uprobe_events"
374
375 /**
376  * @return 0 on success, otherwise -1
377  */
378 static int
379 write_uprobe_event(int fd, int ret_probe, const char *binary,
380                   const char *symbol, unsigned long addr,
381                   char *fetch_args)
382 {
383     char probe_event[128] = {0};
384     char event[64] = {0};
385     struct probe_list *pl;
386     char event_id_path[128] = {0};
387     char *symbol_name;
388     int id_fd, ret;
389
390     /* In case some symbols cannot write to uprobe_events debugfs file */
391     symbol_name = format_symbol_name(symbol);
392
393     if (!fetch_args)
394         fetch_args = " ";
395
396     if (ret_probe) {
397         snprintf(event, 64, "ktap_uprobes_%d/ret_%s",
398                 getpid(), symbol_name);
399         snprintf(probe_event, 128, "r:%s %s:0x%lx %s",
400                 event, binary, addr, fetch_args);
401     } else {
402         snprintf(event, 64, "ktap_uprobes_%d/%s",
403                 getpid(), symbol_name);
404         snprintf(probe_event, 128, "p:%s %s:0x%lx %s",
405                 event, binary, addr, fetch_args);
406     }
407
408     sprintf(event_id_path, "/sys/kernel/debug/tracing/events/%s/id", event);
409     /* if event id already exist, then don't write to uprobes_event again */
410     id_fd = open(event_id_path, O_RDONLY);
411     if (id_fd > 0) {
412         close(id_fd);
413
414         /* remember add event id to ids_array */
415         ret = add_event(event_id_path);
416         if (ret)
417             goto error;
418
419         goto out;
420     }
421

```

```

422     verbose_printf("write uprobe event %s\n", probe_event);
423
424     if (write(fd, probe_event, strlen(probe_event)) <= 0) {
425         fprintf(stderr, "Cannot write %s to %s\n", probe_event,
426             UPROBE_EVENTS_PATH);
427         goto error;
428     }
429
430     /* add to cleanup list */
431     pl = malloc(sizeof(struct probe_list));
432     if (!pl)
433         goto error;
434
435     pl->type = UPROBE_EVENT;
436     pl->next = probe_list_head;
437     memcpy(pl->event, event, 64);
438     probe_list_head = pl;
439
440     ret = add_event(event_id_path);
441     if (ret < 0)
442         goto error;
443
444 out:
445     free(symbol_name);
446     return 0;
447
448 error:
449     free(symbol_name);
450     return -1;
451 }
452
453 /**
454  * TODO: avoid copy-paste stuff
455  *
456  * @return 1 on success, otherwise 0
457  */
458 #ifndef NO_LIBELF
459 static int parse_events_resolve_symbol(int fd, char *event, int type)
460 {
461     char *colon, *binary, *fetch_args;
462     unsigned long symbol_address;
463
464     colon = strchr(event, ':');
465     if (!colon)
466         return -1;
467
468     symbol_address = strtoul(colon + 1 /* skip ":" */, NULL, 0);
469
470     fetch_args = strchr(event, ' ');
471
472     /**
473      * We already have address, no need in resolving.
474      */
475     if (symbol_address) {
476         int ret;
477
478         binary = strdup(event, colon - event);
479         ret = write_uprobe_event(fd, strstr(event, "%return"), binary,
480             "NULL", symbol_address, fetch_args);
481         free(binary);
482         return ret;
483     }
484
485     fprintf(stderr, "error: cannot resolve event \"%s\" without libelf, "
486         "please recompile ktap with NO_LIBELF disabled\n",
487         event);
488     exit(EXIT_FAILURE);
489     return -1;
490 }
491
492 #else
493 static int uprobe_symbol_actor(const char *name, vaddr_t addr, void *arg)
494 {
495     struct probe_cb_base *base = (struct probe_cb_base *)arg;
496     int ret;
497

```

```

498     if (!strglobmatch(name, base->symbol))
499         return 0;
500
501     verbose_printf("uprobe: binary: \"%s\" symbol \"%s\" "
502                  "resolved to 0x%lx\n",
503                  base->binary, base->symbol, (unsigned long)addr);
504
505     ret = write_uprobe_event(base->fd, base->ret_probe, base->binary,
506                             name, addr, base->fetch_args);
507     if (ret)
508         return ret;
509
510     return 0;
511 }
512
513 static int parse_events_resolve_symbol(int fd, char *event, int type)
514 {
515     char *colon, *end;
516     vaddr_t symbol_address;
517     int ret;
518     struct probe_cb_base base = {
519         .fd = fd,
520         .event = event
521     };
522
523     colon = strchr(event, ':');
524     if (!colon)
525         return 0;
526
527     base.ret_probe = strstr(event, "%return");
528     symbol_address = strtoul(colon + 1 /* skip ":" */, NULL, 0);
529     base.binary = strndup(event, colon - event);
530
531     base.fetch_args = strchr(event, ' ');
532
533     /*
534      * We already have address, no need in resolving.
535      */
536     if (symbol_address) {
537         int ret;
538         ret = write_uprobe_event(fd, base.ret_probe, base.binary,
539                                 "NULL", symbol_address,
540                                 base.fetch_args);
541         free(base.binary);
542         return ret;
543     }
544
545     end = strpbrk(event, "% ");
546     if (end)
547         base.symbol = strndup(colon + 1, end - 1 - colon);
548     else
549         base.symbol = strdup(colon + 1);
550
551     ret = parse_dso_symbols(base.binary, type, uprobe_symbol_actor,
552                            (void *)&base);
553     if (!ret) {
554         fprintf(stderr, "error: cannot find symbol %s in binary %s\n",
555                base.symbol, base.binary);
556         ret = -1;
557     } else if (ret > 0) {
558         /* no error found when parse symbols */
559         ret = 0;
560     }
561
562     free(base.binary);
563     free(base.symbol);
564
565     return ret;
566 }
567 #endif
568
569 static int parse_events_add_uprobe(char *old_event, int type)
570 {
571     int ret;
572     int fd;
573

```

```

574 fd = open(UPROBE_EVENTS_PATH, O_WRONLY);
575 if (fd < 0) {
576     fprintf(stderr, "Cannot open %s\n", UPROBE_EVENTS_PATH);
577     return -1;
578 }
579
580 ret = parse_events_resolve_symbol(fd, old_event, type);
581
582 close(fd);
583 return ret;
584 }
585
586 static int parse_events_add_probe(char *old_event)
587 {
588     char *separator;
589
590     separator = strchr(old_event, ':');
591     if (!separator || (separator == old_event))
592         return parse_events_add_kprobe(old_event);
593     else
594         return parse_events_add_uprobe(old_event, FIND_SYMBOL);
595 }
596
597 static int parse_events_add_sdt(char *old_event)
598 {
599     return parse_events_add_uprobe(old_event, FIND_STAPSDT_NOTE);
600 }
601
602 static void strim(char *s)
603 {
604     size_t size;
605     char *end;
606
607     size = strlen(s);
608     if (!size)
609         return;
610
611     end = s + size - 1;
612     while (end >= s && isspace(*end))
613         end--;
614
615     *(end + 1) = '\0';
616 }
617
618 static int get_sys_event_filter_str(char *start,
619                                     char **sys, char **event, char **filter)
620 {
621     char *separator, *separator2, *ptr, *end;
622
623     while (*start == ' ')
624         start++;
625
626     /* find sys */
627     separator = strchr(start, ':');
628     if (!separator || (separator == start)) {
629         return -1;
630     }
631
632     ptr = malloc(separator - start + 1);
633     if (!ptr)
634         return -1;
635
636     strncpy(ptr, start, separator - start);
637     ptr[separator - start] = '\0';
638
639     strim(ptr);
640     *sys = ptr;
641
642     if (!strcmp(*sys, "probe") && (*(separator + 1) == '/')) {
643         /* it's uprobe event */
644         separator2 = strchr(separator + 1, ':');
645         if (!separator2)
646             return -1;
647     } else
648         separator2 = separator;
649

```

```

650     /* find filter */
651     end = start + strlen(start);
652     while (*--end == ' ') {
653     }
654
655     if (*end == '/') {
656         char *filter_start;
657
658         filter_start = strchr(separator2, '/');
659         if (filter_start == end)
660             return -1;
661
662         ptr = malloc(end - filter_start);
663         if (!ptr)
664             return -1;
665
666         memcpy(ptr, filter_start + 1, end - filter_start - 1);
667         ptr[end - filter_start - 1] = '\0';
668
669         *filter = ptr;
670
671         end = filter_start;
672     } else {
673         *filter = NULL;
674         end++;
675     }
676
677     /* find event */
678     ptr = malloc(end - separator);
679     if (!ptr)
680         return -1;
681
682     memcpy(ptr, separator + 1, end - separator - 1);
683     ptr[end - separator - 1] = '\0';
684
685     strim(ptr);
686     *event = ptr;
687
688     return 0;
689 }
690
691 static char *get_next_eventdef(char *str)
692 {
693     char *separator;
694
695     separator = strchr(str, ',');
696     if (!separator)
697         return str + strlen(str);
698
699     *separator = '\0';
700     return separator + 1;
701 }
702
703 ktap_eventdesc_t *kp_parse_events(const char *eventdef)
704 {
705     char *str = strdup(eventdef);
706     char *sys, *event, *filter, *next;
707     ktap_eventdesc_t *evdef_info;
708     int ret;
709
710     idmap_init();
711
712     parse_next_eventdef:
713     next = get_next_eventdef(str);
714
715     if (get_sys_event_filter_str(str, &sys, &event, &filter))
716         goto error;
717
718     verbose_printf("parse_eventdef: sys[%s], event[%s], filter[%s]\n",
719                 sys, event, filter);
720
721     if (!strcmp(sys, "probe"))
722         ret = parse_events_add_probe(event);
723     else if (!strcmp(sys, "sdt"))
724         ret = parse_events_add_sdt(event);
725     else

```

```

726     ret = parse\_events\_add\_tracepoint(sys, event);
727
728     if (ret)
729         goto error;
730
731     /* don't trace ftrace:function when all tracepoints enabled */
732     if (!strcmp(sys, ""))
733         idmap\_clear(1);
734
735
736     if (filter && *next != '\\0') {
737         fprintf(stderr, "Error: eventdef only can append one filter\n");
738         goto error;
739     }
740
741     str = next;
742     if (*next != '\\0')
743         goto parse_next_eventdef;
744
745     evdef_info = malloc(sizeof(*evdef_info));
746     if (!evdef_info)
747         goto error;
748
749     evdef_info->nr = id\_nr;
750     evdef_info->id_arr = get\_id\_array();
751     evdef_info->filter = filter;
752
753     idmap\_free();
754     return evdef_info;
755 error:
756     idmap\_free();
757     cleanup\_event\_resources();
758     return NULL;
759 }
760
761 void cleanup\_event\_resources(void)
762 {
763     struct probe\_list *pl;
764     const char *path;
765     char probe_event[128] = {0};
766     int fd, ret;
767
768     for (pl = probe\_list\_head; pl; pl = pl->next) {
769         if (pl->type == KPROBE_EVENT)
770             path = KPROBE\_EVENTS\_PATH;
771         else if (pl->type == UPROBE_EVENT)
772             path = UPROBE\_EVENTS\_PATH;
773         else {
774             fprintf(stderr, "Cannot cleanup event type %d\n",
775                 pl->type);
776             continue;
777         }
778
779         snprintf(probe_event, 128, "-:%s", pl->event);
780
781         fd = open(path, O_WRONLY);
782         if (fd < 0) {
783             fprintf(stderr, "Cannot open %s\n", UPROBE\_EVENTS\_PATH);
784             continue;
785         }
786
787         ret = write(fd, probe_event, strlen(probe_event));
788         if (ret <= 0) {
789             fprintf(stderr, "Cannot write %s to %s\n", probe_event,
790                 path);
791             close(fd);
792             continue;
793         }
794
795         close(fd);
796     }
797 }
798

```

userspace/kp_reader.c - ktap

Functions defined

- [block_sigint](#)
- [kp_create_reader](#)
- [reader_thread](#)
- [sigfunc](#)

Macros defined

- [MAX_BUFLLEN](#)
- [PATH_MAX](#)
- [handle_error](#)

Source code

```
1  /*
2  * reader.c - ring buffer reader in userspace
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <stdio.h>
23 #include <stdlib.h>
24 #include <string.h>
25 #include <unistd.h>
26 #include <sys/mman.h>
27 #include <sys/stat.h>
28 #include <sys/poll.h>
29 #include <sys/signal.h>
30 #include <fcntl.h>
31 #include <pthread.h>
32
33 #define MAX_BUFLLEN 131072
34 #define PATH_MAX 128
35
36 #define handle_error(str) do { perror(str); exit(-1); } while(0)
37
38 void sigfunc(int signo)
39 {
40     /* should not not reach here */
41 }
42
43 static void block_sigint()
44 {
45     sigset_t mask;
```

```

46     sigemptyset(&mask);
47     sigaddset(&mask, SIGINT);
48
49     pthread_sigmask(SIG_BLOCK, &mask, NULL);
50 }
51
52
53 static void *reader_thread(void *data)
54 {
55     char buf[MAX_BUFLen];
56     char filename[PATH_MAX];
57     const char *output = data;
58     int failed = 0, fd, out_fd, len;
59
60     block_sigint();
61
62     if (output) {
63         out_fd = open(output, O_CREAT | O_WRONLY | O_TRUNC,
64                     S_IRUSR|S_IWUSR);
65         if (out_fd < 0) {
66             fprintf(stderr, "Cannot open output file %s\n", output);
67             return NULL;
68         }
69     } else
70         out_fd = 1;
71
72     sprintf(filename, "/sys/kernel/debug/ktap/trace_pipe_%d", getpid());
73
74     open_again:
75     fd = open(filename, O_RDONLY);
76     if (fd < 0) {
77         usleep(10000);
78
79         if (failed++ == 10) {
80             fprintf(stderr, "Cannot open file %s\n", filename);
81             return NULL;
82         }
83         goto open_again;
84     }
85
86     while ((len = read(fd, buf, sizeof(buf))) > 0)
87         write(out_fd, buf, len);
88
89     close(fd);
90     close(out_fd);
91
92     return NULL;
93 }
94
95 int kp_create_reader(const char *output)
96 {
97     pthread_t reader;
98
99     signal(SIGINT, sigfunc);
100
101     if (pthread_create(&reader, NULL, reader_thread, (void *)output) < 0)
102         handle_error("pthread_create reader_thread failed\n");
103
104     return 0;
105 }
106

```

[One Level Up](#)

[Top Level](#)

userspace/kp_symbol.h - ktap

Data types defined

- [symbol_actor](#)
- [vaddr_t](#)

Macros defined

- [FIND_STAPSDT_NOTE](#)
- [FIND_SYMBOL](#)

Source code

```
1  /*
2  * symbol.h - extract symbols from DSO.
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2013 Azat Khuzhin <a3at.mail@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22
23 #define FIND_SYMBOL 1
24 #define FIND_STAPSDT_NOTE 2
25
26 #ifndef NO_LIBELF
27
28 #include <gelf.h>
29 #include <sys/queue.h>
30
31 typedef GElf_Addr vaddr_t;
32 typedef int (*symbol_actor)(const char *name, vaddr_t addr, void *arg);
33
34 /**
35 * Parse all DSO symbols/sdt notes and all for every of them
36 * an actor.
37 *
38 * @exec - path to DSO
39 * @type - see FIND_*
40 * @symbol_actor - actor to call (callback)
41 * @arg - argument for @actor
42 *
43 * @return
44 * If there have errors, return negative value;
45 * No symbols found, return 0;
46 * Otherwise return number of dso symbols found
47 */
48 int
49 parse_dso_symbols(const char *exec, int type, symbol_actor actor, void *arg);
50 #endif
```

userspace/kp_symbol.c - ktap

Global variables defined

- [dbg_bin_dir](#)
- [dbg_link_name](#)

Functions defined

- [dso_follow_debuglink](#)
- [dso_sdt_notes](#)
- [dso_symbols](#)
- [elf_section_by_name](#)
- [elf_symbols](#)
- [find_load_address](#)
- [parse_dso_symbols](#)
- [sdt_note_addr](#)
- [sdt_note_data](#)
- [sdt_note_name](#)

Macros defined

- [SDT_NOTE_COUNT](#)
- [SDT_NOTE_NAME](#)
- [SDT_NOTE_SCN](#)
- [SDT_NOTE_TYPE](#)

Source code

```
1 /*
2  * symbol.c
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2013 Azat Khuzhin <a3at.mail@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
```

```

22 #include <stdio.h>
23 #include <stdlib.h>
24
25 #include <sys/types.h>
26 #include <sys/stat.h>
27 #include <unistd.h>
28 #include <fcntl.h>
29 #include <string.h>
30 #include <linux/limits.h>
31
32 #include <libelf.h>
33
34 #include "../include/ktap_types.h"
35 #include "kp_symbol.h"
36
37 const char *dbg_link_name = ".gnu_debuglink";
38 const char *dbg_bin_dir = "/usr/lib/debug";
39
40 static Elf_Scn *elf_section_by_name(Elf *elf, GElf_Ehdr *ep,
41                                     GElf_Shdr *shp, const char *name)
42 {
43     Elf_Scn *scn = NULL;
44
45     /* Elf is corrupted/truncated, avoid calling elf_strptr. */
46     if (!elf_rawdata(elf_getscn(elf, ep->e_shstrndx), NULL))
47         return NULL;
48
49     while ((scn = elf_nextscn(elf, scn)) != NULL) {
50         char *str;
51
52         gelf_getshdr(scn, shp);
53         str = elf_strptr(elf, ep->e_shstrndx, shp->sh_name);
54         if (!strcmp(name, str))
55             break;
56     }
57
58     return scn;
59 }
60
61 /**
62  * @return v_addr of "LOAD" program header, that have zero offset.
63  */
64 static int find_load_address(Elf *elf, vaddr_t *load_address)
65 {
66     GElf_Phdr phdr;
67     size_t i, phdrnum;
68
69     if (elf_getphdrnum(elf, &phdrnum))
70         return -1;
71
72     for (i = 0; i < phdrnum; i++) {
73         if (gelf_getphdr(elf, i, &phdr) == NULL)
74             return -1;
75
76         if (phdr.p_type != PT_LOAD || phdr.p_offset != 0)
77             continue;
78
79         *load_address = phdr.p_vaddr;
80         return 0;
81     }
82
83     /* cannot found load address */
84     return -1;
85 }
86
87 static size_t elf_symbols(GElf_Shdr shdr)
88 {
89     return shdr.sh_size / shdr.sh_entsize;
90 }
91
92 static int dso_symbols(Elf *elf, symbol_actor actor, void *arg)
93 {
94     Elf_Data *elf_data = NULL;
95     Elf_Scn *scn = NULL;
96     GElf_Sym sym;
97     GElf_Shdr shdr;

```

```

98     int symbols_count = 0;
99     vaddr_t load_address;
100
101     if (find_load_address(elf, &load_address))
102         return -1;
103
104     while ((scn = elf_nextscn(elf, scn))) {
105         int i;
106
107         gelf_getshdr(scn, &shdr);
108
109         if (shdr.sh_type != SHT_SYMTAB)
110             continue;
111
112         elf_data = elf_getdata(scn, elf_data);
113
114         for (i = 0; i < elf_symbols(shdr); i++) {
115             char *name;
116             vaddr_t addr;
117             int ret;
118
119             gelf_getsym(elf_data, i, &sym);
120
121             if (GELF_ST_TYPE(sym.st_info) != STT_FUNC)
122                 continue;
123
124             name = elf_strptr(elf, shdr.sh_link, sym.st_name);
125             addr = sym.st_value - load_address;
126
127             ret = actor(name, addr, arg);
128             if (ret)
129                 return ret;
130
131             ++symbols_count;
132         }
133     }
134
135     return symbols_count;
136 }
137
138 #define SDT_NOTE_TYPE 3
139 #define SDT_NOTE_COUNT 3
140 #define SDT_NOTE_SCN ".note.stapsdt"
141 #define SDT_NOTE_NAME "stapsdt"
142
143 static vaddr_t sdt_note_addr(Elf *elf, const char *data, size_t len, int type)
144 {
145     vaddr_t vaddr;
146
147     /*
148      * Three addresses need to be obtained :
149      * Marker location, address of base section and semaphore location
150      */
151     union {
152         Elf64_Addr a64[3];
153         Elf32_Addr a32[3];
154     } buf;
155
156     /*
157      * dst and src are required for translation from file to memory
158      * representation
159      */
160     Elf_Data dst = {
161         .d_buf = &buf, .d_type = ELF_T_ADDR, .d_version = EV_CURRENT,
162         .d_size = gelf_fsize(elf, ELF_T_ADDR, SDT_NOTE_COUNT, EV_CURRENT),
163         .d_off = 0, .d_align = 0
164     };
165
166     Elf_Data src = {
167         .d_buf = (void *) data, .d_type = ELF_T_ADDR,
168         .d_version = EV_CURRENT, .d_size = dst.d_size, .d_off = 0,
169         .d_align = 0
170     };
171
172     /* Check the type of each of the notes */
173     if (type != SDT_NOTE_TYPE)

```

```

174     return 0;
175
176     if (len < dst.d_size + SDT_NOTE_COUNT)
177         return 0;
178
179     /* Translation from file representation to memory representation */
180     if (gelf_xlatetom(elf, &dst, &src,
181         elf_getident(elf, NULL)[EI_DATA]) == NULL)
182         return 0; /* TODO */
183
184     memcpy(&vaddr, &buf, sizeof(vaddr));
185
186     return vaddr;
187 }
188
189 static const char *sdt_note_name(Elf *elf, GElf_Nhdr *nhdr, const char *data)
190 {
191     const char *provider = data + gelf_fsize(elf,
192         ELF_T_ADDR, SDT_NOTE_COUNT, EV_CURRENT);
193     const char *name = (const char *)memchr(provider, '\\0',
194         data + nhdr->n_descsz - provider);
195
196     if (name++ == NULL)
197         return NULL;
198
199     return name;
200 }
201
202 static const char *sdt_note_data(const Elf_Data *data, size_t off)
203 {
204     return ((data->d_buf) + off);
205 }
206
207 static int dso_sdt_notes(Elf *elf, symbol_actor actor, void *arg)
208 {
209     GElf_Ehdr ehdr;
210     Elf_Scn *scn = NULL;
211     Elf_Data *data;
212     GElf_Shdr shdr;
213     size_t shstrndx;
214     size_t next;
215     GElf_Nhdr nhdr;
216     size_t name_off, desc_off, offset;
217     vaddr_t vaddr = 0;
218     int symbols_count = 0;
219
220     if (gelf_getehdr(elf, &ehdr) == NULL)
221         return 0;
222     if (elf_getshdrstrndx(elf, &shstrndx) != 0)
223         return 0;
224
225     /*
226      * Look for section type = SHT_NOTE, flags = no SHF_ALLOC
227      * and name = .note.stapsdt
228     */
229     scn = elf_section_by_name(elf, &ehdr, &shdr, SDT_NOTE_SCN);
230     if (!scn)
231         return 0;
232     if (!(shdr.sh_type == SHT_NOTE) || (shdr.sh_flags & SHF_ALLOC))
233         return 0;
234
235     data = elf_getdata(scn, NULL);
236
237     for (offset = 0;
238         (next = gelf_getnote(data, offset, &nhdr, &name_off, &desc_off)) > 0;
239         offset = next) {
240         const char *name;
241         int ret;
242
243         if (nhdr.n_namesz != sizeof(SDT_NOTE_NAME) ||
244             memcmp(data->d_buf + name_off, SDT_NOTE_NAME,
245                 sizeof(SDT_NOTE_NAME)))
246             continue;
247
248         name = sdt_note_name(elf, &nhdr, sdt_note_data(data, desc_off));
249         if (!name)

```

```

250         continue;
251
252     vaddr = sdt\_note\_addr(elf, sdt\_note\_data(data, desc_off),
253                       nhdr.n_descsz, nhdr.n_type);
254     if (!vaddr)
255         continue;
256
257     ret = actor(name, vaddr, arg);
258     if (ret)
259         return ret;
260
261     ++symbols_count;
262 }
263
264 return symbols_count;
265 }
266
267 int dso\_follow\_debuglink(Elf *elf,
268                          const char *orig_exec,
269                          int type,
270                          symbol\_actor actor,
271                          void *arg)
272 {
273     GElf_Ehdr ehdr;
274     size_t shstrndx, orig_exec_dir_len;
275     GElf_Shdr shdr;
276     Elf_Scn *dbg_link_scn;
277     Elf_Data *dbg_link_scn_data;
278     char *dbg_link, *dbg_bin, *last_slash;
279     int symbols_count;
280
281     /* First try to find the .gnu_debuglink section in the binary. */
282     if (gelf_getehdr(elf, &ehdr) == NULL)
283         return 0;
284     if (elf_getshdrstrndx(elf, &shstrndx) != 0)
285         return 0;
286
287     dbg_link_scn = elf\_section\_by\_name(elf, &ehdr, &shdr, dbg\_link\_name);
288     if (dbg_link_scn == NULL)
289         return 0;
290
291     /* Debug link section found, read of the content (only get the first
292        string, no checksum checking atm). This is debug binary file name. */
293     dbg_link_scn_data = elf_getdata(dbg_link_scn, NULL);
294     if (dbg_link_scn_data == NULL ||
295         dbg_link_scn_data->d_size <= 0 ||
296         dbg_link_scn_data->d_buf == NULL)
297         return 0;
298
299     /* Now compose debug executable name */
300     dbg_link = (char *) (dbg_link_scn_data->d_buf);
301     dbg_bin = malloc(strlen(dbg\_bin\_dir) + 1 +
302                    strlen(orig_exec) + 1 +
303                    strlen(dbg_link) + 1);
304     if (!dbg_bin)
305         return 0;
306
307     orig_exec_dir_len = PATH\_MAX;
308     last_slash = strrchr(orig_exec, '/');
309     if (last_slash != NULL)
310         orig_exec_dir_len = last_slash - orig_exec;
311
312     sprintf(dbg_bin, "%s/%.*s/%s",
313           dbg\_bin\_dir, (int)orig_exec_dir_len, orig_exec, dbg_link);
314
315     /* Retry symbol search with the debug binary */
316     symbols_count = parse\_dso\_symbols(dbg_bin, type, actor, arg);
317
318     free(dbg_bin);
319
320     return symbols_count;
321 }
322
323 int parse\_dso\_symbols(const char *exec, int type, symbol\_actor actor, void *arg)
324 {
325     int symbols_count = 0;

```

```

326 Elf *elf;
327 int fd;
328
329 if (elf_version(EV_CURRENT) == EV_NONE)
330     return -1;
331
332 fd = open(exec, O_RDONLY);
333 if (fd < 0)
334     return -1;
335
336 elf = elf_begin(fd, ELF_C_READ, NULL);
337 if (elf) {
338     switch (type) {
339     case FIND_SYMBOL:
340         symbols_count = dso_symbols(elf, actor, arg);
341         if (symbols_count != 0)
342             break;
343         /* If no symbols found, try in the debuglink binary. */
344         symbols_count = dso_follow_debuglink(elf,
345                                             exec,
346                                             type,
347                                             actor,
348                                             arg);
349         break;
350     case FIND_STAPSDT_NOTE:
351         symbols_count = dso_sdt_notes(elf, actor, arg);
352         break;
353     }
354     elf_end(elf);
355 }
356
357 close(fd);
358 return symbols_count;
359 }
360 }

```

[One Level Up](#)

[Top Level](#)

runtime/lib_timer.c - ktap

Global variables defined

- [timer lib funcs](#)

Data types defined

- [ktap_hrtimer](#)

Functions defined

- [do_tick_profile](#)
- [hrtimer_ktap_fn](#)
- [kp_exit_timers](#)
- [kp_lib_init_timer](#)
- [kplib_timer_profile](#)
- [kplib_timer_tick](#)
- [set_profile_timer](#)
- [set_tick_timer](#)

Source code

```
1 /*
2  * lib_timer.c - timer library support for ktap
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <linux/ctype.h>
23 #include <linux/slab.h>
24 #include <linux/delay.h>
25 #include <linux/sched.h>
26 #include "../include/ktap_types.h"
27 #include "ktap.h"
28 #include "kp_obj.h"
29 #include "kp_vm.h"
30 #include "kp_events.h"
31
32 struct ktap_hrtimer {
33     struct hrtimer timer;
34     ktap_state_t *ks;
35     ktap_func_t *fn;
```

```

36     u64 ns;
37     struct list_head list;
38 };
39
40 /*
41  * Currently ktap disallow tracing event in timer callback closure,
42  * that will corrupt ktap_state t and ktap stack, because timer closure
43  * and event closure use same irq percpu ktap_state t and stack.
44  * We can use a different percpu ktap_state t and stack for timer purpose,
45  * but that's don't bring any big value with cost on memory consuming.
46  *
47  * So just simply disable tracing in timer closure,
48  * get_recursion_context()/put_recursion_context() is used for this purpose.
49  */
50 static enum hrtimer_restart hrtimer_ktap_fn(struct hrtimer *timer)
51 {
52     struct ktap_hrtimer *t;
53     ktap_state t *ks;
54     int rctx;
55
56     rcu_read_lock_sched_notrace();
57
58     t = container_of(timer, struct ktap_hrtimer, timer);
59     rctx = get_recursion_context(t->ks);
60
61     ks = kp_vm_new_thread(t->ks, rctx);
62     set_func(ks->top, t->fn);
63     incr_top(ks);
64     kp_vm_call(ks, ks->top - 1, 0);
65     kp_vm_exit_thread(ks);
66
67     hrtimer_add_expires_ns(timer, t->ns);
68
69     put_recursion_context(ks, rctx);
70     rcu_read_unlock_sched_notrace();
71
72     return HRTIMER_RESTART;
73 }
74
75 static int set_tick_timer(ktap_state t *ks, u64 period, ktap_func t *fn)
76 {
77     struct ktap_hrtimer *t;
78
79     t = kp_malloc(ks, sizeof(*t));
80     if (unlikely(!t))
81         return -ENOMEM;
82     t->ks = ks;
83     t->fn = fn;
84     t->ns = period;
85
86     INIT_LIST_HEAD(&t->list);
87     list_add(&t->list, &(G(ks)->timers));
88
89     hrtimer_init(&t->timer, CLOCK_MONOTONIC, HRTIMER_MODE_REL);
90     t->timer.function = hrtimer_ktap_fn;
91     hrtimer_start(&t->timer, ns_to_ktime(period), HRTIMER_MODE_REL);
92
93     return 0;
94 }
95
96 static int set_profile_timer(ktap_state t *ks, u64 period, ktap_func t *fn)
97 {
98     struct perf_event_attr attr;
99
100     memset(&attr, 0, sizeof(attr));
101     attr.type = PERF_TYPE_SOFTWARE;
102     attr.config = PERF_COUNT_SW_CPU_CLOCK;
103     attr.sample_type = PERF_SAMPLE_RAW | PERF_SAMPLE_TIME |
104         PERF_SAMPLE_CPU | PERF_SAMPLE_PERIOD;
105     attr.sample_period = period;
106     attr.size = sizeof(attr);
107     attr.disabled = 0;
108
109     return kp_event_create(ks, &attr, NULL, NULL, fn);
110 }
111

```

```

112 static int do_tick_profile(ktap_state_t *ks, int is_tick)
113 {
114     const char *str = kp_arg_checkstring(ks, 1);
115     ktap_func_t *fn = kp_arg_checkfunction(ks, 2);
116     const char *tmp;
117     char interval_str[32] = {0};
118     char suffix[10] = {0};
119     int i = 0, ret, n;
120     int factor;
121
122     tmp = str;
123     while (isdigit(*tmp))
124         tmp++;
125
126     strncpy(interval_str, str, tmp - str);
127     if (kstrtoint(interval_str, 10, &n))
128         goto error;
129
130     strncpy(suffix, tmp, 9);
131     while (suffix[i] != ' ' && suffix[i] != '\0')
132         i++;
133
134     suffix[i] = '\0';
135
136     if (!strcmp(suffix, "s") || !strcmp(suffix, "sec"))
137         factor = NSEC_PER_SEC;
138     else if (!strcmp(suffix, "ms") || !strcmp(suffix, "msec"))
139         factor = NSEC_PER_MSEC;
140     else if (!strcmp(suffix, "us") || !strcmp(suffix, "usec"))
141         factor = NSEC_PER_USEC;
142     else
143         goto error;
144
145     if (is_tick)
146         ret = set_tick_timer(ks, (u64)factor * n, fn);
147     else
148         ret = set_profile_timer(ks, (u64)factor * n, fn);
149
150     return ret;
151
152 error:
153     kp_error(ks, "cannot parse timer interval: %s\n", str);
154     return -1;
155 }
156
157 /*
158  * tick-n probes fire on only one CPU per interval.
159  * valid time suffixes: sec/s, msec/ms, usec/us
160  */
161 static int kplib_timer_tick(ktap_state_t *ks)
162 {
163     /* timer.tick cannot be called in trace_end state */
164     if (G(ks)->state != KTAP_RUNNING) {
165         kp_error(ks,
166             "timer.tick only can be called in RUNNING state\n");
167         return -1;
168     }
169
170     return do_tick_profile(ks, 1);
171 }
172
173 /*
174  * A profile-n probe fires every fixed interval on every CPU
175  * valid time suffixes: sec/s, msec/ms, usec/us
176  */
177 static int kplib_timer_profile(ktap_state_t *ks)
178 {
179     /* timer.profile cannot be called in trace_end state */
180     if (G(ks)->state != KTAP_RUNNING) {
181         kp_error(ks,
182             "timer.profile only can be called in RUNNING state\n");
183         return -1;
184     }
185
186     return do_tick_profile(ks, 0);
187 }

```

```
188
189 void kp_exit_timers(ktap\_state\_t *ks)
190 {
191     struct ktap\_hrtimer *t, *tmp;
192     struct list_head *timers_list = &(G(ks)->timers);
193
194     list_for_each_entry_safe(t, tmp, timers_list, list) {
195         hrtimer_cancel(&t->timer);
196         kp\_free(ks, t);
197     }
198 }
199
200 static const ktap\_libfunc\_t timer_lib_funcs[] = {
201     {"profile", kplib\_timer\_profile},
202     {"tick", kplib\_timer\_tick},
203     {NULL}
204 };
205
206 int kp_lib_init_timer(ktap\_state\_t *ks)
207 {
208     return kp\_vm\_register\_lib(ks, "timer", timer\_lib\_funcs);
209 }
210
```

[One Level Up](#)

[Top Level](#)

runtime/ffi/ffi_symbol.c - ktap

Functions defined

- [add ffi func to ctable](#)
- [ffi_free_symbols](#)
- [ffi_get_csym_by_id](#)
- [ffi_get_csym_id](#)
- [ffi_set_csym_arr](#)
- [get_csym_arr](#)
- [get_csym_nr](#)
- [get_ffi_ctable](#)
- [set_csym_arr](#)
- [set_csym_nr](#)
- [setup_ffi_ctable](#)
- [setup_ffi_symbol_table](#)

Source code

```
1 /*
2  * ffi_symbol.c - ktapvm kernel module ffi symbol submodule
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22
23 #include "../include/ktap_types.h"
24 #include "../include/ktap_ffi.h"
25 #include "../ktap.h"
26 #include "../kp_vm.h"
27 #include "../kp_obj.h"
28 #include "../kp_str.h"
29 #include "../kp_tab.h"
30
31 static inline csymbol *get_csym_arr(ktap\_state\_t *ks)
32 {
33     return g(ks)->ffis.csym_arr;
34 }
35
36 static inline int get_csym_nr(ktap\_state\_t *ks)
37 {
38     return g(ks)->ffis.csym_nr;
```

```

39 }
40
41 static inline void set_csym_arr(ktap_state_t *ks, csymbol *csym)
42 {
43     G(ks)->ffis.csym_arr = csym;
44 }
45
46 static inline void set_csym_nr(ktap_state_t *ks, int nr)
47 {
48     G(ks)->ffis.csym_nr = nr;
49 }
50
51
52 static inline ktap_tab_t *get_ffi_ctable(ktap_state_t *ks)
53 {
54     return G(ks)->ffis.ctable;
55 }
56
57 static int setup_ffi_ctable(ktap_state_t *ks)
58 {
59     ktap_val_t ffi_name, ffi_lib_name, ffi_mt;
60     const ktap_val_t *gt, *ffit;
61     ktap_tab_t *registry;
62
63     gt = kp_tab_getint(hvalue(&G(ks)->registry), KTAP_RIDX_GLOBALS);
64
65     G(ks)->ffis.ctable = kp_tab_new(ks, 0, 512);
66     if (!G(ks)->ffis.ctable)
67         return -1;
68
69     /* get global["ffi"] */
70     set_string(&ffi_name, kp_str_new(ks, "ffi"));
71     registry = hvalue(gt);
72     ffit = kp_tab_get(ks, registry, &ffi_name);
73     /* insert ffi C table to ffi table */
74     set_table(&ffi_mt, get_ffi_ctable(ks));
75     set_string(&ffi_lib_name, kp_str_new(ks, "C"));
76     registry = hvalue(ffit);
77     kp_tab_setvalue(ks, registry, &ffi_lib_name, &ffi_mt);
78
79     return 0;
80 }
81
82 inline csymbol *ffi_get_csym_by_id(ktap_state_t *ks, int id)
83 {
84     return &(get_csym_arr(ks)[id]);
85 }
86
87 csymbol_id ffi_get_csym_id(ktap_state_t *ks, char *name)
88 {
89     int i;
90
91     for (i = 0; i < get_csym_nr(ks); i++) {
92         if (!strcmp(name, csym_name(ffi_get_csym_by_id(ks, i)))) {
93             return i;
94         }
95     }
96
97     kp_error(ks, "Cannot find csymbol with name %s\n", name);
98     return 0;
99 }
100
101 static void add_ffi_func_to_ctable(ktap_state_t *ks, csymbol_id id)
102 {
103     ktap_val_t func_name, fv;
104     ktap_cdata_t *cd;
105     csymbol *cs;
106
107     /* push cdata to ctable */
108     set_cdata(&fv, kp_obj_newobject(ks, KTAP_TYPE_CDATA, sizeof(ktap_cdata_t),
109                                     NULL));
110     cd = cdvalue(&fv);
111     cd_set_csym_id(cd, id);
112
113     cs = id_to_csym(ks, id);
114     set_string(&func_name, kp_str_new(ks, csym_name(cs)));

```

```

115     kp\_tab\_setvalue(ks, get\_ffi\_ctable(ks), &func_name, &fv);
116 }
117
118 static int setup\_ffi\_symbol\_table(ktap\_state\_t *ks)
119 {
120     int i;
121     csymbol *cs;
122
123     if (setup\_ffi\_ctable(ks))
124         return -1;
125
126     /* push all functions to ctable */
127     for (i = 0; i < get\_csym\_nr(ks); i++) {
128         cs = &get\_csym\_arr(ks)[i];
129         switch (cs->type) {
130             case FFI_FUNC:
131                 kp\_verbose\_printf(ks, "[%d] loading C function %s\n",
132                     i, csym\_name(cs));
133                 add\_ffi\_func\_to\_ctable(ks, i);
134                 kp\_verbose\_printf(ks, "%s loaded\n", csym\_name(cs));
135                 break;
136             case FFI_STRUCT:
137             case FFI_UNION:
138                 break;
139             default:
140                 break;
141         }
142     }
143
144     return 0;
145 }
146
147 void ffi\_free\_symbols(ktap\_state\_t *ks)
148 {
149     int i;
150     csymbol\_id *arg_ids;
151     csymbol *cs;
152
153     if (!get\_csym\_arr(ks))
154         return;
155
156     for (i = 0; i < get\_csym\_nr(ks); i++) {
157         cs = &get\_csym\_arr(ks)[i];
158         switch (csym\_type(cs)) {
159             case FFI_FUNC:
160                 arg_ids = csym\_func\_arg\_ids(cs);
161                 if (arg_ids)
162                     kp\_free(ks, arg_ids);
163                 break;
164             case FFI_STRUCT:
165             case FFI_UNION:
166                 /*@TODO finish this 20.11 2013 (houqp)*/
167                 break;
168             default:
169                 break;
170         }
171     }
172
173     kp\_free(ks, get\_csym\_arr(ks));
174 }
175
176 int ffi\_set\_csym\_arr(ktap\_state\_t *ks, int cs\_nr, csymbol *new_arr)
177 {
178     set\_csym\_nr(ks, cs\_nr);
179     set\_csym\_arr(ks, new_arr);
180     return setup\_ffi\_symbol\_table(ks);
181 }
182

```

[One Level Up](#)

[Top Level](#)

runtime/kp_mempool.c - ktap

Functions defined

- [kp_mempool_alloc](#)
- [kp_mempool_destroy](#)
- [kp_mempool_init](#)

Source code

```
1  /*
2  * kp_mempool.c - ktap memory pool, service for string allocation
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include "../include/ktap_types.h"
23 #include "kp_obj.h"
24 #include "kp_str.h"
25
26 #include <linux/ctype.h>
27 #include <linux/module.h>
28 #include "ktap.h"
29
30
31 /*
32 * allocate memory from mempool, the allocated memory will be free
33 * util ktap exit.
34 * TODO: lock-free allocation
35 */
36 void *kp_mempool_alloc(ktap_state_t *ks, int size)
37 {
38     ktap_global_state_t *g = G(ks);
39     void *mempool = g->mempool;
40     void *freepos = g->mp_freepos;
41     void *addr;
42     unsigned long flags;
43
44     local_irq_save(flags);
45     arch_spin_lock(&g->mp_lock);
46
47     if (unlikely((unsigned long)((char *)freepos + size) >
48                 (unsigned long)((char *)mempool + g->mp_size)) {
49         addr = NULL;
50         goto out;
51     }
52
53     addr = freepos;
54     g->mp_freepos = (char *)freepos + size;
55 out:
56
57     arch_spin_unlock(&g->mp_lock);
58     local_irq_restore(flags);
```

```

59     return addr;
60 }
61
62 /*
63  * destroy mempool.
64  */
65 void kp_mempool_destroy(ktap\_state\_t *ks)
66 {
67     ktap\_global\_state\_t *g = G(ks);
68
69     if (!g->mempool)
70         return;
71
72     vfree(g->mempool);
73     g->mempool = NULL;
74     g->mp_freepos = NULL;
75     g->mp_size = 0;
76 }
77
78 /*
79  * pre-allocate size Kbytes memory pool.
80  */
81 int kp_mempool_init(ktap\_state\_t *ks, int size)
82 {
83     ktap\_global\_state\_t *g = G(ks);
84
85     g->mempool = vmalloc(size * 1024);
86     if (!g->mempool)
87         return -ENOMEM;
88
89     g->mp_freepos = g->mempool;
90     g->mp_size = size * 1024;
91     g->mp_lock = (arch_spinlock_t) __ARCH_SPIN_LOCK_UNLOCKED;
92     return 0;
93 }
94

```

[One Level Up](#)

[Top Level](#)

runtime/lib_base.c - ktap

Global variables defined

- [base lib funcs](#)

Functions defined

- [kp_lib_init_base](#)
- [kplib_arch](#)
- [kplib_curr_taskinfo](#)
- [kplib_delete](#)
- [kplib_exit](#)
- [kplib_gettimeofday_ms](#)
- [kplib_gettimeofday_ns](#)
- [kplib_gettimeofday_s](#)
- [kplib_gettimeofday_us](#)
- [kplib_in_interrupt](#)
- [kplib_in_iowait](#)
- [kplib_ipof](#)
- [kplib_kernel_string](#)
- [kplib_kernel_v](#)
- [kplib_len](#)
- [kplib_num_cpus](#)
- [kplib_pairs](#)
- [kplib_print](#)
- [kplib_print_hist](#)
- [kplib_print_trace_clock](#)
- [kplib_printf](#)
- [kplib_stack](#)
- [kplib_stack](#)
- [kplib_stringof](#)
- [kplib_user_string](#)

Macros defined

- [HISTOGRAM_DEFAULT_TOP_NUM](#)

- [RET VALUE](#)
- [RET VALUE](#)

Source code

```

1  /*
2  * lib_base.c - base library
3  *
4  * Caveat: all kernel funtion called by ktap library have to be lock free,
5  * otherwise system will deadlock.
6  *
7  * This file is part of ktap by Jovi Zhangwei.
8  *
9  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
10 *
11 * ktap is free software; you can redistribute it and/or modify it
12 * under the terms and conditions of the GNU General Public License,
13 * version 2, as published by the Free Software Foundation.
14 *
15 * ktap is distributed in the hope it will be useful, but WITHOUT
16 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
17 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
18 * more details.
19 *
20 * You should have received a copy of the GNU General Public License along with
21 * this program; if not, write to the Free Software Foundation, Inc.,
22 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
23 */
24
25 #include <linux/version.h>
26 #include <linux/hardirq.h>
27 #include <linux/module.h>
28 #include <linux/kallsyms.h>
29 #include <linux/sched.h>
30 #include <linux/uaccess.h>
31 #include <linux/utsname.h>
32 #include <linux/time.h>
33 #include <linux/clocksource.h>
34 #include <linux/ring_buffer.h>
35 #include <linux/stacktrace.h>
36 #include <linux/cred.h>
37 #if LINUX_VERSION_CODE >= KERNEL_VERSION(3, 5, 0)
38 #include <linux/uidgid.h>
39 #endif
40 #include "../include/ktap_types.h"
41 #include "ktap.h"
42 #include "kp_obj.h"
43 #include "kp_str.h"
44 #include "kp_tab.h"
45 #include "kp_transport.h"
46 #include "kp_events.h"
47 #include "kp_vm.h"
48
49 static int kplib_print(ktap_state_t *ks)
50 {
51     int i;
52     int n = kp_arg_nr(ks);
53
54     for (i = 1; i <= n; i++) {
55         ktap_val_t *arg = kp_arg(ks, i);
56         if (i > 1)
57             kp_puts(ks, "\t");
58         kp_obj_show(ks, arg);
59     }
60
61     kp_puts(ks, "\n");
62     return 0;
63 }
64
65 /* don't engage with intern string in printf, use buffer directly */
66 static int kplib_printf(ktap_state_t *ks)
67 {
68     struct trace_seq *seq;

```

```

69     preempt_disable_notraced();
70
71
72     seq = kp_this_cpu_print_buffer(ks);
73     trace_seq_init(seq);
74
75     if (kp_str_fmt(ks, seq))
76         goto out;
77
78     seq->buffer[seq->len] = '\0';
79     kp_transport_write(ks, seq->buffer, seq->len + 1);
80
81 out:
82     preempt_enable_notraced();
83     return 0;
84 }
85
86 #define HISTOGRAM_DEFAULT_TOP_NUM    20
87
88 static int kplib_print_hist(ktap_state_t *ks)
89 {
90     int n ;
91
92     kp_arg_check(ks, 1, KTAP_TTAB);
93     n = kp_arg_checkoptnumber(ks, 2, HISTOGRAM_DEFAULT_TOP_NUM);
94
95     n = min(n, 1000);
96     n = max(n, HISTOGRAM_DEFAULT_TOP_NUM);
97
98     kp_tab_print_hist(ks, hvalue(kp_arg(ks, 1)), n);
99
100    return 0;
101 }
102
103 static int kplib_pairs(ktap_state_t *ks)
104 {
105     kp_arg_check(ks, 1, KTAP_TTAB);
106
107     set_cfunc(ks->top++, (ktap_cfunction)kp_tab_next);
108     set_table(ks->top++, hvalue(kp_arg(ks, 1)));
109     set_nil(ks->top++);
110     return 3;
111 }
112
113 static int kplib_len(ktap_state_t *ks)
114 {
115     int len = kp_obj_len(ks, kp_arg(ks, 1));
116
117     if (len < 0)
118         return -1;
119
120     set_number(ks->top, len);
121     incr_top(ks);
122     return 1;
123 }
124
125 static int kplib_delete(ktap_state_t *ks)
126 {
127     kp_arg_check(ks, 1, KTAP_TTAB);
128     kp_tab_clear(hvalue(kp_arg(ks, 1)));
129     return 0;
130 }
131
132 #ifdef CONFIG_STACKTRACE
133 static int kplib_stack(ktap_state_t *ks)
134 {
135     uint16_t skip, depth = 10;
136
137     depth = kp_arg_checkoptnumber(ks, 1, 10); /* default as 10 */
138     depth = min_t(uint16_t, depth, KP_MAX_STACK_DEPTH);
139     skip = kp_arg_checkoptnumber(ks, 2, 10); /* default as 10 */
140
141     set_kstack(ks->top, depth, skip);
142     incr_top(ks);
143     return 1;
144 }

```

```

145 #else
146 static int kplib_stack(ktap_state_t *ks)
147 {
148     kp_error(ks, "Please enable CONFIG_STACKTRACE before call stack()\n");
149     return -1;
150 }
151 #endif
152
153
154 extern unsigned long long ns2usecs(cycle_t nsec);
155 static int kplib_print_trace_clock(ktap_state_t *ks)
156 {
157     unsigned long long t;
158     unsigned long secs, usec_rem;
159     u64 timestamp;
160
161     /* use ring buffer's timestamp */
162     timestamp = ring_buffer_time_stamp(G(ks)->buffer, smp_processor_id());
163
164     t = ns2usecs(timestamp);
165     usec_rem = do_div(t, USEC_PER_SEC);
166     secs = (unsigned long)t;
167
168     kp_printf(ks, "%5lu.%06lu\n", secs, usec_rem);
169     return 0;
170 }
171
172 static int kplib_num_cpus(ktap_state_t *ks)
173 {
174     set_number(ks->top, num_online_cpus());
175     incr_top(ks);
176     return 1;
177 }
178
179 /* TODO: intern string firstly */
180 static int kplib_arch(ktap_state_t *ks)
181 {
182     ktap_str_t *ts = kp_str_newz(ks, utsname()->machine);
183     if (unlikely(!ts))
184         return -1;
185
186     set_string(ks->top, ts);
187     incr_top(ks);
188     return 1;
189 }
190
191 /* TODO: intern string firstly */
192 static int kplib_kernel_v(ktap_state_t *ks)
193 {
194     ktap_str_t *ts = kp_str_newz(ks, utsname()->release);
195     if (unlikely(!ts))
196         return -1;
197
198     set_string(ks->top, ts);
199     incr_top(ks);
200     return 1;
201 }
202
203 static int kplib_kernel_string(ktap_state_t *ks)
204 {
205     unsigned long addr = kp_arg_checknumber(ks, 1);
206     char str[256] = {0};
207     ktap_str_t *ts;
208     char *ret;
209
210     ret = strncpy((void *)str, (const void *)addr, 256);
211     (void) &ret; /* Silence compiler warning. */
212     str[255] = '\0';
213
214     ts = kp_str_newz(ks, str);
215     if (unlikely(!ts))
216         return -1;
217
218     set_string(ks->top, ts);
219     incr_top(ks);
220     return 1;

```

```

221 }
222
223 static int kplib_user_string(ktap_state_t *ks)
224 {
225     unsigned long addr = kp_arg_checknumber(ks, 1);
226     char str[256] = {0};
227     ktap_str_t *ts;
228     int ret;
229
230     pagefault_disable();
231     ret = __copy_from_user_inatomic((void *)str, (const void *)addr, 256);
232     (void) &ret; /* Silence compiler warning. */
233     pagefault_enable();
234     str[255] = '\0';
235
236     ts = kp_str_newz(ks, str);
237     if (unlikely(!ts))
238         return -1;
239
240     set_string(ks->top, ts);
241     incr_top(ks);
242     return 1;
243 }
244
245 static int kplib_stringof(ktap_state_t *ks)
246 {
247     ktap_val_t *v = kp_arg(ks, 1);
248     const ktap_str_t *ts = NULL;
249
250     if (itype(v) == KTAP_TEVENTSTR) {
251         ts = kp_event_stringify(ks);
252     } else if (itype(v) == KTAP_TKIP) {
253         char str[KSYM_SYMBOL_LEN];
254
255         SPRINT_SYMBOL(str, nvalue(v));
256         ts = kp_str_newz(ks, str);
257     }
258
259     if (unlikely(!ts))
260         return -1;
261
262     set_string(ks->top++, ts);
263     return 1;
264 }
265
266 static int kplib_ipof(ktap_state_t *ks)
267 {
268     unsigned long addr = kp_arg_checknumber(ks, 1);
269
270     set_ip(ks->top++, addr);
271     return 1;
272 }
273
274 static int kplib_gettimeofday_ns(ktap_state_t *ks)
275 {
276     set_number(ks->top, gettimeofday_ns());
277     incr_top(ks);
278
279     return 1;
280 }
281
282 static int kplib_gettimeofday_us(ktap_state_t *ks)
283 {
284     set_number(ks->top, gettimeofday_ns() / NSEC_PER_USEC);
285     incr_top(ks);
286
287     return 1;
288 }
289
290 static int kplib_gettimeofday_ms(ktap_state_t *ks)
291 {
292     set_number(ks->top, gettimeofday_ns() / NSEC_PER_MSEC);
293     incr_top(ks);
294
295     return 1;
296 }

```

```

297
298 static int kplib_gettimeofday_s(ktap_state_t *ks)
299 {
300     set_number(ks->top, gettimeofday_ns() / NSEC_PER_SEC);
301     incr_top(ks);
302
303     return 1;
304 }
305
306 /*
307  * use gdb to get field offset of struct task_struct, for example:
308  *
309  * gdb vmlinux
310  * (gdb)p &(((struct task_struct *)0).prio)
311  */
312 static int kplib_curr_taskinfo(ktap_state_t *ks)
313 {
314     int offset = kp_arg_checknumber(ks, 1);
315     int fetch_bytes = kp_arg_checkoptnumber(ks, 2, 4); /* fetch 4 bytes */
316
317     if (offset >= sizeof(struct task_struct)) {
318         set_nil(ks->top++);
319         kp_error(ks, "access out of bound value of task_struct\n");
320         return 1;
321     }
322
323     #define RET_VALUE ((unsigned long)current + offset)
324
325     switch (fetch_bytes) {
326     case 4:
327         set_number(ks->top, *(unsigned int *)RET_VALUE);
328         break;
329     case 8:
330         set_number(ks->top, *(unsigned long *)RET_VALUE);
331         break;
332     default:
333         kp_error(ks, "unsupported fetch bytes in curr_task_info\n");
334         set_nil(ks->top);
335         break;
336     }
337
338     #undef RET_VALUE
339
340     incr_top(ks);
341     return 1;
342 }
343
344 /*
345  * This built-in function mainly purpose scripts/schedule/schedtimes.kp
346  */
347 static int kplib_in_iowait(ktap_state_t *ks)
348 {
349     set_number(ks->top, current->in_iowait);
350     incr_top(ks);
351
352     return 1;
353 }
354
355 static int kplib_in_interrupt(ktap_state_t *ks)
356 {
357     int ret = in_interrupt();
358
359     set_number(ks->top, ret);
360     incr_top(ks);
361     return 1;
362 }
363
364 static int kplib_exit(ktap_state_t *ks)
365 {
366     kp_vm_try_to_exit(ks);
367
368     /* do not execute bytecode any more in this thread */
369     return -1;
370 }
371
372 static const ktap_libfunc_t base_lib_funcs[] = {

```

```

373 {"print", kplib\_print},
374 {"printf", kplib\_printf},
375 {"print_hist", kplib\_print\_hist},
376
377 {"pairs", kplib\_pairs},
378 {"len", kplib\_len},
379 {"delete", kplib\_delete},
380
381 {"stack", kplib\_stack},
382 {"print_trace_clock", kplib\_print\_trace\_clock},
383
384 {"num_cpus", kplib\_num\_cpus},
385 {"arch", kplib\_arch},
386 {"kernel_v", kplib\_kernel\_v},
387 {"kernel_string", kplib\_kernel\_string},
388 {"user_string", kplib\_user\_string},
389 {"stringof", kplib\_stringof},
390 {"ipof", kplib\_ipof},
391
392 {"gettimeofday_ns", kplib\_gettimeofday\_ns},
393 {"gettimeofday_us", kplib\_gettimeofday\_us},
394 {"gettimeofday_ms", kplib\_gettimeofday\_ms},
395 {"gettimeofday_s", kplib\_gettimeofday\_s},
396
397 {"curr_taskinfo", kplib\_curr\_taskinfo},
398
399 {"in_iowait", kplib\_in\_iowait},
400 {"in_interrupt", kplib\_in\_interrupt},
401
402 {"exit", kplib\_exit},
403 {NULL}
404 };
405
406 int kp\_lib\_init\_base(ktap\_state\_t *ks)
407 {
408     return kp\_vm\_register\_lib(ks, NULL, base\_lib\_funcs);
409 }

```

[One Level Up](#)

[Top Level](#)

runtime/lib_kdebug.c - ktap

Global variables defined

- [kdebug lib funcs](#)

Functions defined

- [kp lib init kdebug](#)
- [kplib kdebug kprobe](#)
- [kplib kdebug trace by id](#)
- [kplib kdebug trace end](#)

Source code

```
1  /*
2  * lib_kdebug.c - kdebug library support for ktap
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <linux/module.h>
23 #include <linux/ctype.h>
24 #include <linux/slab.h>
25 #include <linux/version.h>
26 #include <linux/ftrace_event.h>
27 #include "../include/ktap_types.h"
28 #include "ktap.h"
29 #include "kp_obj.h"
30 #include "kp_str.h"
31 #include "kp_transport.h"
32 #include "kp_vm.h"
33 #include "kp_events.h"
34
35 /**
36 * function kdebug.trace_by_id
37 *
38 * @uaddr: userspace address refer to ktap\_eventdesc t
39 * @closure
40 */
41 static int kplib_kdebug_trace_by_id(ktap_state_t *ks)
42 {
43     unsigned long uaddr = kp_arg_checknumber(ks, 1);
44     ktap_func_t *fn = kp_arg_checkfunction(ks, 2);
45     struct task_struct *task = G(ks)->trace_task;
46     ktap_eventdesc_t eventsdesc;
47     char *filter = NULL;
48     int *id_arr;
49     int ret, i;
50 }
```

```

51 if (G(ks)->mainthread != ks) {
52     kp_error(ks,
53         "kdebug.trace_by_id only can be called in mainthread\n");
54     return -1;
55 }
56
57 /* kdebug.trace_by_id cannot be called in trace_end state */
58 if (G(ks)->state != KTAP_RUNNING) {
59     kp_error(ks,
60         "kdebug.trace_by_id only can be called in RUNNING state\n");
61     return -1;
62 }
63
64 /* copy ktap_eventdesc_t from userspace */
65 ret = copy_from_user(&eventsdesc, (void *)uaddr,
66     sizeof(ktap_eventdesc_t));
67 if (ret < 0)
68     return -1;
69
70 if (eventsdesc.filter) {
71     int len;
72
73     len = strlen_user(eventsdesc.filter);
74     if (len > 0x1000)
75         return -1;
76
77     filter = kmalloc(len + 1, GFP_KERNEL);
78     if (!filter)
79         return -1;
80
81     /* copy filter string from userspace */
82     if (strncpy_from_user(filter, eventsdesc.filter, len) < 0) {
83         kfree(filter);
84         return -1;
85     }
86 }
87
88 id_arr = kmalloc(eventsdesc.nr * sizeof(int), GFP_KERNEL);
89 if (!id_arr) {
90     kfree(filter);
91     return -1;
92 }
93
94 /* copy all event id from userspace */
95 ret = copy_from_user(id_arr, eventsdesc.id_arr,
96     eventsdesc.nr * sizeof(int));
97 if (ret < 0) {
98     kfree(filter);
99     kfree(id_arr);
100    return -1;
101 }
102
103 fn = clvalue(kp_arg(ks, 2));
104
105 for (i = 0; i < eventsdesc.nr; i++) {
106     struct perf_event_attr attr;
107
108     cond_resched();
109
110     if (signal_pending(current)) {
111         flush_signals(current);
112         kfree(filter);
113         kfree(id_arr);
114         return -1;
115     }
116
117     memset(&attr, 0, sizeof(attr));
118     attr.type = PERF_TYPE_TRACEPOINT;
119     attr.config = id_arr[i];
120     attr.sample_type = PERF_SAMPLE_RAW | PERF_SAMPLE_TIME |
121         PERF_SAMPLE_CPU | PERF_SAMPLE_PERIOD;
122     attr.sample_period = 1;
123     attr.size = sizeof(attr);
124     attr.disabled = 0;
125
126     /* register event one by one */

```

```

127     ret = kp\_event\_create(ks, &attr, task, filter, fn);
128     if (ret < 0)
129         break;
130 }
131
132 kfree(filter);
133 kfree(id_arr);
134 return 0;
135 }
136
137 static int kplib\_kdebug\_trace\_end(ktap\_state\_t *ks)
138 {
139     /* trace\_end\_closure will be called when ktap main thread exit */
140     G(ks)->trace\_end\_closure = kp\_arg\_checkfunction(ks, 1);
141     return 0;
142 }
143
144 #if 0
145 static int kplib\_kdebug\_tracepoint(ktap\_state\_t *ks)
146 {
147     const char *event_name = kp\_arg\_checkstring(ks, 1);
148     ktap\_func\_t *fn = kp\_arg\_checkfunction(ks, 2);
149
150     if (G(ks)->mainthread != ks) {
151         kp\_error(ks,
152             "kdebug.tracepoint only can be called in mainthread\n");
153         return -1;
154     }
155
156     /* kdebug.tracepoint cannot be called in trace_end state */
157     if (G(ks)->state != KTAP\_RUNNING) {
158         kp\_error(ks,
159             "kdebug.tracepoint only can be called in RUNNING state\n");
160         return -1;
161     }
162
163     return kp\_event\_create\_tracepoint(ks, event_name, fn);
164 }
165 #endif
166
167 static int kplib\_kdebug\_kprobe(ktap\_state\_t *ks)
168 {
169     const char *event_name = kp\_arg\_checkstring(ks, 1);
170     ktap\_func\_t *fn = kp\_arg\_checkfunction(ks, 2);
171
172     if (G(ks)->mainthread != ks) {
173         kp\_error(ks,
174             "kdebug.kprobe only can be called in mainthread\n");
175         return -1;
176     }
177
178     /* kdebug.kprobe cannot be called in trace_end state */
179     if (G(ks)->state != KTAP\_RUNNING) {
180         kp\_error(ks,
181             "kdebug.kprobe only can be called in RUNNING state\n");
182         return -1;
183     }
184
185     return kp\_event\_create\_kprobe(ks, event_name, fn);
186 }
187 static const ktap\_libfunc\_t kdebug\_lib\_funcs[] = {
188     {"trace_by_id", kplib\_kdebug\_trace\_by\_id},
189     {"trace_end", kplib\_kdebug\_trace\_end},
190
191     #if 0
192     {"tracepoint", kplib\_kdebug\_tracepoint},
193     #endif
194     {"kprobe", kplib\_kdebug\_kprobe},
195     {NULL}
196 };
197
198 int kp\_lib\_init\_kdebug(ktap\_state\_t *ks)
199 {
200     return kp\_vm\_register\_lib(ks, "kdebug", kdebug\_lib\_funcs);
201 }
202

```

[One Level Up](#)

[Top Level](#)

runtime/lib_ansi.c - ktap

Global variables defined

- [ansi lib funcs](#)

Functions defined

- [kp lib init ansi](#)
- [kplib ansi clear screen](#)
- [kplib ansi new line](#)
- [kplib ansi reset color](#)
- [kplib ansi set color](#)
- [kplib ansi set color2](#)
- [kplib ansi set color3](#)

Source code

```
1 /*
2  * lib_ansi.c - ANSI escape sequences library
3  *
4  * http://en.wikipedia.org/wiki/ANSI\_escape\_code
5  *
6  * This file is part of ktap by Jovi Zhangwei.
7  *
8  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
9  *
10 * ktap is free software; you can redistribute it and/or modify it
11 * under the terms and conditions of the GNU General Public License,
12 * version 2, as published by the Free Software Foundation.
13 *
14 * ktap is distributed in the hope it will be useful, but WITHOUT
15 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
16 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
17 * more details.
18 *
19 * You should have received a copy of the GNU General Public License along with
20 * this program; if not, write to the Free Software Foundation, Inc.,
21 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
22 */
23
24 #include "../include/ktap_types.h"
25 #include "ktap.h"
26 #include "kp_vm.h"
27
28 /**
29  * function ansi.clear_screen - Move cursor to top left and clear screen.
30  *
31  * Description: Sends ansi code for moving cursor to top left and then the
32  * ansi code for clearing the screen from the cursor position to the end.
33  */
34
35 static int kplib_ansi_clear_screen(ktap\_state\_t *ks)
36 {
37     kp\_printf(ks, "\033[1;1H\033[J");
38     return 0;
39 }
40
41 /**
42  * function ansi.set_color - Set the ansi Select Graphic Rendition mode.
43  * @fg: Foreground color to set.
```

```

44 *
45 * Description: Sends ansi code for Select Graphic Rendition mode for the
46 * given foreground color. Black (30), Blue (34), Green (32), Cyan (36),
47 * Red (31), Purple (35), Brown (33), Light Gray (37).
48 */
49
50 static int kplib_ansi_set_color(ktap_state_t *ks)
51 {
52     int fg = kp_arg_checknumber(ks, 1);
53
54     kp_printf(ks, "\033[%dm", fg);
55     return 0;
56 }
57
58 /**
59 * function ansi.set_color2 - Set the ansi Select Graphic Rendition mode.
60 * @fg: Foreground color to set.
61 * @bg: Background color to set.
62 *
63 * Description: Sends ansi code for Select Graphic Rendition mode for the
64 * given foreground color, Black (30), Blue (34), Green (32), Cyan (36),
65 * Red (31), Purple (35), Brown (33), Light Gray (37) and the given
66 * background color, Black (40), Red (41), Green (42), Yellow (43),
67 * Blue (44), Magenta (45), Cyan (46), White (47).
68 */
69 static int kplib_ansi_set_color2(ktap_state_t *ks)
70 {
71     int fg = kp_arg_checknumber(ks, 1);
72     int bg = kp_arg_checknumber(ks, 2);
73
74     kp_printf(ks, "\033[%d;%dm", fg, bg);
75     return 0;
76 }
77
78 /**
79 * function ansi.set_color3 - Set the ansi Select Graphic Rendition mode.
80 * @fg: Foreground color to set.
81 * @bg: Background color to set.
82 * @attr: Color attribute to set.
83 *
84 * Description: Sends ansi code for Select Graphic Rendition mode for the
85 * given foreground color, Black (30), Blue (34), Green (32), Cyan (36),
86 * Red (31), Purple (35), Brown (33), Light Gray (37), the given
87 * background color, Black (40), Red (41), Green (42), Yellow (43),
88 * Blue (44), Magenta (45), Cyan (46), White (47) and the color attribute
89 * All attributes off (0), Intensity Bold (1), Underline Single (4),
90 * Blink Slow (5), Blink Rapid (6), Image Negative (7).
91 */
92 static int kplib_ansi_set_color3(ktap_state_t *ks)
93 {
94     int fg = kp_arg_checknumber(ks, 1);
95     int bg = kp_arg_checknumber(ks, 2);
96     int attr = kp_arg_checknumber(ks, 3);
97
98     if (attr)
99         kp_printf(ks, "\033[%d;%d;%dm", fg, bg, attr);
100     else
101         kp_printf(ks, "\033[%d;%dm", fg, bg);
102
103     return 0;
104 }
105
106 /**
107 * function ansi.reset_color - Resets Select Graphic Rendition mode.
108 *
109 * Description: Sends ansi code to reset foreground, background and color
110 * attribute to default values.
111 */
112 static int kplib_ansi_reset_color(ktap_state_t *ks)
113 {
114     kp_printf(ks, "\033[0;0m");
115     return 0;
116 }
117
118 /**
119 * function ansi.new_line - Move cursor to new line.

```

```
120 *
121 * Description: Sends ansi code new line.
122 */
123 static int kplib_ansi_new_line (ktap_state_t *ks)
124 {
125     kp_printf(ks, "\12");
126     return 0;
127 }
128
129 static const ktap_libfunc_t ansi_lib_funcs[] = {
130     {"clear_screen", kplib_ansi_clear_screen},
131     {"set_color", kplib_ansi_set_color},
132     {"set_color2", kplib_ansi_set_color2},
133     {"set_color3", kplib_ansi_set_color3},
134     {"reset_color", kplib_ansi_reset_color},
135     {"new_line", kplib_ansi_new_line},
136     {NULL}
137 };
138
139 int kp_lib_init_ansi(ktap_state_t *ks)
140 {
141     return kp_vm_register_lib(ks, "ansi", ansi_lib_funcs);
142 }
```

[One Level Up](#)

[Top Level](#)

runtime/lib_ffi.c - ktap

Global variables defined

- [ffi lib funcs](#)

Functions defined

- [kp lib init ffi](#)
- [kplib ffi cast](#)
- [kplib ffi new](#)
- [kplib ffi sizeof](#)

Source code

```
1  /*
2  * lib_ffi.c - FFI library
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include "../include/ktap_types.h"
23 #include "../include/ktap_ffi.h"
24 #include "ktap.h"
25 #include "kp_vm.h"
26
27
28 static int kplib_ffi_new(ktap_state_t *ks)
29 {
30     int n = kp_arg_nr(ks);
31     csymbol_id cs_id = kp_arg_checknumber(ks, 1);
32     int array_size = kp_arg_checknumber(ks, 2);
33     int is_array = kp_arg_checknumber(ks, 3);
34     ktap_cdata_t *cd;
35
36     if (unlikely(n != 3)) {
37         /* this is not likely to happen since ffi.new arguments are
38          * generated by compiler */
39         set_nil(ks->top++);
40         kp_error(ks, "wrong number of arguments\n");
41         return 1;
42     }
43
44     if (unlikely(cs_id > max_csym_id(ks)))
45         kp_error(ks, "invalid csymbol id\n");
46
47     kp_verbose_printf(ks, "ffi.new symbol %s with length %d\n",
48                     id_to_csym(ks, cs_id)->name, array_size);
49
50     if (is_array)
```

```

51     cd = kp\_cdata\_new\_ptr(ks, NULL, array_size, cs_id, 1);
52     else
53         cd = kp\_cdata\_new\_by\_id(ks, NULL, cs_id);
54     set\_cdata(ks->top, cd);
55     incr\_top(ks);
56
57     return 1;
58 }
59
60 static int kplib\_ffi\_cast(ktap\_state\_t *ks)
61 {
62     int n = kp\_arg\_nr(ks);
63     csymbol\_id cs_id = kp\_arg\_checknumber(ks, 1);
64     unsigned long addr = kp\_arg\_checknumber(ks, 2);
65     ktap\_cdata\_t *cd;
66
67     if (unlikely(n != 2)) {
68         /* this is not likely to happen since ffi.cast arguments are
69          * generated by compiler */
70         set\_nil(ks->top++);
71         kp\_error(ks, "wrong number of arguments\n");
72         return 1;
73     }
74
75     if (unlikely(cs_id > max\_csym\_id(ks)))
76         kp\_error(ks, "invalid csymbol id\n");
77
78     cd = kp\_cdata\_new\_record(ks, (void *)addr, cs_id);
79     set\_cdata(ks->top, cd);
80     incr\_top(ks);
81     return 1;
82 }
83
84 static int kplib\_ffi\_sizeof(ktap\_state\_t *ks)
85 {
86     /*@TODO finish this 08.11 2013 (houqp)*/
87     return 0;
88 }
89
90 static const ktap\_libfunc\_t ffi\_lib\_funcs[] = {
91     {"sizeof", kplib\_ffi\_sizeof},
92     {"new", kplib\_ffi\_new},
93     {"cast", kplib\_ffi\_cast},
94     {NULL}
95 };
96
97 int kp\_lib\_init\_ffi(ktap\_state\_t *ks)
98 {
99     return kp\_vm\_register\_lib(ks, "ffi", ffi\_lib\_funcs);
100 }

```

[One Level Up](#)

[Top Level](#)

runtime/ffi/cdata.c - ktap

Functions defined

- [kp_cdata_dump](#)
- [kp_cdata_init](#)
- [kp_cdata_new](#)
- [kp_cdata_new_by_id](#)
- [kp_cdata_new_number](#)
- [kp_cdata_new_ptr](#)
- [kp_cdata_new_record](#)
- [kp_cdata_pack](#)
- [kp_cdata_ptr_get](#)
- [kp_cdata_ptr_set](#)
- [kp_cdata_record_get](#)
- [kp_cdata_record_set](#)
- [kp_cdata_type_match](#)
- [kp_cdata_unpack](#)
- [kp_cdata_value](#)

Source code

```
1  /*
2  * cdata.c - support functions for ktap\_cdata\_t
3  *
4  * This file is part of ktap by Jovi Zhangwei
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22
23 #include "../include/ktap_types.h"
24 #include "../include/ktap_ffi.h"
25 #include "../kp_obj.h"
26
27 ktap\_cdata\_t *kp\_cdata\_new(ktap\_state\_t *ks, csymbol\_id id)
28 {
29     ktap\_cdata\_t *cd;
30
31     cd = &kp_obj_newobject(ks, KTAP_TYPE_CDATA, sizeof(ktap\_cdata\_t), NULL)->cd;
```

```

32     cd_set_csym_id(cd, id);
33
34     return cd;
35 }
36
37 ktap_cdata_t *kp_cdata_new_number(ktap_state_t *ks, void *val, csymbol_id id)
38 {
39     ktap_cdata_t *cd;
40
41     cd = kp_cdata_new(ks, id);
42     cd_int(cd) = (cdata_number)val;
43
44     return cd;
45 }
46
47 /* argument nmemb here indicates the length of array that is pointed to,
48 * -1 for unknown */
49 ktap_cdata_t *kp_cdata_new_ptr(ktap_state_t *ks, void *addr,
50                               int nmemb, csymbol_id id, int to_allocate)
51 {
52     ktap_cdata_t *cd;
53     size_t memb_size;
54     csymbol_id deref_id;
55
56     cd = kp_cdata_new(ks, id);
57
58     if (to_allocate) {
59         /* allocate new empty space */
60         deref_id = csym_ptr_deref_id(id to csym(ks, id));
61         memb_size = csym_size(ks, id to csym(ks, deref_id));
62         cd_ptr(cd) = kp_rawobj_alloc(ks, memb_size * nmemb);
63     } else {
64         cd_ptr(cd) = addr;
65     }
66
67     cd_ptr_nmemb(cd) = nmemb;
68
69     return cd;
70 }
71
72 ktap_cdata_t *kp_cdata_new_record(ktap_state_t *ks, void *val, csymbol_id id)
73 {
74     ktap_cdata_t *cd;
75     size_t size;
76
77     cd = kp_cdata_new(ks, id);
78
79     /* if val == NULL, allocate new empty space */
80     if (val == NULL) {
81         size = csym_size(ks, id to csym(ks, id));
82         cd_record(cd) = kp_rawobj_alloc(ks, size);
83     } else
84         cd_record(cd) = val;
85
86     return cd;
87 }
88
89 ktap_cdata_t *kp_cdata_new_by_id(ktap_state_t *ks, void *val, csymbol_id id)
90 {
91     csymbol *cs = id to csym(ks, id);
92
93     switch (csym_type(cs)) {
94     case FFI_VOID:
95         kp_error(ks, "Error: Cannot new a void type\n");
96         return NULL;
97     case FFI_UINT8:
98     case FFI_INT8:
99     case FFI_UINT16:
100    case FFI_INT16:
101    case FFI_UINT32:
102    case FFI_INT32:
103    case FFI_UINT64:
104    case FFI_INT64:
105        return kp_cdata_new_number(ks, val, id);
106    case FFI_PTR:
107        return kp_cdata_new_ptr(ks, val, 0, id, 0);

```

```

108 case FFI_STRUCT:
109 case FFI_UNION:
110     return kp_cdata_new_record(ks, val, id);
111 case FFI_FUNC:
112     kp_error(ks, "Error: Cannot new a function type\n");
113     return NULL;
114 case FFI_UNKNOWN:
115 default:
116     kp_error(ks, "Error: unknown csymbol type %s\n", csym_name(cs));
117     return NULL;
118 }
119 }
120
121 void kp_cdata_dump(ktap_state_t *ks, ktap_cdata_t *cd)
122 {
123     switch (cd_type(ks, cd)) {
124     case FFI_UINT8:     case FFI_INT8:
125         kp_printf(ks, "c int(0x%01x)", cd_int(cd));
126         break;
127     case FFI_UINT16:   case FFI_INT16:
128         kp_printf(ks, "c int(0x%02x)", cd_int(cd));
129         break;
130     case FFI_UINT32:   case FFI_INT32:
131         kp_printf(ks, "c int(0x%04x)", cd_int(cd));
132         break;
133     case FFI_UINT64:   case FFI_INT64:
134         kp_printf(ks, "c int(0x%08x)", cd_int(cd));
135         break;
136     case FFI_PTR:
137         kp_printf(ks, "c pointer(0x%p)", cd_ptr(cd));
138         break;
139     case FFI_STRUCT:
140         kp_printf(ks, "c struct(0x%p)", cd_struct(cd));
141         break;
142     case FFI_UNION:
143         kp_printf(ks, "c union(0x%p)", cd_union(cd));
144         break;
145     default:
146         kp_printf(ks, "unsupported cdata type %d!\n", cd_type(ks, cd));
147     }
148 }
149
150 /* Notice: Even if the types are matched, there may exist the lost of
151 * data in the unpack process due to precision */
152 int kp_cdata_type_match(ktap_state_t *ks, csymbol *cs, ktap_val_t *val)
153 {
154     ffi_type type;
155
156     type = csym_type(cs);
157     if (type == FFI_FUNC)
158         goto error;
159
160     switch (ttypenv(val)) {
161     case KTAP_TYPE_LIGHTUSERDATA:
162         if (type != FFI_PTR) goto error;
163         break;
164     case KTAP_TYPE_BOOLEAN:
165     case KTAP_TYPE_NUMBER:
166         if (type != FFI_UINT8 && type != FFI_INT8
167             && type != FFI_UINT16 && type != FFI_INT16
168             && type != FFI_UINT32 && type != FFI_INT32
169             && type != FFI_UINT64 && type != FFI_INT64)
170             goto error;
171         break;
172     case KTAP_TYPE_STRING:
173         if (type != FFI_PTR && type != FFI_UINT8 && type != FFI_INT8)
174             goto error;
175         break;
176     case KTAP_TYPE_CDATA:
177         if (cs != cd_csym(ks, cdvalue(val)))
178             goto error;
179         break;
180     default:
181         goto error;
182     }
183     return 0;

```

```

184
185 error:
186     return -1;
187 }
188
189 static void kp_cdata_value(ktap_state_t *ks, ktap_val_t *val, void **out_addr,
190                          size_t *out_size, void **temp)
191 {
192     ktap_cdata_t *cd;
193     csymbol *cs;
194     ffi_type type;
195
196     switch (ttypenv(val)) {
197     case KTAP_TYPE_BOOLEAN:
198         *out_addr = &bvalue(val);
199         *out_size = sizeof(int);
200         return;
201     case KTAP_TYPE_LIGHTUSERDATA:
202         *out_addr = pvalue(val);
203         *out_size = sizeof(void *);
204         return;
205     case KTAP_TYPE_NUMBER:
206         *out_addr = &nvalue(val);
207         *out_size = sizeof(ktap_number);
208         return;
209     case KTAP_TYPE_STRING:
210         *temp = (void *)svalue(val);
211         *out_addr = temp;
212         *out_size = sizeof(void *);
213         return;
214     }
215
216     cd = cdvalue(val);
217     cs = cd_csym(ks, cd);
218     type = csym_type(cs);
219     *out_size = csym_size(ks, cs);
220     switch (type) {
221     case FFI_VOID:
222         kp_error(ks, "Error: Cannot copy data from void type\n");
223         return;
224     case FFI_UINT8:
225     case FFI_INT8:
226     case FFI_UINT16:
227     case FFI_INT16:
228     case FFI_UINT32:
229     case FFI_INT32:
230     case FFI_UINT64:
231     case FFI_INT64:
232         *out_addr = &cd_int(cd);
233         return;
234     case FFI_PTR:
235         *out_addr = &cd_ptr(cd);
236         return;
237     case FFI_STRUCT:
238     case FFI_UNION:
239         *out_addr = cd_record(cd);
240         return;
241     case FFI_FUNC:
242     case FFI_UNKNOWN:
243         kp_error(ks, "Error: internal error for csymbol %s\n",
244                 csym_name(cs));
245         return;
246     }
247 }
248
249 /* Check whether or not type is matched before unpacking */
250 void kp_cdata_unpack(ktap_state_t *ks, char *dst, csymbol *cs, ktap_val_t *val)
251 {
252     size_t size = csym_size(ks, cs), val_size;
253     void *val_addr, *temp;
254
255     kp_cdata_value(ks, val, &val_addr, &val_size, &temp);
256     if (val_size > size)
257         val_size = size;
258     memmove(dst, val_addr, val_size);
259     memset(dst + val_size, 0, size - val_size);

```

```

260 }
261
262 /* Check whether or not type is matched before packing */
263 void kp_cdata_pack(ktap_state_t *ks, ktap_val_t *val, char *src, csymbol *cs)
264 {
265     size_t size = csym_size(ks, cs), val_size;
266     void *val_addr, *temp;
267
268     kp_cdata_value(ks, val, &val_addr, &val_size, &temp);
269     if (size > val_size)
270         size = val_size;
271     memmove(val_addr, src, size);
272     memset(val_addr + size, 0, val_size - size);
273 }
274
275 /* Init its cdata type, but not its actual value */
276 static void kp_cdata_init(ktap_state_t *ks, ktap_val_t *val, void *addr, int len,
277     csymbol_id id)
278 {
279     ffi_type type = csym_type(id to csym(ks, id));
280
281     switch (type) {
282     case FFI_PTR:
283         set_cdata(val, kp_cdata_new_ptr(ks, addr, len, id, 0));
284         break;
285     case FFI_STRUCT:
286     case FFI_UNION:
287         set_cdata(val, kp_cdata_new_record(ks, addr, id));
288         break;
289     case FFI_UINT8:
290     case FFI_INT8:
291     case FFI_UINT16:
292     case FFI_INT16:
293     case FFI_UINT32:
294     case FFI_INT32:
295     case FFI_UINT64:
296     case FFI_INT64:
297         /* set all these value into ktap_number(long) */
298         set_number(val, 0);
299         break;
300     default:
301         set_cdata(val, kp_cdata_new(ks, id));
302         break;
303     }
304 }
305
306 void kp_cdata_ptr_set(ktap_state_t *ks, ktap_cdata_t *cd,
307     ktap_val_t *key, ktap_val_t *val)
308 {
309     ktap_number idx;
310     csymbol *cs;
311     size_t size;
312     char *addr;
313
314     if (!is_number(key)) {
315         kp_error(ks, "array index should be number\n");
316         return;
317     }
318     idx = nvalue(key);
319     if (unlikely(idx < 0 || (cd_ptr_nmemb(cd) >= 0
320         && idx >= cd_ptr_nmemb(cd)))) {
321         kp_error(ks, "array index out of bound\n");
322         return;
323     }
324
325     cs = csym_ptr_deref(ks, cd_csym(ks, cd));
326     if (kp_cdata_type_match(ks, cs, val)) {
327         kp_error(ks, "array member should be %s type\n", csym_name(cs));
328         return;
329     }
330     size = csym_size(ks, cs);
331     addr = cd_ptr(cd);
332     addr += size * idx;
333     kp_cdata_unpack(ks, addr, cs, val);
334 }
335

```

```

336 void kp_cdata_ptr_get(ktap_state_t *ks, ktap_cdata_t *cd,
337                      ktap_val_t *key, ktap_val_t *val)
338 {
339     ktap_number idx;
340     csymbol *cs;
341     size_t size;
342     char *addr;
343     csymbol_id cs_id;
344
345     if (!is_number(key)) {
346         kp_error(ks, "array index should be number\n");
347         return;
348     }
349     idx = nvalue(key);
350     if (unlikely(idx < 0 || (cd_ptr_nmemb(cd) >= 0
351                            && idx >= cd_ptr_nmemb(cd)))) {
352         kp_error(ks, "array index out of bound\n");
353         return;
354     }
355
356     cs_id = csym_ptr_deref_id(cd_csym(ks, cd));
357     cs = id_to_csym(ks, cs_id);
358     size = csym_size(ks, cs);
359     addr = cd_ptr(cd);
360     addr += size * idx;
361
362     kp_cdata_init(ks, val, addr, -1, cs_id);
363     kp_cdata_pack(ks, val, addr, cs);
364 }
365
366 void kp_cdata_record_set(ktap_state_t *ks, ktap_cdata_t *cd,
367                         ktap_val_t *key, ktap_val_t *val)
368 {
369     const char *mb_name;
370     csymbol *cs, *mb_cs;
371     csymbol_struct *csst;
372     struct_member *mb;
373     char *addr;
374
375     if (!is_shrstring(key)) {
376         kp_error(ks, "struct member name should be string\n");
377         return;
378     }
379     mb_name = svalue(key);
380     cs = cd_csym(ks, cd);
381     csst = csym_struct(cs);
382     mb = csymst_mb_by_name(ks, csst, mb_name);
383     if (mb == NULL) {
384         kp_error(ks, "struct member %s doesn't exist\n", mb_name);
385         return;
386     }
387
388     mb_cs = id_to_csym(ks, mb->id);
389     if (kp_cdata_type_match(ks, mb_cs, val)) {
390         kp_error(ks, "struct member should be %s type\n",
391                csym_name(mb_cs));
392         return;
393     }
394
395     addr = cd_record(cd);
396     addr += csym_record_mb_offset_by_name(ks, cs, mb_name);
397     kp_cdata_unpack(ks, addr, mb_cs, val);
398 }
399
400 void kp_cdata_record_get(ktap_state_t *ks, ktap_cdata_t *cd,
401                         ktap_val_t *key, ktap_val_t *val)
402 {
403     const char *mb_name;
404     csymbol *cs, *mb_cs;
405     csymbol_struct *csst;
406     struct_member *mb;
407     char *addr;
408     csymbol_id mb_cs_id;
409
410     if (!is_shrstring(key)) {
411         kp_error(ks, "struct member name should be string\n");

```

```
412     return;
413 }
414
415 mb_name = svalue(key);
416 cs = cd_csym(ks, cd);
417 csst = csym_struct(cs);
418 mb = csymst_mb_by_name(ks, csst, mb_name);
419 if (mb == NULL) {
420     kp_error(ks, "struct member %s doesn't exist\n", mb_name);
421     return;
422 }
423
424 mb_cs_id = mb->id;
425 mb_cs = id_to_csym(ks, mb_cs_id);
426 addr = cd_record(cd);
427 addr += csym_record_mb_offset_by_name(ks, cs, mb_name);
428
429 kp_cdata_init(ks, val, addr, mb->len, mb_cs_id);
430 if (mb->len < 0)
431     kp_cdata_pack(ks, val, addr, mb_cs);
432 }
433
```

[One Level Up](#)

[Top Level](#)

runtime/ffi/ffi_util.c - ktap

Functions defined

- [csym_align](#)
- [csym_record_mb_offset_by_name](#)
- [csym_size](#)
- [csymst_mb_by_name](#)
- [init_csym_struct](#)
- [init_csym_union](#)

Source code

```
1 /*
2  * ffi_util.c - utility function for ffi module
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22
23 #include "../include/ktap_types.h"
24 #include "../include/ktap_ffi.h"
25 #include "../ktap.h"
26
27 static void init_csym_struct(ktap_state_t *ks, csymbol_struct *csst)
28 {
29     int nr = csymst_mb_nr(csst);
30     size_t size = 0;
31     size_t align = 1;
32     int i, var_len;
33     size_t var_size, var_align;
34
35     if (csymst_mb_nr(csst) < 0) {
36         kp_error(ks, "Find size of undefined struct");
37         return;
38     }
39
40     for (i = 0; i < nr; i++) {
41         csymbol *var_cs = csymst_mb_csym(ks, csst, i);
42         var_len = csymst_mb_len(csst, i);
43         if (var_len < 0) {
44             var_len = 1;
45         } else {
46             var_cs = csym_ptr_deref(ks, var_cs);
47         }
48         var_size = csym_size(ks, var_cs) * var_len;
49         var_align = csym_align(ks, var_cs);
50         size = ALIGN(size, var_align);
51         size += var_size;
```

```

52     align = align > var_align ? align : var_align;
53 }
54 size = ALIGN(size, align);
55 csst->size = size;
56 csst->align = align;
57 }
58
59 static void init_csym_union(ktap_state_t *ks, csymbol_struct *csst)
60 {
61     int nr = csymst_mb_nr(csst);
62     size_t size = 0;
63     size_t align = 1;
64     int i, var_len;
65     size_t var_size, var_align;
66
67     if (csymst_mb_nr(csst) < 0) {
68         kp_error(ks, "Find size of undefined struct");
69         return;
70     }
71
72     for (i = 0; i < nr; i++) {
73         csymbol *var_cs = csymst_mb_csym(ks, csst, i);
74         var_len = csymst_mb_len(csst, i);
75         if (var_len < 0) {
76             var_len = 1;
77         } else {
78             var_cs = csym_ptr_deref(ks, var_cs);
79         }
80         var_size = csym_size(ks, var_cs) * var_len;
81         var_align = csym_align(ks, var_cs);
82         size = size > var_size ? size : var_size;
83         align = align > var_align ? align : var_align;
84     }
85     csst->size = size;
86     csst->align = align;
87 }
88
89
90 size_t csym_size(ktap_state_t *ks, csymbol *cs)
91 {
92     ffi_type type = csym_type(cs);
93     switch(type) {
94     case FFI_STRUCT:
95         if (csym_struct(cs)->align == 0)
96             init_csym_struct(ks, csym_struct(cs));
97         return csym_struct(cs)->size;
98     case FFI_UNION:
99         if (csym_struct(cs)->align == 0)
100             init_csym_union(ks, csym_struct(cs));
101         return csym_struct(cs)->size;
102     default:
103         return ffi_type_size(type);
104     }
105 }
106
107 size_t csym_align(ktap_state_t *ks, csymbol *cs)
108 {
109     ffi_type type = csym_type(cs);
110     switch(type) {
111     case FFI_STRUCT:
112         if (csym_struct(cs)->align == 0)
113             init_csym_struct(ks, csym_struct(cs));
114         return csym_struct(cs)->align;
115     case FFI_UNION:
116         if (csym_struct(cs)->align == 0)
117             init_csym_union(ks, csym_struct(cs));
118         return csym_struct(cs)->align;
119     default:
120         return ffi_type_align(type);
121     }
122 }
123
124 size_t csym_record_mb_offset_by_name(ktap_state_t *ks,
125     csymbol *cs, const char *name)
126 {
127     csymbol_struct *csst = csym_struct(cs);

```

```

128     int nr = csymst\_mb\_nr(csst);
129     size_t off = 0, sub_off;
130     size_t align = 1;
131     int i, var_len;
132     size_t var_size, var_align;
133
134     if (nr < 0) {
135         kp\_error(ks, "Find size of undefined struct");
136         return 0;
137     }
138
139     for (i = 0; i < nr; i++) {
140         csymbol *var_cs = csymst\_mb\_csym(ks, csst, i);
141         var_len = csymst\_mb\_len(csst, i);
142         if (var_len < 0) {
143             var_len = 1;
144         } else {
145             var_cs = csym\_ptr\_deref(ks, var_cs);
146         }
147         var_size = csym\_size(ks, var_cs) * var_len;
148         var_align = csym\_align(ks, var_cs);
149         off = ALIGN(off, var_align);
150         if (!strcmp(name, csymst\_mb\_name(csst, i)))
151             return off;
152         if (!strcmp("", csymst\_mb\_name(csst, i))) {
153             if (csym\_type(var_cs) != FFI_STRUCT &&
154                 csym\_type(var_cs) != FFI_UNION) {
155                 kp\_error(ks, "Parse error: non-record type without name");
156                 return -1;
157             }
158             sub_off = csym\_record\_mb\_offset\_by\_name(ks,
159                 var_cs, name);
160             if (sub_off >= 0)
161                 return off + sub_off;
162         }
163         if (csym\_type(cs) == FFI_STRUCT)
164             off += var_size;
165         else
166             off = 0;
167         align = align > var_align ? align : var_align;
168     }
169     return -1;
170 }
171
172 struct member *csymst\_mb\_by\_name(ktap\_state\_t *ks,
173     csymbol\_struct *csst, const char *name)
174 {
175     int nr = csymst\_mb\_nr(csst);
176     int i;
177     struct member *memb;
178     csymbol *cs = NULL;
179
180     if (nr < 0) {
181         kp\_error(ks, "Find size of undefined struct");
182         return NULL;
183     }
184
185     for (i = 0; i < nr; i++) {
186         if (!strcmp(name, csymst\_mb\_name(csst, i)))
187             return csymst\_mb(csst, i);
188         if (!strcmp("", csymst\_mb\_name(csst, i))) {
189             cs = csymst\_mb\_csym(ks, csst, i);
190             if (csym\_type(cs) != FFI_STRUCT && csym\_type(cs) != FFI_UNION) {
191                 kp\_error(ks, "Parse error: non-record type without name");
192                 return NULL;
193             }
194             memb = csymst\_mb\_by\_name(ks, csym\_struct(cs), name);
195             if (memb != NULL)
196                 return memb;
197         }
198     }
199     return NULL;
200 }

```

runtime/lib_table.c - ktap

Global variables defined

- [table lib funcs](#)

Functions defined

- [kp lib init table](#)
- [kplib table new](#)

Source code

```
1 /*
2  * lib_table.c - Table library
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <linux/ctype.h>
23 #include <linux/slab.h>
24 #include <linux/delay.h>
25 #include <linux/sched.h>
26 #include "../include/ktap_types.h"
27 #include "ktap.h"
28 #include "kp_obj.h"
29 #include "kp_vm.h"
30 #include "kp_tab.h"
31
32 static int kplib_table_new(ktap_state_t *ks)
33 {
34     int narr = kp_arg_checkoptnumber(ks, 1, 0);
35     int nrec = kp_arg_checkoptnumber(ks, 2, 0);
36     ktap_tab_t *h;
37
38     h = kp_tab_new_ah(ks, narr, nrec);
39     if (!h) {
40         set_nil(ks->top);
41     } else {
42         set_table(ks->top, h);
43     }
44
45     incr_top(ks);
46     return 1;
47 }
48
49 static const ktap_libfunc_t table_lib_funcs[] = {
50     {"new", kplib_table_new},
51     {NULL}
52 };
53
54 int kp_lib_init_table(ktap_state_t *ks)
```

```
55 {  
56   return kp\_vm\_register\_lib(ks, "table", table\_lib\_funcs);  
57 }  
58
```

[One Level Up](#)

[Top Level](#)

runtime/lib_net.c - ktap

Global variables defined

- [net lib funcs](#)

Functions defined

- [kp lib init net](#)
- [kplib net format ip addr](#)
- [kplib net ip sock daddr](#)
- [kplib net ip sock saddr](#)

Source code

```
1  /*
2  * lib_base.c - base library
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22 #include <net/inet_sock.h>
23 #include "../include/ktap_types.h"
24 #include "ktap.h"
25 #include "kp_obj.h"
26 #include "kp_str.h"
27 #include "kp_vm.h"
28
29 /**
30 * Return the source IP address for a given sock
31 */
32 static int kplib_net_ip_sock_saddr(ktap_state_t *ks)
33 {
34     struct inet_sock *isk;
35     int family;
36
37     /* TODO: need to validate the address firstly */
38
39     isk = (struct inet_sock *)kp_arg_checknumber(ks, 1);
40     family = isk->sk.__sk_common.skc_family;
41
42     if (family == AF_INET) {
43         set_number(ks->top, isk->inet_rcv_saddr);
44     } else {
45         kp_error(ks, "ip_sock_saddr only support ipv4 now\n");
46         set_nil(ks->top);
47     }
48
49     incr_top(ks);
50     return 1;

```

```

51 }
52
53 /**
54  * Return the destination IP address for a given sock
55  */
56 static int kplib_net_ip_sock_daddr(ktap_state_t *ks)
57 {
58     struct inet_sock *isk;
59     int family;
60
61     /* TODO: need to validate the address firstly */
62
63     isk = (struct inet_sock *)kp_arg_checknumber(ks, 1);
64     family = isk->sk.__sk_common.skc_family;
65
66     if (family == AF_INET) {
67         set_number(ks->top, isk->inet_daddr);
68     } else {
69         kp_error(ks, "ip_sock_daddr only support ipv4 now\n");
70         set_nil(ks->top);
71     }
72
73     incr_top(ks);
74     return 1;
75 }
76
77
78 /**
79  * Returns a string representation for an IP address
80  */
81 static int kplib_net_format_ip_addr(ktap_state_t *ks)
82 {
83     __be32 ip = (__be32)kp_arg_checknumber(ks, 1);
84     ktap_str_t *ts;
85     char ipstr[32];
86
87     snprintf(ipstr, 32, "%pI4", &ip);
88     ts = kp_str_newz(ks, ipstr);
89     if (ts) {
90         set_string(ks->top, kp_str_newz(ks, ipstr));
91         incr_top(ks);
92         return 1;
93     } else
94         return -1;
95 }
96
97 static const ktap_libfunc_t net_lib_funcs[] = {
98     {"ip_sock_saddr", kplib_net_ip_sock_saddr},
99     {"ip_sock_daddr", kplib_net_ip_sock_daddr},
100    {"format_ip_addr", kplib_net_format_ip_addr},
101    {NULL}
102 };
103
104 int kp_lib_init_net(ktap_state_t *ks)
105 {
106     return kp_vm_register_lib(ks, "net", net_lib_funcs);
107 }

```

[One Level Up](#)

[Top Level](#)

userspace/kp_bcwrite.c - ktap

Global variables defined

- [bc_mode](#)
- [bc_names](#)
- [function_nr](#)

Data types defined

- [BCWriteCtx](#)
- [BCWriteCtx](#)

Functions defined

- [bcwrite_bytecode](#)
- [bcwrite_footer](#)
- [bcwrite_header](#)
- [bcwrite_kgc](#)
- [bcwrite_knum](#)
- [bcwrite_ktab](#)
- [bcwrite_ktabk](#)
- [bcwrite_proto](#)
- [bcwrite_uint32](#)
- [dump_bytecode](#)
- [kp_bcwrite](#)
- [kp_dump_proto](#)

Macros defined

- [BCNAME](#)
- [BCNAME](#)

Source code

```
1  /*
2  * Bytecode writer
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * Copyright (C) 1994-2013 Lua.org, PUC-Rio.
9  * - The part of code in this file is copied from lua initially.
10 * - lua's MIT license is compatible with GPL.
11 *
```

```

12 * ktap is free software; you can redistribute it and/or modify it
13 * under the terms and conditions of the GNU General Public License,
14 * version 2, as published by the Free Software Foundation.
15 *
16 * ktap is distributed in the hope it will be useful, but WITHOUT
17 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
18 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
19 * more details.
20 *
21 * You should have received a copy of the GNU General Public License along with
22 * this program; if not, write to the Free Software Foundation, Inc.,
23 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
24 */
25
26 #include <stdio.h>
27 #include <stdlib.h>
28 #include <string.h>
29
30 #include "../include/ktap_types.h"
31 #include "cparser.h"
32 #include "kp_util.h"
33
34 /* Context for bytecode writer. */
35 typedef struct BCWriteCtx {
36     SBuf sb;          /* Output buffer. */
37     ktap_proto_t *pt; /* Root prototype. */
38     ktap_writer wfunc; /* Writer callback. */
39     void *wdata;      /* Writer callback data. */
40     int strip;        /* Strip debug info. */
41     int status;       /* Status from writer callback. */
42 } BCWriteCtx;
43
44
45 static char *bcwrite_uint32(char *p, uint32_t v)
46 {
47     memcpy(p, &v, sizeof(uint32_t));
48     p += sizeof(uint32_t);
49     return p;
50 }
51
52 /* -- Bytecode writer ----- */
53
54 /* Write a single constant key/value of a template table. */
55 static void bcwrite_ktab(BCWriteCtx *ctx, const ktap_val_t *o, int narrow)
56 {
57     char *p = kp_buf_more(&ctx->sb, 1+10);
58     if (is_string(o)) {
59         const ktap_str_t *str = rawtsvalue(o);
60         int len = str->len;
61         p = kp_buf_more(&ctx->sb, 5+len);
62         p = bcwrite_uint32(p, BCDUMP_KTAB_STR+len);
63         p = kp_buf_wmem(p, getstr(str), len);
64     } else if (is_number(o)) {
65         p = bcwrite_uint32(p, BCDUMP_KTAB_NUM);
66         p = kp_buf_wmem(p, &nvalue(o), sizeof(ktap_number));
67     } else {
68         kp_assert(tvispri(o));
69         p = bcwrite_uint32(p, BCDUMP_KTAB_NIL+~itype(o));
70     }
71     setsbufP(&ctx->sb, p);
72 }
73
74 /* Write a template table. */
75 static void bcwrite_ktab(BCWriteCtx *ctx, char *p, const ktap_tab_t *t)
76 {
77     int narray = 0, nhash = 0;
78     if (t->asize > 0) { /* Determine max. length of array part. */
79         ptrdiff_t i;
80         ktap_val_t *array = t->array;
81         for (i = (ptrdiff_t)t->asize-1; i >= 0; i--)
82             if (!is_nil(&array[i]))
83                 break;
84         narray = (int)(i+1);
85     }
86     if (t->hmask > 0) { /* Count number of used hash slots. */
87         int i, hmask = t->hmask;

```

```

88     ktap_node_t *node = t->node;
89     for (i = 0; i <= hmask; i++)
90         nhash += !is_nil(&node[i].val);
91 }
92 /* Write number of array slots and hash slots. */
93 p = bcwrite_uint32(p, narray);
94 p = bcwrite_uint32(p, nhash);
95 setsbufP(&ctx->sb, p);
96 if (narray) { /* Write array entries (may contain nil). */
97     int i;
98     ktap_val_t *o = t->array;
99     for (i = 0; i < narray; i++, o++)
100         bcwrite_ktabk(ctx, o, 1);
101 }
102 if (nhash) { /* Write hash entries. */
103     int i = nhash;
104     ktap_node_t *node = t->node + t->hmask;
105     for (;;) node--
106         if (!is_nil(&node->val)) {
107             bcwrite_ktabk(ctx, &node->key, 0);
108             bcwrite_ktabk(ctx, &node->val, 1);
109             if (--i == 0)
110                 break;
111         }
112 }
113 }
114
115 /* Write GC constants of a prototype. */
116 static void bcwrite_kgc(BCWriteCtx *ctx, ktap_proto_t *pt)
117 {
118     int i, sizekgc = pt->sizekgc;
119     ktap_obj_t **kr = (ktap_obj_t **)pt->k - (ptrdiff_t)sizekgc;
120
121     for (i = 0; i < sizekgc; i++, kr++) {
122         ktap_obj_t *o = *kr;
123         int tp, need = 1;
124         char *p;
125
126         /* Determine constant type and needed size. */
127         if (o->gch.gct == ~KTAP_TSTR) {
128             tp = BCDUMP_KGC_STR + ((ktap_str_t *)o)->len;
129             need = 5 + ((ktap_str_t *)o)->len;
130         } else if (o->gch.gct == ~KTAP_TPROTO) {
131             kp_assert((pt->flags & PROTO_CHILD));
132             tp = BCDUMP_KGC_CHILD;
133         } else {
134             kp_assert(o->gch.gct == ~KTAP_TTAB);
135             tp = BCDUMP_KGC_TAB;
136             need = 1+2*5;
137         }
138
139         /* Write constant type. */
140         p = kp_buf_more(&ctx->sb, need);
141         p = bcwrite_uint32(p, tp);
142         /* Write constant data (if any). */
143         if (tp >= BCDUMP_KGC_STR) {
144             p = kp_buf_wmem(p, getstr((ktap_str_t *)o),
145                 ((ktap_str_t *)o)->len);
146         } else if (tp == BCDUMP_KGC_TAB) {
147             bcwrite_ktab(ctx, p, (ktap_tab_t *)o);
148             continue;
149         }
150         setsbufP(&ctx->sb, p);
151     }
152 }
153
154 /* Write number constants of a prototype. */
155 static void bcwrite_knum(BCWriteCtx *ctx, ktap_proto_t *pt)
156 {
157     int i, sizekn = pt->sizekn;
158     const ktap_val_t *o = (ktap_val_t *)pt->k;
159     char *p = kp_buf_more(&ctx->sb, 10*sizekn);
160
161     for (i = 0; i < sizekn; i++, o++) {
162         if (is_number(o))
163             p = kp_buf_wmem(p, &nvalue(o), sizeof(ktap_number));

```

```

164     }
165     setsbufP(&ctx->sb, p);
166 }
167
168 /* Write bytecode instructions. */
169 static char *bcwrite_bytecode(BCWriteCtx *ctx, char *p, ktap\_proto\_t *pt)
170 {
171     int nbc = pt->sizebc-1; /* Omit the [JI]FUNC* header. */
172
173     p = kp\_buf\_wmem(p, proto\_bc(pt)+1, nbc*(int)sizeof(BCIns));
174     return p;
175 }
176
177 /* Write prototype. */
178 static void bcwrite_proto(BCWriteCtx *ctx, ktap\_proto\_t *pt)
179 {
180     int sizedbg = 0;
181     char *p;
182
183     /* Recursively write children of prototype. */
184     if (pt->flags & PROTO\_CHILD) {
185         ptrdiff\_t i, n = pt->sizekgc;
186         ktap\_obj\_t **kr = (ktap\_obj\_t **)pt->k - 1;
187         for (i = 0; i < n; i++, kr--) {
188             ktap\_obj\_t *o = *kr;
189             if (o->gch.gct == -KTAP\_TPROTO)
190                 bcwrite\_proto(ctx, (ktap\_proto\_t *)o);
191         }
192     }
193
194     /* Start writing the prototype info to a buffer. */
195     p = kp\_buf\_need(&ctx->sb,
196         5+4+6*5+(pt->sizebc-1)*(int)sizeof(BCIns)+pt->sizeuv*2);
197     p += 4; /* Leave room for final size. */
198
199     /* Write prototype header. */
200     *p++ = (pt->flags & (PROTO\_CHILD|PROTO\_VARARG|PROTO\_FFI));
201     *p++ = pt->numparams;
202     *p++ = pt->framesize;
203     *p++ = pt->sizeuv;
204     p = bcwrite\_uint32(p, pt->sizekgc);
205     p = bcwrite\_uint32(p, pt->sizekn);
206     p = bcwrite\_uint32(p, pt->sizebc-1);
207     if (!ctx->strip) {
208         if (proto\_lineinfo(pt))
209             sizedbg = pt->sizept -
210                 (int)((char *)proto\_lineinfo(pt) - (char *)pt);
211         p = bcwrite\_uint32(p, sizedbg);
212         if (sizedbg) {
213             p = bcwrite\_uint32(p, pt->firstline);
214             p = bcwrite\_uint32(p, pt->numline);
215         }
216     }
217
218     /* Write bytecode instructions and upvalue refs. */
219     p = bcwrite\_bytecode(ctx, p, pt);
220     p = kp\_buf\_wmem(p, proto\_uv(pt), pt->sizeuv*2);
221     setsbufP(&ctx->sb, p);
222
223     /* Write constants. */
224     bcwrite\_kgc(ctx, pt);
225     bcwrite\_knum(ctx, pt);
226
227     /* Write debug info, if not stripped. */
228     if (sizedbg) {
229         p = kp\_buf\_more(&ctx->sb, sizedbg);
230         p = kp\_buf\_wmem(p, proto\_lineinfo(pt), sizedbg);
231         setsbufP(&ctx->sb, p);
232     }
233
234     /* Pass buffer to writer function. */
235     if (ctx->status == 0) {
236         int n = sbufllen(&ctx->sb) - 4;
237         char *q = sbufB(&ctx->sb);
238         p = bcwrite\_uint32(q, n); /* Fill in final size. */
239         kp\_assert(p == sbufB(&ctx->sb) + 4);

```

```

240     ctx->status = ctx->wfunc(q, n + 4, ctx->wdata);
241 }
242 }
243
244 /* Write header of bytecode dump. */
245 static void bcwrite_header(BCWriteCtx *ctx)
246 {
247     ktap_str_t *chunkname = proto_chunkname(ctx->pt);
248     const char *name = getstr(chunkname);
249     int len = chunkname->len;
250     char *p = kp_buf_need(&ctx->sb, 5+5+len);
251     *p++ = BCDUMP_HEAD1;
252     *p++ = BCDUMP_HEAD2;
253     *p++ = BCDUMP_HEAD3;
254     *p++ = BCDUMP_VERSION;
255     *p++ = (ctx->strip ? BCDUMP_F_STRIP : 0) + (KP_BE ? BCDUMP_F_BE : 0);
256
257     if (!ctx->strip) {
258         p = bcwrite_uint32(p, len);
259         p = kp_buf_wmem(p, name, len);
260     }
261     ctx->status = ctx->wfunc(sbufB(&ctx->sb),
262         (int)(p - sbufB(&ctx->sb)), ctx->wdata);
263 }
264
265 /* Write footer of bytecode dump. */
266 static void bcwrite_footer(BCWriteCtx *ctx)
267 {
268     if (ctx->status == 0) {
269         uint8_t zero = 0;
270         ctx->status = ctx->wfunc(&zero, 1, ctx->wdata);
271     }
272 }
273
274 /* Write bytecode for a prototype. */
275 int kp_bcwrite(ktap_proto_t *pt, ktap_writer writer, void *data, int strip)
276 {
277     BCWriteCtx ctx;
278
279     ctx.pt = pt;
280     ctx.wfunc = writer;
281     ctx.wdata = data;
282     ctx.strip = strip;
283     ctx.status = 0;
284
285     kp_buf_init(&ctx.sb);
286     kp_buf_need(&ctx.sb, 1024); /* Avoids resize for most prototypes. */
287     bcwrite_header(&ctx);
288     bcwrite_proto(&ctx, ctx.pt);
289     bcwrite_footer(&ctx);
290
291     kp_buf_free(&ctx.sb);
292     return ctx.status;
293 }
294
295 /* -- Bytecode dump ----- */
296
297 static const char * const bc_names[] = {
298 #define BCNAME(name, ma, mb, mc, mt)      #name,
299     BCDEF(BCNAME)
300 #undef BCNAME
301     NULL
302 };
303
304 static const uint16_t bc_mode[] = {
305     BCDEF(BCMODE)
306 };
307
308 static void dump_bytecode(ktap_proto_t *pt)
309 {
310     int nbc = pt->sizebc - 1; /* Omit the FUNC* header. */
311     BCIns *ins = proto_bc(pt) + 1;
312     ktap_obj_t **kbase = pt->k;
313     int i;
314
315     printf("-- BYTECODE -- %s:%d-%d\n", getstr(pt->chunkname),

```

```

316     pt->firstline, pt->firstline + pt->numline);
317
318     for (i = 0; i < nbc; i++, ins++) {
319         int op = bc\_op(*ins);
320
321         printf("%04d\t%s", i + 1, bc\_names[op]);
322
323         printf("\t%d", bc\_a(*ins));
324         if (bcmode\_b(op) != BCMnone)
325             printf("\t%d", bc\_b(*ins));
326
327         if (bcmode\_hasd(op))
328             printf("\t%d", bc\_d(*ins));
329         else
330             printf("\t%d", bc\_c(*ins));
331
332         if (bcmode\_b(op) == BCMstr || bcmode\_c(op) == BCMstr) {
333             printf("\t ; ");
334             if (bcmode\_d(op) == BCMstr) {
335                 int idx = ~bc\_d(*ins);
336                 printf("\t%s", getstr((ktap\_str\_t *)kbase[idx]));
337             }
338         }
339         printf("\n");
340     }
341 }
342
343 static int function_nr = 0;
344
345 void kp\_dump\_proto(ktap\_proto\_t *pt)
346 {
347     printf("\n-----\n");
348     printf("function proto %d:\n", function\_nr++);
349     printf("numparams: %d\n", pt->numparams);
350     printf("framesize: %d\n", pt->framesize);
351     printf("sizebc: %d\n", pt->sizebc);
352     printf("sizekgc: %d\n", pt->sizekgc);
353     printf("sizekn: %d\n", pt->sizekn);
354     printf("sizept: %d\n", pt->sizept);
355     printf("sizeuv: %d\n", pt->sizeuv);
356     printf("firstline: %d\n", pt->firstline);
357     printf("numline: %d\n", pt->numline);
358
359     printf("has child proto: %d\n", pt->flags & PROTO\_CHILD);
360     printf("has vararg: %d\n", pt->flags & PROTO\_VARARG);
361     printf("has ILOOP: %d\n", pt->flags & PROTO\_ILOOP);
362
363     dump\_bytecode(pt);
364
365     /* Recursively dump children of prototype. */
366     if (pt->flags & PROTO\_CHILD) {
367         ptrdiff\_t i, n = pt->sizekgc;
368         ktap\_obj\_t **kr = (ktap\_obj\_t **)pt->k - 1;
369         for (i = 0; i < n; i++, kr--) {
370             ktap\_obj\_t *o = *kr;
371             if (o->gch.gct == ~KTAP\_TPROTO)
372                 kp\_dump\_proto((ktap\_proto\_t *)o);
373         }
374     }
375 }
376

```

[One Level Up](#)

[Top Level](#)

test/ffi_test/ktap_ffi_test.c - ktap

Global variables defined

- [ffi_test_array](#)
- [ffi_test_exit](#)
- [ffi_test_init](#)
- [ffi_test_int1](#)
- [ffi_test_int2](#)
- [ffi_test_pointer1](#)
- [ffi_test_sched_clock](#)
- [ffi_test_struct](#)
- [ffi_test_struct_array](#)
- [ffi_test_struct_loop](#)
- [ffi_test_struct_noname](#)
- [ffi_test_union](#)
- [ffi_test_var_arg](#)
- [ffi_test_void](#)

Data types defined

- [ffi_struct](#)
- [ffi_struct_array](#)
- [ffi_struct_loop](#)
- [ffi_struct_noname](#)
- [ffi_union](#)

Functions defined

- [ffi_test_array](#)
- [ffi_test_exit](#)
- [ffi_test_init](#)
- [ffi_test_int1](#)
- [ffi_test_int2](#)
- [ffi_test_pointer1](#)
- [ffi_test_sched_clock](#)
- [ffi_test_struct](#)

- [ffi_test_struct_array](#)
- [ffi_test_struct_loop](#)
- [ffi_test_struct_noname](#)
- [ffi_test_union](#)
- [ffi_test_var_arg](#)
- [ffi_test_void](#)

Source code

```

1  #include <linux/module.h>
2  #include <linux/kernel.h>
3  #include <linux/init.h>
4  #include <linux/sched.h>
5
6  void ffi_test_void(void)
7  {
8  }
9  EXPORT_SYMBOL(ffi_test_void);
10
11 int ffi_test_int1(unsigned char a, char b, unsigned short c, short d)
12 {
13     return a + b + c + d;
14 }
15 EXPORT_SYMBOL(ffi_test_int1);
16
17 long long ffi_test_int2(unsigned int a, int b, unsigned long c, long d,
18     unsigned long long e, long long f, long long g)
19 {
20     return a + b + c + d + e + f + g;
21 }
22 EXPORT_SYMBOL(ffi_test_int2);
23
24 void *ffi_test_pointer1(char *a) {
25     return a;
26 }
27 EXPORT_SYMBOL(ffi_test_pointer1);
28
29 long long ffi_test_var_arg(int n, ...) {
30     va_list ap;
31     int i;
32     long long sum = 0;
33     va_start(ap, n);
34     for (i = 0; i < n; i++) {
35         sum += va_arg(ap, long long);
36     }
37     va_end(ap);
38     return sum;
39 }
40 EXPORT_SYMBOL(ffi_test_var_arg);
41
42 unsigned long long ffi_test_sched_clock(void)
43 {
44     return sched_clock();
45 }
46 EXPORT_SYMBOL(ffi_test_sched_clock);
47
48 int ffi_test_array(int *arr, int idx)
49 {
50     return arr[idx];
51 }
52 EXPORT_SYMBOL(ffi_test_array);
53
54 struct ffi_struct {
55     int val;
56 };
57
58 int ffi_test_struct(struct ffi_struct *s)
59 {

```

```

60     return s->val;
61 }
62 EXPORT_SYMBOL(fffi_test_struct);
63
64 struct fffi_struct_loop {
65     int val;
66     struct fffi_struct_loop *next;
67 };
68
69 int fffi_test_struct_loop(struct fffi_struct_loop *s)
70 {
71     if (s == s->next)
72         return s->val;
73     return 0;
74 }
75 EXPORT_SYMBOL(fffi_test_struct_loop);
76
77 union fffi_union {
78     int val;
79     int val2;
80     int val3;
81 };
82
83 int fffi_test_union(union fffi_union *s)
84 {
85     return s->val;
86 }
87 EXPORT_SYMBOL(fffi_test_union);
88
89 struct fffi_struct_array {
90     int val[20];
91 };
92
93 int fffi_test_struct_array(struct fffi_struct_array *s)
94 {
95     int sum = 0, i;
96     for (i = 0; i < 20; i++)
97         sum += s->val[i];
98     return sum;
99 }
100 EXPORT_SYMBOL(fffi_test_struct_array);
101
102 struct fffi_struct_noname {
103     struct {
104         int pad;
105         int val;
106     };
107 };
108
109 int fffi_test_struct_noname(struct fffi_struct_noname *s)
110 {
111     return s->val;
112 }
113 EXPORT_SYMBOL(fffi_test_struct_noname);
114
115
116 static int __init fffi_test_init(void)
117 {
118     return 0;
119 }
120
121 static void __exit fffi_test_exit(void)
122 {
123 }
124
125
126 MODULE_DESCRIPTION("ktap fffi test module");
127 MODULE_LICENSE("GPL");
128
129 module_init(fffi_test_init);
130 module_exit(fffi_test_exit);

```

[One Level Up](#)

[Top Level](#)

test/lib/ - ktap

- [Test/](#)

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

test/lib/Test/ - ktap

- [ktap/](#)
- [ktap.pm](#)

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

test/lib/Test/ktap/ - ktap

- [features.pm](#)

[One Level Up](#)

[Top Level](#)

test/lib/Test/ktap/features.pm - ktap

Data types defined

- [has_ffi](#)

Source code

```
1 package Test::ktap::features;
2
3 use strict;
4 use warnings;
5
6 use Exporter 'import';
7 our @EXPORT_OK = qw( &has_ffi );
8
9 sub has_ffi {
10     return `./ktap -V 2>&1 | grep -c FFI` == 1;
11 }
12
13 1;
```

test/lib/Test/ktap.pm - ktap

Data types defined

- [bail_out](#)
- [parse_cmd](#)
- [run_test](#)
- [run_tests](#)

Source code

```
1  # Copyright (C) Yichun Zhang (agentzh)
2
3  package Test::ktap;
4
5  use Test::Base -Base;
6  use POSIX ();
7  use IPC::Run ();
8
9  our @EXPORT = qw( run_tests );
10
11 sub run_tests () {
12     for my $block (Test::Base::blocks()) {
13         run_test($block);
14     }
15 }
16
17 sub bail_out (@) {
18     Test::More::BAIL_OUT(@_);
19 }
20
21 sub parse_cmd ($) {
22     my $cmd = shift;
23     my @cmd;
24     while (1) {
25         if ($cmd =~ /\G\s*"(.*)"\/gmsc) {
26             push @cmd, $1;
27         } elsif ($cmd =~ /\G\s*'(.*)'\/gmsc) {
28             push @cmd, $1;
29         } elsif ($cmd =~ /\G\s*(\S+)\/gmsc) {
30             push @cmd, $1;
31         } else {
32             last;
33         }
34     }
35     return @cmd;
36 }
37
38 sub run_test ($) {
39     my $block = shift;
40     my $name = $block->name;
41
42     my $timeout = $block->timeout() || 10;
43     my $opts = $block->opts;
44     my $args = $block->args;
45
46     my $cmd = "./ktap";
47
48     if (defined $opts) {
49         $cmd .= " $opts";
50     }
51
52     my $kpf;
53     if (defined $block->src) {
```

```

57     $kpfiler = POSIX::tmpnam() . ".kp";
58     open my $out, ">$kpfiler" or
59         bail_out("cannot open $kpfiler for writing: $!");
60     print $out ($block->src);
61     close $out;
62     $cmd .= " $kpfiler"
63 }
64
65 if (defined $args) {
66     $cmd .= " $args";
67 }
68
69 #warn "CMD: $cmd\n";
70
71 my @cmd = parse_cmd($cmd);
72
73 my ($out, $err);
74
75 eval {
76     IPC::Run::run(\@cmd, \undef, \$out, \$err,
77                 IPC::Run::timeout($timeout));
78 };
79 if ($?) {
80     # timed out
81     if ($? =~ /timeout/) {
82         if (!defined $block->expect_timeout) {
83             fail("$name: ktap process timed out");
84         }
85     } else {
86         fail("$name: failed to run command [$cmd]: $?");
87     }
88 }
89
90 my $ret = ($? >> 8);
91
92 if (defined $kpfiler) {
93     unlink $kpfiler;
94 }
95
96 if (defined $block->out) {
97     is $out, $block->out, "$name - stdout eq okay";
98 }
99
100 my $regex = $block->out_like;
101 if (defined $regex) {
102     if (!ref $regex) {
103         $regex = qr/$regex/ms;
104     }
105     like $out, $regex, "$name - stdout like okay";
106 }
107
108 if (defined $block->err) {
109     is $err, $block->err, "$name - stderr eq okay";
110 }
111
112 $regex = $block->err_like;
113 if (defined $regex) {
114     if (!ref $regex) {
115         $regex = qr/$regex/ms;
116     }
117     like $err, $regex, "$name - stderr like okay";
118 }
119
120 my $exp_ret = $block->ret;
121 if (!defined $exp_ret) {
122     $exp_ret = 0;
123 }
124 is $ret, $exp_ret, "$name - exit code okay";
125 }
126
127 1;
128 # vi: et

```

[One Level Up](#)

[Top Level](#)

test/util/ - ktap

- [reindex](#)

[One Level Up](#)

[Top Level](#)

test/util/reindex - ktap

Data types defined

- [reindex](#)

Source code

```
1  #!/usr/bin/env perl
2
3  # reindex
4  # reindex .t files for Test::Base based test files
5  # Copyright (C) Yichun Zhang (agentzh)
6
7  use strict;
8  use warnings;
9
10 use Getopt::Std;
11
12 my %opts;
13 getopts('hb:', \%opts);
14 if (@opts{h} or ! @ARGV) {
15     die "Usage: reindex [-b 0] t/*.t\n";
16 }
17
18 my $init = $opts{b};
19 $init = 1 if not defined $init;
20
21 my @files = map glob, @ARGV;
22 for my $file (@files) {
23     next if -d $file or $file !~ /\..t_?$/;
24     reindex($file);
25 }
26
27 sub reindex {
28     my $file = $_[0];
29     open my $in, $file or
30         die "Can't open $file for reading: $!";
31     my @lines;
32     my $counter = $init;
33     my $changed;
34     while (<$in>) {
35         s/\r$//;
36         my $num;
37         s/ ^ === \s+ TEST \s+ (\d+)/$num=$1; "=== TEST " . $counter++/xie;
38         next if !defined $num;
39         if ($num != $counter-1) {
40             $changed++;
41         }
42     } continue {
43         push @lines, $_;
44     }
45     close $in;
46     my $text = join '', @lines;
47     $text =~ s/(?x) \n+ === \s+ TEST/\n\n\n\n=== TEST/ixsg;
48     $text =~ s/__(DATA|END)__\n+=== TEST/__$1__\n\n=== TEST/;
49     # $text =~ s/\n+$/\n\n/s;
50     if (! $changed and $text eq join '', @lines) {
51         warn "reindex: $file:\t\tskipped.\n";
52         return;
53     }
54     open my $out, "> $file" or
55         die "Can't open $file for writing: $!";
56     binmode $out;
57     print $out $text;
58     close $out;
59
60     warn "reindex: $file:\tdone.\n";
61 }
```


test/benchmark/cmp_neq.sh - ktap

```
1  #!/bin/sh
2
3  # This script compare number equality performance between ktap and stap.
4  # It also compare different ktap tracing interfaces.
5  #
6  # 1. ktap -e 'trace syscalls:sys_enter_futex {}'
7  # 2. ktap -e 'kdebug.tracepoint("sys_enter_futex", function () {})'
8  # 3. ktap -e 'trace probe:SyS_futex uaddr=%di {}'
9  # 4. ktap -e 'kdebug.kprobe("SyS_futex", function () {})'
10 # 5. stap -e 'probe syscall.futex {}'
11 # 6. ktap -d -e 'trace syscalls:sys_enter_futex {}'
12 # 7. ktap -d -e 'kdebug.tracepoint("sys_enter_futex", function () {})'
13 # 8. ktap -e 'trace syscalls:sys_enter_futex /kernel_buildin_filter/ {}'
14
15 #Result:
16 #ktap number computation and comparsion overhead is bigger than stap,
17 #nearly 10+% (4 vs. 5 in above)), ktap is not very slow.
18 #
19 #Perf backend tracing overhead is big, because it need copy temp buffer, and
20 #code path is very long than direct callback(1 vs. 4 in above).
21
22 gcc -o sembench sembench.c -O2 -lpthread
23
24 COMMAND="./sembench -t 200 -w 20 -r 30 -o 2"
25
26 #-----#
27
28 echo -e "without tracing:"
29 # $COMMAND; $COMMAND; $COMMAND
30
31 #-----#
32
33 ../../ktap -q -e 'trace syscalls:sys_enter_futex {
34     var uaddr = arg2
35     if (uaddr == 0x100 || uaddr == 0x200 || uaddr == 0x300 ||
36         uaddr == 0x400 || uaddr == 0x500 || uaddr == 0x600 ||
37         uaddr == 0x700 || uaddr == 0x800 || uaddr == 0x900 ||
38         uaddr == 0x1000) {
39         printf("%x %x\n", arg1, arg2)
40     }
41 }' &
42 echo -e "\nktap tracing: trace syscalls:sys_enter_futex { if (arg2 == 0x100 || arg2 == 0x200 ... }"
43 $COMMAND; $COMMAND; $COMMAND
44 pid=`pidof ktap`
45 disown $pid; kill -9 $pid; sleep 1
46
47 #-----#
48
49 ../../ktap -q -e 'kdebug.tracepoint("sys_enter_futex", function () {
50     var arg = arg2
51     if (arg == 0x100 || arg == 0x200 || arg == 0x300 || arg == 0x400 ||
52         arg == 0x500 || arg == 0x600 || arg == 0x700 || arg == 0x800 ||
53         arg == 0x900 || arg == 0x1000) {
54         printf("%x %x\n", arg1, arg2)
55     }
56 })' &
57 echo -e "\nktap tracing: kdebug.tracepoint("sys_enter_futex", function (xxx) {})"
58 $COMMAND; $COMMAND; $COMMAND
59 pid=`pidof ktap`
60 disown $pid; kill -9 $pid; sleep 1
61
62 #-----#
63
64 ../../ktap -q -e 'trace probe:SyS_futex uaddr=%di {
65     var arg = arg1
66     if (arg == 0x100 || arg == 0x200 || arg == 0x300 || arg == 0x400 ||
67         arg == 0x500 || arg == 0x600 || arg == 0x700 || arg == 0x800 ||
68         arg == 0x900 || arg == 0x1000) {
69         printf("%x\n", arg1)
70     }
71 }' &
72 echo -e "\nktap tracing: trace probe:SyS_futex uaddr=%di {...}"
```

```

72 $COMMAND; $COMMAND; $COMMAND
73 pid=`pidof ktap`
74 disown $pid; kill -9 $pid; sleep 1
75
76
77 #-----#
78 ../../ktap -q -e 'kdebug.kprobe("Sys_futex", function () {
79     var uaddr = 1
80     if (uaddr == 0x100 || uaddr == 0x200 || uaddr == 0x300 ||
81         uaddr == 0x400 || uaddr == 0x500 || uaddr == 0x600 ||
82         uaddr == 0x700 || uaddr == 0x800 || uaddr == 0x900 ||
83         uaddr == 0x1000) {
84         printf("%x\n", uaddr)
85     }}' &
86 echo -e '\nktap tracing: kdebug.kprobe("Sys_futex", function () {})'
87 $COMMAND; $COMMAND; $COMMAND
88 pid=`pidof ktap`
89 disown $pid; kill -9 $pid; sleep 1
90
91 #-----#
92
93 stap -e 'probe syscall.futex {
94     uaddr = $uaddr
95     if (uaddr == 0x100 || uaddr == 0x200 || uaddr == 0x300 ||
96         uaddr == 0x400 || uaddr == 0x500 || uaddr == 0x600 ||
97         uaddr == 0x700 || uaddr == 0x800 || uaddr == 0x900 ||
98         uaddr == 0x1000) {
99         printf("%x\n", uaddr)
100     }}' &
101
102 echo -e "\nstap tracing: probe syscall.futex { if (uaddr == 0x100 || addr == 0x200 ... }"
103 $COMMAND; $COMMAND; $COMMAND
104 pid=`pidof stap`
105 disown $pid; kill -9 $pid; sleep 1
106
107 #-----#
108
109
110 ../../ktap -d -q -e 'trace syscalls:sys_enter_futex {
111     var uaddr = arg2
112     if (uaddr == 0x100 || uaddr == 0x200 || uaddr == 0x300 ||
113         uaddr == 0x400 || uaddr == 0x500 || uaddr == 0x600 ||
114         uaddr == 0x700 || uaddr == 0x800 || uaddr == 0x900 ||
115         uaddr == 0x1000) {
116         printf("%x %x\n", arg1, arg2)
117     }}' &
118
119 echo -e "\nktap tracing dry-run: trace syscalls:sys_enter_futex { if (arg2 == 0x100 || arg2 == 0x200 ... }"
120 $COMMAND; $COMMAND; $COMMAND
121 pid=`pidof ktap`
122 disown $pid; kill -9 $pid; sleep 1
123
124
125 #-----#
126
127 ../../ktap -d -q -e 'kdebug.tracepoint("sys_enter_futex", function () {
128     var arg = arg2
129     if (arg == 0x100 || arg == 0x200 || arg == 0x300 || arg == 0x400 ||
130         arg == 0x500 || arg == 0x600 || arg == 0x700 || arg == 0x800 ||
131         arg == 0x900 || arg == 0x1000) {
132         printf("%x %x\n", arg1, arg2)
133     }}' &
134
135 echo -e '\nktap tracing dry-run: kdebug.tracepoint("sys_enter_futex", function (xxx) {})'
136 $COMMAND; $COMMAND; $COMMAND
137 pid=`pidof ktap`
138 disown $pid; kill -9 $pid; sleep 1
139
140
141 #-----#
142
143 ../../ktap -q -e 'trace syscalls:sys_enter_futex /
144     uaddr == 0x100 || uaddr == 0x200 || uaddr == 0x300 || uaddr == 0x400 ||
145     uaddr == 0x500 || uaddr == 0x600 || uaddr == 0x700 || uaddr == 0x800 ||
146     uaddr == 0x900 || uaddr == 0x1000/ {
147     printf("%x %x\n", arg1, arg2)

```

```
148     }' &
149
150 echo -e "\nktap tracing: trace syscalls:sys_enter_futex /uaddr == 0x100 || uaddr == 0x200 .../ {"
151 $COMMAND; $COMMAND; $COMMAND
152 pid=`pidof ktap`
153 disown $pid; kill -9 $pid; sleep 1
154
155 #-----#
156
157 rm -rf ./sembench
158
```

[One Level Up](#)

[Top Level](#)

test/benchmark/cmp_profile.sh - ktap

```
1 #!/bin/sh
2
3 # This script compare stack profiling performance between ktap and stap.
4 #
5 # 1. ktap -e 'profile-1000us { s[stack(-1, 12)] += 1 }'
6 # 2. stap -e 'probe timer.profile { s[backtrace()] += 1 }'
7 # 3. stap -e 'probe timer.profile { s[backtrace()] <<< 1 }'
8
9 #Result:
10 #Currently the stack profiling overhead is nearly same between ktap and stap.
11 #
12 #ktap resolve kernel stack to string in runtime, which is very time consuming,
13 #optimize it in future.
14
15
16 gcc -o sembench sembench.c -O2 -lpthread
17
18 COMMAND="./sembench -t 200 -w 20 -r 30 -o 2"
19
20 #-----#
21
22 echo -e "without tracing:"
23 $COMMAND; $COMMAND; $COMMAND
24
25 #-----#
26
27 ../../ktap -q -e 'var s = table.new(0, 20000) profile-1000us { s[stack(-1, 12)] += 1 }' &
28
29 echo -e "\nktap tracing: profile-1000us { s[stack(-1, 12)] += 1 }"
30 $COMMAND; $COMMAND; $COMMAND
31 pid=`pidof ktap`
32 disown $pid; kill -9 $pid; sleep 1
33
34 #-----#
35
36 stap -o /dev/null -e 'global s[20000]; probe timer.profile { s[backtrace()] += 1 }' &
37
38 echo -e "\nstap tracing: probe timer.profile { s[backtrace()] += 1 }"
39 $COMMAND; $COMMAND; $COMMAND
40 pkill stap
41
42 #-----#
43
44 stap -o /dev/null -e 'global s[20000]; probe timer.profile { s[backtrace()] <<< 1 }' &
45
46 echo -e "\nstap tracing: probe timer.profile { s[backtrace()] <<< 1 }"
47 $COMMAND; $COMMAND; $COMMAND
48 pkill stap
49
50 #-----#
51
52
53 rm -rf ./sembench
54
```

test/benchmark/cmp_table.sh - ktap

```
1  #!/bin/sh
2
3  # This script compare table performance between ktap and stap.
4  #
5  # 1. ktap -e 'trace syscalls:sys_enter_futex { s[execname] += 1 }'
6  # 2. ktap -e 'kdebug.tracepoint("sys_enter_futex", function () { s[execname] += 1 })'
7  # 3. ktap -e 'kdebug.kprobe("Sys_futex", function () { s[execname] += 1 })'
8  # 4. stap -e 'probe syscall.futex { s[execname()] += 1 }'
9  # 5. ktap -e 'kdebug.kprobe("Sys_futex", function () { s[probename] += 1 })'
10 # 6. stap -e 'probe syscall.futex { s[name] += 1 }'
11 # 7. ktap -e 'kdebug.kprobe("Sys_futex", function () { s["constant_string_key"] += 1 })'
12 # 8. stap -e 'probe syscall.futex { s["constant_string_key"] += 1 }'
13
14 #Result:
15 #Currently ktap table operation overhead is smaller than stap.
16
17
18 gcc -o sembench sembench.c -O2 -lpthread
19
20 COMMAND="./sembench -t 200 -w 20 -r 30 -o 2"
21
22 #-----#
23
24 echo -e "without tracing:"
25 $COMMAND; $COMMAND; $COMMAND
26
27 #-----#
28
29 ../../ktap -q -e 'var s = {} trace syscalls:sys_enter_futex { s[execname] += 1 }' &
30
31 echo -e "\nktap tracing: trace syscalls:sys_enter_futex { s[execname] += 1 }"
32 $COMMAND; $COMMAND; $COMMAND
33 pid=`pidof ktap`
34 disown $pid; kill -9 $pid; sleep 1
35
36 #-----#
37
38 ../../ktap -q -e 'var s = {} kdebug.tracepoint("sys_enter_futex", function () {
39     s[execname] += 1 })' &
40
41 echo -e "\nktap tracing: kdebug.tracepoint("sys_enter_futex", function () { s[execname] += 1 })"
42 $COMMAND; $COMMAND; $COMMAND
43 pid=`pidof ktap`
44 disown $pid; kill -9 $pid; sleep 1
45
46 #-----#
47
48 ../../ktap -q -e 'var s = {} kdebug.kprobe("Sys_futex", function () {
49     s[execname] += 1 })' &
50
51 echo -e "\nktap tracing: kdebug.kprobe("Sys_futex", function () { s[execname] += 1 })"
52 $COMMAND; $COMMAND; $COMMAND
53 pid=`pidof ktap`
54 disown $pid; kill -9 $pid; sleep 1
55
56 #-----#
57
58 stap -e 'global s; probe syscall.futex { s[execname()] += 1 }' &
59
60 echo -e "\nstap tracing: probe syscall.futex { s[execname()] += 1 }"
61 $COMMAND; $COMMAND; $COMMAND
62 pkill stap
63
64 #-----#
65
66 ../../ktap -q -e 'var s = {} kdebug.kprobe("Sys_futex", function () {
67     s[probename] += 1 })' &
68
69 echo -e "\nktap tracing: kdebug.kprobe("Sys_futex", function () { s[probename] += 1 })"
70 $COMMAND; $COMMAND; $COMMAND
71 pid=`pidof ktap`
```

```

72 disown $pid; kill -9 $pid; sleep 1
73
74 #-----#
75
76 stap -e 'global s; probe syscall.futex { s[name] += 1 }' &
77
78 echo -e "\nstap tracing: probe syscall.futex { s[name] += 1 }"
79 $COMMAND; $COMMAND; $COMMAND
80 pkill stap
81
82 #-----#
83
84 ../../ktap -q -e 'var s = {} s["const_string_key"] = 0 kdebug.kprobe("SyS_futex", function () {
85     s["const_string_key"] += 1 })' &
86
87 echo -e '\nktap tracing: kdebug.kprobe("SyS_futex", function () { s["const_string_key"] += 1 })'
88 $COMMAND; $COMMAND; $COMMAND
89 pid=`pidof ktap`
90 disown $pid; kill -9 $pid; sleep 1
91
92 #-----#
93
94 stap -e 'global s; probe syscall.futex { s["const_string_key"] += 1 }' &
95
96 echo -e "\nstap tracing: probe syscall.futex { s["const_string_key"] += 1 }"
97 $COMMAND; $COMMAND; $COMMAND
98 pkill stap
99
100 #-----#
101
102 stap -o /dev/null -e 'global s; probe syscall.futex { s["const_string_key"] <<< 1 }' &
103
104 echo -e "\nstap tracing: probe syscall.futex { s["const_string_key"] <<< 1 }"
105 $COMMAND; $COMMAND; $COMMAND
106 pkill stap
107
108 #-----#
109
110
111 rm -rf ./sembench
112

```

[One Level Up](#)

[Top Level](#)

test/benchmark/test.sh - ktap

```
1 #!/bin/sh
2
3 gcc -o sembench sembench.c -O2 -lpthread
4
5 COMMAND="./sembench -t 200 -w 20 -r 30 -o 2"
6
7 #-----#
8
9 echo -e "without tracing:"
10 $COMMAND
11 $COMMAND
12 $COMMAND
13
14 #-----#
15
16 ../../ktap -d -q -e 'trace syscalls:sys_enter_futex {}' &
17 echo -e "\nktap tracing in dry-run mode: trace syscalls:sys_enter_futex {}"
18 $COMMAND
19 $COMMAND
20 $COMMAND
21 pid=`pidof ktap`
22 disown $pid; kill -9 $pid; sleep 1
23
24 #-----#
25
26 ../../ktap -q -e 'trace syscalls:sys_enter_futex {}' &
27 echo -e "\nktap tracing: trace syscalls:sys_enter_futex {}"
28 $COMMAND
29 $COMMAND
30 $COMMAND
31 pid=`pidof ktap`
32 disown $pid; kill -9 $pid; sleep 1
33
34 #-----#
35
36 ../../ktap -q -e 'kdebug.tracepoint("sys_enter_futex", function () {})' &
37 echo -e '\nktap tracing: kdebug.tracepoint("sys_enter_futex", function () {})'
38 $COMMAND
39 $COMMAND
40 $COMMAND
41 pid=`pidof ktap`
42 disown $pid; kill -9 $pid; sleep 1
43
44 #-----#
45
46 ../../ktap -q -e 'kdebug.tracepoint("sys_enter_futex", function () {
47     var arg = 1
48     if (arg == 0x100 || arg == 0x200 || arg == 0x300 || arg == 0x400 ||
49         arg == 0x500 || arg == 0x600 || arg == 0x700 || arg == 0x800 ||
50         arg == 0x900 || arg == 0x1000) {
51         printf("%x %x\n", arg1, arg2)
52     }
53 })' &
54 echo -e '\nktap tracing: kdebug.tracepoint("sys_enter_futex", function (xxx) {})'
55 $COMMAND
56 $COMMAND
57 $COMMAND
58 pid=`pidof ktap`
59 disown $pid; kill -9 $pid; sleep 1
60
61
62 #-----#
63
64 perf record -e syscalls:sys_enter_futex -a &
65 echo -e "\nperf tracing: perf record -e syscalls:sys_enter_futex -a"
66 $COMMAND
67 $COMMAND
68 $COMMAND
69 pid=`pidof perf`
70 disown $pid; kill -9 $pid; sleep 1; rm -rf perf.data
71
```

```

72 #-----#
73
74 #ltnng cannot compiled successful in my box.
75
76 #ltnng create
77 #ltnng enable-event -k syscalls:sys_enter_futex &
78 #echo -e "\nltnng tracing: ltnng enable-event -k syscalls:sys_enter_futex"
79 # $COMMAND
80 # $COMMAND
81 # $COMMAND
82 #pid=`pidof perf`
83 #disown $pid; kill -9 $pid; sleep 1;
84 #ltnng destroy
85
86
87
88 #-----#
89
90 ../../ktap -q -e 'trace syscalls:sys_enter_futex {
91     if (arg2 == 0x100 || arg2 == 0x200 || arg2 == 0x300 || arg2 == 0x400 ||
92         arg2 == 0x500 || arg2 == 0x600 || arg2 == 0x700 || arg2 == 0x800 ||
93         arg2 == 0x900 || arg2 == 0x1000) {
94         printf("%x %x\n", arg1, arg2)
95     }
96 }' &
97 echo -e "\nktap tracing: trace syscalls:sys_enter_futex { if (arg2 == 0x100 || arg2 == 0x200 ... }"
98 $COMMAND
99 $COMMAND
100 $COMMAND
101 pid=`pidof ktap`
102 disown $pid; kill -9 $pid; sleep 1
103
104 #-----#
105
106 ../../ktap -q -e 'trace syscalls:sys_enter_futex /
107     uaddr == 0x100 || uaddr == 0x200 || uaddr == 0x300 || uaddr == 0x400 ||
108     uaddr == 0x500 || uaddr == 0x600 || uaddr == 0x700 || uaddr == 0x800 ||
109     uaddr == 0x900 || uaddr == 0x1000/ {
110     printf("%x %x\n", arg1, arg2)
111 }' &
112
113 echo -e "\nktap tracing: trace syscalls:sys_enter_futex /uaddr == 0x100 || uaddr == 0x200 .../ {"
114 $COMMAND
115 $COMMAND
116 $COMMAND
117 pid=`pidof ktap`
118 disown $pid; kill -9 $pid; sleep 1
119
120 #-----#
121 #ltnng don't support kernal event filter now
122 #-----#
123
124 #-----#
125 #systemtap compiled failure in my box
126 #-----#
127
128 rm -rf ./sembench

```

[One Level Up](#)

[Top Level](#)

runtime/ffi/ffi_call.c - ktap

Data types defined

- [arg_status](#)

Functions defined

- [ffi_call](#)
- [ffi_call_unsupported](#)
- [ffi_call_x86_64](#)
- [ffi_get_arg_csym](#)
- [ffi_set_return](#)
- [ffi_type_check](#)
- [ffi_unpack](#)

Macros defined

- [ALIGN_STACK](#)
- [GPR_SIZE](#)
- [MAX_GPR](#)
- [MAX_GPR_SIZE](#)
- [NEWSTACK_SIZE](#)
- [REDZONE_SIZE](#)
- [STACK_ALIGNMENT](#)
- [ffi_call](#)
- [ffi_call_arch](#)

Source code

```
1 /*
2  * ffi\_call.c - foreign function calling library support for ktap
3  *
4  * This file is part of ktap by Jovi Zhangwei.
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
```

```

20 */
21
22 #include <linux/ctype.h>
23 #include <linux/slab.h>
24 #include "../include/ktap_types.h"
25 #include "../include/ktap_ffi.h"
26 #include "../ktap.h"
27 #include "../kp_vm.h"
28 #include "../kp_obj.h"
29
30 static int ffi_type_check(ktap_state_t *ks, csymbol_func *csf, int idx)
31 {
32     StkId arg;
33     csymbol *cs;
34
35     if (idx >= csymf_arg_nr(csf))
36         return 0;
37     arg = kp_arg(ks, idx + 1);
38     cs = csymf_arg(ks, csf, idx);
39
40     if (!kp_cdata_type_match(ks, cs, arg))
41         return 0;
42     else {
43         kp_error(ks, "Cannot convert to csymbol %s for arg %d\n",
44                 csym_name(cs), idx);
45         return -1;
46     }
47 }
48
49 static csymbol *ffi_get_arg_csym(ktap_state_t *ks, csymbol_func *csf, int idx)
50 {
51     StkId arg;
52     csymbol *cs;
53
54     if (idx < csymf_arg_nr(csf))
55         return csymf_arg(ks, csf, idx);
56
57     arg = kp_arg(ks, idx + 1);
58     cs = id_to_csym(ks, ffi_get_csym_id(ks, "void *"));
59     switch (ttypenv(arg)) {
60     case KTAP_TYPE_LIGHTUSERDATA:
61     case KTAP_TYPE_BOOLEAN:
62     case KTAP_TYPE_NUMBER:
63     case KTAP_TYPE_STRING:
64         return cs;
65     case KTAP_TYPE_CDATA:
66         return cd_csym(ks, cdvalue(arg));
67     default:
68         kp_error(ks, "Error: Cannot get type for arg %d\n", idx);
69         return cs;
70     }
71 }
72
73 static void ffi_unpack(ktap_state_t *ks, char *dst, csymbol_func *csf,
74                       int idx, int align)
75 {
76     StkId arg = kp_arg(ks, idx + 1);
77     csymbol *cs = ffi_get_arg_csym(ks, csf, idx);
78     size_t size = csym_size(ks, cs);
79
80     /* initialize the destination section */
81     memset(dst, 0, ALIGN(size, align));
82
83     kp_cdata_unpack(ks, dst, cs, arg);
84 }
85
86 #ifdef __x86_64
87
88 enum arg_status {
89     IN_REGISTER,
90     IN_MEMORY,
91     IN_STACK,
92 };
93
94 #define ALIGN_STACK(v, a) ((void *) (ALIGN(((uint64_t)v), a)))
95 #define STACK_ALIGNMENT 8

```

```

96 #define REDZONE_SIZE 128
97 #define GPR_SIZE (sizeof(void *))
98 #define MAX_GPR 6
99 #define MAX_GPR_SIZE (MAX_GPR * GPR_SIZE)
100 #define NEWSTACK_SIZE 512
101
102 #define ffi_call_arch(ks, cf, rvalue) ffi_call_x86_64(ks, cf, rvalue)
103
104 extern void ffi_call_assem_x86_64(void *stack, void *temp_stack,
105     void *func_addr, void *rvalue, ffi_type rtype);
106
107 static void ffi_call_x86_64(ktap_state_t *ks, csymbol_func *csf, void *rvalue)
108 {
109     int i;
110     int gpr_nr;
111     int arg_bytes; /* total bytes needed for exceeded args in stack */
112     int mem_bytes; /* total bytes needed for memory storage */
113     char *stack, *stack_p, *gpr_p, *arg_p, *mem_p, *tmp_p;
114     int arg_nr;
115     csymbol *rsym;
116     ffi_type rtype;
117     size_t rsize;
118     bool ret_in_memory;
119     /* New stack to call C function */
120     char space[NEWSTACK_SIZE];
121
122     arg_nr = kp_arg_nr(ks);
123     rsym = csymf_ret(ks, csf);
124     rtype = csym_type(rsym);
125     rsize = csym_size(ks, rsym);
126     ret_in_memory = false;
127     if (rtype == FFI_STRUCT || rtype == FFI_UNION) {
128         if (rsize > 16) {
129             rvalue = kp_malloc(ks, rsize);
130             rtype = FFI_VOID;
131             ret_in_memory = true;
132         } else {
133             /* much easier to always copy 16 bytes from registers */
134             rvalue = kp_malloc(ks, 16);
135         }
136     }
137
138     gpr_nr = 0;
139     arg_bytes = mem_bytes = 0;
140     if (ret_in_memory)
141         gpr_nr++;
142     /* calculate bytes needed for stack */
143     for (i = 0; i < arg_nr; i++) {
144         csymbol *cs = ffi_get_arg_csym(ks, csf, i);
145         size_t size = csym_size(ks, cs);
146         size_t align = csym_align(ks, cs);
147         enum arg_status st = IN_REGISTER;
148         int n_gpr_nr = 0;
149         if (size > 32) {
150             st = IN_MEMORY;
151             n_gpr_nr = 1;
152         } else if (size > 16)
153             st = IN_STACK;
154         else
155             n_gpr_nr = ALIGN(size, GPR_SIZE) / GPR_SIZE;
156
157         if (gpr_nr + n_gpr_nr > MAX_GPR) {
158             if (st == IN_MEMORY)
159                 arg_bytes += GPR_SIZE;
160             else
161                 st = IN_STACK;
162         } else
163             gpr_nr += n_gpr_nr;
164         if (st == IN_STACK) {
165             arg_bytes = ALIGN(arg_bytes, align);
166             arg_bytes += size;
167             arg_bytes = ALIGN(arg_bytes, STACK_ALIGNMENT);
168         }
169         if (st == IN_MEMORY) {
170             mem_bytes = ALIGN(mem_bytes, align);
171             mem_bytes += size;

```

```

172     mem_bytes = ALIGN(mem_bytes, STACK_ALIGNMENT);
173     }
174 }
175
176 /* apply space to fake stack for C function call */
177 if (16 + REDZONE_SIZE + MAX_GPR_SIZE + arg_bytes +
178     mem_bytes + 6 * 8 >= NEWSTACK_SIZE) {
179     kp_error(ks, "Unable to handle that many arguments by now\n");
180     return;
181 }
182 stack = space;
183 /* 128 bytes below %rsp is red zone */
184 /* stack should be 16-bytes aligned */
185 stack_p = ALIGN_STACK(stack + REDZONE_SIZE, 16);
186 /* save general purpose registers here */
187 gpr_p = stack_p;
188 memset(gpr_p, 0, MAX_GPR_SIZE);
189 /* save arguments in stack here */
190 arg_p = gpr_p + MAX_GPR_SIZE;
191 /* save arguments in memory here */
192 mem_p = arg_p + arg_bytes;
193 /* set additional space as temporary space */
194 tmp_p = mem_p + mem_bytes;
195
196 /* copy arguments here */
197 gpr_nr = 0;
198 if (ret_in_memory) {
199     memcpy(gpr_p, &rvalue, GPR_SIZE);
200     gpr_p += GPR_SIZE;
201     gpr_nr++;
202 }
203 for (i = 0; i < arg_nr; i++) {
204     csymbol *cs = ffi_get_arg_csym(ks, csf, i);
205     size_t size = csym_size(ks, cs);
206     size_t align = csym_align(ks, cs);
207     enum arg_status st = IN_REGISTER;
208     int n_gpr_nr = 0;
209     if (size > 32) {
210         st = IN_MEMORY;
211         n_gpr_nr = 1;
212     } else if (size > 16)
213         st = IN_STACK;
214     else
215         n_gpr_nr = ALIGN(size, GPR_SIZE) / GPR_SIZE;
216
217     if (st == IN_MEMORY)
218         mem_p = ALIGN_STACK(mem_p, align);
219     /* Tricky way about storing it above mem_p. It won't overflow
220     * because temp region can be temporarily used if necesseary. */
221     ffi_unpack(ks, mem_p, csf, i, GPR_SIZE);
222     if (gpr_nr + n_gpr_nr > MAX_GPR) {
223         if (st == IN_MEMORY) {
224             memcpy(arg_p, &mem_p, GPR_SIZE);
225             arg_p += GPR_SIZE;
226         } else
227             st = IN_STACK;
228     } else {
229         memcpy(gpr_p, mem_p, n_gpr_nr * GPR_SIZE);
230         gpr_p += n_gpr_nr * GPR_SIZE;
231         gpr_nr += n_gpr_nr;
232     }
233     if (st == IN_STACK) {
234         arg_p = ALIGN_STACK(arg_p, align);
235         memcpy(arg_p, mem_p, size);
236         arg_p += size;
237         arg_p = ALIGN_STACK(arg_p, STACK_ALIGNMENT);
238     }
239     if (st == IN_MEMORY) {
240         mem_p += size;
241         mem_p = ALIGN_STACK(mem_p, STACK_ALIGNMENT);
242     }
243 }
244
245 kp_verbose_printf(ks, "Stack location: %p -redzone- %p -general purpose "
246     "register used- %p -zero- %p -stack for argument- %p"
247     " -memory for argument- %p -temp stack-\n",

```

```

248     stack, stack_p, gpr_p, stack_p + MAX_GPR_SIZE,
249     arg_p, mem_p);
250     kp_verbose_printf(ks, "GPR number: %d; arg in stack: %d; "
251     "arg in mem: %d\n",
252     gpr_nr, arg_bytes, mem_bytes);
253     kp_verbose_printf(ks, "Return: address %p type %d\n", rvalue, rtype);
254     kp_verbose_printf(ks, "Number of register used: %d\n", gpr_nr);
255     kp_verbose_printf(ks, "Start FFI call on %p\n", csf->addr);
256     ffi_call_assem_x86_64(stack_p, tmp_p, csf->addr, rvalue, rtype);
257 }
258
259 /* non-supported platform */
260
261 #define ffi_call(ks, cf, rvalue) ffi_call_unsupported(ks, cf, rvalue)
262
263 static void ffi_call_unsupported(ktap_state_t *ks,
264     csymbol_func *csf, void *rvalue)
265 {
266     kp_error(ks, "unsupported architecture.\n");
267 }
268
269 #endif /* end for platform-specific setting */
270
271
272 static int ffi_set_return(ktap_state_t *ks, void *rvalue, csymbol_id ret_id)
273 {
274     ktap_cdata_t *cd;
275     ffi_type type = csym_type(id_to_csym(ks, ret_id));
276
277     /* push return value to ktap stack */
278     switch (type) {
279     case FFI_VOID:
280         return 0;
281     case FFI_UINT8:
282     case FFI_INT8:
283     case FFI_UINT16:
284     case FFI_INT16:
285     case FFI_UINT32:
286     case FFI_INT32:
287     case FFI_UINT64:
288     case FFI_INT64:
289         set_number(ks->top, (ktap_number)rvalue);
290         break;
291     case FFI_PTR:
292         cd = kp_cdata_new_ptr(ks, rvalue, -1, ret_id, 0);
293         set_cdata(ks->top, cd);
294         break;
295     case FFI_STRUCT:
296     case FFI_UNION:
297         cd = kp_cdata_new_record(ks, rvalue, ret_id);
298         set_cdata(ks->top, cd);
299         break;
300     case FFI_FUNC:
301     case FFI_UNKNOWN:
302         kp_error(ks, "Error: Have not support ffi_type %s\n",
303         ffi_type_name(type));
304         return 0;
305     }
306     incr_top(ks);
307     return 1;
308 }
309
310 /*
311 * Call C into function
312 * First argument should be function symbol address, argument types
313 * and return type.
314 * Left arguments should be arguments for calling the C function.
315 * Types between Ktap and C are converted automatically.
316 * Only support x86_64 function call by now
317 */
318 int ffi_call(ktap_state_t *ks, csymbol_func *csf)
319 {
320     int i;
321     int expected_arg_nr, arg_nr;
322     ktap_closure_t *cl;
323     void *rvalue;

```

```

324 expected_arg_nr = csymf\_arg\_nr(csf);
325 arg_nr = kp\_arg\_nr(ks);
326
327 /* check stack status for C call */
328 if (!csf->has_var_arg && expected_arg_nr != arg_nr) {
329     kp\_error(ks, "wrong argument number %d, which should be %d\n",
330             arg_nr, expected_arg_nr);
331     goto out;
332 }
333
334 if (csf->has_var_arg && expected_arg_nr > arg_nr) {
335     kp\_error(ks, "argument number %d, which should be bigger than %d\n",
336             arg_nr, expected_arg_nr);
337     goto out;
338 }
339
340 /* maybe useful later, leave it here first */
341 cl = clvalue(kp\_arg(ks, arg_nr + 1));
342
343 /* check the argument types */
344 for (i = 0; i < arg_nr; i++) {
345     if (ffi\_type\_check(ks, csf, i) < 0)
346         goto out;
347 }
348
349 /* platform-specific calling workflow */
350 ffi\_call\_arch(ks, csf, &rvalue);
351 kp\_verbose\_printf(ks, "Finish FFI call\n");
352
353 out:
354 return ffi\_set\_return(ks, rvalue, csymf\_ret\_id(csf));
355 }

```

[One Level Up](#)

[Top Level](#)

runtime/kp_bcread.h - ktap

Macros defined

- [__KTAP_BCREAD_H__](#)

Source code

```
1 #ifndef \_\_KTAP\_BCREAD\_H\_\_
2 #define \_\_KTAP\_BCREAD\_H\_\_
3
4 ktap\_proto\_t *kp\_bcread(ktap\_state\_t *ks, unsigned char *buff, int len);
5
6 #endif /* \_\_KTAP\_BCREAD\_H\_\_ */
```

runtime/kp_mempool.h - ktap

Macros defined

- [__KTAP_MEMPOOL_H__](#)

Source code

```
1 #ifndef \_\_KTAP\_MEMPOOL\_H\_\_
2 #define \_\_KTAP\_MEMPOOL\_H\_\_
3
4 void *kp\_mempool\_alloc(ktap\_state\_t *ks, int size);
5 void kp\_mempool\_destroy(ktap\_state\_t *ks);
6 int kp\_mempool\_init(ktap\_state\_t *ks, int size);
7
8 #endif /* \_\_KTAP\_MEMPOOL\_H\_\_ */
```

runtime/kp_transport.h - ktap

Macros defined

- [__KTAP_TRANSPORT_H__](#)

Source code

```
1 #ifndef \_\_KTAP\_TRANSPORT\_H\_\_
2 #define \_\_KTAP\_TRANSPORT\_H\_\_
3
4 void kp\_transport\_write(ktap\_state\_t *ks, const void *data, size_t length);
5 void kp\_transport\_event\_write(ktap\_state\_t *ks, struct ktap\_event\_data *e);
6 void kp\_transport\_print\_kstack(ktap\_state\_t *ks, uint16_t depth, uint16_t skip);
7 void *kp\_transport\_reserve(ktap\_state\_t *ks, size_t length);
8 void kp\_transport\_exit(ktap\_state\_t *ks);
9 int kp\_transport\_init(ktap\_state\_t *ks, struct dentry *dir);
10
11 int \_trace\_seq\_puts(struct trace_seq *s, const char *str);
12
13 #endif /* \_\_KTAP\_TRANSPORT\_H\_\_ */
```

runtime/ffi/call_x86_64.S - ktap

```
1  /*
2  * call_x86_64.S - assembly code to call C function and handle return value
3  *
4  * This file is part of ktap by Jovi Zhangwei
5  *
6  * Copyright (C) 2012-2013 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
7  *
8  * ktap is free software; you can redistribute it and/or modify it
9  * under the terms and conditions of the GNU General Public License,
10 * version 2, as published by the Free Software Foundation.
11 *
12 * ktap is distributed in the hope it will be useful, but WITHOUT
13 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
14 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
15 * more details.
16 *
17 * You should have received a copy of the GNU General Public License along with
18 * this program; if not, write to the Free Software Foundation, Inc.,
19 * 51 Franklin St - Fifth Floor, Boston, MA 02110-1301 USA.
20 */
21
22
23 #ifdef __x86_64
24
25     .file "call_x86_64.S"
26     .text
27
28 /*     ffi_call_assem_x86_64(void *stack, void *temp_stack,
29 *         void *rvalue, void *func_addr, ffi_type rftype)
30 *     @stack: base address of register values and new stack
31 *     @temp_stack: stack to store temporary values
32 *     @func_addr: Function address
33 *     @rvalue: where to put return value
34 *     @rftype: FFI type of return value
35 */
36     .align 2
37     .globl ffi_call_assem_x86_64
38     .type ffi_call_assem_x86_64,@function
39
40 ffi_call_assem_x86_64:
41     movq    (%rsp), %rax    /* save return address */
42     /* move stuffs to temp memory region(void *temp_stack) */
43     movq    %rcx, (%rsi)   /* save pointer to return value */
44     movq    %r8, 8(%rsi)   /* save return_ffi_type */
45     movq    %rbp, 16(%rsi) /* save %rbp */
46     movq    %rax, 24(%rsi) /* save return address */
47     movq    %rsp, 32(%rsi) /* save %rsp */
48     movq    %rsi, %rbp    /* point %rbp to temp memory region */
49
50     movq    %rdx, %r11    /* move function address to %r11 */
51
52     movq    %rdi, %r10    /* set %r10 point to register region */
53     movq    (%r10), %rdi  /* load registers */
54     movq    8(%r10), %rsi
55     movq    16(%r10), %rdx
56     movq    24(%r10), %rcx
57     movq    32(%r10), %r8
58     movq    40(%r10), %r9
59     xorq    %rax, %rax
60
61     leaq    48(%r10), %rsp
62
63     callq   *%r11
64
65     movq    32(%rbp), %rsp /* restore %rsp */
66     movq    24(%rbp), %rcx /* restore return address */
67     movq    %rcx, (%rsp)
68
69     movq    (%rbp), %rcx   /* get pointer to return value */
70     movq    8(%rbp), %r8   /* get return_ffi_type */
71     movq    16(%rbp), %rbp /* restore rbp */
```

```

72     leaq    .Lreturn_table(%rip), %r11    /* start address of return_table */
73     movslq  (%r11, %r8, 8), %r11        /* fetch target address from table */
74     jmpq    *%r11                       /* jump according to value in table */
75
76
77     .align 8
78 .Lreturn_table:
79     .quad   .Lreturn_void              /* FFI_VOID */
80     .quad   .Lreturn_uint8            /* FFI_UINT8 */
81     .quad   .Lreturn_int8             /* FFI_INT8 */
82     .quad   .Lreturn_uint16           /* FFI_UINT16 */
83     .quad   .Lreturn_int16            /* FFI_INT16 */
84     .quad   .Lreturn_uint32           /* FFI_UINT32 */
85     .quad   .Lreturn_int32            /* FFI_INT32 */
86     .quad   .Lreturn_uint64           /* FFI_UINT64 */
87     .quad   .Lreturn_int64            /* FFI_INT64 */
88     .quad   .Lreturn_ptr              /* FFI_PTR */
89     .quad   .Lreturn_func             /* FFI_FUNC */
90     .quad   .Lreturn_struct           /* FFI_STRUCT */
91     .quad   .Lreturn_unknown         /* FFI_UNKNOWN */
92
93     .align 8
94 .Lreturn_void:
95 .Lreturn_func:
96 .Lreturn_unknown:
97     retq
98     .align 8
99 .Lreturn_uint8:
100    movzbq  %al, %rax
101    movq    %rax, (%rcx)
102    retq
103    .align 8
104 .Lreturn_int8:
105    movsbq  %al, %rax
106    movq    %rax, (%rcx)
107    retq
108    .align 8
109 .Lreturn_uint16:
110    movzwbq %ax, %rax
111    movq    %rax, (%rcx)
112    retq
113    .align 8
114 .Lreturn_int16:
115    movswq  %ax, %rax
116    movq    %rax, (%rcx)
117    retq
118    .align 8
119 .Lreturn_uint32:
120    movl    %eax, %eax
121    movq    %rax, (%rcx)
122    retq
123    .align 8
124 .Lreturn_int32:
125    movslq  %eax, %rax
126    movq    %rax, (%rcx)
127    retq
128    .align 8
129 .Lreturn_uint64:
130 .Lreturn_int64:
131 .Lreturn_ptr:
132    movq    %rax, (%rcx)
133    retq
134 /* Struct type indicates that struct is put into at most two registers,
135 * and 16 bytes space is always available
136 */
137     .align 8
138 .Lreturn_struct:
139    movq    %rax, (%rcx)
140    movq    %rdx, 8(%rcx)
141    retq
142
143 #endif /* end for __x86_64 */

```

[One Level Up](#)

[Top Level](#)

userspace/kp_parse.h - ktap

```
1
2 ktap\_proto\_t *kp\_parse(LexState *ls);
3 ktap\_str\_t *kp\_parse\_keepstr(LexState *ls, const char *str, size_t l);
4
```

[One Level Up](#)

[Top Level](#)

include/ktap_errmsg.h - ktap

Macros defined

- [ERRDEF](#)

Source code

```
1 /*
2  * VM error messages.
3  * Copyright (C) 2012-2014 Jovi Zhangwei <jovi.zhangwei@gmail.com>.
4  * Copyright (C) 2005-2014 Mike Pall.
5  */
6
7 /* Basic error handling. */
8 ERRDEF(ERRMEM, "not enough memory")
9 ERRDEF(ERRERR, "error in error handling")
10
11 /* Allocations. */
12 ERRDEF(STROV, "string length overflow")
13 ERRDEF(UDATAOV, "userdata length overflow")
14 ERRDEF(STKOV, "stack overflow")
15 ERRDEF(STKOV, "stack overflow (%s)")
16 ERRDEF(TABOV, "table overflow")
17
18 /* Table indexing. */
19 ERRDEF(NANIDX, "table index is NaN")
20 ERRDEF(NILIDX, "table index is nil")
21 ERRDEF(NEXTIDX, "invalid key to " KTAP\_QL("next"))
22
23 /* Metamethod resolving. */
24 ERRDEF(BADCALL, "attempt to call a %s value")
25 ERRDEF(BADOPRT, "attempt to %s %s " KTAP\_QL("next") " (a %s value)")
26 ERRDEF(BADOPRV, "attempt to %s a %s value")
27 ERRDEF(BADCMP, "attempt to compare %s with %s")
28 ERRDEF(BADCMPV, "attempt to compare two %s values")
29 ERRDEF(GETLOOP, "loop in gettable")
30 ERRDEF(SETLOOP, "loop in settable")
31 ERRDEF(OPCALL, "call")
32 ERRDEF(OPINDEX, "index")
33 ERRDEF(OPARITH, "perform arithmetic on")
34 ERRDEF(OPCAT, "concatenate")
35 ERRDEF(OPLEN, "get length of")
36
37 /* Type checks. */
38 ERRDEF(BADSELF, "calling " KTAP\_QL("next") " on bad self (%s)")
39 ERRDEF(BADARG, "bad argument #%d to " KTAP\_QL("next") " (%s)")
40 ERRDEF(BADTYPE, "%s expected, got %s")
41 ERRDEF(BADVAL, "invalid value")
42 ERRDEF(NOVAL, "value expected")
43 ERRDEF(NOCORO, "coroutine expected")
44 ERRDEF(NOTABN, "nil or table expected")
45 ERRDEF(NOLFUNC, "ktap function expected")
46 ERRDEF(NOFUNCL, "function or level expected")
47 ERRDEF(NOSFT, "string/function/table expected")
48 ERRDEF(NOPROXY, "boolean or proxy expected")
49 ERRDEF(FORINIT, KTAP\_QL("for") " initial value must be a number")
50 ERRDEF(FORLIM, KTAP\_QL("for") " limit must be a number")
51 ERRDEF(FORSTEP, KTAP\_QL("for") " step must be a number")
52
53 /* C API checks. */
54 ERRDEF(NOENV, "no calling environment")
55 ERRDEF(CYIELD, "attempt to yield across C-call boundary")
56 ERRDEF(BADLU, "bad light userdata pointer")
57
58 /* Standard library function errors. */
59 ERRDEF(ASSERT, "assertion failed!")
60 ERRDEF(PROTMT, "cannot change a protected metatable")
61 ERRDEF(UNPACK, "too many results to unpack")
62 ERRDEF(RDRSTR, "reader function must return a string")
```

```

63 ERRDEF(PRTOSTR, KTAP\_QL("tostring") " must return a string to " KTAP\_QL("print"))
64 ERRDEF(IDXRNG, "index out of range")
65 ERRDEF(BASERNG, "base out of range")
66 ERRDEF(LVLRNG, "level out of range")
67 ERRDEF(INVLVL, "invalid level")
68 ERRDEF(INVOPT, "invalid option")
69 ERRDEF(INVOPTM, "invalid option " KTAP\_QS)
70 ERRDEF(INVFMT, "invalid format")
71 ERRDEF(SETFENV, KTAP\_QL("setfenv") " cannot change environment of given object")
72 ERRDEF(CORUN, "cannot resume running coroutine")
73 ERRDEF(CODEAD, "cannot resume dead coroutine")
74 ERRDEF(COSUSP, "cannot resume non-suspended coroutine")
75 ERRDEF(TABINS, "wrong number of arguments to " KTAP\_QL("insert"))
76 ERRDEF(TABCAT, "invalid value (%s) at index %d in table for " KTAP\_QL("concat"))
77 ERRDEF(TABSORT, "invalid order function for sorting")
78 ERRDEF(IOCLFL, "attempt to use a closed file")
79 ERRDEF(IOSTDCL, "standard file is closed")
80 ERRDEF(OSUNIQF, "unable to generate a unique filename")
81 ERRDEF(OSDATEF, "field " KTAP\_QS " missing in date table")
82 ERRDEF(STRDUMP, "unable to dump given function")
83 ERRDEF(STRSLC, "string slice too long")
84 ERRDEF(STRPATB, "missing " KTAP\_QL("[" " after " KTAP\_QL("%" " in pattern")
85 ERRDEF(STRPATC, "invalid pattern capture")
86 ERRDEF(STRPATE, "malformed pattern (ends with " KTAP\_QL("%" ")")
87 ERRDEF(STRPATM, "malformed pattern (missing " KTAP\_QL("]" ")")
88 ERRDEF(STRPATU, "unbalanced pattern")
89 ERRDEF(STRPATX, "pattern too complex")
90 ERRDEF(STRCAPI, "invalid capture index")
91 ERRDEF(STRCAPN, "too many captures")
92 ERRDEF(STRCAPU, "unfinished capture")
93 ERRDEF(STRFMT, "invalid option " KTAP\_QS " to " KTAP\_QL("format"))
94 ERRDEF(STRGSRV, "invalid replacement value (a %s)")
95 ERRDEF(BADMODN, "name conflict for module " KTAP\_QS)
96 ERRDEF(JITOPT, "unknown or malformed optimization flag " KTAP\_QS)
97
98 /* Lexer/parser errors. */
99 ERRDEF(XMODE, "attempt to load chunk with wrong mode")
100 ERRDEF(XNEAR, "%s near " KTAP\_QS)
101 ERRDEF(XLINES, "chunk has too many lines")
102 ERRDEF(XLEVELS, "chunk has too many syntax levels")
103 ERRDEF(XNUMBER, "malformed number")
104 ERRDEF(XLSTR, "unfinished long string")
105 ERRDEF(XLCOM, "unfinished long comment")
106 ERRDEF(XSTR, "unfinished string")
107 ERRDEF(XESC, "invalid escape sequence")
108 ERRDEF(XLDELIM, "invalid long string delimiter")
109 ERRDEF(XTOKEN, KTAP\_QS " expected")
110 ERRDEF(XJUMP, "control structure too long")
111 ERRDEF(XSLOTS, "function or expression too complex")
112 ERRDEF(XLIMC, "chunk has more than %d local variables")
113 ERRDEF(XLIMM, "main function has more than %d %s")
114 ERRDEF(XLIMF, "function at line %d has more than %d %s")
115 ERRDEF(XMATCH, KTAP\_QS " expected (to close " KTAP\_QS " at line %d)")
116 ERRDEF(XFIXUP, "function too long for return fixup")
117 ERRDEF(XPARAM, "<name> or " KTAP\_QL("...") " expected")
118 ERRDEF(XAMBIG, "ambiguous syntax (function call x new statement)")
119 ERRDEF(XFUNARG, "function arguments expected")
120 ERRDEF(XSYMBOL, "unexpected symbol")
121 ERRDEF(XDOTS, "cannot use " KTAP\_QL("...") " outside a vararg function")
122 ERRDEF(XSYNTAX, "syntax error")
123 ERRDEF(XFOR, KTAP\_QL("=") " or " KTAP\_QL("in") " expected")
124 ERRDEF(XBREAK, "no loop to break")
125 ERRDEF(XLUNDEF, "undefined label " KTAP\_QS)
126 ERRDEF(XLDUP, "duplicate label " KTAP\_QS)
127 ERRDEF(XGSCOPE, "<goto %s> jumps into the scope of local " KTAP\_QS)
128 ERRDEF(XEVENTDEF, "cannot parse eventdef " KTAP\_QS)
129
130 /* Bytecode reader errors. */
131 ERRDEF(BCFMT, "cannot load incompatible bytecode")
132 ERRDEF(BCBAD, "cannot load malformed bytecode")
133
134 #undef ERRDEF
135

```

ktapvm.mod.c - ktap

Global variables defined

- [__this_module](#)
- [__used](#)

Source code

```
1 #include <linux/module.h>
2 #include <linux/vermagic.h>
3 #include <linux/compiler.h>
4
5 MODULE_INFO(vermagic, VERMAGIC_STRING);
6
7 struct module __this_module
8 __attribute__((section(".gnu.linkonce.this_module"))) = {
9     .name = KBUILD_MODNAME,
10    .init = init_module,
11    #ifdef CONFIG_MODULE_UNLOAD
12    .exit = cleanup_module,
13    #endif
14    .arch = MODULE_ARCH_INIT,
15 };
16
17 static const char __module_depends[]
18 __used
19 __attribute__((section(".modinfo"))) =
20 "depends=";
21
```

[One Level Up](#)

[Top Level](#)

vim/ - ktap

- [ftdetect/](#)
- [syntax/](#)

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

vim/ftdetect/ - ktap

- [ktap.vim](#)

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

vim/ftdetect/ktap.vim - ktap

```
1 augroup filetype
2   au BufNewFile,BufRead *.kp    set filetype=ktap
3 augroup end
```

[One Level Up](#)

[Top Level](#)

[One Level Up](#)

[Top Level](#)

vim/syntax/ - ktap

- [ktap.vim](#)

[One Level Up](#)

[Top Level](#)

vim/syntax/ktap.vim - ktap

Global variables defined

- [b:current_syntax](#)
- [did_lua_syntax_inits](#)

Source code

```
1  " Vim syntax file
2  " Language:      ktap
3  " Maintainer:   Jovi Zhangwei <jovi.zhangwei@gmail.com>
4  " First Author: Jovi Zhangwei <jovi.zhangwei@gmail.com>
5  " Last Change:  2013 Dec 19
6
7  " For version 5.x: Clear all syntax items
8  " For version 6.x: Quit when a syntax file was already loaded
9  if version < 600
10     syn clear
11     elseif exists("b:current_syntax")
12         finish
13     endif
14
15     setlocal iskeyword=@,48-57,_,$,
16
17     syn keyword ktapStatement break continue return
18     syn keyword ktapRepeat while for in
19     syn keyword ktapConditional if else elseif
20     syn keyword ktapDeclaration trace trace_end
21     syn keyword ktapIdentifier var
22     syn keyword ktapFunction function
23     syn match   ktapBraces "[{}]\[\]"
24     syn match   ktapParens "[()]"
25     syn keyword ktapReserved argstr probename arg0 arg1 arg2 arg3 arg4 arg5 arg6 arg7 arg8 arg9
26     syn keyword ktapReserved cpu pid tid uid execname
27
28
29     syn region ktapTraceDec start="\<trace\>"lc=5 end="{me=s-1 contains=ktapString,ktapNumber
30     syn region ktapTraceDec start="\<trace_end\>"lc=9 end="{me=s-1 contains=ktapString,ktapNumber
31     syn match ktapTrace contained "\<lw\+\>" containedin=ktapTraceDec
32
33     syn region ktapFuncDec start="\<function\>"lc=8 end=":\|("me=s-1 contains=ktapString,ktapNumber
34     syn match ktapFuncCall contained "\<lw\+\ze\(\s\|\n\)*("
35     syn match ktapFunc contained "\<lw\+\>" containedin=ktapFuncDec,ktapFuncCall
36
37     syn match ktapStat contained "@\<lw\+\ze\(\s\|\n\)*("
38
39     " decimal number
40     syn match ktapNumber "\<d\+\>"
41     " octal number
42     syn match ktapNumber "\<0\o\+\>" contains=ktapOctalZero
43     " Flag the first zero of an octal number as something special
44     syn match ktapOctalZero contained "\<0"
45     " flag an octal number with wrong digits
46     syn match ktapOctalError "\<0\o*\[89]\d*"
47     " hex number
48     syn match ktapNumber "\<0x\x\+\>"
49     " numeric arguments
50     syn match ktapNumber "\<\$d\+\>"
51     syn match ktapNumber "\<\$#"
52
53     syn region ktapString oneline start="+ " skip="+ \" end="+ "
54     " string arguments
55     syn match ktapString "@d\+\>"
56     syn match ktapString "@#"
57     syn region ktapString2 matchgroup=ktapString start="\[\z(=*\\)\[" end="]\z1\]" contains=@Spell
58
59     " syn keyword ktapTodo contained TODO FIXME XXX
60
```

```

61 syn match ktapComment "#.*"
62
63 " treat ^#! as special
64 syn match ktapSharpBang "^#!.*"
65
66
67 syn keyword ktapFunc printf print print_hist stack
68 syn keyword ktapFunc gettimeofday_us
69 syn keyword ktapFunc pairs
70
71 syn match ktapFunc /\<ffi\.cdef\>/
72 syn match ktapFunc /\<ffi\.new\>/
73 syn match ktapFunc /\<ffi\.free\>/
74 syn match ktapFunc /\<ffi\.C\>/
75
76
77
78 " Define the default highlighting.
79 " For version 5.7 and earlier: only when not done already
80 " For version 5.8 and later: only when an item doesn't have highlighting yet
81 if version >= 508 || !exists("did_lua_syntax_inits")
82   if version < 508
83     let did_lua_syntax_inits = 1
84     command -nargs=+ HiLink hi link <args>
85   else
86     command -nargs=+ HiLink hi def link <args>
87   endif
88
89   HiLink ktapNumber      Number
90   HiLink ktapOctalZero  PreProc " c.vim does it this way...
91   HiLink ktapOctalError Error
92   HiLink ktapString     String
93   HiLink ktapString2    String
94   HiLink ktapTodo       Todo
95   HiLink ktapComment    Comment
96   HiLink ktapSharpBang  PreProc
97   HiLink ktapStatement  Statement
98   HiLink ktapConditional Conditional
99   HiLink ktapRepeat     Repeat
100  HiLink ktapTrace       Function
101  HiLink ktapFunc        Function
102  HiLink ktapStat        Function
103  HiLink ktapFunction    Function
104  HiLink ktapBraces      Function
105  HiLink ktapDeclaration Typedef
106  HiLink ktapIdentifier  Identifier
107  HiLink ktapReserved    Keyword
108
109  delcommand HiLink
110 endif
111
112 let b:current_syntax = "ktap"

```

[One Level Up](#)

[Top Level](#)