

August 1981

Report. No. SUN-CS-81-879

*Also numbered:
HPP-81-14*

Interlisp-VAX: A Report

by

Larry M. Masinter

Department of Computer Science

Stanford University
Stanford, CA 94305



Interlisp-VAX: A Report

Larry M. Masinter

August 1, 1981

Contents:

- I. Introduction
- II. Interlisp-VAX: **Overview** and Status
- III. What will Interlisp-VAX be like?
- IV. Conclusions

The views expressed in this report are those of the author. They do not necessarily reflect those of the Xerox Corporation, Stanford University, or the University of Southern California.

This study was funded in part through the SUMEX Computer Project at Stanford University under grant RR-00785 from the Biotechnology Resources Program of the National Institutes of Health.

I. INTRODUCTION

Since November 1979, a group at the Information Sciences Institute of the University of Southern California has been working on an implementation of Interlisp for the DEC VAX-series' computers. This report is a description of the current status, future prospects, and estimated character of that Interlisp-VAX implementation. It is the result of several days of discussion with those at ISI involved with the implementation (Dave Dyer, Hans Koomen, Ray Bates, Dan Lynch): with John L. White of MIT, who is working on an implementation of another Lisp for the VAX (NIL); with the implementors of Interlisp-Jericho at BBN (Alice Hartley, Norton Greenfeld, Martin Yonke, John Vittal, Frank Zdybel, Jeff Gibbons, Daryl Lewis); with the implementors of Franz Lisp and Berkeley Unix² at U.C. Berkeley (Richard Fateman, Bill Joy, Keith Sklower, John Foderaro); and with my colleagues at Xerox PARC.

An earlier draft of this report was circulated to the parties involved in the Interlisp-VAX discussions. This document has been revised as a result of comments received.

Why Interlisp-VAX?

In early April 1981, a meeting of ARPA-sponsored or related Lisp users was held at SRI, to discuss the status and future of Lisp. Those of the community who were current Interlisp users felt strongly that: (1) there was a need for Interlisp to continue to be a viable programming environment in the 1980's, strongly standardized among all implementations; and (2) the most important new implementation of Interlisp would be for the VAX. There were several reasons for the choice of both the VAX and Interlisp.

Why VAX? The primary reason is that many sites already have VAX's which are and will continue to be used not only for Lisp and AI research, but also for use as general purpose, time-shared computing resources, for running FORTRAN, business computing, etc. The VAX is considered to be the most important "technology transfer" vehicle for Interlisp AI programs in the early through mid '80s. It is already spread widely throughout industry and industrial laboratories, and it is very widespread among ARPA's military clientele. It is unlikely that researchers who develop application-oriented AI systems in Interlisp will want to re-implement them in some other language. It is also unlikely that these institutions (private and military) will buy machines specifically for AI programs if those programs constitute only an occasional part of their computing needs. The VAX is believed to be the most likely vehicle for transferring applications to those institutions.

In addition, the VAX (for better or for worse) appears to be the machine that many computer science departments around the country have chosen for their "next generation" machine. Insofar as there is a need to spread the concepts and software technologies developed in Interlisp to these departments, it is believed that there is a need to have Interlisp running on the VAX.

Why Interlisp? The Interlisp programming environment has been in wide use in the Artificial Intelligence community for a number of years. It is a powerful, integrated environment, having evolved over the years into a stable system. The availability of multiple, compatible implementations on a number of machines means that researchers can easily transport their programs from any implementation to another.

Why this **report**? Because of the perceived importance of Interlisp-VAX to the community, and because of my experience with Interlisp and its implementations, I was asked by Stanford and ISI to evaluate the status of the on-going project at ISI, and to estimate the magnitude of the tasks remaining, expected performance and character of the resulting product. Many ongoing research institutions are making plans for their future computational requirements, and many of the decisions about choice of programming language and hardware hinge on the prospects for Interlisp-VAX. In light of the large amount of confusion in the community about the future availability of Interlisp on the VAX, it was thought important to have an outside assessment of the future of the project.

II. INTERLISP-VAX: OVERVIEW AND STATUS

A. Project definition and history

While the Interlisp-VAX project started in November 1979, most of the first year was taken up with project startup and training of personnel (None of the project members were originally familiar with the C programming language, Unix, or Interlisp, either as a programming environment or its overall implementation). Thus, most of the work on the implementation to date has been accomplished since November 1980.

The goal of the project has been to produce a version of Interlisp which runs on the VAX, which:

Is as **compatible** with **Interlisp-10** as practical.

While it is difficult to give a metric for compatibility, the goal of the project is that most Interlisp programs in the community will run in Interlisp-VAX merely by recompiling them. At a **bare** minimum, the Interlisp-VAX code must run the standard Interlisp packages such as the Interlisp Editor, Masterscope, Break Package, Record Package, DWIM, CLisp, and so forth.

Uses **the extended** virtual address space of the VAX.

One of the primary motivations for investing in the VAX is that the VAX potentially has a large (at least 2^{30} bit) address space. Many Interlisp- users long ago ran out of address space and spend much of their time trying to squeeze programs into available address space.

Has adequate performance.

The Lisp produced is expected to make reasonable use of the hardware. It is difficult to give a single number which describes the performance of a system (because some things will run faster and some slower), but the average performance of Interlisp-VAX must be within a factor of 2 of other Lisps which run on the VAX (c.g., Franz and NIL). In addition, Interlisp-VAX must be competitive in price/performance to a DEC-20 for the size of programs which it is now able to run, and also be able to handle larger programs.

B. Summary of the Interlisp-VAX architecture

Interlisp-VAX is a non-microcoded implementation more similar in architecture to Interlisp- than Interlisp-D or Interlisp-Jericho. This is appropriate for the VAX, which has a powerful "native" instruction set and is time-shared between a number of users, not all of whom would be running Interlisp. Interlisp-VAX is intended to run on top of the Berkeley Unix operating system. Unlike

Interlisp-10, in which the kernel is written in assembly language, the kernel of Interlisp-VAX is written in the high-level systems implementation language C. This might well simplify the transportation of Interlisp-VAX to another machine which has a C compiler and similar characteristics (byte addressable memory, Unix, 32-bit registers).

Without going into great detail, the important aspects of the Interlisp-VAX architecture are as follows: deep binding; full implementation of spaghetti stacks; compilation to VAX native code (with no "block compiler"); memory allocation in 64 KByte "sectors" with sector-table giving type per sector; no CDR-coding (CONS cells take 64 bits); a "stop and copy" garbage collector; and 31-bit immediate integers (with plans, but no implementation, of "bignums"). These design choices seem reasonable for the VAX, with exceptions noted below in the section on performance.

Interlisp-VAX has the following component pieces:

1. *Machine-independent "higher-level" Interlisp code.* This includes, for example, the Interlisp editor, file package, and the Master-scope program analyzer. This code is shared, intact, with Interlisp-10, Interlisp-D and Interlisp-Jericho.
2. *Interlisp-D code.* This is Lisp code which, although shared with Interlisp-D (and Interlisp-Jericho), is not used in Interlisp-10. For example, the implementation of Terminal Tables and Read Tables may be shared with the other Interlisp implementations.
3. *VAX-specific Lisp code.* This code is necessary to interface to the C kernel and perform other VAX-specific operations. For example, the implementation of the DATATYPE package, while in Lisp, must satisfy constraints placed by the Interlisp-VAX garbage collector. Thus, it has some essential differences from the version of the DATATYPE package for other Interlisp implementations.
4. *C kernel.* The C kernel handles memory management, garbage collection, the interpreter, spaghetti stack support (including FUNARGs, RESUME for processes/coroutine support), bootstrapping and interface to the operating system.
5. *Lisp/C interface.* A small amount of VAX machine code is necessary for the interface between Lisp and code generated by the C compiler. Primarily of interest here is the code which is part of function call and return.
6. *VAX code generator.* The VAX native-code generator takes the output of the Interlisp-D ByteCompiler and generates VAX native code. The Byte-Compiler [Masinter & Deutsch 1980] is a machine-independent optimizing compiler which produces intermediate "linearized lisp" code for an abstract stack machine.

C. Current implementation status of Interlisp-VAX: what's been done?

Many of the major design decisions for Interlisp-VAX have been made, including layout of memory, important code sequences (e.g., function call and return for all of the various cases), representations of pointers and system data types, and many parts of the interface to the operating system. In addition, the following tasks have been accomplished:

1. *Higher-level Inter-lisp software.* The "shared" Interlisp software has been examined, and a few problems identified and fixed. The rest will run in Interlisp-VAX with little change.
2. *Inter-lisp-D code.* An initial pass has been made over the Interlisp-D code, identifying which portions can be shared.
3. *VAX-specific Lisp code.* The major pieces which have been written are a version of the DATATYPE package, an array package, and the compiled code loader and parser.

4. **C kernel.** Most of the C kernel has been completed, in the sense that the code is there and has passed preliminary tests.
5. **Lisp/C interface.** This has been completed.
6. **VAX code generator.** A first version of the VAX code generator has been produced and, to a great extent, debugged. The important design decisions have been made about function call sequences, as well as some of the important open-coding sequences (e.g., CAR and CDR).

D. Tasks remaining in existing code

1. **Higher-level Interlisp software.** Problems may arise in implementing Interlisp's notions of files, versions, and dates under Unix. If so, it may be necessary to fix those portions of the Interlisp higher-level software to be more implementation-independent.
2. **In terlisp- D code.** Unfortunately, Interlisp-D is a "moving target," and it is difficult to rely on the sources staying compatible. Since code is shared between Interlisp-D and Interlisp-Jericho, the same code will most likely run under Interlisp-VAX. Problems may arise insofar as the lower levels of Interlisp-VAX differ.
3. **Vax-specific Lisp code.** The DATATYPE implementation requires some work. The array package seems to be relatively complete, although the program has not been extensively tested. The compiled code loader/parser has been completed and tested in "cross-compilation" mode, while running in Interlisp-10.

The VAX/Unix I/O package still requires much work. The interface between Interlisp and Unix is to be accomplished via (1) the Interlisp-D FILE10 package, which gives an interface to buffered, random access files from higher-level Interlisp software, (2) some VAX-specific Lisp code, which then interfaces to (3) some pieces of the C kernel. The interfaces between many of these pieces are being designed, but some of the pieces have not been written.

There is a body of the Interlisp environment which, although nominally not part of the "core" of Interlisp, forms a useful part of most of its implementations. For example, the DIRECTORY package and GETFILEINFO are Interlisp- facilities which, while not part of the Moore VM document, can be implemented in Interlisp-VAX. They are part of Interlisp-10, Interlisp-D and Interlisp-Jericho, and are used by Inter-lisp application programs.

Interface with Unix's notion of terminals and interrupts has been considered, but the final details have not been worked out. Initial versions of Interlisp-VAX will have a very simple notion of interrupts.

4. **C kernel.** Future changes will likely be required depending on the needs of the Lisp-level T/O package, interrupts, and a new version of Unix which will allow Interlisp to use the high end of memory. The C kernel contains some especially "tricky" areas: interpreter, stack management and garbage collection. These were not completed as of June 1981. Experience with other Lisp/Interlisp implementations has been that debugging and complete testing are difficult. Bugs often are found in the handling of obscure and rare cases, as the code interacts with many other parts of the system. I expect Interlisp-VAX to have its share of problems in these areas.
5. **Lisp/C interface.** Changing the Lisp/C interface will only be necessary in response to fixing some of the expected "performance bugs" of Interlisp-VAX; e.g., free-variable-pointer-caching (discussed below) may require changes in the function call sequences.
6. **VAX code generator.** My examination of the Vax code generator uncovered a few minor problems due to a misunderstanding of conventions required by the ByteCompiler. Undoubtedly, a few more will surface.

More importantly, the current code generator for VAX native code will (as' planned) require much work to bring it to the point where it generates production-quality code. In particular:

- a. A register-allocating version of the code generator (in some ways a complete rewrite) would significantly improve performance on the VAX.
- b. A "peephole optimizer" for VAX instructions would enable Interlisp-VAX to take advantage of the VAX's complex instruction repertoire.
- c. More "open" compilation of frequently-used routines will be necessary in many circumstances. Although many open-coding sequences have been incorporated, adding more will of course require additional time and effort.
- d. Modification of the ByteCompiler to suppress boxing of intermediate results would pay off in speed for integer calculations and space for floating arithmetic.

E. Other areas requiring work '

In addition to the areas outlined above, a number of other areas need attention:

1. **Free variable pointer caching.** There is a very serious performance problem in Interlisp-VAX, the correction of which will require major changes to the Interlisp-VAX system. Interlisp-VAX uses deep binding. While deep binding is a reasonable choice for Interlisp-D (because of microcoded free variable lookup), it may be a source of a large performance penalty in Interlisp-VAX, especially in interpreted code. In any case, there is currently no mechanism for "caching" free variable pointers: so free variables are "looked up" at every reference, even within an inner loop. This is unacceptable. A design needs to be worked out and integrated into the compiler and stack access mechanism. No one scheme is clearly optimal, although whatever scheme is chosen will require changes to the compiler, interpreter, garbage collector and stack manipulation routines.
2. **Bootstrapping.** Bootstrapping is as complicated in Interlisp-VAX as it is in other Interlisp implementations, for a variety of reasons. For example, debugging "low level" pieces of the system is made more difficult because bootstrap-load order requirements are difficult to detect without running the (time-consuming) bootstrap process. Traditionally, this is merely a source of frustration rather than an insurmountable barrier.

Documentation. Documentation of Interlisp-VAX is needed discussing its differences from other Interlisps and areas such as interface to Unix. There is some intention to participate in the upcoming major revision of the Interlisp Reference Manual!

Access to Unix facilities. Interface from Lisp to Unix facilities such as pipes, processes, and shell programs, will greatly increase the utility of Interlisp-VAX. These facilities are **not** necessary for running current Interlisp- programs, except to the extent they replace Interlisp-10 facilities (e.g., SUBSYS).

SYSOUT. The current Interlisp-VAX SYSOUT facility dumps the entire allocated virtual memory of the Lisp system (currently, without any of the "shared" Interlisp code over 1 MByte). At some future date, Berkeley Unix will provide a mechanism which will allow writing out individual pages and a page map, making SYSOUT files more manageable.

Porting to other VAX operating systems. Many sites do not run the Berkeley Unix operating system, instead choosing VMS (the DEC-supplied operating system for the VAX), or EUNICE (a Unix compatibility package developed at SRI). These are candidates for "other implementations" of Interlisp-VAX. Because of Interlisp's heavy use of the operating system's memory management facilities, porting Interlisp-VAX to these other operating systems will likely prove quite difficult.

III. WHAT WILL INTERLISP-VAX BE LIKE?

Assuming the above tasks are completed, the question remains: what will it be like? There are two issues: in what way will Interlisp-VAX differ from other Lisp implementations, and what performance can be expected?

A. Comparison of Interlisp-VAX to other Interlisps

Full Interlisp-VAX is intended to be highly compatible with Intclisp-10, to the point where many complex programs would move gracefully between it and other Interlisp implementations. The only areas of incompatibility are those which are necessarily not shared between any implementations: access to machine code within Lisp routines, etc. In addition, there are currently no plans for "linked" function calls in Interlisp-VAX, nor for a "block" compiler. These are minor difficulties.

Interlisp-VAX will be able to access some of the facilities of the Unix environment to good effect; e.g., one might imagine using it as an interactive "shell" programming language.

Interlisp-VAX will not have any particular capabilities for bit-mapped graphics.

Interlisp-VAX will have a larger "small" arithmetic range.

B. Performance

There are two major factors in the performance of Interlisp on the VAX. First is in the actual CPU time to complete various operations, and second is in the amount of time spent paging.

1. CPU performance

The performance profile of a Lisp system is complex, and there are many areas where Interlisp-VAX's relative performance to other Interlisp implementations will vary over a wide range. There seem to be a few areas of critical performance to any program: function call, variable reference, data structure access, arithmetic, and garbage collection. An appropriate weighted average of performance in those areas is a good overall measure of total system performance.

One important way of estimating performance of Interlisp-VAX is to use as a comparison the code in other Lisp implementations for the same task, taking into account the differences in the various code sequences. Comparisons are made between Intclisp-VAX and Franz, NIL, and Interlisp-10.

a. Function call and return

A function call for Interlisp-VAX will be at least twice as slow as a similar function call in Franz Lisp, partly because of language requirements (Franz does not check that the number of arguments passed matches the number of arguments expected), and partly because of the design of the Interlisp-VAX stack format (Variable names are pushed as well as the values).

In Franz Lisp, a minimal call/return takes 17 microseconds (VAX-111780). Call/return in Interlisp-VAX may be as high as 100 microseconds, although the average will most likely be nearer to 40 microseconds.

In Interlisp- on a DEC 2060, a block-internal call takes on the order of a microsecond (PUSHJ, POPJ), and the minimal (non-block) call/return takes 57 instructions (roughly 25 microseconds), while some functions, because of the Interlisp Swapper, may take more than 200 instructions for call/return (100 microseconds). The variation in function call time will apparently be high for Interlisp-VAX and Intclisp-10. For some functions, Interlisp-VAX function call will be slightly faster. For calls which in Interlisp- would be block internal, an Interlisp-VAX call might be 50 times slower. Note that benchmarks which purport to make comparisons with Interlisp- should explicitly control for the possibly enormous variation in Intclisp-10 function call time.

b. Variable reference

Performance on local variable reference in Franz and Interlisp-VAX will be similar if Interlisp-VAX delivers its optimizing, register allocation code generator. Currently, variable reference will often be slightly slower. More importantly, free variable access will be very significantly slower in Interlisp-VAX, even after a variable caching scheme is implemented, because of the cost of variable lookup when using deep binding.

c. Garbage collection

The "stop and copy" variety of garbage collection, while compacting the address space and thus reducing the working set of subsequent computations, is more expensive in CPU time and memory usage than the "mark and sweep" variety by a nominal factor of two. Garbage collections for large address space systems can be expensive, even using mark and sweep. A full VAXSYMA garbage collection is reported to take on the order of 3 seconds of CPU time. A garbage collection of an 8 MByte address space in ELISP (a Lisp for DEC 2060's which uses the extended addressing feature [Hedrick]) took between 20 and 40 seconds of CPU time. This figure includes some charge for paging overhead. It seems likely that (1) garbage collection is swap limited, and (2) the respective operating systems used to gather those times do not do a particularly good job of filtering out swap overhead from CPU time. It is not unreasonable to expect, however, that an Interlisp-VAX garbage collection will take twice as long as a Franz Lisp collection, because of the intrinsic overhead of "stop and copy" over "mark and sweep."

An alternative computation can be made as follows. Assuming an Interlisp-VAX system uses 4 MBytes of memory, then with a compacting garbage collection but no other memory localization algorithms, I believe that most user programs would "dirty" at least $\frac{1}{4}$ of all system pages (i.e., 1 MByte) within a relatively small amount of time. Let us suppose a garbage collection occurs after a user has allocated the equivalent 40K CONS cells, or 0.32 MByte of storage. This would involve referencing 1.6 MBytes of memory. This would mean that a garbage collection would take, at a minimum, between 2 and 20 seconds of CPU time on a VAX-111780.

2. Paging Performance and Real Memory Requirements

I spent a considerable amount of time trying to estimate the number of users or sizes of Lisp systems that some typical VAX configurations might support. I believe that this is one of the most important factors in Interlisp-VAX performance, because of the predicted large virtual address spaces of Interlisp-VAX programs (one of the main reasons for going to Interlisp-VAX in the first place).

a. Operating system considerations

Interlisp-VAX will be implemented on top of the Berkeley Unix operating system. Another possible candidate for a host operating system is a Unix compatibility package written at SRI by the name of EUNICE, which runs under the DEC-supplied operating system VMS. There is some controversy over the relative performance and functionality of VMS vs. Unix. A fairly comprehensive set of benchmarks [Kashtan] showed that VMS out-performed Unix in a variety of paging configurations. It is claimed by the Berkeley Unix implementors that (a) many of the benchmarks were atypical of real computations, and (b) tests were run on an early version of Berkeley Unix and performance has improved considerably since then. I believe the choice of operating system can be made on grounds other than predicted performance for running Interlisp: reliability, maintenance, cost, etc. Further, converting Interlisp-VAX to run under EUNICE rather than Berkeley Unix will be a relatively minor job compared to the magnitude of the Interlisp-VAX implementation itself. It seems that the difference between operating systems makes for only a relatively small factor in the overall performance, if the real memory available is too small to hold the "working set" of programs attempting to run at any one time.

b. Real memory requirements of Interlisp-VAX

There are a variety of ways of estimating memory needs. The best estimates seem to come from: (1) comparison with MACSYMA in Franz Lisp (VAXSYMA), and (2) comparison with Interlisp- and Interlisp-D.

1) Virtual Address space (minimum). Many current Interlisp- programs run with a virtual address space of 2 MBytes (2 full "forks"). A similar system in Interlisp-VAX will probably require a 4 MByte address space because: (1) there is expansion for 32 rather than 18 bit addresses (no CDR coding); (2) the copying garbage collector, when it runs, will require twice the allocated space; and (3) Interlisp-VAX allocates storage in quanta of 64 KByte sectors rather than a 2 KByte "page" as in Inter-lisp-lo, giving more "breakage" per datatype. This figure is consistent with numbers extrapolated from Interlisp-D.

2) Working set. In current Interlisp-10, the "working set" of many programs is 0.5 MBytes or more (that is, the amount of real memory outside of the "system" necessary to- keep the program from spending more than half of its time paging). Extrapolating, using the same figures as above, the working set of a "typical" Interlisp-VAX application will be over 1 MByte. This figure is consistent with memory requirements extrapolated from Interlisp-D.

3) Calculation of real memory requirements. If there are i users, j of whom are active, they will need $i \cdot 31$ KBytes of page table (31 KBytes = 4 MBytes/128), plus $0.75 \cdot j$ MByte bytes for their working set. For example, 5 users, 2 of whom are actively running at any one time, would require less than 2 MBytes of real memory (outside of i/o buffers, etc.).

However, if systems increase in *allocated* space (independent of the working set) because more programming or data is contained in their virtual address space, one might imagine a situation where the virtual address spaces were in the 20-30 MByte range (Many users do not believe that a 2^{24} byte virtual address space, 16 MByte, is big enough for applications they plan in the near future). In such a situation, **each** such Interlisp process would require as much as 0.2 MBytes of real memory for its page table, independent of its activity. This might severely limit the number of users who could be active on the system at any one time.

c. Problem areas

There are some problem areas, both with Unix and with VMS, which will have to be resolved:

1) Sharing. VMS currently has more flexibility in allowing sharing of space among users in a piecemeal fashion. In the current Interlisp-VAX design, only 0.1 MBytes of the address space are "pure" in the sense that Berkeley Unix would allow it to be shared among multiple users. Insofar as multiple users have the same large virtual address space (e.g., they are running the same program with a large, fairly static "knowledge base"), sharing is important to improving the number of users allowable at any one time.

2) Problems with large virtual address space. VMS requires disk/swap space to be pre-allocated, at system generation time, for the maximum allowable in the system. With multiple users with large address space programs, this adds considerably to the amount of disk space required on the system (even if most of those users are inactive). In addition, VMS requires an additional pre-allocated swap file which contains $J \cdot W$ pages, where J is the maximum number of processes with independent address spaces (100 would not be an unreasonable figure for a machine used by many users for editing, background processing, etc.), while W is the maximum "working set" of a single process (which, for large address space processes, should be at least 2 MBytes).

On the other hand, Berkeley Unix currently requires the page tables of all processes to be locked down, which may be a significant drain for very large address space programs where the data in the address space is in fact infrequently referenced.

IV. CONCLUSIONS: WHITHER INTERLISP-VAX

A. There aren't any good alternatives

Given the requirements of technology transfer to university, industrial and military sites, there are few other options. Even though Interlisp-VAX will probably not be cost-effective for intensive Lisp users, it may be for those whose requirements are for casual and occasional use of Interlisp or tools developed in it. There are a few alternatives which could benefit from further exploration:

In terlisp-3 70

There is a version of Interlisp for IBM/370 machines, originally developed at Uppsala University and modified at the Weizmann Institute [Raim]. Interlisp- might be a possibility for some sites, although reports from several sources are that the Interlisp-370 is incomplete, not particularly compatible with other Interlisps, and has serious performance and reliability problems. However, I believe that this alternative should be more seriously explored.

Implementing Interlisp on top of NIL or Franz

This might have been a reasonable way to approach the initial Interlisp-VAX implementation, but it does not seem cost-effective at this point.

Emulating Interlisp-D on a VAX

An alternative not presently explored in any detail would be to write an Interpreter for Interlisp-D byte codes and run Interlisp-D on a VAX (cf. [Rowan]). Performance would be poor (perhaps a factor of 4-5 slower than currently projected), but code would be more easily transportable.

Automatic conversion of Interlisp programs to other VAX Lisps

This is an approach which has rarely succeeded. Programs which convert between language dialects are heuristic at best, and require considerable hand-holding. For any particular program, converting to another language might be cost-effective, but on the whole it is not.

B. Performance: mixed results

Performance in the Lisp community is often measured in DEC KA-10 or KL-10 equivalents, e.g., "1/4 of the speed as on a KL-10." One would like to be able to draw the inference that, if a KL-10 adequately supports 40 users with 8 actively computing (the rest editing, reading mail, etc.), 1/4 of that would amount to 10 users with 2 actively computing. Unfortunately, these performance figures can be misleading: (1) because of the wide variation in Interlisp-10 speeds on the same problem, and (2) because timings on small benchmarks do not give an accurate picture of the number of active users who can be supported in a working environment.

More reasonable estimations of performance can be drawn from experience with VAX's running Franz Lisp or VAXSYMA. While no exact figures are available, experience has been that a VAX-11/780 with 4 MBytes real memory can support 30 users, of whom 3 are actively using VAXSYMA. Interlisp working-set and virtual address space requirements will exceed those of VAXSYMA.

Although the VAX is purported to be quite cost-effective for FORTRAN, the instruction set is not particularly effective for Lisp, and even less so for Interlisp. The "CALLS" instruction, which is intended to be used for function calls in high-level languages, assumes a model of the stack which does not match Interlisp's. While the Interlisp-VAX design takes advantage of "CALLS" in a clever way, function call is still relatively more expensive than it is on microcoded machines which can have an Intel-lisp-specific function call instruction.

Virtual address space and real memory

Although the VAX is a large virtual address space machine, the address space may not be particularly usable on configurations typical in many installations. For example, the following configurations were proposed as "typical" VAX installations:

VAX-11/750 with 2 MBytes real memory (maximum for 750)

VAX-11/780 with 4 MBytes real memory

VAX-11/780 with 8 MBytes real memory (requires additional memory controller)

Also proposed are configurations not currently available: "single-user" VAX machines with memory in the 1-2 MByte range, or 750's and 780's with more memory (requiring 64K RAM chips).

Because of Interlisp-VAX's large virtual address space and working set, a machine with only 2 MBytes of real memory might be able to support at most one or two large address space active users at a time. Generous amounts of disk, swapping space, and real memory will be required -- more so than in Interlisp-10 to support the same users, and much more so than in Interlisp-D or Interlisp- Jericho.

Very few time-sharing systems have adequately dealt with giant address spaces for multiple users. The success of very-large-address-space Interlisp-VAX will depend on the cooperation and support of the Berkeley Unix implementors.

C. There is much left to do

There is an unfortunate tendency to underestimate the magnitude of the task of transporting a system the size and complexity of Interlisp. Interlisp is not merely an interpreter and a few utility routines. It is a rich and complex programming environment with facilities which were heavily influenced by Tenex, its original host operating system. Porting it to another machine and continuing to upgrade it is a major undertaking. I cannot stress this enough.

The publication of the Interlisp Virtual Machine Specification [Moore] was an important step forward in the creation of transportable Interlisp, in that it identified a major portion of what the "higher-level" Interlisp support software required in order to run. Unhappily, as complete and well-written as that document was, it is *not* an accurate guide for the construction of a useful Interlisp implementation, in that many areas are designated as being left to the implementor while many Interlisp applications require exact compatibility with Interlisp-10. The VM is also not a good measure of the magnitude of implementing Interlisp. For example, the VM mentions the compiler only in passing; however, providing a reasonable Interlisp compiler is a major portion of the task of transporting Interlisp to a new (non-microcoded) machine.

Transporting Interlisp is harder than merely implementing "some" Lisp dialect. It is much more difficult to be strictly compatible while using the underlying power of the machine to the fullest. Compatibility makes the implementation harder because there is an existing standard against which the implementation can be judged. For a "new" Lisp, it is always possible to declare oneself "done" at almost any point. The necessity of emulating exactly the behavior of another system is what makes the task more difficult.

How much is left to do?

It is difficult to give a "man-month" figure for Interlisp-VAX for several reasons. First, of course, the notion of "man-month" independent of implementor is a well-known paradox; start-up time and personnel training can delay a project for many months (as in the early months of the Interlisp-VAX project).

Second, there are several tasks ahead which will undoubtedly encounter unforeseen problems. "System shakedown" is a catch-all phrase which can cover many months of discovering problems or previously undetected system requirements. Software completion is not measured well by proportion of lines-of-code written.

Finally, there is a wide range of variation of what is meant by "Interlisp-VAX." On the one hand, an initial version may be available relatively soon. This version will likely have serious performance

problems (mainly because of free variable **reference** and non-tuned code generation), and will likely not be fully functional or compatible with Interlisp-10. The task of bringing Interlisp-VAX to the level of functionality, performance and reliability of Interlisp- and Interlisp-D remains awesome.

Unfortunately, there is not a good perception in the Interlisp user community of the amount of work between the first release and a system which will be acceptable to current Interlisp users. For this reason, the recent "prc-announcement" message [Dyer] was at best misleading for those trying to make plans based on Interlisp-VAX availability. While this initial version might in fact be a reasonable alternative to, say, converting a large Interlisp program to Franz Lisp (because the conversion cost would be higher than the performance difference would warrant), it will not be comparable with most other Interlisp implementations (-10, -D, -Jericho).

The Interlisp-VAX project is and has been from the beginning drastically undermanned. The initial proposal for implementation of Interlisp-VAX in one year with no existing personnel was at best wishful thinking. Hans Koomen will be leaving within the near future. This is a serious, although possibly not fatal, blow to the continuation of the project, even with the addition of more staff members (Ray Bates and Don Voreck).

The project needs a team of implementors who are committed to its goals, are qualified to carry it out, and will stick with the project once the initial release has been made. If Interlisp-VAX is to be viable, there needs to be a long-term (3-4 year) commitment to its maintenance and support by a team of qualified personnel. This level of support or greater has been required by every other serious implementation of Lisp that I know of, including Inter-lisp-10, Interlisp-D, Interlisp-Jericho, and Lisp Machine Lisp. There is no reason why anyone should imagine that Interlisp-VAX would be different.

NOTES

¹VAX is a trademark of Digital Equipment Corporation.

²Unix is a trademark of Bell Laboratories.

BIBLIOGRAPHY

- Burton, R.R., *et al.* "Interlisp-D: Overview and Status." In ***Papers on Interlisp-D***, Xerox Palo Alto Research Center, CIS-5 (SSL-80-4), 1980.
Describes the Interlisp-D implementation effort, including some words of wisdom on why implementing Interlisp is hard.
- Dyer, D., *et al.* ***INTERLISP-VAX***. Message-ID: <[USC-ISIB]17-Jul-81 15:10:10.MILLAR>
This was the "official pre-announcement of the availability of Interlisp-VAX."
- Hedrick, C. ***Some Tests of Rig Core Images***. [message file] Rutgers University.
30 May 81 0447-EDT.
Discusses ELISP implementation on an extended address DEC 2060.
- Kashtan, D. Unix ***and VMS: Some Performance Comparisons***. [message file] SRI International.
Compares performance of VAX/VMS version 1.6 and VM Unix Berkeley version 2.1.
- Masinter, L. M. and Deutsch, L. P. "Local Optimization in a Compiler for Stack-based Lisp Machines." In ***Papers on Interlisp-D***, Xerox Palo Alto Research Center, CIS-5 (SSL-80-4), 1980.
Describes the byte compiler.
- Moore, J. ***The Inter-lisp Virtual Machine Specification***. Xerox Palo Alto Research Center, CSL 76-5, revised March 1979.
- Raim, M. Personal communication.
- Teitelman, W. and Masinter, L. ***The Inter-lisp Programming Environment***. IEEE Computer, April 1981, pp. 25-33.
Overview of Interlisp.