

decrease in mental effort, or increase in comprehensibility, appears to motivate the purification of algorithms.

REFERENCES

- [1] N. Bulut and M. H. Halstead, "Impurities found in algorithm implementations," *ACM SIGPLAN Notices*, vol. 9, pp. 9-12, Mar. 1974.
- [2] L. J. Chmura and H. F. Ledgard, *Cobol With Style: Programming Proverbs*. Rochelle Park, NJ: Hayden, 1976.
- [3] O.-J. Dahl, E. W. Dijkstra, and C. A. R. Hoare, *Structured Programming*. New York: Academic, 1972.
- [4] J. L. Elshoff, "Measuring commercial PL/I programs using Halstead's criteria," *ACM SIGPLAN Notices*, vol. 11, pp. 38-46, May 1976.
- [5] M. H. Halstead, "Natural laws controlling algorithm structure?" *ACM SIGPLAN Notices*, vol. 7, pp. 19-26, Feb. 1972.
- [6] —, *Elements of Software Science*. New York: Elsevier, 1977.
- [7] I. D. Hill, R. S. Scowen, and B. A. Wichmann, "Writing algorithms in ALGOL 60," *Software Practice and Experience*, vol. 5, pp. 223-224, July-Sept. 1975.
- [8] B. W. Kernighan and P. J. Plauger, *The Elements of Programming Style*. New York: McGraw-Hill, 1974.
- [9] —, "Programming style: Examples and counterexamples," *ACM Computing Surveys*, vol. 6, pp. 303-319, Dec. 1974.
- [10] D. E. Knuth, "A review of 'structured programming,'" *Dep. Comput. Sci., Stanford Univ., Stanford, CA, Tech. Rep. 371*, June 1973.
- [11] D. E. Knuth and R. W. Floyd, "Notes on avoiding GO TO-statements," *Inform. Process. Lett.*, vol. 1, pp. 23-31, Feb. 1971.
- [12] J. M. Yohe, "An overview of programming practices," *ACM Computing Surveys*, vol. 6, pp. 221-245, Dec. 1974.

Ronald D. Gordon, for a photograph and biography, see this issue, p. 90.

Designing Software for Ease of Extension and Contraction

DAVID L. PARNAS

Abstract—Designing software to be extensible and easily contracted is discussed as a special case of design for change. A number of ways that extension and contraction problems manifest themselves in current software are explained. Four steps in the design of software that is more flexible are then discussed. The most critical step is the design of a software structure called the "uses" relation. Some criteria for design decisions are given and illustrated using a small example. It is shown that the identification of *minimal* subsets and *minimal* extensions can lead to software that can be tailored to the needs of a broad variety of users.

Index Terms—Contractibility, extensibility, modularity, software engineering, subsets, supersets.

Manuscript received June 7, 1978; revised October 26, 1978. The earliest work in this paper was supported by NV Phillips Computer Industrie, Apeldoorn, The Netherlands. This work was also supported by the National Science Foundation and the German Federal Ministry for Research and Technology (BMFT). This paper was presented at the Third International Conference on Software Engineering, Atlanta, GA, May 1978.

The author is with the Department of Computer Science, University of North Carolina, Chapel Hill, NC 27514. He is also with the Information Systems Staff, Communications Sciences Division, Naval Research Laboratory, Washington, DC.

I. INTRODUCTION

THIS paper is being written because the following complaints about software systems are so common.

1) "We were behind schedule and wanted to deliver an early release with only a <proper subset of intended capabilities>, but found that that subset would not work until everything worked."

2) "We wanted to add <simple capability>, but to do so would have meant rewriting all or most of the current code."

3) "We wanted to simplify and speed up the system by removing the <unneeded capability>, but to take advantage of this simplification we would have had to rewrite major sections of the code."

4) "Our SYSGEN was intended to allow us to tailor a system to our customers' needs but it was not flexible enough to suit us."

After studying a number of such systems, I have identified some simple concepts that can help programmers to design software so that subsets and extensions are more easily obtained. These concepts are simple if you think about software in the way suggested by this paper. Programmers do not commonly do so.

II. SOFTWARE AS A FAMILY OF PROGRAMS

When we were first taught how to program, we were given a specific problem and told to write one program to do that job. Later we compared our program to others, considering such issues as space and time utilization, but still assuming that we were producing a single product. Even the most recent literature on programming methodology is written on that basis. Dijkstra's *A Discipline of Programming* [1] uses predicate transformers to specify *the* task to be performed by *the* program to be written. The use of the definite article implies that there is a unique problem to be solved and but one program to write.

Today, the software designer should be aware that he is not designing a single program but a family of programs. As discussed in an earlier paper [2], we consider a set of programs to be a program family if they have so much in common that it pays to study their common aspects before looking at the aspects that differentiate them. This rather pragmatic definition does not tell us what pays, but it does explain the motivation for designing program families. We want to exploit the commonalities, share code, and reduce maintenance costs.

Some of the ways that the members of a program family may differ are listed below.

- 1) They may run on different hardware configurations.
- 2) They may perform the same functions but differ in the format of the input and output data.
- 3) They may differ in certain data structures or algorithms because of differences in the available resources.
- 4) They may differ in some data structures or algorithms because of differences in the size of the input data sets or the relative frequency of certain events.
- 5) Some users may require only a subset of the services or features that other users need. These "less demanding" users may demand that they not be forced to pay for the resources consumed by the unneeded features.

Engineers are taught that they must try to anticipate the changes that may be made, and are shown how to achieve designs that can easily be altered when these anticipated changes occur. For example, an electrical engineer will be advised that the world has not standardized the 60-cycle 110-V current. Television designers are fully aware of the differing transmission conventions that exist in the world. It is standard practice to design products that are easily changed in those aspects. Unfortunately, there is no magic technique for handling unanticipated changes. The makers of conventional watches have no difficulty altering a watch that shows the day so that it displays "MER" instead of "WED," but I would expect a long delay for redesign were the world to switch to a ten day week.

Software engineers have not been trained to design for change. The usual programming courses neither mention the need to anticipate changes nor do they offer techniques for designing programs in which changes are easy. Because programs are abstract mathematical objects, the software engineers' techniques for responding to anticipated changes are more subtle and more difficult to grasp than the techniques used by designers of physical objects. Further, we have been led astray by the other designers of abstract objects—mathematicians who state

and prove theorems. When a mathematician becomes aware of the need for a set of closely related theorems, he responds by proving a more general theorem. For mathematicians, a more general result is always superior to a more specialized product. The engineering analogy to the mathematician's approach would be to design television sets containing variable transformers and tuners that are capable of detecting several types of signals. Except for the U.S. armed forces stationed overseas, there is little market for such a product. Few of us consider relocations so likely that we are willing to pay to have the generality present in the product. My guess is that the market for calendar watches for a variable length week is even smaller than the market for the television sets just described.

In [2] I have treated the subject of the design of program families rather generally and in terms of text in a programming language. In this paper I focus on the fifth situation described above; families of programs in which some members are subsets of other family members or several family members share a common subset. I discuss an earlier stage of design, the stage when one identifies the major components of the system and defines relations between those components. We focus on this early stage because the problems described in the introduction result from failure to consider early design decisions carefully.

III. HOW DOES THE LACK OF SUBSETS AND EXTENSIONS MANIFEST ITSELF?

Although we often speak of programs that are "not subsettable" or "not extensible," we must recognize that phrase as inaccurate. It is always possible to remove code from a program and have a runnable result. Any software system can be extended (TSO proves that). The problem is that the subsets and extensions are not the programs that we would have designed if we had set out to design just that product. Further, the amount of work needed to obtain the product seems all out of proportion to the nature of the change. The obstacles commonly encountered in trying to extend or shrink systems fall into four classes.

A. Excessive Information Distribution

A system may be hard to extend or contract if too many programs were written assuming that a given feature is present or not present. This was illustrated by an operating system in which an early design decision was that the system would support three conversational languages. There were many sections of the system where knowledge of this decision was used. For example, error message tables had room for exactly three entries. An extension to allow four languages would have required that a great deal of code be rewritten. More surprisingly, it would have been difficult to reduce the system to one that efficiently supported only two of the languages. One could remove the third language, but to regain the table space, one would have had to rewrite the same sections of code that would be rewritten to add a language.

B. A Chain of Data Transforming Components

Many programs are structured as a chain of components, each receiving data from the previous component, processing it

(and changing the format), before sending the data to the next program in the chain. If one component in this chain is not needed, that code is often hard to remove because the output of its predecessor is not compatible with the input requirements of its successor. A program that does nothing but change the format must be substituted. One illustration would be a payroll program that assumed unsorted input. One of the components of the system accepts the unsorted input and produces output that is sorted by some key. If the firm adopts an office procedure that results in sorted input, this phase of the processing is unnecessary. To eliminate that program, one may have to add a program that transfers data from a file in the input format to a file in the format appropriate for the next phase. It may be almost as efficient to allow the original SORT component to sort the sorted input.

C. Components That Perform More Than One Function

Another common error is to combine two simple functions into one component because the functions seem too simple to separate. For example, one might be tempted to combine synchronization with message sending and acknowledgment in building an operating system. The two functions seem closely related; one might expect that for the sake of reliability one should insist on a "handshake" with each exchange of synchronization signals. If one later encounters an application in which synchronization is needed very frequently, one may find that there is no simple way to strip the message sending out of the synchronization routines. Another example is the inclusion of run-time type-checking in the basic subroutine call mechanism. In applications where compile-time checking or verification eliminates the need for the run-time type-check, another subroutine call mechanism will be needed. The irony of these situations is that the "more powerful" mechanism could have been built separately from, but *using*, simpler mechanisms. Separation would result in a system in which the simpler mechanism was available for use where it sufficed.

D. Loops in the "Uses" Relation

In many software design projects, the decisions about what other component programs to use are left to individual systems programmers. If a programmer knows of a program in another module, and feels that it would be useful in his program, he includes a call on that program in his text. Programmers are encouraged to use the work of other programmers as much as possible because, when each programmer writes his own routines to perform common functions, we end up with a system that is much larger than it need be.

Unfortunately, there are two sides to the question of program usage. Unless some restraint is exercised, one may end up with a system in which nothing works until everything works. For example, while it may seem wise to have an operating system scheduler use the file system to store its data (rather than use its own disk routines), the result will be that the file system must be present and working before any task scheduling is possible. There are users for whom an operating system subset without a file system would be useful. Even if

one has no such users, the subset would be useful during development and testing.

IV. STEPS TOWARDS A BETTER STRUCTURE

This section discusses four parts of a methodology that I believe will help the software engineer to build systems that do not evidence the problems discussed above.

A. Requirements Definition: Identifying the Subsets First

One of the clearest morals in the earlier discussion about "design for change" as it is taught in other areas of engineering is that one must anticipate changes before one begins the design. At a past conference [3] many of the papers exhorted the audience to spend more time identifying the actual requirements before starting on a design. I do not want to repeat such exhortations, but I do want to point out that the identification of the possible subsets is part of identifying the requirements. Treating the easy availability of certain subsets as an operational requirement is especially important to government officials who purchase software. Many officials despair of placing strict controls on the production methods used by their contractors because they are forbidden by law to tell the contractor how to perform his job. They may tell him what they require, but not how to build it. Fortunately, the availability of subsets may be construed as an operational property of the software.

On the other hand, the identification of the required subsets is not a simple matter of asking potential users what they could do without. First, users tend to overstate their requirements. Second, the answer will not characterize the set of subsets that might be wanted in the future. In my experience, identification of the potentially desirable subsets is a demanding intellectual exercise in which one first searches for the *minimal* subset that might conceivably perform a useful service and then searches for a set of *minimal* increments to the system. Each increment is small—sometimes so small that it seems trivial. The emphasis on minimality stems from our desire to avoid components that perform more than one function (as discussed in Section III-C). Identifying the minimal subset is difficult because the minimal system is not usually a program that anyone would ask for. If we are going to build the software family, the minimal subset is useful; it is not usually worth building by itself. Similarly, the maximum flexibility is obtained by looking for the smallest possible increments in capability: often these are smaller increments than a user would think of. Whether or not he would think of them before system development, he is likely to want that flexibility later.

The search for a minimal subset and minimal extensions can best be shown by an example. One example of a minimal subset is given in [4]. Another example will be given later in this paper.

B. Information Hiding: Interface and Module Definition

In an earlier section we touched upon the difference between the mathematician's concept of generality and an engineer's

approach to design flexibility. Where the mathematician wants his product, a theorem or method of proof, to be as general as possible, i.e., applicable, without change, in as many situations as possible, an engineer often must tailor his product to the situation actually at hand. Lack of generality is necessary to make the program as efficient or inexpensive as possible. If he must develop a family of products, he tries to isolate the changeable parts in modules and to develop an interface between the module and the rest of the product that remains valid for all versions. The crucial steps are as follows.

1) Identification of the items that are likely to change. These items are termed "secrets."

2) Location of the specialized components in separate modules.

3) Designing intermodule interfaces that are insensitive to the anticipated changes. The changeable aspects or "secrets" of the modules are not revealed by the interface.

It is exactly this that the concept of information hiding [5], encapsulation, or abstraction [6] is intended to do for software. Because software is an abstract or mathematical product, the modules may not have any easily recognized physical identity. They are not necessarily separately compilable or coincident with memory overlay units. The interface must be general but the contents should not be. Specialization is necessary for economy and efficiency.

The concept of information hiding is very general and is applicable in many software change situations—not just the issue of subsets and extensions that we address in this paper. The ideas have also been extensively discussed in the literature [5]–[9]. The special implications for our problem are simply that, as far as possible, even the presence or absence of a component should be hidden from other components. If one program uses another directly, the presence of the second program cannot be fully hidden from its user. However, there is never any reason for a component to "know" how many other programs use it. All data structures that reveal the presence or number of certain components should be included in separate information hiding modules with abstract interfaces [10]. Space and other considerations make it impossible to discuss this concept further in this paper; it will be illustrated in the example. Readers for whom this concept is new are advised to read some of the articles mentioned above.

C. The Virtual Machine (VM) Concept

To avoid the problems that we have described as "a chain of data transforming components," it is necessary to stop thinking of systems in terms of components that correspond to steps in the processing. This way of thinking dies hard. It is almost certain that your first introduction to programming was in terms of a series of statements intended to be executed in the order that they were explained to you. We are goal oriented; we know what we start with and what we want to produce. It is natural to think in terms of steps progressing towards that goal. It is the fact that we are designing a family of systems that makes this "natural" approach the wrong one.

The viewpoint that seems most appropriate to designing soft-

ware families is often termed the virtual machine approach. Rather than write programs that perform the transformation from input data to output data, we design software machine extensions that will be useful in writing many such programs. Where our hardware machine provides us with a set of instructions that operate on a small set of data types, the extended or virtual machine will have additional data types as well as "software instructions" that operate on those data types. These added features will be tailored to the class of programs that we are building. While the VM instructions are designed to be generally useful, they can be left out of a final product if the user's programs do not use them. The programmer writing programs for the virtual machine should not need to distinguish between instructions that are implemented in software and those that are hardware implemented. To achieve a true virtual machine, the hardware resources that are used in implementing the extended instruction set must be unavailable to the user of the virtual machine. The designer has traded these resources for the new data elements and instructions. Any attempt to use those resources again will invalidate the concept of virtual machine and lead to complications. Failure to provide for isolation of resources is one of the reasons for the failure of some attempts to use macros to provide a virtual machine. The macro user must be careful not to use the resources used in the code generated by the macros.

There is no reason to accomplish the transformation from the hardware machine to a virtual machine with all of the desired features in a single leap. Instead we will use the machine at hand to implement a few new instructions. At each step we take advantage of the newly introduced features. Such a step-by-step approach turns a large problem into a set of small ones and, as we will see later, eases the problem of finding the appropriate subsets. Each element in this series of virtual machines is a useful subset of the system.

D. Designing the "Uses" Structure

The concept of an abstract machine is an intuitive way of thinking about design. A precise description of the concept comes through a discussion of the relation "uses" [11], [12].

1) *The relation "uses"*: We consider a system to be divided into a set of programs that can be invoked either by the normal flow of control mechanisms, by an interrupt, or by an exception handling mechanism. Each of these programs is assumed to have a specification that defines exactly the effect that an invocation of the program should have.

We say of two programs A and B that A *uses* B if correct execution of B may be necessary for A to complete the task described in its specification. That is, A *uses* B if there exist situations in which the correct functioning of A depends upon the availability of a correct implementation of B. Note that to decide whether A *uses* B or not, one must examine both the implementation *and* the specification of A.

The "uses" relation and "invokes" very often coincide, but *uses* differs from *invokes* in two ways:

a) Certain invocations may not be instances of "uses." If A's specification requires only that A *invoke* B when certain

conditions occur, then A has fulfilled its specification when it has generated a correct call to B. A is correct even if B is incorrect or absent. A proof of correctness of A need only make assumptions about the way to invoke B.

b) A program A may use B even though it never invokes it. The best illustration of this is interrupt handling. Most programs in a computer system are only correct on the assumption that the interrupt handling routine will correctly handle the interrupts (leave the processor in an acceptable state). Such programs use the interrupt handling routines even though they never call them. "Uses" can also be formulated as "requires the presence of a correct version of."

Systems that have achieved a certain "elegance" (e.g., T.H.E. [5], Venus [6]) have done so by having parts of the system "use" other parts in such a way that the "user" programs were simplified. For example, the transport stream mechanism in T.H.E. uses the segmenting mechanism to great advantage. In contrast, many large and complex operating systems achieve their size and complexity by having "independent" parts. For example, there are many systems in which "spooling," virtual memory management, and the file system all perform their own backup store operations. Code to perform these functions is present in each of the components. Whenever such components must share a single device, complex interfaces exist.

The disadvantage of unrestrained "usage" of each others facilities is that the system parts become highly interdependent. Often there are no subsets of the system that can be used before the whole system is complete. In practice, some duplication of effort seems preferable to a system in which nothing runs unless everything runs.

2) *The uses hierarchy*: By restricting the relation "uses" so that its graph is loop free we can retain the primary advantages of having system parts "use" each other while eliminating the problems. In that case it is possible to assign the programs to the levels of a hierarchy by the following rules:

- a) level 0 is the set of all programs that use no other program;
- b) level i ($i \geq 1$) is the set of all programs that use at least one program on level $i - 1$ and no program at a level higher than $i - 1$.

If such a hierarchical ordering exists, then each level offers a testable and usable subset of the system. In fact, one can get additional subsets by including only parts of a level. The easy availability of these subsets is very valuable for the construction of any software systems and is vital for developing a broad family of systems.

The design of the "uses" hierarchy should be one of the major milestones in a design effort. The division of the system into independently callable subprograms has to go on in parallel with the decisions about uses, because they influence each other.

3) *The criteria to be used in allowing one program to use another*: We propose to allow A "uses" B when all of the following conditions hold:

- a) A is essentially simpler because it uses B;
- b) B is not substantially more complex because it is not allowed to use A;

- c) there is a useful subset containing B and not A;
- d) there is no conceivably useful subset containing A but not B.

During the process of designing the "uses" relation, we often find ourselves in a situation where two programs could obviously benefit from using each other and the conditions above cannot be satisfied. In such situations, we resolve the apparent conflicts by a technique that we call "sandwiching." One of the programs is "sliced" into two parts in a way that allows the programs to "use" each other and still satisfy the above conditions. If we find ourselves in a position where A would benefit from using B, but B can also benefit from using A, we may split B into two programs: B1 and B2. We then allow A to use B2 and B1 to use A. The result would appear to be a sandwich with B as the bread and A as the filling. Often, we then go on to split A. We start with a few levels and end up with many.

An earlier report [11] introduced many of the ideas that are in this paper and illustrated them by proposing a "uses" relation for a family of operating systems. It contains several examples of situations where "sandwiching" led us from a "T.H.E.-like structure" [14] to a structure with more than twice as many levels. For example, the virtual memory mechanism was split into address translation and dynamic allocation of memory areas to segments.

The most frequent instances of splitting and sandwiching came because initially we were assuming that a "level" would be a "module" in the sense of Section IV-B. We will discuss this in the final part of this paper.

4) *Use of the word "convenience"*: It will trouble some readers that it is usual to use the word "convenience" to describe a reason for introducing a certain facility at a given level of the hierarchy. A more substantial basis would seem more scientific.

As discussed in [11] and [13], we must assume that the hardware itself is capable of performing all necessary functions. As one goes higher in the levels, one can lose capabilities (as resources are consumed)—not gain them. On the other hand, at the higher levels the new functions can be implemented with simpler programs because of the additional programs that can be used. We speak of "convenience" to make it clear that one could implement any functions on a lower level, but the availability of the additional programs at the higher level is useful. For each function we give the lowest level at which the features that are useful for implementing that function (with the stated restrictions) are available. In each case, we see no functions available at the next higher level that would be useful for implementing the functions as described. If we implemented the program one level lower we would have to duplicate programs that become available at that level.

V. EXAMPLE: AN ADDRESS PROCESSING SUBSYSTEM

As an example of designing for extensibility and subsets, we consider a set of programs to read in, store, and write out lists of addresses. This example has also been used, to illustrate a different point, in [10] and has been used in several classroom experiments to demonstrate module interchangeability. This

The following items of information will be found in the addresses to be processed and constitute the only items of relevance to the application programs:

- Last name
- Given names (first name and possible middle names)
- Organization (Command or Activity)
- Internal identifier (Branch or Code)
- Street address or P.O. box
- City or mail unit identifier
- State
- Zip code
- Title
- Branch of service if military
- GS grade if civil service

Each of the above will be strings of characters in the standard ANSI alphabet, and each of the above may be empty or blank.

Fig. 1.

example is intended as an integral part of this paper; several statements in the final summation are supported only in this section.

A. Our Basic Assumptions

- 1) The information items discussed in Fig. 1 will be the items to be processed by all application programs.
- 2) The input formats of the addresses are subject to change.
- 3) The output formats of the addresses are subject to change.
- 4) Some systems will use a single fixed format for input and output. Other systems will need the ability to choose from several input or output formats at run-time. Some systems will be required in which the user can specify the format using a format definition language.
- 5) The representation of addresses in main storage will vary from system to system.
- 6) In most systems, only a subset of the total set of addresses stored in the system need be in main storage at any one time. The number of addresses needed may vary from system to system, and in some systems the number of addresses to be kept in main memory may vary at run-time.

B. We Propose the Following Design Decisions

- 1) The input and output programs will be table driven: the table will specify the format to be used for input and output. The contents and organization of these format tables will be the "secrets" of the input and output modules.
- 2) The representation of addresses in core will be the "secret" of an address storage module (ASM). The implementation chosen for this module will be such that the operations of changing a portion of an address will be relatively inexpensive, compared to making the address table larger or smaller.
- 3) When the number of addresses to be stored exceeds the capacity of an ASM, programs will use an address file module (AFM). An AFM can be made upward compatible with an ASM; programs that were written to use ASM's could operate using an AFM in the same way. The AFM provides additional

commands to allow more efficient usage by programs that do not assume the random access properties of an ASM. These programs are described below.

4) Our implementation of an AFM would use an ASM as a submodule as well as another submodule that we will call block file module (BFM). The BFM stores blocks of data that are sufficiently large to represent an address, but the BFM is not specialized to the handling of addresses. An ASM that is used within an AFM may be said to have two interfaces. In the "normal interface" that an ASM presents to an outside user, an address is a set of fields and the access functions hide or abstract from the representation. Fig. 2 is a list of the access programs that comprise this interface. In the second interface, the ASM deals with blocks of contiguous storage and abstract from the contents. There are commands for the ASM to input and output "addresses" but the operands are storage blocks whose interpretation as addresses is known only within the ASM. The AFM makes assumptions about the association between blocks and addresses but not about the way that an address's components are represented as blocks. The BFM is completely independent of the fact that the blocks contain address information. The BFM might, in fact, be a manufacturer supplied access method.

C. Component Programs

1) Module: Address Input

- INAD: Reads in an address that is assumed to be in a format specified by a format table and calls ASM or AFM functions to store it.
- INFSL: Selects a format from an existing set of format tables. The selected format is the one that will be used by INAD. There is always a format selected.
- INFCR: Adds a new format to the tables used by INFSL. The format is specified in a "format language." Selection is *not* changed (i.e., INAD still uses the same format table).
- INTABEXT: Adds a blank table to the set of input format tables.
- INTABCHG: Rewrites a table in the input format tables using a description in a format language. Selection is not changed.
- INFDEL: Deletes a table from the set of format tables. The selected format cannot be deleted.
- INADSEL: Reads in an address using one of a set of formats. Choice is specified by an integer parameter.
- INADFO: Reads in an address in a format specified as one of its parameters (a string in the format definition language). The format is selected and added to the tables and subsequent addresses could be read in using INAD.

2) Module: Address Output

- OUTAD: Prints an address in a format specified by a format table. The information to be printed

Access Functions for "Normal Interface"

MODULE: ASM		NAME OF ACCESS PROGRAM*						INPUT PARAMETERS		OUTPUT	
*ADDTIT:	asm	X	integer	X	string	→	asm	*			
ADDGN:	asm	X	integer	X	string	→	asm	*			
ADDLN:	asm	X	integer	X	string	→	asm	*			
ADDSERV:	asm	X	integer	X	string	→	asm	*			
ADDBORC:	asm	X	integer	X	string	→	asm	*			
ADDCORA:	asm	X	integer	X	string	→	asm	*			
ADDSORP:	asm	X	integer	X	string	→	asm	*			
ADDCITY:	asm	X	integer	X	string	→	asm	*			
ADDSTATE:	asm	X	integer	X	string	→	asm	*			
ADDZIP:	asm	X	integer	X	string	→	asm	*			
ADDGSL:	asm	X	integer	X	string	→	asm	*			
SETNUM:	asm	X	integer	→	asm	*					
FETTIT:	asm	X	integer	→	string						
FETGN:	asm	X	integer	→	string						
FETLN:	asm	X	integer	→	string						
FETSERV:	asm	X	integer	→	string						
FETBORC:	asm	X	integer	→	string						
FETCORA:	asm	X	integer	→	string						
FETSORP:	asm	X	integer	→	string						
FETCITY:	asm	X	integer	→	string						
FETSTATE:	asm	X	integer	→	string						
FETZIP:	asm	X	integer	→	string						
FETGSL:	asm	X	integer	→	string						
FETNUM:	asm	→	integer								

*These are abbreviations: ADDTIT = ADD TITLE; ADDGN = ADD GIVEN NAME, etc.

Fig. 2. Syntax of ASM functions.

is assumed to be in an ASM and identified by its position in an ASM.

- OUTFSL: Selects a format table from an existing set of output format tables. The selected format is the one that will be used by OUTAD.
- OUTTABEXT: Adds a "blank" table to the set of output format tables.
- OUTTABCHG: Rewrites the contents of a format table using information in a format language.
- OUTFCR: Adds a new format to the set of formats that can be selected by OUTFSL in a format description language.
- OUTFDEL: Deletes a table from the set of format tables that can be selected by OUTFSL.
- OUTADSEL: Prints out an address using one of a set of formats.
- OUTADFO: Prints out an address in a format specified in a format definition language string, which is one of the actual parameters. The format is added to the tables and selected.

3) Module: Address Storage (ASM)

FET: (Component Name): This is a set of functions used to read information from an address

store. Returns a string as a value. See Fig. 2.

- ADD: (Component Name): This is a set of functions used to write information in an address store. Each takes a string and an integer as parameters. The integer specifies an address within the ASM. See Fig. 2.
- 0BLOCK: Takes an integer parameter, returns a storage block as a value.
- 1BLOCK: Accepts a storage block and integer as parameters. Its effect is to change the contents of an address store—which is reflected by a change in the values of the FET programs.
- ASMEXT: Extends an address store by appending a new address with empty components at the end of the address store.
- ASMSHR: "Shrinks" the address store.
- ASMCR: Creates a new address store. The parameter specifies the number of components. All components are initially empty.
- ASMDL: Deletes an existing address store.

4) Module: Block File Module

BLFET: Accepts an integer as a parameter and returns a "block."

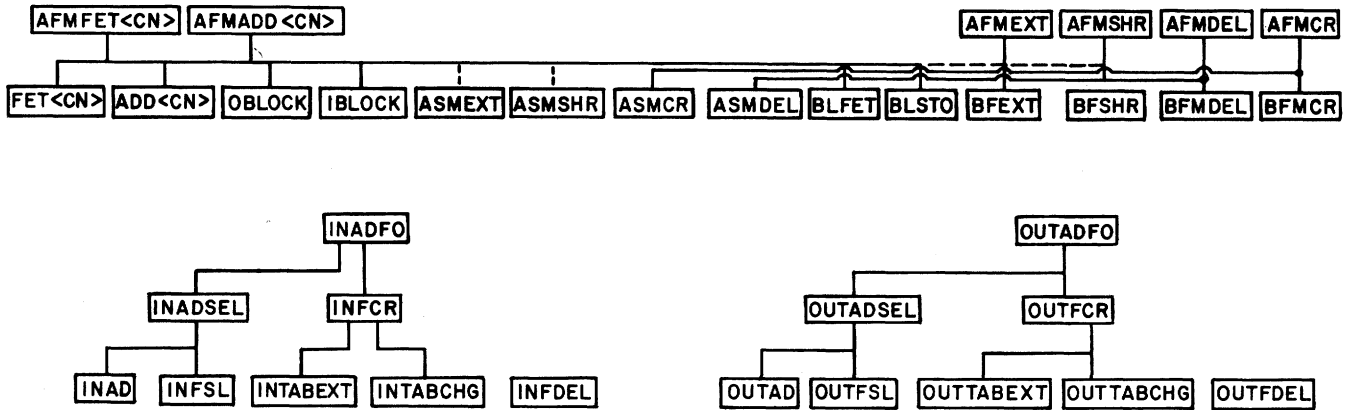


Fig. 3.

- BLSTO: Accepts a block and an integer and stores the block.
- BFEXT: Extends BFM by adding additional blocks to its capacity.
- BFSHR: Reduces the size of the BFM by removing some blocks.
- BFMCR: Creates a file of blocks.
- BFMDEL: Deletes an existing file of blocks.

5) Module: Address File Module

This module includes implementations of all of the ASM functions except 0BLOCK and 1BLOCK. To avoid confusion in the diagram showing the uses hierarchy we have changed the names to:

- AFMADD (Component Name) defined as in Fig. 2
- AFMFET (Component Name) defined as in Fig. 1
- AFMEXT defined as in BFM above
- AFMSHR defined as in BFM above
- AFMCR defined as in BFM above
- AFMDEL defined as in BFM above.

D. Uses Relation

Fig. 3 shows the *uses* relation between the component programs. It is important to note that we are now discussing the implementation of those programs, not just their specifications. The *uses* relation is characterized by the fact that there are a large number of relatively simple, *single-purpose* programs on the lowest level. The upper level programs are implemented by means of these lower level programs so that they too are quite simple. This *uses* relation diagram characterizes the set of possible subsets.

E. Discussion

To pick a subset, one identifies the set of upper level programs that the user needs and includes only those programs that those programs use (directly or indirectly). For example, a user who uses addresses in a single format does not need the component programs that interpret format description lan-

guages. Systems that work with a small set of addresses can be built without any BFM components. A program that works as a query system and never prints out a complete address would not need any Address Output components.

The system is also easily extended. For example, one could add a capability to read in addresses with self-defining files. If the first record on a file was a description of the format in something equivalent to the format description language, one could write a program that would be able to read in that record, use INTABCHG to build a new format table, and then read in the addresses. Programs that do things with addresses (such as print out "personalized" form letters) can also be added using these programs and selecting only those capabilities that they actually need.

One other observation that can be made is that the upper level programs can be used to "generate" lower level versions. For example, the format description languages can be used to generate the tables used for the fixed format versions. There is no need for a separate SYSGEN program.

We will elaborate on this observation in the conclusion.

VI. SOME REMARKS ON OPERATING SYSTEMS: WHY GENERALS ARE SUPERIOR TO COLONELS

An earlier report [11] discusses the design of a "uses" hierarchy for operating systems. Although there have been some refinements to the proposals of that report, its basic contents are consistent with the present proposals. This section compares the approach outlined in this paper and the "kernel" approach or "nucleus" approach to OS design [18]-[20]. It is tempting to say that the suggestions in this paper do not conflict with the "kernel" approach. These proposals can be viewed as a refinement of the nucleus approach. The first few levels of our system could be labeled "kernel," and one could conclude that we are just discussing a fine structure within the kernel.

To yield to that temptation would be to ignore an essential difference between the approaches suggested in this paper and the kernel approach. The system kernels known to me are

such that some desirable subsets cannot be obtained without major surgery. It was assumed that the nucleus must be in every system family member. In the RC4000 system the inability to separate synchronization from message passing has led some users to bypass the kernel to perform teletype handling functions. In Hydra as originally proposed [19], "type checking" was so intrinsic to the call mechanism that it appeared impossible to disable it when it was not needed or affordable.¹

Drawing a line between "kernel" and the rest of the system, and putting "essential" services of "critical programs" in the nucleus yields a system in which kernel features cannot be removed and certain extensions are impractical. Looking for a *minimal* subset and a set of *minimal* independent incremental function leads to a system in which one can trim away unneeded features. I know of no feature that is always needed. When we say that two functions are *almost* always used together, we should remember that "almost" is a euphemism for "not."

VII. SUMMATION

This paper describes an approach to software intended to result in systems that can be tailored to fit the needs of a broad variety of users. The points most worthy of emphasis are as follows.

1) *The Requirements Include Subsets and Extensions:* It is essential to recognize the identification of useable subsets as part of the preliminaries to software design. Flexibility cannot be an afterthought. Subsetability is needed, not just to meet a variety of customers' needs, but to provide a fail-safe way of handling schedule slippage.

2) *Advantages of the Virtual Machine Approach:* Designing software as a set of virtual machines has definite advantages over the conventional (flowchart) approach to system design. The virtual machine "instructions" provide facilities that are useful for purposes beyond those originally conceived. These instructions can easily be omitted from a system if they are not needed. Remove a jammer box from a flowchart and there is often a need to "fill the hole" with conversion programs.

3) *On the Difference Between Software Generality and Software Flexibility:* Software can be considered "general" if it can be used, *without change*, in a variety of situations. Software can be considered flexible, if it is *easily changed* to be used in a variety of situations. It appears unavoidable that there is a run-time cost to be paid for generality. Clever designers can achieve flexibility without significant run-time cost, but there is a design-time cost. One should incur the design-time cost only if one expects to recover it when changes are made.

Some organizations may choose to pay the run-time cost for generality. They build general software rather than flexible software because of the maintenance problems associated with maintaining several different versions. Factors influencing this decision include a) the availability of extra computer resources,

b) the facilities for program change and maintenance available at each installation, and c) the extent to which design techniques ease the task of applying the same change to many versions of a program.

No one can tell a designer how much flexibility and generality should be built into a product, but the decision should be a conscious one. Often, it just happens.

4) *On the Distinction Between Modules, Subprograms, and Levels:* Several systems and at least one dissertation [14]-[17] have, in my opinion, blurred the distinction between modules, subprograms, and levels. Conventional programming techniques consider a subroutine or other callable program to be a module. If one wants the modules to include all programs that must be designed together and changed together, then, as our example illustrates, one will usually include many small subprograms in a single module. If does not matter what word we use; the point is that the unit of change is not a single callable subprogram.

In several systems, modules and levels have coincided [14], [15]. This had led to the phrase "level of abstraction." Each of the modules in the example abstract from some detail that is assumed likely to change. In our approach there is no correspondence between modules and levels. Further, I have not found a relation, "more abstract than," that would allow me to define an abstraction hierarchy [12]. Although I am myself guilty of using it, in most cases the phrase "levels of abstraction" is an abuse of language.

Janson has suggested that a design such as this one (or the one discussed in [11]) contain "soft modules" that can represent a breach of security principles. Obviously an error in any program in one of our modules can violate the integrity of that module. All module programs that will be included in a given subset must be considered in proving the correctness of that module. However, I see no way that allowing the component programs to be on different levels of a "uses" hierarchy makes this process more difficult or makes the system less secure. The boundaries of our modules are quite firm and clearly identified.

The essential difference between this paper and other discussions of hierarchically structured designs is the emphasis on subsets and extensions. My search for a criterion to be used in designing the *uses* hierarchy has convinced me that if one does not care about the existence of subsets, it does not really matter what hierarchy one uses. Any design can be bent until it works. It is only in the ease of change that they differ.

5) *On Avoiding Duplication:* Some earlier work [21] has suggested that one needs to have duplicate or near duplicate modules in a hierarchically structured system. For example, they suggest that one needs one implementation of processes to give a fixed number of processes at a low level and another to provide for a varying number of processes at a user's level. Similar ideas have appeared elsewhere. Were such duplication to be necessary, it would be a sound argument against the use of "structured" approaches. One can avoid such duplication if one allows the programs that vary the size of a data structure to be on a higher level than the other programs that operate on that data structure. For example, in an operating system, the programs to create and delete processes need not be on the

¹Accurate reports on the current status and performance of that system are not available to me.

same level as the more frequently used scheduling operations. In designing software, I regard the need to code similar functions in two separate programs as an indication of a fundamental error in my thinking.

6) *Designing for Subsets and Extensions Can Reduce the Need for Support Software:* We have already mentioned that this design approach can eliminate the need for separate SYSGEN programs. We can also eliminate the need for special-purpose compilers. The price of the convenience features offered by such languages is often a compiler and run-time package distinctly larger than the system being built. In our approach, each level provides a "language extension" available to the programmer of the next level. We never build a compiler; we just build our system, but we get convenience features anyway.

7) *Extension at Run-Time Versus Extension During SYSGEN:* At a later stage in the design we will have to choose data structures and take the difference between run-time extension and SYSGEN extension into consideration. Certain data structures are more easily accessed but harder to extend while the program is running; others are easily extended but at the expense of a higher access cost. These differences do not affect our early design decisions because they are hidden in modules.

8) *On the Value of a Model:* My work on this example and similar ones has gone much faster because I have learned to exploit a pattern that I first noticed in the design discussed in [11]. Low level operations assume the existence of a fixed data structure of some type. The operations on the next level allow the swapping of a data element with others from a fixed set of similar elements. The high level programs allow the creation and deletion of such data elements. This pattern appears several times in both designs. Although I have not designed your system for you, I believe that you can take advantage of a similar pattern. If so, this paper has served its purpose.


ACKNOWLEDGMENT

The ideas presented in this paper have been developed over a lengthy period and with the cooperation and help of many collaborators. I am grateful to numerous Philips employees for thought provoking comments and questions. Price's collaboration was invaluable at Carnegie-Mellon University. The help of W. Bartussek, G. Handzel, and H. Wuerges at the Technische Hochschule Darmstadt led to substantial improvements. Heninger, Weiss, and J. Shore at the Naval Research Laboratory helped me to understand the application of the concepts in areas other than operating systems. B. Trombka and J. Guttg both helped in the design of pilots of the address process system. Discussions with P. J. Courtois have helped me to better understand the relation between software structure and run-time characteristics of computer systems. Dr. E. Britton, H. Rettenmaier, L. Belady, Dr. D. Stanat, G. Fran, and Dr. W. Wright made many helpful suggestions about an earlier draft of this paper. If you find portions of this paper helpful, these people deserve your thanks.

REFERENCES

- [1] E. W. Dijkstra, *A Discipline of Programming*. Englewood Cliffs, NJ: Prentice-Hall, 1976.

- [2] D. L. Parnas, "On the design and development of program families," *IEEE Trans. Software Eng.*, vol. SE-2, pp. 1-9, Mar. 1976.
- [3] 2nd Int. Conf. Software Engineering, Oct. 13-15, 1976; also, *IEEE Trans. Software Eng.*, (Special Issue), vol. SE-2, Dec. 1976.
- [4] D. L. Parnas, G. Handzel, and H. Wuerges, "Design and specification of the minimal subset of an operating system family," presented at the 2nd Int. Conf. Software Engineering, Oct. 13-15, 1976; also, *IEEE Trans. Software Eng.*, (Special Issue), vol. SE-2, pp. 301-307, Dec. 1976.
- [5] D. L. Parnas, "On the criteria to be used in decomposing systems into modules," *Commun. Ass. Comput. Mach.*, Dec. 1972.
- [6] T. A. Linden, "The use of abstract data types to simplify program modifications," in *Proc. Conf. Data: Abstraction, Definition and Structure*, Mar. 22-24, 1976; also, *ACM SIGPLAN Notices* (Special Issue), vol. II, 1976.
- [7] D. L. Parnas, "A technique for software module specification with examples," *Commun. Ass. Comput. Mach.*, May 1972.
- [8] —, "Information distribution aspects of design methodology," in *1971 Proc. IFIP Congr.* Amsterdam, The Netherlands: North-Holland, 1971.
- [9] —, "The use of precise specifications in the development of software," in *1977 Proc. IFIP Congr.* Amsterdam, The Netherlands: North-Holland, 1977.
- [10] —, "Use of abstract interfaces in the development of software for embedded computer systems," Naval Res. Lab., Washington, DC, NRL Rep. 8047, June 1977.
- [11] —, "Some hypotheses about the 'uses' hierarchy for operating systems," Technische Hochschule Darmstadt, Darmstadt, West Germany, Tech. Rep., Mar. 1976.
- [12] —, "On a 'buzzword': Hierarchical structure," in *1974 Proc. IFIP Congr.* Amsterdam, The Netherlands: North-Holland, 1974.
- [13] D. L. Parnas and D. L. Siewiorek, "Use of the concept of transparency in the design of hierarchically structured systems," *Commun. Ass. Comput. Mach.*, vol. 18, July 1975.
- [14] E. W. Dijkstra, "The structure of the 'THE'-multiprogramming system," *Commun. Ass. Comput. Mach.*, vol. 11, pp. 341-346, May 1968.
- [15] B. Liskov, "The design of the Venus operating system," *Commun. Ass. Comput. Mach.*, vol. 15, pp. 144-149, Mar. 1972.
- [16] P. A. Janson, "Using type extension to organize virtual memory mechanisms," Lab. for Comput. Sci., M.I.T., Cambridge, MA, MIT-LCS-TR167, Sept. 1976.
- [17] —, "Using type-extension to organize virtual memory mechanisms," IBM Zurich Res. Lab., Switzerland, Res. Rep. RZ 858 (#28909), August 31, 1977.
- [18] P. Brinch Hansen, "The nucleus of the multiprogramming system," *Commun. Ass. Comput. Mach.*, vol. 13, pp. 238-241, 250, Apr. 1970.
- [19] W. Wulf, E. Cohen, A. Jones, R. Lewin, C. Pierson, and F. Pollack, "HYDRA: The kernel of a multiprocessor operating system," *Commun. Ass. Comput. Mach.*, vol. 17, pp. 337-345, June 1974.
- [20] G. J. Popek and C. S. Kline, "The design of a verified protection system," in *Proc. Int. Workshop Prot. In Oper. Syst.*, IRIA, pp. 183-196.
- [21] A. R. Saxena and T. H. Bredt, "A structured specification of a hierarchical operating system," in *Proc. 1975 Int. Conf. Reliable Software*.



David L. Parnas received the B.S., M.S., and Ph.D. degrees in electrical engineering—systems and communications sciences from the Carnegie Institute of Technology, Pittsburgh, PA.

He held the position of Assistant Professor of Computer Science at the University of Maryland and at Carnegie-Mellon University. During the period 1969-1970 he was employed by Philips-Electrologica, Apeldoorn, The Netherlands, and at the MBL Research Laboratory, Brussels, Belgium. He then returned to Carnegie-Mellon

University where he held the rank of Associate Professor until 1973. In June of 1973 he was appointed Professor and Head of the Research Group on Operating Systems I at the Technical University of Darmstadt, Germany, where he remained through August 1976. He is presently Professor in the Department of Computer Science, University of North Carolina, Chapel Hill. He is also with the Information Systems Staff,

Communications Sciences Division, at the Naval Research Laboratory, Washington, DC. He has published papers in the areas of computer design languages and simulation techniques. His current interests are in the field of software engineering methods, computer system design, abstract specification for programs, verification that a program meets its specifications, and cooperating sequential processes.

The Design of a Message Switching System: An Application and Evaluation of Modula

GREGORY R. ANDREWS

Abstract—Modula is a new programming language for implementing dedicated, parallel systems. Following a systematic design technique, this paper illustrates the use of Modula for the design of a message switching communication system. A message switching system poses a number of interesting problems: a high degree of concurrent activity exists, a variety of I/O devices need to be controlled, messages can have multiple destinations, and messages can be preempted. The strengths and weaknesses of Modula with respect to these specific problems and its utility as a general purpose language are evaluated.

Index Terms—Concurrent systems, message switching, Modula, modular design, monitors, processes, software design, structured multiprogramming.

I. INTRODUCTION

THE unmistakable trend in recent years has been toward the use of high-level languages for systems programming. In an effort to improve upon available tools, three new languages have recently been designed: Concurrent Pascal [2] and Modula [8] aid in the design and implementation of multiprogramming systems while Euclid [4] is intended for implementing verifiable systems such as compilers or operating system nuclei. All three borrow heavily from the work of Wirth in the design of Pascal [7]. Although intended primarily for the development of small operating systems, both Concurrent Pascal and Modula are applicable to parallel systems in general.

In this paper, the design of one specific example, a message switching communication system, is developed using Modula. Our purposes are: 1) to present a system design technique; 2) to illustrate the use of Modula as a design, documentation, and implementation language; and 3) to evaluate Modula's

Manuscript received March 2, 1978; revised June 26, 1978. This work was supported in part by the Battelle Columbus Laboratories under Scientific Services Program Contract 0562 and in part under National Science Foundation Grant MCS 77-07554.

The author is with the Department of Computer Science, Cornell University, Ithaca, NY 14853.

utility. A message switching communication system was chosen as the application because it shows the range of Modula's applicability and presents a number of interesting implementation problems. Modula was chosen as the target language because it is specifically intended for dedicated multiprogramming systems, provides much needed facilities for controlling input/output, and appears to be very efficiently implemented [9], [10].

The specific design technique used here is described in Section II. The communication system itself is then developed in Sections III-V. Section III specifies the functions and external interfaces of the system. Section IV summarizes the major features of Modula and presents the system organization, information and control flow, and block interfaces in terms of Modula components. Section V refines the organization by giving outlines of Modula programs for the most interesting parts of the system. Finally, Section VI evaluates the utility of Modula for the design of parallel systems by reflecting on aspects of the communication system.

Modula is used here as a specification language since its compiler is not generally available. Even without a compiler though, we feel that programming any system in a structured, high-level language such as Modula is a valuable prelude to actual implementation. It serves as an intermediate step between specification and coding that helps one develop and reason about the implementation. It also provides meaningful documentation when used as comments in whatever implementation language is eventually employed.

II. DESIGN TECHNIQUE

The design described here was developed in three major steps: system specification, system organization, and program implementation. Each major step consists of a number of parts. The first step involves specifying the major functions of the system and the specific formats of input/output messages. This characterizes both what and how information is processed; hence it completely characterizes the user's