From UNCOL to ANDF: Progress in Standard Intermediate Languages

Stavros Macrakis
Open Software Foundation

1. Introduction

ANDF is an architecture- and language- neutral distribution format resembling a compiler intermediate language. Architecture- and language-neutral intermediate languages were first discussed in the UNCOL project of 1958-1962. UNCOL was never fully defined nor implemented. This paper takes a look at the history of UNCOL and shows why ANDF can succeed where UNCOL failed. The paper presupposes an understanding of comparative programming language semantics and implementation, and is addressed to programming language implementors.

The UNCOL section describes the history and nature of UNCOL. UNCOL was a very ambitious project for its day, and would have required innovations in many areas. But later work on compilers and intermediate languages in fact fulfilled many of the promises of the UNCOL work. Those innovations, together with the widespread use of portable software and open systems made a universal distribution format economically worthwhile. But the technical groundwork had been laid by the compiler community.

The Fortran, C, and Lisp sections discuss a different approach to a universal language—the use of existing programming languages. They have been used as an intermediate language for Ada, Modula-3, and other compilers.

An examination of these compilers shows the strengths and weaknesses of this approach.

The final section discusses the relationship between portability and a universal distribution format: ANDF cannot make programs portable; it provides the mechanisms necessary to support portable programs.

2. UNCOL

What was UNCOL?

UNCOL, a **Un**iversal **C**omputer-**O**riented **L**anguage, is the general name for a group of proposals for a universal intermediate language for compilers.

UNCOL's goal was quick and economic production of compilers for a full range of languages on new machine architectures. Steel presents the economic justification clearly:

> Th[e] capital investment in a translator [*i.e.* compiler] would be well advised if our concern were for a single machine and a single problem oriented language…. In our firm insistence on many machines and many languages, however, the requisite investment is increased multiplicatively…. A primary objective of the effort whose direction is outlined in the following pages is a reduction in the time and money required to live in this changing environment.
>
> …if one is presented with M problem languages and N machine languages, M + N translation programs are required in the UNCOL scheme of things, while M × N translation programs are necessary in the traditional mode…. [Steel61, p. 371]

UNCOL, then, was a means to an end: compiler production. UNCOL was never intended to be a mechanism for porting programs; indeed the very notion of a portable program was hardly mentioned in the UNCOL literature!

According to [Share58], the concept of a universal intermediate language had "been discussed by many independent persons as long ago as 1954. It might not be difficult to prove that 'this was well-known to Babbage,' so no effort has been made to give credit to the originator, if indeed there was a unique originator." The most usual bibliographic reference is to Steel's papers [Steel60, 61] and the [Share58] report.

Various UNCOLs were proposed, but none in detail. Some concepts were utopian:

> [the data description] language includes the first order predicate calculus with identity…. In such languages non-constructable [*sic*] items occur…. In order to circumvent trouble from this source, a set of rules for generator writers must be given which limits the complexity and character of the definitions; rules which the careful framer of definitions will follow instinctively. [Steel61, p. 374]

UNCOL was only one approach proposed to solve the compiler construction problem. An alternative was the compiler-compiler, of which Steel says:

> A single program is written, *exactly once*, which takes as input descriptions of a problem oriented language and a machine language and then proceeds to produce as output a compiler that translates the given problem language statements into the given machine language. The trouble with this approach is that nobody has even the vaguest idea about how to do it despite occasional statements to the contrary. [Steel60, p. 20]

> The compiler-compiler is a more dignified solution…. [It] will be a program that accepts as input a description of a problem oriented language, a description of a compilation machine and a description of an object machine and produces as output a program which runs on the first machine and translates the problem language into the language of the second machine. This is all very simple conceptually. The trouble is that nobody knows how to do it. One of the hoped-for byproducts of the UNCOL effort is sufficient insight into the mechanisms of translation to permit a beginning on [the compiler-compiler]. Until then UNCOL will have to do. [Steel61, p. 344]

But Steel's prediction was incorrect. In fact, progress on intermediate languages and on compiler-compilers has gone hand in hand.

Computer Technology in 1958

A striking feature of the UNCOL papers is the number of innovations that would have been needed to make it work.

Since there was no standard character code, UNCOL would have had to define one. One proposal included over 500 characters thought useful for mathematical notation (including, *e.g.*, black-letter subscripts). On the other hand, no thought was given to national language support.

Bootstrapping was a novel technique. Many critics of UNCOL could not understand how it would be possible to bring an UNCOL compiler up on a new machine. Thus, much of the text of UNCOL papers describes bootstrapping and defends its feasibility. [Steel62]

Indeed, the very notion of an intermediate language for compilers (as opposed to *ad hoc* data formats to pass information between compiler passes) was rather novel.

At the same time, programming language technology was in its infancy. Concepts that we now take for granted were novel or even non-existent, such as records, pointer types, and for that matter data types in general.

Developments in Compiler Technology

Compiler technology has been an active area of research from the 1960's to the present. It is useful to review the chronology of developments to put universal intermediate languages into context.

late 1950'	Inventon of regular languages and finite-state automata—basis for lexicall analysis
	Invention of context-free languages, pushdown automata, and Backus form—basis for syntactic analysis.
1960's:	Development of lexing and parsing algorithms.
	Elaboration of efficient run-time structures for Algol.
1970's:	Development of theory of semantics
	First table-driven code generators.
	Broader use of ILs.
	First compiler-compilers.
1980's:	Application of semantic theory to front-end construction (experimental).

In the open literature, compiler intermediate languages are presented as novel well into the 1970's. This is misleading, since many proprietary compilers used this technique. But it does show that ILs were not fully assimilated into the computer science culture.

Automatic generation of back ends for compilers began in the 1970's and is now the dominant mode of production of compiler back ends.

Front ends are currently not as automated. Although lexical and syntax analysers are routinely generated from formal specifications (regular expressions and BNF), semantics is usually treated by explicit program.

But even for semantics, innovations of the 1970's and 1980's now make it possible to produce a full compiler-compiler: given the syntax (in BNF) and semantics (in some denotational formalism) of the language on the one hand, and the machine description on the other, a compiler can be produced completely automatically. However, runtime performance is currently poor in this fully automatic denotational approach.

Compiler-compilers are routinely used today[1]. Their intermediate languages, though not standard, are constant for a particular compiler-compiler, with minor variants for new machines or languages.

These ILs meet the primary requirement for an UNCOL (to make compiler writing easier), but not all the requirements for ANDF. A design process driven by the specific requirements of ANDF would incorporate IL portability and machine and architecture neutrality from the beginning. This is precisely what the UK Defence Research Agency has done with TDF, the technology selected by OSF as the basis of ANDF. Thus the danger of continual incremental modifications for new languages and machines is averted.

Using existing languages for distribution

A different approach to intermediate languages for distribution is the use of existing programming languages.

The advantage is that they already have compilers on many machines, and if the source language is the same as the chosen distribution language, then the front end becomes trivial.

However, there are disadvantages both on the producer (front-end or preprocessor) and installer (back-end) sides.

1. although the term 'compiler-compiler' is outmoded

The front end must translate source language semantics into the intermediate language. Sometimes, this is straightforward. But more often, there are hidden complications. For instance, languages' treatment of loop termination conditions may be different. More subtly, languages' treatment of variable values after error conditions may be different. Often, a language feature available in the source language is not available in the target, and must thus be emulated. Such an emulation may mean premature implementation decisions. Also, machine-dependent implementation techniques not used by the IL's compiler are not available to the source language translator. These complications make producer design more difficult, reduce the quality of the installed code, and favor the particular language over others.

The back end must translate the intermediate language into machine language. For the result to be predictable, the intermediate language must be fully specified in a portable way. This means not only specifying the semantics independently of the machine, but also providing a complete and portable set of environment inquiries, which can be used either at installation time or at runtime. Full specification and complete environment inquiries are found in few languages.

Many compilers do not handle program-generated code well. Program-generated code often contains peculiar constructs and often exceeds compilers' capacity limits. For that matter, compilers are tuned to provide good code for typical hand-written programs, and not necessarily for translations of programs written in other languages.

Another important objection to using existing languages is the problem of reverse engineering. If protection of proprietary programming is required, the front end is no longer trivial even if the source language and the intermediate language are the same, since it must obscure or 'shroud' the source code.

3. Fortran

Fortran has often been used in the past as an intermediate language because of its simple semantics, relatively good standardization, and wide availability. Nowadays it would usually not be considered because of its

lack of such fundamental features as records and pointers and weak support for character manipulation.

But writing portable Fortran is difficult. Indeed, such major producers of portable Fortran as IMSL and NAG have extensive software toolkits to support the process (cf. [Boyle77]). Depending on such elaborate preprocessing defeats the purpose of a standard intermediate language.

4. C

C as an Intermediate Language

Many language implementations generate C as an IL, both for C extensions (*e.g.* C++) and other languages (Ada [Meridian], Cedar, Eiffel, Modula-3,[1] Pascal [Bothe89], Sather [Lim91], Standard ML [Tarditi91]).

C is widely available and generally compiles into efficient code. Moreover, C's low-level pointer constructs and weak typing allow easy emulation of many other languages' constructs (*e.g.* passing parameters by reference).

C is a particularly appropriate intermediate language for prototyping (in the absence of a standard intermediate language). In prototype implementations, portability and standardization are often not issues. But experimentation shows that good efficiency requires careful choice of idiom in C to translate source language features.[2]

Disadvantages of C as IL

For an intermediate language, many areas of C semantics are implementation-dependent, that is, underdefined. For instance, there is no

1. *cf.* the discussion about the use of C as an intermediate language for Modula-3 compilers on the comp.compilers newsgroup. David Chase's remarks <1990Aug14.163258.2094@esegue.segue.-boston.ma.us> are particularly pertinent.
2. The Sather work shows this quite clearly.

way of specifying integer or floating-point precision, nor for that matter of querying them.

Other areas are overdefined; for instance, parameter passage is always by value, although reference passage or copy-in/copy-out may be more appropriate on some architectures. It is possible to implement these other mechanisms explicitly, but that precludes the installer from choosing an implementation strategy as a function of the target architecture.

C's data definition mechanisms are weak. There is no way, for instance, to declare a variant record with discriminant nor a variable-size array. They can of course be emulated, but this means committing to a particular implementation which may not be appropriate for the target machine.

Many constructs needed for other languages are missing from C. Lexical scope is missing. Exceptions are missing. Safe pointers (for garbage collection) are missing. All of these can be emulated in C, but only by committing to a particular runtime model.[1] Such overspecification reduces efficiency. Also, invariants preserved by the emulation are unknown to the C compiler and thus cannot benefit optimization.

For all of the above reasons, although C has been useful for prototyping extensions to itself (C++) and for producing code rapidly for other languages, C is not an ideal intermediate language.

Discussion: C as a Distribution Format

As an intermediate language for languages other than itself, C is not very strong. C does have a specific advantages as a distribution format, however: if source code is distributed, the mechanisms for pre-processing and machine-dependent static quantities are handled with no additional effort.

Still, many of C's disadvantages as an IL carry over to C as a distribution format.

1. Consider the programming conventions required to support safe pointers in GNU Emacs Lisp.

In addition, the complete semantics of a C program are not specified by the C language; parameters to linkers and loaders may change symbol interpretations.

In principle, many of these objections could be overcome by re-engineering of compilers, careful definition of standard special-purpose datatypes (int_12, int_13, …), standardization of linkers, and more flexible optimizers. But if standard C compilers cannot be used, why use the standard C language for a purpose for which it was not intended?

5. How about Lisp?

Sometimes Lisp is suggested as a universal intermediate language. Indeed, ANDF has some superficial resemblances to Lisp.

Which one?

There are apparently three different things meant when people speak of using Lisp as an intermediate language:

- Using a fully parenthesized notation.
- Using traditional dialects of Lisp.
- Using Scheme as an object-oriented language.

Fully parenthesized notation

The major classes of intermediate languages are linear (triples, quadruples, tuples), tree-structured, and graph-structured. In all cases, a key decision is the actual operators used and their exact semantics.

ANDF is in fact a tree-structured language. Its operators have been carefully chosen to be architecture- and language- neutral.

Lisp is tree-structured as well, but its particular set of operations is specific to Lisp semantics. Replacing the operators with another set of operators would result in a different tree-structured language, not Lisp.

Using traditional dialects of Lisp

Lisp, as a programming language developed for human use, shares several of C's shortcomings as an intermediate language, in particular underspecification (implementation-dependence) for some constructs, and overspecification for others.

An example of underspecification is the lack of a machine-independent way of specifying the range of integers. As for overspecification, Lisp parameter passing is by value for certain primitive types and by sharing for composite types. Other languages may require different semantics. On a more practical level, no commercial Lisp compiler allocates composite objects (records, arrays) on the stack.

So using an existing dialect of Lisp would require radical reworking both of the language and of its compilers.

Scheme-like object orientation

Scheme used as an object-oriented language has several good properties: it is well-defined, it is abstract and thus uncommitted to a particular implementation style, it has very general control structure.

However, it has two fatal flaws: lack of concrete data typing facilities and requirement for garbage collection. Also, none of its implementations are as efficient as implementations of traditional sequential programming languages.

6. Portability and Distribution

Portability is best designed into a software product, and not added on later. Sometimes ANDF is thought of as a portability tool which will mechanically turn a non-portable program into a portable one. If ANDF claimed this, it would be justly criticized for trying to solve an unsolvable problem.

But ANDF is a distribution format for portable software. ANDF is not a tool for making non-portable software into portable software, but a tool for

distributing portable software, which must thus provide *mechanisms* needed by portable software.

Barriers to portability

Portable software is defined as software that can easily be moved from one execution environment to another. Software can be portable as source or as binary of some kind (object, load module, etc.).

At the source level, portability can be blocked by language incompatibility (supersets and subsets), compiler incompatibility (bugs or extensions), non-portable assumptions about the runtime environment (*e.g.* precision of numbers), or use of unavailable or incompatible interfaces.

Language and compiler incompatibility is usually solved by writing in a subset of the language known to be correctly implemented by a wide range of compilers. This subset is typically more restrictive than the official standard, and is determined empirically (!). Sometimes conditional compilation is used to avoid compilers' weak spots.

Environment assumptions are treated by parameterizing the software (preferably at compile time), often using macros. Many languages do not provide standard environment queries, so programs must be explicitly parameterized by installation machine or determine environmental parameters through *ad hoc* mechanisms.

Availability of interfaces can be dealt with either by restricting oneself to the least common denominator interfaces or by providing a platform-dependent layer which translates to native mechanisms. When it is not the interface that varies, but the functionality itself (*e.g.* shared memory or threads), software must sometimes be extensively parameterized or even redesigned.

At the binary level, software is of course only portable among machines of the same instruction set architecture. Moreover, the object or load format must be compatible, as must all system interfaces, including device registers (when accessible to the software), *etc.* When these interfaces are ill-defined, binary compatibility can be a major restraint on architectural innovation, as witness MS-DOS.

When software is distributed as a linkable module, global names must be guaranteed to be unambiguous. This can be a particular problem when several large systems are linked together.

ANDF support for portability

A distribution format such as ANDF cannot make software portable. ANDF provides the *mechanisms* needed for writing and distributing portable software.

Language and compiler incompatibility are avoided by using the same translator (front-end) for all machines. This has the important side-effect of allowing software implementors to use the full power of their programming language, rather than restricting themselves to the subset known to be correctly implemented by all compilers. Back-end validation is of course a critical link in this chain, but since the intermediate language is better-defined than most programming languages, this is feasible.

ANDF provides a full set of environment inquiries to support compile- and run-time parameterization. ANDF also provides explicit specification of numerical precision.

ANDF provides install-time conditions and parameterization of both executable code and of data definitions, thus supporting explicit parameterization when necessary.

ANDF's modular structure also allows for target-specific libraries of code, data definitions, and macros which can present a standard interface for varied low-level interfaces. Of course, ANDF cannot create new functionality where it does not exist, but it can conditionally use different mechanisms in different environments.

ANDF completely eliminates the problems associated with distribution of programs in binary. Naturally, it is independent of instruction-set and load-format. But it also guarantees uniqueness of names for library entry points.

7. Conclusion

A universal intermediate language has been a dream for many years.

UNCOL was an ambitious effort for the early 1960's. An attempt to solve the compiler-writing problem, it ultimately failed because language and compiler technology were not yet mature.

In the 1970's, compiler-compilers ultimately contributed to solving the problem that UNCOL set itself: the economical production of compilers for new languages and new machines. Although their intermediate languages were fairly invariant by language and machine, this was not their principal goal.

With the growth of open systems, distribution of portable programs became more important. Neither UNCOL nor compiler-compiler technology had addressed this issue. But a standard intermediate language would permit true open systems, where programmers could choose their language independently of the implementation platform, and hardware vendors could choose their hardware implementation independently of the installed base of program binaries.

Programming languages solved part of the problem: Fortran and C were often used as intermediate languages. But they did not solve it all: neither was really suitable as an intermediate language.

A closer analysis of the specific needs of portable programs showed that particular mechanisms were essential. By designing an intermediate language from scratch which took account of these requirements, ANDF succeeds as a universal intermediate language.

UNCOL was the first step in three decades' work in software portability and compiler design whose culmination is ANDF.

8. Bibliography

UNCOL

[Conway58] Melvin E. Conway, "Proposal for an UNCOL," *Commun. ACM* 1:3:5 (1958).

Suggests using a single-address model for UNCOL. Sketchy.

[Sammet69] Jean E. Sammet, *Programming Languages: History and Fundamentals,* Prentice-Hall, 1969. Chapter X.2: UNCOL (Significant Unimplemented Concepts), p. 708.

[Share58] Share Ad-Hoc Committee on Universal Languages (J. Strong, J. Olsztyn, J. Wegstein, O. Mock, A. Tritter, T. Steel), "The Problem of Programming Communication with Changing Machines," *Commun. ACM* **1**:8:12 (August 1958) and **1**:9:9 (September 1958).

Describes the requirements for UNCOL, the general concept, and the idea of bootstrapping. No proposal for the language itself.

[Steel60] T. B. Steel, Jr., "UNCOL: Universal Computer Oriented Language Revisited," *Datamation* (January/February 1960), p. 18.

Presents the UNCOL concept, but no technical details.

[Steel61] T.B. Steel, Jr., "A First Version of UNCOL," *Proc. Western Joint Computer Conference* **19** (Los Angeles, May 9-11, 1961), p. 371.

Proposes an elaborate character set, the first-order predicate calculus for a kind of syntactic data description, a notation for indirect addressing and indexing, a set of machine-level primitives, and a vague description of macros and declaratives.

[Steel61AP] T. B. Steel, Jr., "UNCOL: The Myth and the Fact," *Ann. Rev. in Automatic Programming* **2** (1961), p. 325.

Discusses primarily the economics of UNCOL and the concept of bootstrapping.

Intermediate Languages

[Brosgol80] Benjamin M. Brosgol, "$TCOL_{Ada}$ and the 'Middle End' of the PQCC Ada Compiler," Proc. ACM-Sigplan Symposium on the

Ada Programming Language (Boston, December 9-11, 1980) in *Sigplan Notices* **15**:11:101 (November 1980).

[Brown72] P. J. Brown, "Levels of Language for Portable Software," *Commun. ACM* **15**:12:1059 (December 1972).

High and low level macro languages as intermediate languages.

[Chow83] Frederick C. Chow and Mahadevan Ganapathi, "Intermediate Languages in Compiler Construction—A Bibliography," *Sigplan Notices* **18**:11:21 (November 1983).

See also [Ottenstein84].

[Coleman73] S. S. Coleman, P. C. Poole, and W. M. Waite, "The Mobile Programming System, Janus," *Software—Practice and Experience* **5**:5 (1974).

[Ganapathi84] Mahadevan Ganapathi and Charles N. Fischer, "Attributed Linear Intermediate Representations for Retargetable Code Generators," *Software—Practice and Experience* **14**:4:347 (April 1984).

[Griswold77] Ralph E. Griswold, "An Alternative to SIL," in [Brown77], p. 291.

[Haddon78] B. K. Haddon and W. M. Waite, "Experience with the Universal Intermediate Language Janus," *Software—Practice and Experience* **8**:601 (1978).

[Kornerup80] Peter Kornerup, Bent Bruun Kristensen, and Ole Lehrmann Madsen, "Interpretation and Code Generation Based on Intermediate Languages," *Software—Practice and Experience* **10**:635 (1980).

[Lamb87] David Alex Lamb, "IDL: Sharing Intermediate Representations," *ACM Trans. Prog. Lang. and Sys.* **9**:3:297 (July 1987).

A formalism for describing data structures. Has been used to describe intermediate languages (*cf.* TCOL.Ada).

[Nelson79] Philip A. Nelson, "A Comparison of PASCAL Intermediate Languages," Proc. Sigplan Symposium on Compiler Construction (Denver, August 6-10, 1979) in *Sigplan Notices* **14**:8:208.

[Ottenstein84] Karl J. Ottenstein, "Intermediate Program Representations in Compiler Construction: A Supplemental Bibliography," *Sigplan Notices* **19**:7:25 (July 1984).

A supplement to [Chow83].

[Teller80] J. Teller, “Intermediate Languages,” unnumbered Siemens technical report (November 1980).

[Waite76] W.M. Waite, “Intermediate Languages: Current Status,” Portability of Numerical Software (Oak Brook, 1976) in *Lecture Notes in Computer Science* **57**: 269.

Portable Compilers

[Elsworth78] E. F. Elsworth, “Compilation via an Intermediate Language,” *Computer J.* **22**:3:226 (1978).

[Hennessy86] John Hennessy and Mahadevan Ganapathi, “Advances in Compiler Technology,” *Ann. Rev. Computer Science* **1**:83 (1986), p. 83.

[Johnson78] S. C. Johnson, “A Portable Compiler: Theory and Practice,” *Proc. Fifth Annual ACM Symposium on Principles of Programming Languages* (Tucson, January 23-25, 1978).

The Portable C Compiler (pcc).

[Lecarme78] Olivier Lecarme and Marie-Claude Peyrolle-Thomas, “Self-compiling Compilers: An Appraisal of their Implementation and Portability,” *Software—Practice and Experience* **8**:149 (1978).

[Leverett80] Bruce W. Leverett *et al.*, “An Overview of the Production-Quality Compiler-Compiler Project,” *Computer (IEEE)* (August 1980), p. 38.

[Richards71] M. Richards, “The Portability of the BCPL Compiler,” *Software—Practice and Experience* **1**:135 (1971).

[Richards77] M. Richards, “The Implementation of BCPL,” in [Brown77], p. 192.

[Tanenbaum83] Andrew S. Tanenbaum, Hans van Staveren, E.G. Keizer, and Johan W. Stevenson, “A Practical Tool Kit for Making Portable Compilers,” *Commun. ACM* **26**:9:654 (September 1983).

The Amsterdam Compiler Kit: perhaps the best-known compiler-compiler.

[Waite77T] W. M. Waite, “Theory,” in [Brown77], p. 7.

[Waite77J] W. M. Waite, “Janus,” in [Brown77], p. 277.

[Wirth71] N. Wirth, “The Design of a Pascal Compiler,” *Software—Practice and Experience* **1**:309 (1971).

Portability

[Boyle77] James M. Boyle, “Mathematical Software Transportability Systems—Have the Variations a Theme?,” in [Brown77], p. 357.

[Brown77] P.J. Brown, ed., *Software Portability: An Advanced Course*, Cambridge University Press (1977), p. 357.

[Brown77B] P. J. Brown, “Basic Implementation Concepts,” in [Brown77], p. 20.

[Newey72] M. C. Newey, P. C. Poole, and W. M. Waite, “Abstract Machine Modelling to Produce Portable Software—A Review and Evaluation,” *Software—Practice and Experience* **2**: 107 (1972).

[Waite75] W. M. Waite, “Hints on Distributing Portable Software,” *Software—Practice and Experience* **5**: 295 (1975).

Mostly about distribution media and character codes.

C as an Intermediate Language

[Bothe89] Klaus Bothe, Christian Horn, “Übersetzung zwischen höheren Programmiersprachen: eine Lösung des UNCOL-Problems?,” *Angewandte Informatik* (September 1989), p. 283.

Proposes C as UNCOL.

[Lim91] Chu-Cheow Lim, Andreas Stolcke, “Sather Language Design and Performance Evaluation,” Computer Science Division, U.C. Berkeley TR-91-034 (May 1991).

Implementation of an object-oriented language via a C intermediate language. Sather programs pay a small penalty compared to C source programs, but it is not clear how much of this is due to Sather’s features and how much to going through C.

[Tarditi91] David Tarditi, Anurag Acharya, Peter Lee, “No Assembly Required: Compiling Standard ML to C,” Unpublished technical report, School of Computer Science, Carnegie-Mellon University.

Using C as an intermediate language, ML is successfully implemented with a factor of 2 loss in speed, while preserving portability and true tail-recursion.

For further information please contact:

Stavros Macrakis
macrakis@osf.org
(617) 621-7356

Copyright 1993 by Open Software Foundation, Inc.

All Rights Reserved

Permission to reproduce this document without fee is hereby granted, provided that the copyright notice and this permission notice appear in all copies or derivative works. OSF MAKES NO WARRRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. OSF shall not be liable for errors contained herein or for any direct or indirect, incidental, special or consequential damages in connection with the furnishing, performance, or use of this material.