
Integration, extensibility, and ease of modification made Interlisp unique and powerful. Its adaptations will enhance the power of the coming world of personal computing and advanced displays.

The Interlisp Programming Environment

Warren Teitelman and Larry Masinter
Xerox Palo Alto Research Center



Interlisp is a programming environment based on the Lisp programming language.^{1,2} In widespread use in the artificial intelligence community, Interlisp has an extensive set of user facilities, including syntax extension, uniform error handling, automatic error correction, an integrated structure-based editor, a sophisticated debugger, a compiler, and a filing system. Its most popular implementation is Interlisp-10, which runs under both the Tenex and Tops-20 operating systems for the DEC PDP-10 family. Interlisp-10 now has approximately 300 users at 20 different sites (mostly universities) in the US and abroad. It is an extremely well documented and well maintained system.

Interlisp has been used to develop and implement a wide variety of large application systems. Examples include the Mycin system for infectious disease diagnosis,³ the Boyer-Moore theorem prover,⁴ and the BBN speech understanding system.⁵

This article describes the Interlisp environment, the facilities available in it, and some of the reasons why Interlisp developed as it has.

Overview

From its inception, the focus of the Interlisp project has been not so much on the programming language as on the programming environment. An early paper on Interlisp states, "In normal usage, the word 'environment' refers to the aggregate of social and cultural conditions that influence the life of an individual. The programmer's environment influences, and to a large extent determines, what sort of problems he can (and will want to) tackle, how far he can go, and how fast. If the environment is cooperative and helpful (the anthropomorphism is deliberate), the programmer can be more ambitious and productive. If not, he will spend most of his time and

energy fighting a system that at times seems bent on frustrating his best efforts."⁶

The environmental considerations were greatly influenced by the perceived user community and the style of programming in that community: first, typical Lisp users were engaged in experimental rather than production programming; second, they were willing to expend computer resources to improve human productivity; third, we believed users would prefer sophisticated tools, even at the expense of simplicity.

Experimental programming and structured growth. The original architects of the Interlisp system were interested in large artificial intelligence applications programs. Examples of such programs are theorem provers, sophisticated game-playing programs, and speech and other pattern recognition systems. These programs are characterized by the fact that they often cannot be completely specified in advance because the problems—to say nothing of their solutions—are simply not well enough understood. Instead, a program must evolve as a series of experiments, in which the results of each step suggest the direction of the next. During the course of its evolution, a program may undergo drastic revisions as the problem is better understood. One goal of Interlisp was to support this style of program development, which Erik Sandewall* has termed structured growth: "An initial program with a pure and simple structure is written, tested, and then allowed to grow by increasing the ambition of its modules. The process continues recursively as each

*Sandewall's excellent survey article, "Programming in an Interactive Environment: The Lisp Experience," gives an overview of existing programming methodology in the Lisp user's environment, emphasizing methods for interactive program development. It includes a comprehensive description and analysis of current Lisp programming environments in general, and Interlisp and MacLisp in particular.

module is rewritten. The principle applies not only to input/output routines but also to the flexibility of the data handled by the program, sophistication of deduction, the number and versatility of the services provided by the system, etc. The growth can occur both 'horizontally' through the addition of more facilities, and 'vertically' through a deepening of existing facilities and making them more powerful in some sense.''⁷

Computer costs vs. people costs. The second major influence in Interlisp's development was a willingness to "let the machine do it." The developers were willing to expend computer resources to save people resources because computer costs were expected to continue to drop. This perspective sometimes led to tools which were ahead of their time in respect to the available computer resources.

The Advanced Research Projects Administration of the Department of Defense sponsored much of this effort. As a result, we had a fair amount of freedom in the development of Interlisp, i.e., we did not have to justify the cost-effectiveness of our research to a profit-oriented management. ARPA's willingness to make Interlisp available at a number of sites on the Arpanet justified and motivated the extra effort it took to turn a research project into a real system. These Arpanet sites also provided an active and creative user community from which we obtained many valuable suggestions and much-needed feedback.

Interlisp is for experts. The incremental, evolutionary way in which Interlisp developed was not especially conducive to simple interfaces. It was inappropriate to spend a lot of time and effort trying to design the right interface to a new, experimental capability whose utility had not yet been proven. Would the users like automatic error correction? Was the programmer's assistant really a good idea? The inherent complexity of the interactions among some of the more sophisticated tools, such as Masterscope, DWIM, and the programmer's assistant, made it very difficult to provide simple interfaces. In many cases, unification and simplification came only after considerable experience.

Further complexity stemmed from the commitment to accommodate a wide variety of programming styles and to enable the tools to be tailored for many applications. Given the choice of sophistication and generality of tools or simplicity of design, we chose the former, under the assumption that the system was primarily for expert programmers. As a result, mastery of all of the tools and facilities of Interlisp has become quite difficult and initial learning time fairly long. We accept this as part of the price for the system's power and productivity.

Background

Programming environments have been built for a number of languages, on top of a number of operating systems, and for a variety of user communities. Each of these factors can influence the path taken in the development of a programming environment. In the case of Inter-

lisp, the Lisp language itself and the sociological factors in effect during its early development were both important.

The Lisp language. The Lisp language is conducive to the development of sophisticated programming tools because it is easy to write programs that manipulate other programs. The core syntax for the Lisp language is simple, and Lisp programs are naturally represented in simple Lisp data structures in a way that reflects the structure of the program. Since Lisp requires no declarations, programs can be built up incrementally; this is more difficult in declarative languages. This means that Lisp supports the structured growth style of program building.

Early sociology of Interlisp. One unusual historical aspect of the development of Interlisp is that from the very beginning those interested in programming environments were in a position to strongly influence the development of the language system. We were not constrained to live within the language and operating system we were given, as is usually the case. Most of the additions or extensions to the underlying Lisp language performed under the Interlisp project were in response to perceived environmental needs. For example, Interlisp permits accessing the control stack at an unusually detailed level. Capabilities such as this were added to Interlisp to enable development of sophisticated and intelligent debugging facilities. Similarly, uniform error handling was added to the Lisp base in order to permit experimentation with automatic error correction.

Some representative Interlisp facilities

File package. Interactive program development consists alternately of testing of program parts and editing to correct errors discovered during the tests and/or extend the program. Interlisp, unlike many other interactive programming systems, supports both the testing and editing operations. The user talks exclusively to the Interlisp system during the interactive session. During this process, the primary copy of the program (the copy that is changed during editing operations) resides in the programming system as a data structure; editing is performed by modifying this data structure. For this reason, Interlisp is called a residential system.⁷

In a residential system, it is important to be able to take procedures represented by data structures and print them on text files in an input-compatible format for use as backup, for transporting programs from one environment to another, and to provide hard-copy listings. In Interlisp, the file package is a set of functions, conventions, and interfaces with other system packages. The role of the file package is to automate the bookkeeping necessary for a large system consisting of many source files and their compiled counterparts. The file package removes from the user the burden of keeping track of where things are and what things have changed. For example, the file package keeps track of which file contains a particular datum, e.g., a function definition or record declaration. In many cases, it automatically retrieves the necessary

datum, if it is not already in the user's working environment. The file package also keeps track of which files have been in some way modified and need to be dumped, which files have been dumped but still need to be recompiled, etc.⁸

Once the user agrees to operate in the residential mode, it becomes possible to design and implement such powerful tools as DWIM and Masterscope to assist in program development. The file package makes this mode attractive to the user.

The significant breakthrough occurred with the emergence of the idea of having the system notice when a datum was changed and associate this fact with the file containing the datum.

The history of the file package is instructive, as it is a paradigm for the development of user facilities that has frequently been followed in Interlisp. The file package was *not* designed in a coherent, integrated way; nobody sat down and said, "We need a file package." Instead, it evolved gradually. Originally, there was only a very limited facility for symbolically saving state at the end of a session in a form that could be loaded into a Lisp system to restore that state: the PrettyDef function. This took as its arguments a list of function names, a list of variable names, and a file name. PrettyDef wrote ("prettyprinted") the definitions of the named functions and the values of the named variables onto the indicated file. PrettyDef was soon extended to take a set of commands, which could indicate not only the functions and variables to be saved, but properties on property lists, values in arrays, definitions of new editor commands, and record declarations, among others. Finally, PrettyDef was extended to allow the user to augment this simple command language by defining his own filing commands (usually in terms of existing ones).

Concurrently, as the contents of source files became more complicated, the ability to interrogate files as to their contents (e.g., which file contained a particular datum and what functions were contained in a particular file) became more important. This required the system to be able to enumerate all of the user's source files, which was accomplished by adding the file to a global list so it would be noticed when it was first loaded or dumped.

The significant breakthrough occurred with the emergence of the idea (probably through some user saying, "Wouldn't it be nice if the system . . .") of having the system notice when a datum was changed, e.g., defined for the first time, edited, redefined, or reset, and associate this fact with the file containing the datum. This enhancement was relatively straightforward since the ability to decompose and interpret the commands that described the contents of a particular file was already available. It was implemented by a function that took the name of a datum and its type (function, variable, record definition, etc.) and marked the datum as changed and therefore in need of dumping. This function, MarkAsChanged, operated by discovering which file(s) contained the datum and

associating with each such file the name of the object that had changed. Calls to MarkAsChanged were then inserted in all of the parts of the Interlisp system that changed objects—the editor, the DEFINE function, the facility for (re)declaring records, and DWIM (which can modify a function when it performs a spelling correction or some other transformation inside the function).

With this change, the file package assumed a degree of autonomy, often operating automatically and behind the scenes. Furthermore, since it was no longer a function that the user called, but included "tendrils" into many parts of the system, we began to think of it as a package. In this light, a number of extensions became apparent. For example, the "Cleanup" function provided the capability to enumerate all of the user's files and write out those that contain data that had changed. Next, an automatic warning was added to Cleanup, in case an object not associated with any file changed or was newly defined. Then a filing capability was added, which enabled the user to add a datum to a file by automatically modifying the commands for that file. Finally, Cleanup was extended to prompt the user about unfiled objects and to allow the user to specify the destinations of these objects.

By this point, the Interlisp user did not have to worry about maintaining his source files, save for occasionally calling Cleanup. The file package, however, had become smart with respect to its built-in commands, but if a user defined a new type of command, the file package would not necessarily be able to support operations such as adding an object of that type to a file, deleting or renaming an object, or obtaining the "definition" of an object of a particular type from a file. The user was placed in the position of having to choose between using an automatic facility that did just what he wanted—provided he stuck to a predefined class of file objects—or extending this facility to print out his own types of file objects, which meant returning to manual bookkeeping of symbolic files.

Thus, the next extension to the file package identified and exposed its primitive operations and allowed the user to define or change these operations. This resulted in a more complicated interface to the file package than that used directly by most, but it enabled builders of systems within Interlisp to enjoy the same privileges of defining file package operations that the original implementors enjoyed. In fact, we were able to express the semantics of all of the built-in file package commands and types in terms of the above interface. We thus eliminated all distinction between built-in operations and those defined by the user (a good test of the completeness of this lower-level interface) and permitted the user to redefine the way these operations are performed.

The file package supports the abstraction that the user is truly manipulating his program as data and that the file is merely one particular external representation of a collection of program pieces. During a session, the user manipulates the pieces with a variety of tools, occasionally saving what he has done by calling the Cleanup function. The user can also operate in a mode where programs are treated as residing in a data base, i.e., the external file

system, with a variety of sophisticated retrieval tools at his disposal.

Note the evolution of the file package. It started as an isolated facility that was explicitly invoked by the user to perform a particular and limited action. More and more capabilities were added, increasing the range of applicability of the tool. At the same time, the tool was integrated into the system to produce a semi-autonomous configuration in which the tool is invoked automatically in a number of contexts. Finally, the utility of the tool became so great that a form of user extensibility, to adapt the tool to accommodate unforeseen situations, became imperative.

The file package also illustrates one of the principal design criteria of the Interlisp system, the accommodation of a wide range of styles and applications. The user is not forced to choose between using a facility that is powerful and attractive but forces adherence to its prescribed conventions, abandoning the tool, or even creating a personal, renegade version whenever he needs a capability the tool does not provide. In other words, if a particular tool handles 95 percent of the user's applications correctly, he should be able to extend the tool in a prescribed and "blessed" manner to accommodate the remaining five percent without undue effort.

Masterscope, like the file package, has its roots in an extremely simple program.

Masterscope. As the size of systems built within Interlisp grew larger and larger, it became increasingly difficult for a user to predict the effect of a proposed change. It was also growing difficult to effect a pervasive change, for example, to change the calling convention of a low-level procedure and be sure that all of the relevant places in programs would be found and modified. Masterscope is an interactive program for analyzing and cross-referencing user programs that addresses this problem. It contains facilities for analyzing user programs to determine which functions are called, how and where variables are bound, set, or referenced, which functions use particular record declarations, etc.

Masterscope maintains a data base of the results of the analyses it performs. The user can interrogate the data base explicitly (e.g., WHO USES FOO FREELY), or have Masterscope call the editor on all functions that contain expressions that satisfy certain relations (EDIT WHERE ANY FUNCTION USES THE RECORD DICTENTRY).

Masterscope, like the file package, has its roots in an extremely simple program. Called PrintStructure, this program analyzed function definitions and printed out the tree structure of their calls. It was first extended to include the names of the arguments for each function it analyzed, and then to include more information about variable usage within each function. However, as PrintStructure presented more and more information about larger and more complicated program configurations, it became increasingly difficult for the user to extract particular information from this massive output. It became

clear that the user wanted access to specific information rather than a complete listing. This led to the idea of separating the analysis of the program from the interrogation of the data base.

The next stage was integration with other parts of the system. As in the case of the file package, the utility of Masterscope increased greatly when the burden of remembering what had changed, and therefore needed re-analysis, was lifted from the user and carried out automatically behind the scenes. The next phase of the evolution of Masterscope was to permit the user to extend Masterscope's built-in information on analysis of special Lisp forms, such as PROG, SETQ, and LAMBDA expressions. This was accomplished through use of Masterscope "templates," which are essentially patterns for evaluation of functions. Finally, all built-in information was removed from Masterscope and replaced by templates, both to test the completeness of the interface and to expose this information to users so they could change it.

DWIM. According to Sandewall, "One of the most impressive features in the Interlisp system is the DWIM (Do What I Mean) facility, which is invoked when the basic system detects an error and which attempts to guess what the user might have intended."⁷

The most visible part of DWIM⁹ is the spelling corrector, which is invoked from many places in the system, including the file package, Lisp editor, and the Lisp interpreter itself. When an unrecognized file package command, edit command, Lisp function, etc., is encountered, the spelling corrector is invoked. The spelling corrector attempts to find the closest match within a list of relevant items. If an edit command is misspelled, for example, the list of valid edit commands is searched; if the name of a function is misspelled, the corrector scans a list of the functions the user has recently been working with. If the spelling correction is successful, the cause of the error is also repaired, so subsequent corrections will not be necessary. For example, when DWIM corrects a user's misspelled function name in one of his programs, it actually modifies the user's program to contain the correct spelling (and notifies the file package of the change).

Although most users think of DWIM as a single identifiable package, it embodies a pervasive philosophy of user interface design: at the user interface level, system facilities should make reasonable interpretations when given unrecognized input. Spelling correction is only one example of such an interpretation. Depending on how far off the input is, a facility might make the transformation silently and automatically, without seeking user approval. For example, a function expecting a list of items will normally interpret an argument that is a single atom as a list made up of that single atom. In this case, the package in question probably would not even indicate to the user that it had made this correction, and in fact the user might view the package as expecting either a list or an atom. Similarly, the style of interface used throughout Interlisp allows the user to omit various parameters and have these default to reasonable values, such as "the last thing this package operated upon."

DWIM is an embodiment of the idea that the user is interacting with an agent who attempts to interpret the

user's request from contextual information. Since we want the user to feel that he is conversing with the system, he should not be stopped and forced to correct himself or give additional information in situations where the correction or information is obvious.

The iterative expression. The various forms of the Interlisp iterative expression permit the user to specify complicated loops in a straightforward and visible manner. In one sense, the iterative expression represents a language extension, but by its design, implementation, and in particular its extensibility, it more naturally falls into the same category as other Interlisp tools.

An iterative expression in Interlisp consists of a sequence of operators, indicated by keywords, followed by one or more operands; many different operators can be combined in the same iterative statement. For example, (for X in L sum X) iterates the variable X over the elements of the list L, returning the sum of each value seen. The iterative could be further embellished by including "when (GREATERP X 30)" to only sum elements greater than 30, or "while (LESSP \$VAL 50)" to terminate the iteration when the sum exceeds 50. Other operators can be used to specify different ranges. For example, iteration can take place over a range of numbers instead of over the elements of a list: (for i from 1 to 10 . . .). Operators can also specify the value returned by the iterative expression. For example, (for X in L collect (ADD1 X)) would return a new list, consisting of the elements in L, each incremented by 1.

The iterative expression currently understands approximately two dozen operators. Furthermore, new iterative operators can be defined simply. One group, experimenting with a relational data base system, provided access to that data base merely by defining a new iterative operator called "matching." This matching operator can be used in conjunction with all of the other iterative constructs, as in "(for Records matching (payment (> 30) *) sum Record: 3)," which would find all payment records in the data base and sum their third component. Such language extensions are quite difficult in most programming languages.

Programmer's assistant. The central idea of the programmer's assistant is that the user is not talking to a passive executive that merely responds to each input and waits for the next, but is addressing an active intermediary.¹⁰ The programmer's assistant records, in a data structure called the history list, the user's input, a description of the side effects of the operation,⁸ and the result of the operation.

The programmer's assistant also responds to commands that manipulate the history list. For example, the REDO command allows the user to repeat a particular operation or sequence of operations, the FIX command allows the user to invoke the Interlisp editor on the specified events and then re-execute the modified operations, the USE command performs a substitution before re-executing a specified event (e.g., USE PRINT FOR READ), and the UNDO command cancels the effect of the specified operations. In addition to the obvious use of recovering information lost through typing errors, UN-

DO is often used to selectively flip back and forth between two states. For example, the user might make some changes to his program and/or data structures, run an experiment, undo the changes, rerun the experiment, undo the undo, and so on.

The various replay commands, such as REDO and FIX, permit the user to construct complex console operations out of simpler ones, in much the same fashion as programs are constructed. That is, simple operations can be first checked and then combined into large ones. The system always remembers what the programmer has typed, so that keyboard input can be reused in response to an afterthought.

The programmer's assistant has been implemented for use in contexts besides the handling of inputs to the Lisp "listen" loop. For example, the Interlisp editor also uses the programmer's assistant for storing operations on the history list and thereby provides all the history commands for use in an editing session. Similarly, user programs can take advantage of the history facility. A system for natural language queries of a data base of lunar rock samples provides one example of how this facility can be used. After a complicated query regarding the percentage of cobalt in a sample, a user could say USE MANGANESE FOR COBALT to repeat the query with a different parameter.

What makes Interlisp unique?

The Interlisp programming environment has been characterized as friendly, cooperative, and forgiving. While these qualities are desirable, they are not unique to Interlisp. The two attributes that set it apart are the degree to which the system is integrated and the degree to which facilities in the environment can be tailored, modified, or extended.

Integration. Interlisp is not merely a collection of independent programming tools, but an integrated system. By integration, we mean that there need not be any explicit context switch when switching between tasks or programming tools, in switching, for example, from debugging to editing to interrogating Masterscope about the program. Thus, having called the editor from inside the debugger, the user can examine the current run-time state from within the editor or ask Masterscope a question without losing the context of the editing session. Also, the various facilities themselves can use each other in important ways, since they all coexist in the same address space. For example, the editor can directly invoke DWIM, or Masterscope commands can be used to drive the editor. The integration of facilities increases their power.

Integrated programming tools such as these are not feasible without a large virtual address space. Where the size of the programming environment is constrained, it is unreasonable to have a large variety of resident tools that can all interact with the user's run-time environment and with each other. Interlisp-10's large virtual address space of 256K 36-bit words (large, at least, for the early 1970's) made it possible to add new features to the programming

