

~~Handwritten signature~~

Introduction

9

The Modular Programming System (MPS) is a facility for the development 12
of large software systems. Our goal in creating it was a substantial 13
improvement in the ease with which such systems are written, or rather the
ease with which they evolve over a long period. We have not attempted much 15
innovation in language features, since for the most part we were able to
borrow satisfactory control and data handling facilities from Bliss [*1], 16
Pascal [*2] and SUE [*3]; the few exceptions are discussed in the section 17
on data structures below. Instead, we have concentrated on the more global 18
parts of a programming system: control transfers between modules, binding
of external references, specifying the execution context of a module. 19
Following such diverse examples as assembly language and BBN Lisp [*4], we 20
have also sought to define a logically complete virtual machine, every 21
aspect of which is accessible to programs executing within it. The rest of 22
this section outlines the essential features of MPS and comments on their
precursors in other systems. ~~The body of the paper discusses these ideas~~ 23
~~in detail and explains what we have done to realize them.~~

Programming generality is perhaps the most important principle on which 25
the system is based. MPS programs or computations are built out of 26
modules. Each of these is a black box with a set of labelled ports 27
emerging from it. The label on a port consists of a name, which is an 28
arbitrary string, and a type. Modules can be used to build a computation 29
by connecting ports which have compatible types. Programming generality 30
requires that

(1) connections can be made without any changes to the modules; 32

(2) connections are not restricted by any properties of the modules 34
except the port types; 35

(3) a configuration of modules with some connections established can 37
be packaged, again without changes to the components, into a new 38
module whose ports are the unconnected ports of the configuration,
perhaps with new names. Such a configuration is functionally 39
indistinguishable from an atomic module with the same ports. Of 40
course this operation can be applied recursively.

An (atomic) module consists of code and context. The code consists of 43
a body of instructions for the base machine on which MPS programs run,
together with a symbol table and a context specifier which is a list of 44
record types (as in Pascal [*2]; PL/I or Algol 68 call them structures). 45
An example will explain this: in the Algol program 46

```
begin integer a; b; 48
```

```
L: begin Boolean c; 50
```

```
... 52
```

```
end; 54
```

```
... 56
```

```
end 58
```

the block labeled L would have the context specifier ((a: integer, b: 60
integer), (c: Boolean)). The context consists of a list of record 61
generators, each of which is (conceptually at least) a procedure which

returns a record whose type matches that of the specifier. The concepts of 63
local variable and of record component are thus unified, and decisions
about when new instances of variables are to be created can be programmed 64
explicitly in the language. Readers familiar with Simula [*6] will 65
recognize the ancestry of this idea. Of course it is possible to create 66
any number of modules from the same code.

Control transfers between modules are handled by a uniform mechanism 68
which subsumes the ideas of procedure call and return and coroutine 69
linkage; these turn out to differ only in their conventions about when
storage is allocated and released. Every control transfer passes a single 71
argument record. The concepts of formal and actual argument list and of 72
record are thus unified, so that all the machinery described below for 73
handling records of variable type, for creating a record with default
components, and for manipulating a record as a whole extends to parameter 74
lists. It requires no special features to pass a variable number of 75
arguments to a function, to return multiple result values or to intercept a 76
function call and display the actual arguments. The syntax for ordinary 77
procedure calls is unchanged, but it is now viewed as a convenient
abbreviation for a sequence of more primitive actions: create an argument 78
record, transfer control, receive control back, disassemble a result
record.

When a module is created, its ports appear as variables in its context 80
which are initialized to a special "unconnected" value. The code body 82
contains a list of the ports, their types and the symbol table entries
which contain their names. The process of connecting a port is called 83

binding; it is roughly analogous to the linking performed by a conventional loader. In MPS however, binding is entirely under control of the user. 85
The creator of the module, or any one else who has access to it, can 86
connect port P to port Q with the primitive bind (P,Q), or break the 87
connector with bind (P, nil). This can be done at any time, whether or not 88
execution has already begun or control been passed over the port. It is 89
therefore possible to splice in tracing or advising procedures [*7:
Warren], or to replace a module with a new version, without disturbing the 90
state of a computation. There is also a set of standard binding procedures 91
for binding things in standard ways, e.g. by matching names in a tree- 92
structured naming hierarchy [*8: Multics].

An attempt to use an unbound port is translated into a call on a system 94
primitive called the control fault handler; this idea is also taken from 95
[*8]. The handler, however, takes no action of its own, but instead 96
generates a called "control fault" signal. This is a standard control 97
transfer preceded by a special mechanism which binds the port over which
the transfer is going. The binding is done by following the chain of 98
return links looking for one which lies in a scope which has enabled the
signal. Associated with the enable is a catch-phrase, which declares a 99
procedure to be called by the signal. The context of the catch-phrase 100
includes the context of the scope in which it was enabled. The catch- 101
phrase can return a value to the module which generated the signal, reject
the signal and cause the binder to resume its search, or unwind control 102
back to the module which enabled the signal. Other examples of signals 103
might be "out of storage" or "end of file." An especially interesting 104
example is "unwind" itself, which destroys each module which rejects the

signal; by enabling "unwind" a module can get control to clear itself up 105
before being destroyed. Signals are a synthesis of PL/I ON-conditions [*9] 106
and Snobol failure returns [*10], and provide a very attractive (and cheap) 107
mechanism for handling unusual events. The way in which the binding is 108
done means that modules which are not interested in a signal remain
completely unaware of its existence, 109

New types of data structures are built up from primitive types such as 111
integer (m...n) using the type constructors array, record, pointer, and 112
union in the style of Algol 68 [*10] or Pascal [*2]. A type is described 113
by a record of type "type descriptor" which can be manipulated by programs
like any other record. Hence programs such as binders or debuggers can 114
interpret type descriptors at run-time to check the validity of port
connections or to display data structures. It is important to have only 116
one copy of the source text for a type, no matter how many modules use it.
To make this convenient, source programs can "include" declarations from 117
other files. In addition a type descriptor carries information which 118
identifies its source text, so it is possible to compare two types for
structural or for literal identity. The major innovations in the type 120
facility are a systematic treatment of generic functions, which subsumes
operators like plus, assignment or subscripting as special cases, and 121
complete user control over permissible coercions. This latter point is 122
especially important for the binder, which needs to be able to insert
coercions to obtain agreement of the types of two ports which it is trying 123
to connect.

A final point, which has been mentioned several times above, is our 125

insistence that all of the data structures and procedures which form the 126
infrastructure of MPS be accessible to programs. Type descriptors, symbol 127
tables, frames allocated for local variables, ports and arguments are all
MPS data structures whose type descriptors can be used to access and, when 128
appropriate, modify them. The procedures which are invoked for binding, 129
establishing context, allocating storage, signaling and passing control can 130
all be replaced or advised. We have tried very hard not to provide any 131
facilities which cannot be duplicated by ordinary MPS programs. As a 132
result, most of the system facilities are in fact written as standard
programs, as is the debugger.