

CHUGENS, CHUBGRAPHS, CHUGINS: 3 TIERS FOR EXTENDING CHUCK

Spencer Salazar

Ge Wang

Center for Computer Research in Music and Acoustics
Stanford University
{spencer, ge}@ccrma.stanford.edu

ABSTRACT

The Chuck programming language lacks straightforward mechanisms for extension beyond its built-in programming and processing facilities. Chugens address this issue by allowing programmers to live-code new unit generators in Chuck in real-time. Chubgraphs also allow new unit generators to be built in Chuck, by defining specific arrangements of existing unit generators. ChuGins allow a wide array of high-performance unit generators and general functionality to be exposed in Chuck by providing a dynamic binding between Chuck and native C/C++-based compiled code. Performance and code analysis shows that the most suitable approach for extending Chuck is situation-dependent.

1. INTRODUCTION

Since its introduction, the Chuck programming language [14] has become a popular tool for computer music composers, educators, and application software developers. However, to date, its catalogue of audio processing unit generators and general programming functionality has been largely limited to those that are built-in when the Chuck binary executable is compiled. Adding new unit generators mandates recompilation of the entirety of Chuck, requiring a level of expertise and motivation reserved for an elite group of power-users. Furthermore, community-based development efforts are hampered by this centralization of functionality, as developers of new unit generators have no easy way to share their work.

The aim of the work described herein is to provide Chuck with multiple levels of extensibility, each essential and appropriate to specific tasks and levels of user expertise. On the one hand, Chuck's pervasive ethos of on-the-fly development creates the desire to design and implement new audio processors in Chuck itself in real-time, working down to the per-sample level if necessary. Furthermore, implementing these components in Chuck allows their use on any operating system Chuck supports with no additional effort from the developer. For these cases we have developed Chugens and Chubgraphs.

On the other hand, real-time performance requirements often mandate the use of compiled native machine code for complex audio-rate processing. There also exists a wealth of C/C++-based software libraries for audio synthesis and effects, such as FluidSynth [4] and Faust [9]. These situations can be straightforwardly handled given portable bindings between Chuck and

native compiled code, which is precisely the intent of ChuGins.

2. RELATED WORK

Extensibility is a primary concern of music software of all varieties. The popular audio programming environments Max/MSP [15], Pure Data [10], SuperCollider [8], and Csound [4] all provide mechanisms for developing C/C++-based compiled sound processing functions. Max also allows control-rate functionality to be encapsulated in-situ in the form of Javascript code snippets. Max's Gen facility dynamically compiles audio-rate processors, implemented as either data-flows or textual code, to native machine code. Csound allows the execution of Tcl, Lua, and Python code for control-rate and/or audio-rate manipulation and synthesis. The Impromptu environment supports live-coding of audio-rate processors in a Scheme-like language [3], and LuaAV does the same for Lua [13]. Pure allows programmers to dynamically write and execute audio signal processors using Faust [5]. A thriving ecosystem revolves around extensions to popular digital audio workstation software, in the form of VST, RTAS, AudioUnits, and LADSPA plugins developed primarily in C and C++. For general-purpose computing, JNI provides a highly flexible binding between native machine code and the Java virtual machine-based run-time environment [7]. RubyGems is a complete plugin and package management tool for both Ruby-based plugins and C/C++ libraries compiled with Ruby bindings [11].

3. CHUGENS, CHUBGRAPHS, CHUGINS

3.1. Chugens

The goal of Chugens (pronounced "chyoo-jen") is to facilitate rapid prototyping of audio synthesis and processing algorithms. Additionally, Chugens provide a basic framework for extending Chuck's built-in audio processing functionality. Using the Chugen system, a programmer can implement sample-rate audio algorithms within the Chuck development environment, utilizing the full array of programming facilities provided by Chuck. These processing units can be naturally integrated into standard Chuck programs, even in the same script file, providing seamless control of audio-rate processing, control-rate manipulation, and higher-level compositional organization.

A Chugen is created first by subclassing the built-in `Chugen` class. This subclass is required to implement a

`tick` function, which accepts a single floating-point argument (the input sample) and returns a single floating-point value (the output sample). For example, this code uses a Chugen to synthesize a sinusoid using the cosine function:

```
class MyCosine extends Chugen
{
  0 => int p;
  440 => float f;
  second/samp => float SRATE;

  fun float tick(float in)
  {
    return Math.cos(p++*2*pi*f/SRATE);
  }
}
```

In the case of an audio synthesizer that does not process an existing signal, the input sample may be ignored.

A Chugen defined so may be integrated into audio graphs like any standard Chuck ugen. Since the `tick` function is just a standard Chuck class member function, it can be as simple or as elaborate as required. Standard library calls, file I/O, multiprocessing (using `spork`), and other general Chuck programming structures can be integrated into the `tick` function and supporting code. For performance reasons, its important to consider that the `tick` function will be called for every sample of audio, so simple `tick` functions will typically perform better. Moreover, the intrinsic overhead of Chuck's virtual machine architecture will cause Chugens to underperform compared to a native C/C++ implementation. Lastly, since Chugens are fundamentally a specialization of a Chuck class, it may define functions to provide structured access to whichever parameters it wishes to expose to the programmer.

3.2. Chubgraphs

Chubgraphs (pronounced "chub-graph") provide a way to construct new unit generators by composition, arranging multiple existing ugens into a single unit. In this way, common arrangements of existing unit generators can be defined and instantiated. Furthermore, Chubgraph parameters can be exposed in a structured manner via class member functions.

A Chubgraph is defined by extending the `Chubgraph` class. The `Chubgraph` class has member variables named `inlet` and `outlet`; `inlet` is a ugen that represents the input signal to the Chubgraph, and `outlet` is the output signal. The Chubgraph's internal audio processing graph is created by spanning a sequence of ugens between `inlet` and `outlet`. The following Chubgraph implements a basic feedback echo processor:

```
class Feedback extends Chubgraph
{
  inlet => Gain dry => outlet;
  dry => Delay delay => outlet;

  delay => Gain feedback => delay;

  0.8 => feedback.gain;
```

```
  1::second => delay.delay;
}
```

(Chubgraphs that do not wish to process an input signal, such as audio synthesizing algorithms, may omit the connection from `inlet`.)

Compared to Chugens, Chubgraphs have obvious performance advantages, as primary audio-rate processing still occurs in the native machine code underlying its component ugens. However Chubgraphs are limited to audio algorithms that can be expressed as combinations of existing unit generators; implementing, for example, intricate mathematical formulae or conditional logic in the form of a ugen graph is possible but, in our experience, fraught with hazard.

3.3. ChuGins

ChuGins (pronounced "chug-in") allow near limitless possibilities for expansion of Chuck's capabilities. A ChuGin is a distributable dynamic library, typically written in C or C++ compiled to native machine code, which Chuck can be instructed to load at runtime. When loaded, the ChuGin defines one or more classes that are subsequently available to Chuck programs. These classes may define new unit generators or provide general programmatic functionality beyond that built in to Chuck. Since these classes are normal Chuck classes implemented with native code, member functions and variables can be used to provide an interface to control parameters.

ChuGins are best suited for audio algorithms that are reasonably well understood and stand to gain from the performance of compiled machine code. The "write-compile-run" development cycle and C/C++-based programming mandated by ChuGins makes implementation of audio processors require comparatively more effort than the Chubgraph or Chugen approaches. However for ugens a programmer intends to use over an extended period of time, the effort to implement a ChuGin will quickly pay off in the form of lower CPU usage.

An additional advantage of ChuGins is that they may provide functionality far outside the intrinsic capabilities of Chuck. Complex synthesis C/C++ based synthesis packages can be imported wholesale into Chuck, opening up an abundance of sonic possibilities. For example, ChuGins have been implemented to bring audio processing programs from the Faust programming language into Chuck. Similarly, the SoundFont renderer FluidSynth has been packaged as a ChuGin. This functionality is not limited to audio processing; a serial port input/output ChuGin is under development, as are other general purpose programming libraries.

Development of a ChuGin is somewhat more complex than Chubgraphs or Chugens, and does not lend itself to explicit presentation of code herein. Using a set of convenience macros, a ChuGin developer first defines a `query` function, which Chuck calls upon first loading the ChuGin. Chuck provides the `query` function with routines with which to define unit generators and classes, specify what member variables and functions

are associated with these, and indicate a `tick` function in the case of unit generators. These functions are then defined in C/C++, using predefined macros for interactions with the upper-level ChuckK runtime, such as retrieving function arguments, getting and setting member variables, and handling input/output samples. This code is then compiled into a dynamic library using the standard facilities for doing so on the target computing platform (`gcc` for Mac OS X and Linux systems, Visual C++ for Windows systems). Additional C/C++ code or libraries, such as STK, may be compiled into the dynamic library using the mechanisms standard for those operations on the target platform. Furthermore, Faust DSP code can be automatically converted to ChuGin form using FaucK, a Faust architecture file and an accompanying compilation script.

4. PERFORMANCE CASE STUDIES

Execution speed of audio software is typically a vital metric for computer musicians, as real-time audio synthesis requires timely production of tens of thousands of samples per second. To evaluate the performance of our mechanisms for extending ChuckK, we designed and implemented several reference unit generators using each method. For each method, we measured the time required to offline-render 5 minutes of audio using the resulting unit generator.

In many situations not only is CPU time at a premium, but a programmer's time to implement a specified application is also limited. Therefore, often "less complex" programs are more desirable than "more complex" programs because they can be developed faster and be better understood by other programmers. While quantifying code complexity is a nuanced and volatile field of inquiry, we have chosen to represent it in the form of lines of code required to implement the unit generator using each method. In addition to characterizing code complexity, lines of code gives an approximation of much time the extension may have taken to develop.

To construct an overall picture of performance vs. development effort, we have created a composite statistic formed by the product of these two metrics, named "aptness." A lower aptness indicates that a particular technique is more appropriate for a particular processing application than a technique with a higher aptness.

4.1. CombFilter

For this case study we implemented a simple comb filter, the "plucked string filter" described in [12]. For reference, the comparatively concise Chubgraph version is as follows:

```
public class CombFilter extends Chubgraph
{
    inlet => Delay d => outlet;
    d => OneZero oz => d;

    -1 => oz.zero;
    100::samp => d.delay;
    0.75 => oz.gain;
}
```

}

As can be seen in Table 1, the ChuGin variant of CombFilter performs best, but the Chubgraph version is not far behind. The Chugen version performs poorly as a result of requiring execution of Chuck code at audio-rate, and also requires more code than the Chubgraph, as the delay line and moving average filter needed to be implemented from scratch. Given the vastly superior lines of code and aptness metrics, the Chubgraph is probably the best approach to use in this case. Chubgraph excels here because CombFilter is essentially a particular arrangement of existing unit generator primitives, i.e. Delay and OneZero, allowing a concise and efficient implementation. A ChuGin may be more appropriate if speed is the utmost concern, but its inferior code complexity makes it largely undesirable in this class of audio processing algorithm.

CombFilter	Time (s)	Lines of Code	Aptness
Chugen	10.944	18	197
Chubgraph	2.841	9	26
ChuGin	2.001	65	130

Table 1. Performance and complexity measurements for CombFilter.

4.2. Bitcrusher

This audio processor performs sample rate reduction by decimation and destructive sample-width compression to recreate its input at a specific sample rate and sample width. The result typically sounds as if it was synthesized by vintage, low-resolution audio or video game hardware such as the Casiotone family of consumer keyboards or the Nintendo Entertainment System. For reference, the shortest version, a Chugen, is listed here:

```
public class Bitcrusher extends Chugen
{
    8 => int bits;
    4 => int downsample;
    float sample;
    int count;
    Math.pow(2,31) => float fINT_MAX;

    fun float tick(float in)
    {
        if(count++ % downsample == 0)
            Math.min(1,Math.max(-1,in))=> sample;

        (sample * fINT_MAX) $ int => int q32;
        32-bits => int shift;
        ((q32 >> shift) << shift) => q32;

        return q32 / fINT_MAX;
    }
}
```

The results in Table 2 once again show the ChuGin as the performance leader. Chugen is the most concise, and, unsurprisingly, the poorest performing contender. Chubgraph performs better than Chugen, but much worse than ChuGin, and is far more complex in this case. The built-in unit generators at the disposal of

Chubgraphs are ill suited to tasks involving conditional logic and intricate non-linear arithmetic manipulation. Therefore, this algorithm would be ideal to develop and prototype as a Chugen. If a computer musician finds him or herself using this Chugen often, or requires faster execution time, reimplementing as a ChuGin would be appropriate.

Bitcrusher	Time (s)	Lines of Code	Aptness
Chugen	10.600	21	223
Chubgraph	6.544	41	268
ChuGin	1.965	84	165

Table 2. Performance and complexity measurements for Bitcrusher.

5. FUTURE WORK

Further developments we are pursuing include mechanisms for creating general purpose programming libraries in ChuckK. Additionally we are investigating a unified ChuGin repository and distribution system, similar to Debian APT [1] or RubyGems, to simplify the task of finding and installing ChuckK extensions.

Our current system of importing ChuGins does not scale well, as ChuGins cannot be loaded on demand. Rather, ChuckK will load every ChuGin that is installed on a system, which may take a noticeable amount of time if there are many ChuGins present. In the future we wish to more intelligently load ChuGins only when their component ugens and classes are invoked by active ChuckK code.

Chugens generate audio by executing ChuckK virtual machine code at audio-rate; dynamically compiling Chugens instead to native machine code may introduce significant performance gains. Environments such as LuaAV, Gen, and Pure do just this, compiling high-level programs to native code for audio-rate execution.

Additional improvements may be required in the form of process and memory safety in ChuGins. Currently, ChuGins execute in the same process space as ChuckK itself, which means that buggy ChuGins can crash ChuckK outright. Technological approaches can alleviate these problems somewhat, but organizational solutions may also be desirable, such as having registries of vetted and thoroughly tested ChuGins.

6. ACKNOWLEDGEMENTS

Special thanks to Kassen Oud, Casper Schipper, Jorge Herrera, and Hongchan Choi for their invaluable beta testing, bug finding, and feature requesting. Additional support for this research was made possible by a National Science Foundation Creative IT Grant, No. IIS-0855758.

7. REFERENCES

[1] Apt. <http://wiki.debian.org/Apt>. Accessed February 10, 2012.

[2] Cook, P., and Scavone, G. 1999. "STK: The Synthesis Toolkit." In *Proceedings of the International Computer Music Conference*. Beijing, China.

[3] Brown, A. R. and Sorenson, A. 2007. "Dynamic Media Arts Programming in impromptu." In *Proceedings of the 6th ACM SIGCHI conference on Creativity & Cognition*. Washington, DC, USA.

[4] FluidSynth. <http://www.fluidsynth.org/>. Accessed February 11, 2012.

[5] Gräf, A. 2011. "Functional Signal Processing with Pure and Faust Using the LLVM Toolkit." In *Proceedings of the Sound and Music Computing Conference*. Padua, Italy.

[6] Lazzarini, V. 2005. "Extensions to the Csound Language: from User- Defined to Plugin-Opcodes and Beyond." In *Proceedings of the Linux Audio Conference*. Karlsruhe, Germany.

[7] Liang, S. 1996. *The Java Native Interface: Programmers Guide and Specification*. Addison-Wesley, Reading, MA.

[8] McCartney, J. 2002. "Rethinking the Computer Music Language: SuperCollider," *Computer Music Journal*, 26(4): 61-68.

[9] Orlarey, Y., Foer, D., and Letz, S. 2009. "Faust: an Efficient Functional Approach to DSP Programming," *New Computational Paradigms for Computer Music*. Edition Delatour, France.

[10] Puckette, M. S. 1997. "Pure data." In *Proceedings of the International Computer Music Conference*. Thessaloniki, Greece.

[11] RubyGems. <http://rubygems.org/>. Accessed February 10, 2012.

[12] Steiglitz, K. 1996. *Digital Signal Processing Primer*. Addison-Wesley, Reading, MA.

[13] Wakefield, G., Smith, W., and Roberts, C. 2010. "LuaAV: Extensibility and Heterogeneity for Audiovisual Computing." In *Proceedings of the Linux Audio Conference*. Utrecht, the Netherlands.

[14] Wang, G. 2008. *The ChuckK Audio Programming Language: A Strongly-timed, On-the-fly Environmentality*. PhD Thesis. Princeton University.

[15] Zicarelli, D. 1998. "An Extensible Real-Time Signal Processing Environment for MAX." In *Proceedings of the International Computer Music Conference*. Ann Arbor, MI.