A History of UNIX[®] before Berkeley: UNIX[™] Evolution, 1975–1984

Copyright © 1984 Ian F. Darwin

Darwin Open Systems

Geoffrey Collyer

University of Toronto*

ABSTRACT

This article traces some of the intermediate history of the UNIX Operating System, from the mid nineteen-seventies to the early eighties. It is slightly updated from an article that appeared as "The Evolution of UNIX from 1974 to the Present, Part 1" in *Microsystems* [Darw1984a]. It was intended as part 1 of 3; unfortunately that issue was also the last issue of *Microsystems*. This article discusses "Research UNIX": V6, V7 and V8; and tells the tale of many programs and subsystems that are today part of 4BSD, System V or both. Subsequent articles were planned to discuss in more detail the history of Berkeley UNIX, System V, and commercialised UNIXes. We have not written those other articles; this article is being submitted to DaemonNews in hopes that those who have written other histories of other parts of UNIX's history will come forward.

1. Introduction

Nobody needs to be told that UNIX is popular today. In this article we will show you a little of where it was yesterday and over the past decade. And, without meaning in the least to minimise the incredible contributions of Ken Thompson and Dennis Ritchie, we will bring to light many of the others who worked on early UNIX versions, and try to show where some of the key ideas came from, and how they got into the UNIX of today.

Our title says we are talking about UNIX evolution. Evolution means different things to different people. We use the term loosely, to describe the change over time among the many different UNIX variants in use both inside and outside Bell Labs. Ideas, code, and useful programs seem to have made their way back and forth – like mutant genes – among all the many UNIXes living in the phone company over the decade in question.

Part One looks at some of the major components of the current UNIX system – the text-formatting tools, the compilers and program development tools, and so on. Most of the work described in Part One took place at "Research", a part of Bell Laboratories (once Bell Telephone Laboratories, then AT&T Bell Laboratories, now a part of Lucent Technologies; then as now "the Labs") located primarily at Murray Hill, NJ and Holmdel, NJ, and the ancestral home of UNIX. In planned (but not written) later parts, we would have looked at some of the myriad versions of UNIX – there are far more than one might suspect. This includes a look at Columbus, PWB (Programmer's Workbench) and USG (UNIX Support Group) UNIXes and at U. C. Berkeley (BSD) UNIXes: CBUNIX, PWB/UNIX, UNIX/TS, and BSD UNIX respectively. We use the term 'USG UNIX' somewhat loosely in this paper to refer to PWB and USG UNIXes and System III and System V, since these (especially later System V) are the merger of the major internal (to the Bell System) variants of UNIX and 'USG' has long been the designator for these systems outside of the

^{*} now with Bell Laboratories, Murray Hill, New Jersey 07974, a part of Lucent Technologies

Bell System, probably in part due to #ifdef USG in source code from these systems. You'll begin to get a glimpse inside the history of the UNIXverse.

2. Basic Sources

Since we can't say *everything* about UNIX in this article, we'll give some pointers for those who want to read more. Full acknowledgements will be found at the tail end of each installment. But some basic sources must be mentioned.

It is a truism that the final source of information about UNIX is UNIX itself. And this, of course, requires that you have a source license. And to get a source license, you must sign in blood that you will not divulge the source code of UNIX in any way, shape or form. So, in preparing this article, we have stayed clear of looking at source code. But there are times when you want to do so, not so much to find out how some feature evolved as to see how it really works. (Since this was written in 1984, the famed Lions book [Lion1977a] exegesis of the UNIX kernel has been published openly. We are saddened to note the passing of John Lions in late 1998. And, of course, the source code of the three free Berkeley UNIX derivatives is readily available.)

The UNIX Manuals are a prime source of information about UNIX. It's trendy to deride the UNIX manuals, but for their intended purpose and audience the Bell Laboratories man pages generally exemplify good technical writing: concise, accurate, and to the point. The *Programmer's Manual* or *User's Manual* as it is variously called, more colloquially known as "Volume 1", summarises in a standard format each command, system call, library function, and many special files (in the technical sense!), system file formats, games, miscellany, and maintenance information. Comparing a series of manuals of different vintages offers the student of UNIX evolution a good view on changing conditions.

Volume Two of the Manual set (beginning with PWB and V7; before that it all fit in one binder) is a series of short papers. These range from notes on installing the system to reference manuals on compilers to introductory tutorials. These, too, are typically well-written but occasionally incomplete. They are concise and to the point; some people find this obscure. Remember the audience and the background. The papers are written for the benefit of someone with the source code and with some knowledge of the system. It was always assumed that you would have somebody around to help you – a wizard. Or you would go to the conferences and ask others about problems. A careful reading of the manuals was (and is) required to become a wizard, along with hands-on time spent *using* (and eventually modifying) the system, learning by doing.

These papers, and others written at Research, established an interesting tradition, so counter to mainstream computerdom: *You write the program, you write the documentation*. (Later examples of this include the Literate Programming project and the Java language document generator *javadoc*.) In almost every case, the authors of the program are the authors of the paper describing its details. And in almost every case, acknowledgement is made to those who contributed significant ideas, advice or moral support to the project. This, of course, has made our work in this paper easier. It also speaks volumes about management and about programmers – both those programmers who write effective summaries of their programs, and those who don't condescend to do so.

And the UNIX manuals are sometimes derided for the "BUGS" section. This is the place where the author(s) of a program list its design limitations. One UNIX critic said of this policy: "If they know about the bugs, why don't they fix them?" The point is that the early UNIX authors established the beneficial habit of documenting limits to the program, rather than always letting the end user find them. Dennis Ritchie comments: "Every other manual has bugs sections; they just aren't published." Many of the BUGS sections were intended as pointers for further development of the programs, rather than as warnings to the user. Ritchie adds: "Our habit of trying to document bugs and limitations visibly was enormously useful to the system. As we put out each edition, the presence of these sections shamed us into fixing innumerable things rather than exhibiting them in public. I remember clearly adding or editing many of these sections, then saying to myself "I can't write this," and fixing the code instead." [Ritchie, personal correspondence]

After the manuals, another important series of papers in a similar vein appeared in the *Bell System Technical Journal*, July-August 1978. This special issue – part 2 of the July-August 1978 issue [1978a] – is

often referred to as "the blue book" for its blue binding (though all issues of the BSTJ from this time period were blue). The magazine is now called *Bell Labs Technical Journal* and did another special issue on UNIX in October, 1984[1984a]. This issue is must reading for the serious UNIXophile.

Many of the technical reports from Research are published as Computing Science Technical Reports (CSTRs); those still in print are available from Bell Labs and some of them can be found on the World-Wide Web, at http://www.cs.bell-labs.com/cm/cs/cstr.html currently.

Brian Kernighan has co-written several books containing interesting historical details. We quote later from *Software Tools* [Kern1976a], a book he wrote with P. J. Plauger, and *The UNIX Programming Environment* [Kern1984a] *written with Rob Pike*.

Finally, access to a nearly-complete collection of back issues of *;login:*, the journal of the USENIX Association, has been invaluable.

The only way to keep abreast of ongoing development work in the UNIX community is to attend, or at least read the proceedings of, the USENIX meetings held twice a year. Everyone doing serious technical work presents it here. Other conferences are more introductory or marketing-oriented.

3. Text Processing Tools

One of the guiding lights of the UNIX utilities or software tools has been the deeply-felt conviction that text should be stored in as simple, as general a format as possible so that any program can easily process it. This idea (sometimes attributed to the late Joseph Ossanna; see below) has had the widest impact possible on UNIX in all its varieties. However, there has been a regrettable tendency to move away from it in recent times, especially among commercial software developers.

We have rather arbitrarily divided the software tools into text processing tools and program development tools. Remember that UNIX makes no distinction between text files, program files and data files. Many of the same techniques can be applied to all three. But more on this later. First, an outline of the major tools and their development.

3.1. An old editor made new

The original standard UNIX line-mode text editor *ed* has a lineage longer than many of us do. As early as 1969, the first assembly-language version of *ed* was in place. Although later rewritten in C, the editor is fundamentally the same program as used then. As Kernighan and Plauger wrote in 1976,

The earliest traceable version of the editor presented here is TECO, written for the first PDP-1 timesharing system at MIT. It was subsequently implemented on the SDS-940 as the "quick editor" QED by L. P. Deutsch and B. W. Lampson; see "An online editor,", *CACM* December, 1967 [Deut1967a]. K. L. Thompson adapted QED for CTSS on the IBM 7090 at MIT, and later D. M. Ritchie wrote a version for the GE-635 (now HIS-6070) at Bell Labs.

The latest version is *ed*, a simplified form of QED for the PDP-11, written by Ritchie and Thompson. Our editor closely resembles *ed*, at least in outward appearance. [Software Tools, page 217]

['CACM' is the Communications of the Association for Computing Machinery.]

This is not to say that *ed* is the same as the TECO found on some DEC computers. For one thing, TECO is character-oriented while *ed* is line-oriented. It seems rather a case of "common ancestry."

During the 1970's, the editor went through countless revisions. Nearly every university had its own modified versions of *ed* and *qed*; some had several modified versions. Jay Michlin of Bell Labs wrote (in IBM assembler) a QED for IBM's mainframe TSO; this was released to Universities in the mid-70's. This was, in fact, one of my (Darwin) earlier exposures to the UNIX philosophy; around 1975, I heard about a "spiffy new editor" for TSO (the IBM mainframe Time Sharing Option), so ordered and installed it on the TSO system at the University of Toronto. And, of course, *ed* would visit Berkeley and, while there, mutate into *ex* and *vi*.

Did this wide variety of editor versions lead to massive confusion? Not really. For although most of the editors added new commands and features, they seldom deleted them. The result was that you could – and this is still true – learn a basic set of *ed* commands and special characters usable on every version.

Today the Seventh Edition, 4.?BSD and System III/V versions of *ed* are all sufficiently similar that one can move freely amongst them with only minor inconvenience. The manual pages for every current version of *ed* are all recognizably derived from the Sixth Edition document. System III/V extends the 'u' (undo) command, but most of the other commands are constant. If you've used *ed*, you've used an editor with a long history, and probably a long future.

3.2. roff

Having a good text editor is only half the text processing battle. Having entered your text, you still must format it neatly for presentation. That's the function of a text-formatting program. The earliest UNIX formatter known to man is *roff*, a line-command formatter. Like *ed*, *roff* is part of a large and diverse family, one that includes the *runoff* package found on Digital Equipment computers (the latest release is called DSR, for DEC Standard Runoff). The earliest Runoff program is attributed by Kernighan and Plauger to J. Saltzer, who wrote it for CTSS. Runoff also is an ancestor of the Script programs available on IBM mainframe systems; that descent would be equally interesting for IBMers to trace (no doubt we'll get letters from those with information to SHARE with us).

Roff was written by M. D. McIlroy, at Research. Like *ed*, *roff* was well in place by the First Edition of UNIX. It was considered static by the time of the Sixth Edition, regarded as obsolescent by the time of the Seventh, and dropped altogether by the time of System III.

3.3. nroff and troff – The assembler of text

Computerists are never satisfied. So after *roff* came "New Roff", or *nroff*, written by the late Joseph Ossanna, who throughout his career was concerned with improving the way text was handled. Ossanna's *nroff*, as Kernighan and Pike relate,

"was much more ambitious [than *roff*]. Rather than trying to provide every style of document that users might ever want, Ossanna made *nroff* programmable, so that many formatting tasks were handled by programming in the *nroff* language.

"When a small typesetter was acquired in 1973, *nroff* was extended to handle the multiple sizes and fonts and the richer character set that the typesetter provided. The new program was called *troff* (which by analogy to "en-roff" is pronounced "tee-roff"). *nroff* and *troff* are basically the same program..." with divergent processing appropriate to the differences in output device. [UNIX Programming Environment page 289]

They point out that *troff* is tremendously flexible, and indeed many computer books have been typeset using it. But it can be complex to use. As a result, most everyone uses one or another "macro package" – a series of pre-programmed commands – and optionally one of the preprocessors (such as *eqn*, *tbl*, *refer*, and more recently *pic*, *grap* and *ideal*). *Troff* was originally written in assembler, but was redone in C in 1975. Joseph Ossanna wrote both versions and maintained it until his death in 1977.

3.4. Macro Packages

The earliest macro package to come into wide use was "ms", for "manuscript" [Lesk1977a]. Written by Mike Lesk, the *ms* macros provide a powerful but easy-to-learn (by comparison with bare *nroff*) approach to document formatting. The ms macros were distributed with the Sixth and Seventh Edition UNIX and most subsequent releases. The package was picked up and extended at Berkeley.

The USG versions of UNIX include a macro package called 'mm', for 'memorandum macros' [Dolo1978a]. These do most of the same things as ms, in slightly different ways, with the addition of numbered lists and a few other bells and whistles, but are about half again as big as ms. The startup time is such with mm that USG in 1979 had to resort to a compacted form of the macro packages; this made it into System III.

There are two versions of the 'man' macro package used to format the manual pages in Volume 1 of the UNIX Manual Set. One was used on V6, and the other from V7 on. If you see a manual page beginning with '.th' instead of '.TH', it's from V6. System III has an (undocumented) command *mancvt*, and 4.1BSD had *trman*, to convert files from the old to the new format.

There is also the 'mv' macros for producing viewgraph or slide presentations. This is a USG product, and versions of the USG manuals from PWB 1 up to just before System III carried the now-famous line:

The PWB/UNIX document entitled: PWB/UNIX View Graph and Slide Macros is not yet available.

System III manuals appeared with scarcely a mention of mv, and (finally) it was documented with the System V manuals.

3.5. tbl, eqn, refer

One view of *troff* is as an assembler language for text processing. If this be true, then *eqn*, *tbl*, *refer* and later preprocessors are the high-level compilers that go with it.

Mathematics has always been an inconvenience to traditional typesetting. This observation led Brian Kernighan and Lorinda Cherry to develop *eqn* for UNIX, and would later lead Donald Knuth to write his TeX typesetting package with math capabilities built in.

The *eqn* program reads an entire *nroff/troff* input file and passes it unchanged except for "equation specifications" delimited by .EQ and .EN requests (or pairs of delimiter characters for in-line equations). Material inside these requests is used to construct equations of considerable complexity from simple input. English words such as 'sum', 'x sub i', and 'infinity' produce the expected results (a large Sigma, x with a subscript i, and the infinity symbol respectively). In most cases no typesetter wizardry is required. A list of some forty extra character definitions lives in the file */usr/pub/eqnchar*; these can be copied, extended, or altered by the knowledgeable user.

Eqn was written by Brian Kernighan and Lorinda Cherry. The "new graphic symbols" in */usr/pub/eqnchar* are the work of Carmela L'Hommedieu (formerly Scrocca) at Bell Labs. The first public write-up of *eqn* appears to be a paper by Kernighan and Cherry in the *CACM*, March 1975 [Kern1975a]. The software was included in the Sixth Edition UNIX.

Like *eqn*, *tbl* is a preprocessor that passes over a formatter input file looking for special requests (here .TS and .TE) [Lesk1976a]. The material between these requests is expected to be a series of special commands to *tbl* and some tabular data. To greatly over-simplify how this program works, *tbl* replaces tab characters with explicit horizontal and vertical moves to make the rows and columns in the table align exactly under control of the table specification. It is invaluable for putting tabular material of any kind into documents.

Mike Lesk wrote *tbl* at Research; the idea for it came from an earlier table formatting program by J. F. Gimpel. *Tbl* first appeared outside the Labs with the V6 release of UNIX. It appeared in its present form on the "Phototypesetter Version 7" (interim V7) and PWB tapes, in Seventh Edition UNIX distributions, and in all systems since then.

Lesk also wrote *refer* [Lesk1978a], a bibliographic citation and reference package, which first appeared in V7. Regrettably, not all UNIX distributions picked it up (like Lesk's *tbl*, *refer* contained many bugs, some of which manifested as core dumps on various architectures; that probably didn't help, though there have also been strong rumours that AT&T didn't want to distribute any form of software incorporating indices, notably the V7 *dbm* library, in order to leave a market for third-party data base companies). It incorporates knowledge of how to format citations and can produce from an approximate citation a footnote (or reference) mark and a formal citation as a footnote or reference at the end. Once you build up (or copy) a large collection of references, this is very handy. If your system doesn't come with *refer*, *grefer* can be found in the *groff* distribution.

3.6. Typesetter Independent troff

New! Improved! Yet again! That's right. *troff* is infinitely perfectible. In 1979, Brian Kernighan at Research set out to re-write *troff*. Rather than rewrite it completely and be incompatible with the tens or hundreds of thousands of documents in existence, he chose to "clean up" *troff*. It turned out to be rather more akin to cleaning the Augean Stables than he had imagined, but resolve did not desert Kernighan. Finally he emerged with a tape for the Typesetter-Independent (later dubbed "Device-Independent", to

avoid a sexually-charged acronym) *troff* [Kern1982a], along with revised *tbl/eqn/refer* and two new preprocessors, *pic* and Chris Van Wyk's *ideal* [Kern1981a]. The new *ditroff* produced output that is readable ASCII containing a low-level description of the commands needed to drive a typesetter; this 'intermediate language' is fed to a particular postprocessor to generate the particular commands needed to drive a typesetter. *Pic* (as the name implies) does pictures. It is useful for drawing ''flow-chart''-like drawings, but there is much more to it than that. *Ideal* also draws pictures, but is somewhat more mathematical in usage than is *pic*.

This set of products formed the basis of the commercialised "Documentor's Workbench" package from AT&T. And work continues, of course. Kernighan has been working on cleaning up the appearance of *eqn* output. Recently, Kernighan and John Bentley have written *grap*, a graph plotting preprocessor for *pic*. The program was released with Release 2 of Documentor's WorkBench, released in early 1986. A report on *grap* is also available from Bell Laboratories as a CSTR [Bent1984a].

3.7. Troff in the outside world

Several Universities have further developed troff. Since these are changing all the time, it is not feasible to mention them all here. A future issue of the USENIX journal *;login:* will contain a summary of this work.

Many companies picked up the commercial (binary sublicense) release of Documentor's WorkBench. As a result one can now purchase this product in binary-only form for most any UNIX computer. Most of these *``troff* houses'' have added postprocessors for specific output devices; few have made substantial change to the main program. One of the authors (Darwin) was employed at SoftQuad in the mid-1980's, then believed to be the only company doing serious development on *troff* internals (AT&T had no developers assigned to the product). SoftQuad had drastically changed the intermediate language devised by Kernighan. It is now more readable and more amenable to additional processing with the UNIX software tools than is the standard *ditroff*. Most *troff* houses support the common output devices (in 1986 this means the Apple LaserWriterTM and other PostScriptTM printers, Imagen imPRESSTM printers, the Hewlett-Packard LaserJetTM, and other devices; by the mid 1990's it means PostScript and HP).

A few attempts have been made to port Documentor's WorkBench to non-UNIX environments. Both SoftQuad and Elan Systems have an MS-DOS port of complete DWB. Other companies (including Image Network) have ported the product to Digital Equipment's VMS system. But the main usage of *troff* continues to be on UNIX.

Because troff required a commercial license, it had to be cloned for free UNIX. The *groff* formatter suite used on most free BSD systems was cloned by James Clark in the 1980's from UNIX *troff*, with ideas from SoftQuad and other extended versions of *troff*.

3.8. Of Mice and Blits

Tired of typing at a dull, boring 24×80 screen, Rob Pike and Bart Locanthi had a better idea. Integrating the ideas of the Alto project and related work at Xerox PARC (Palo Alto Research Centre) with the UNIX approach to things, they built a special terminal called the Blit [Pike1985a, Pike1984a, Carg1984a] (not an acronym) with a Motorola 68000 processor, high-resolution screen, good keyboard, a fabulous mouse, **and** software split between the host and the terminal. Their particular combination of these ingredients makes possible a form of interaction with the computer that was years ahead of its time. Pike, in addition to being a radical advocate of the UNIX approach to software development, is quite visionary. Some of his work was described at recent (1980's) USENIX conferences. And large parts of the audience didn't seem to grasp the essentials of what he is saying. Some of the same ideas are of course found in other bitmapped screen environments, such as the SUN workstation and The X Window System, in somewhat different forms.

Blits were available only inside AT&T, mostly in Bell Labs, although a few have been released to selected Universities. The Blit has been commercialised by AT&T/Teletype, and was sold (with a different microprocessor) as the AT&T DMD 5620 terminal.

(More recently, much of this work has been consolidated into an operating system called Plan 9 From Bell Laboratories.) [Pike1990a, Pike1995a]

3.9. Style, Diction, Writer's Workbench

One of the authors (Darwin) has long been interested in the computerised processing of text, a term I take to mean more than is commonly included as "word processing." So I was quite interested to read a paper by L. E. McMahon, Lorinda L. Cherry and R. Morris entitled "Statistical Text Processing" in the 1978 special BSTJ issue on UNIX [McMa1978a]. I would later use several of the techniques mentioned in the paper.

At the end of the paper, Ms. Cherry describes a program *parts* for finding parts of speech in English text. This was written to be the first pass of a system to add inflection to the *speak* program written by Doug McIlroy, but the person doing the stress part left the company. Ms. Cherry wasn't interested in the stress assignment, so she documented the work done so far and went on to other things.

In the spring of 1979, W. Vesterman of Rutgers approached Doug McIlroy at Research about computerizing one of the techniques Vesterman used in teaching writing. The students had to count surface features in their text and in a sample of text written by a professional writer. That summer, Ms. Cherry expanded *parts* considerably, and added the code that turned it into *style*, a program to analyse the readability and other characteristics of a textual document. She also developed *diction* to check for awkward word uses, overused words, and other problems facing everyone who composes text for others to read. She also modified *deroff* to find the real text in a document. Vesterman consulted on this work.

And when the 4.1BSD release of the system came out, I was pleasantly surprised to see that *style* and *diction* were present [Cher1981a]. Bell Labs has a policy of sometimes releasing software to educational institutions; this probably explains the release at Berkeley.

While this was going on, the Human Factors group at Piscataway (now at Summit) was getting interested in automating document review, and Nina Macdonald of that group called Ms. Cherry about using *parts*. She had worked at Murray Hill in a Linguistics group and was familiar with the program. Ms. Macdonald took *style* and *diction*, and WWB evolved from there. Writer's Workbench (WWB) consists of *style*, *diction* and a dozen or so related programs for finding problems in written work. The "chattiness" level of the programs is set for the beginning user, but can easily be adjusted by the advanced user. The ideas for this work came from the Piscataway group, the Murray Hill group, and from Colorado State University, where extensive use of the Writer's Workbench (described at USENIX, Toronto, July 1983) currently puts several thousand undergraduates on WWB each year. The use of WWB is perceived to improve significantly the students' writing skills.

Many writers will be thankful to all who contributed, because these programs have proven themselves useful many times over. If buying a 4.1 or 4.2BSD system, insist on *style* and *diction*. If you get a System V UNIX, consider getting the WWB add-on if you'll be doing any document preparation. Writer's Workbench is one product that should survive and prosper as UNIX continues to evolve. The next major release of WWB (3.0) was scheduled for the spring of 1985.

4. Compilers, Languages, Tools

What is an operating system without languages and utilities? Despite the limited support for FOR-TRAN, UNIX has always been known for the diversity of languages and tools provided. Some of these are well-known, others are less-well-known than perhaps they ought to be.

4.1. The C Programming Language

The early evolution of the C language has been described elsewhere (see Dennis Ritchie's paper in *Microsystems*, October, 1983. Dennis is rather modest, and doesn't tell you that the UNIX world has named the C compiler described there "the Ritchie compiler" to distinguish it from other C translators). As we pick up the threads of the story, Fifth Edition UNIX has been in the field for some time. It is May, 1975, and the new improved Sixth Edition [Thom1975a] is about to be released. Ritchie has added some support for a new data type, "long integers" referred to with the keyword long, but this will not be documented. Not all the runtime support has been installed, and the tape goes out without it. Later Ken Thompson will announce that the support for 'longs', limited though it be, was there all along in V6. There is no support for "short s".

There follows a succession of releases of the C compiler. The PWB 1.0 release of UNIX, the first

outside the labs of a non-Research UNIX from Bell, goes out in 1977. And a special-release tape known only as "Phototypesetter Version 7" includes a new release of *troff* as well as the C compiler, assembler, loader, archiver and bits of the C library, including the first release of 'stdio'. These compilers seem to be from about the same vintage. Both compilers support another new data type modifier, unsigned, which causes all bits of an integer to be treated as magnitude (on the PDP-11, for example, signed ints run from -32,768 to +32,767 while unsigned ints are from 0 to 65,535). These compilers add typedefs, which allow you to generate your own names for existing data types, for a degree of independence from the machine data types. One of these compilers is somewhat buggy – the concept of "cast" is in the code but doesn't work properly. Bit fields exist but are buggy; this is documented. The Phototypesetter Version 7 was primarily a release of *troff*; the compiler was included because it was necessary for *troff* (very convenient, since Research wanted to get the latest C out into the field anyway).

In 1978, *The C Programming Language* is published [Kern1978a] and serves as the standard for the C language for over a decade, until the ANSI C standard [Amer1989a, ISO1990a] appears, at long last, in December 1989.

Finally the Seventh Edition of UNIX [Thom1978a], the first portable UNIX, is released in June or July of 1979 and announced at the Summer USENIX conference in Toronto, along with 32V (the DEC VAX port of V7 that all later VAX implementations derive from, including 3BSD, 4BSD, System III and System V). This is the last Research UNIX that the outside world, with the exception of a few academic institutions, will ever see, though there will be three more editions (V8 through V10, all for the VAX) [Pres1986a, Pres1985a, Labo1986a, Comp1990a] incorporating interesting ideas before Plan 9 takes over as the operating system under active development by Research in the early 1990's. Many afficionados today regard V7 as the pinnacle of UNIX; superior to all its ancestors and descendants. There is some truth in this, though V7 lacked several things that later became important: interactive networking (though *uucp* provides batch networking), a network filesystem, and bitmapped graphics. These will be added in V8 in superior form to anything seen outside AT&T until Plan 9. The networking infrastructure will seep out in heavily-modified form as System V STREAMS (a heavy-weight form of Research Streams). The graphics will become available in the AT&T DMD 5620 terminal described above, which is unfortunately much more expensive than competing terminals and PCs, but nevertheless develops a cult following.

Of course, V7 has another C compiler. This one, for the most part, is a "shaken down" version of the "Phototypesetter C" compiler. It is a lot more solid, although bit fields are still broken and now the bug is not documented. And there is a special kludge for *uucp* whereby casting an expression involving a character pointer to type unsigned treats the character referenced by the pointer as an 'unsigned character', a concept not yet in the compiler or language. (It will later appear as the type unsigned char.) This kludge will be quietly withdrawn later. The compiler has a bug in that it treats the right, not the left, side of an assignment as the value of the assignment expression. The V7 *stdio* exploits these two bugs, thereby making it non-portable. The semantics of casts will remain unsettled until slightly after V7.

Along with the Ritchie C compiler, V7 of UNIX includes the first release outside Bell Labs of a second C compiler, bearing the impressive name of the "Portable C compiler." [John1978a] Written by S. C. Johnson, this compiler has been in development since 1975 and uses the program development tools *yacc* but not *lex*. (A part of the *yacc* grammar for this compiler was published in the C manual with PWB in 1977.)

The portable C compiler turns out to be not as portable as desired, so a second version is developed over the next few years, called *pcc2*. At the ACM National Conference in 1983, Steve Johnson describes *pcc2* in some detail, and shows an example of its portability. Of the many "back ends" (assembler or machine code generators) for it, one compiles a C language algorithm into the commands necessary to drive a VLSI fabrication process. So your program (if you work in the right part of the Labs!) can be compiled into a custom microprocessor, optimised to execute your program and nothing else! That sure out-classes the EPROM versions of Intel and Motorola microprocessors. *pcc2* has only recently been released; it is the C compiler for System V Release 2 (?) and the Software Generation System cross-compiler. Later versions of the C compiler are described in a paper [Kris1986a] by David M. Kristol of AT&T.

One immediate beneficiary of the two-pass nature of *pcc* was the Fortran compiler, to which we will return shortly. But a second major fallout from *pcc* is a program called *lint*, which does partial compilation of C programs with much greater error checking. Like *pcc*, *lint* first appears with V7. We continue to

recommend the use of *lint* to provide some reassurance of program correctness and portability.

Berkeley has taken the C language in some new directions. They have relaxed some restrictions on compiled programs. Most notably, variables can be almost any length and need not be unique in the first seven or eight characters. While this sounds handy, it is a major annoyance to the rest of the world, which has to change programs written with such 'features' in order even to compile them. Berkeley programmers also tend to rely to an unprecedented extent on the 'asm' keyword, which allows you to interpolate assembler language code into the middle of the C program. 'asm' buys an increase in micro-efficiency but with a tremendous loss of portability. To preserve portability, the programmer should use #ifdef to include in the source code both an assembler version and a portable C version. But the latter is often omitted. A fine example was shown in a talk by Mike Tilson, then of Toronto's Human Computing Resources (reprinted in the November, 1984 final issue of *Microsystems*) [Tils1983a]. Here's the code:

```
to = bp->b_ptr;
asm("movc3 r8,(r11),(r7)");
bp->b ptr += put;
```

What it does is left as an exercise to the reader. The writer of this code left no clues to how his mayhem works. As Mike says: "The variable "to" is one of the registers used in this VAX assembly instruction. You guess which." Oh, we almost forgot. The three lines above are Copyright © 1980 by the Regents of the University of California.

Meanwhile, back at Research, B. Stroustrup has been busy adding "classes" to C[Stro1981a]. Classes (nothing to do with going back to school) are the interesting part of Simula 67. They provide for orderly interchange of data between modules, with no possibility of hidden dependencies. A class consists of data (normally inaccessible from outside the class) and functions which are normally accessible from outside but which may be declared as inaccessible. One typically defines a class and publishes the names of the accessible functions. Functions outside the class cannot reference the data within that class except by calling the class's publicly-accessible functions. This enforces modularity by hiding the details of a particular class's internals from other routines. Classes can be nested, of course, so you can develop such things as queues and stacks of objects. The compiler encompassing all recent developments, including classes, declaration of function parameters for type checking, and other recent developments is given the name "C++" [Stro1984a, Stro1984b]. C programmers will recognize the pun; for others, it simply means "an incremented (augmented?) form of the C language, which retains the value of the old language''. C++ has been in use for some time within the Labs, and was recently (1985) released for external sale. [C++ hasn't worked out as well as we'd hoped in 1984 [Stro1989a]. If the differences between Algol 60 and Simula 67, edited with taste to remove complexity such as the built-in text class, had been added to C, the result might have been an interesting language (C+=simula-algol60 rather than C++), but C++ is now too big, redundant and confused for our tastes. Luckily the ANSI C++ committee has put something of a brake on further bloat (though failed ideas from languages like Lisp have been tossed in), but it's too late already. Java is in a sense a reaction to C++: a distilled C++ without concern for C compatibility.]

In addition to a C compiler, one needs a series of library routines to do Input/Output and some 'extra-linguistic' operations such as setexit()/reset() in V6 or setjmp()/longjmp() in V7. The first 'portable C library' [Lesk1975a] was written by Mike Lesk, and was implemented on PDP-11 UNIX, the IBM 370, and the Honeywell 6000 with the GCOS operating system. It set the style for subsequent development, and in Version 7 there was a 'new portable I/O library' [Kern1979a] written by Ritchie. This has become known as ''stdio'' (pronounced ''stuh-DYE-oh'') for the name of its header file. This is the I/O library distributed with all real UNIXes today.

The current C compilers for the PDP-11 continue to derive from the Ritchie versions. *Pcc* for the PDP-11 never worked as well as the Ritchie compiler. Most other machines use *pcc*-based C translators since the Ritchie compiler only works for PDP-11's (though Amdahl bravely ported it to the IBM 370 for their UTS [Wals1982a] UNIX port). Many systems integrators wait earnestly for the release of *pcc2* since porting *pcc* takes a non-trivial amount of work.

The future of the C language is not primarily in the hands of people like Dennis Ritchie and Steve Johnson and B. Stroustrup. Rather, it is in the hands of the ANSI C Language standards committee. But in the final sense, it is in the hands of programmers everywhere. Partly because ANSI is a democratic agency,

and any member of the committee has as much voice as a Dennis Ritchie or a Steve Johnson. And also because C is a powerful language, and like all powerful tools it can be used or abused. This is not the place for a tutorial on C style, but the interested reader can refer to the paper by Tilson cited above and our own USENIX paper [Darw1985a]. Good use of C leads to rapid development of high-level code; poor use of the language leads to code that looks like it was written in assembler. As we have seen, in a few cases it has been.

4.2. fortran

Since UNIX comes from a Computer Science research background, it is perhaps natural that Fortran, that octogenarian, reptilian but ubiquitous language should be the object of some disdain among UNIXophiles. Indeed, the V6 *how to get started* document says that "No debugger is much help for Fortran." And the Sixth Edition manual set included the *C Reference* and the *C Tutorial*, but nothing on Fortran proper beyond the manual page for fc(1), a compiler for a slight variant of the ANSI Fortran-66 standard. (But see the discussion of Ratfor below.)

fc produced executable programs that used threaded code and floating-point instructions heavily; thus it ran slowly on machines without a floating-point processor, on which floating point had to be interpreted.

The prime mover behind the next Fortran compiler was Stuart Feldman, who had been interested in compilers for some time. In 1976 he released a CSTR on *'Fortlex* – A General Purpose Lexical Analyzer for Fortran.' This program [Feld1976a] reads a Fortran program and breaks it up into lexical tokens of the appropriate type, a non-trivial job. *Fortlex* was used in the construction of various Fortran programming aids, such as a program to change all double precision variables, functions and library calls to single precision. The paper also includes the *yacc* grammar for a Fortran scanner to be used with *Fortlex*; not a complete compiler, but possibly a basis for one.

And the Fortran weakness of UNIX was remedied with a vengeance for the Seventh Edition. A compiler for the full ANSI Fortran-77 standard was included, apparently the first implementation of the 1977 standard on any system anywhere, along with a paper detailing its use and implementation. The back-end of this compiler was the same back end as the Portable C Compiler, so that it would be easy to adapt to new computers.

Although the Fortran compiler is part of all standard UNIX systems, most suppliers of 68000-based UNIX boxes do not include f77. Whether this is so they can charge extra for it, or because they couldn't figure out how to port it, is unclear. But commercial implementations are available for most micro-based UNIXes.

4.3. yacc+lex

One of the major tools used in compiler development is the *yacc* (yet another compiler compiler) [John1975a] program by Steve Johnson. When this program was developed in the early 1970's, compiler generators were being generated by many universities and other research institutes. As Kernighan and Pike remark, Johnson's choice of name for his program is ironic in that his has endured while most of the others have now been retired.

yacc reads in the specification of a language, and generates a program which parses that language. Note that this is not limited to "programming languages", but can be applied to any input that is structured. In addition to the *nrws* nroff-to-Wordstar program used to translate some articles for *Microsystems*, many applications are mentioned in the *yacc* paper. That manual mentions "compilers for C, APL, Pascal, RAT-FOR, etc, ..., a phototypesetter system, several desk calculators, a document retrieval system, and a Fortran debugging system" as programs which have been written in *yacc*. More recently, Cobol and Ada compilers have been constructed.

But syntax analysis is only one part of compilation. Another part is lexical analysis, or scanning of the input looking for certain kinds of tokens. For this, too, UNIX has an answer. The *lex* program by Mike Lesk and E. Schmidt [Lesk1975b] provides this function. Since it is part of the UNIX tradition, of course, *lex* uses many of the same conventions. In particular, *lex* uses the same notation for "regular expressions" as is used in the editors and elsewhere to describe the patterns to be looked for. If you've mastered commands like /[hH]e/ in the editor, you already know how to construct expressions for *lex*. And of course

it works with *yacc*. The naming conventions of these two programs are such that the output of both can be loaded together to form a working unit. Indeed many programs consist of *yacc* and *lex* outputs compiled and loaded together.

Yacc was present in Version 6 (the manual page is dated late 1974); *lex* first appeared outside the labs in the PWB 1.0 release.

4.4. make

It's hard to imagine UNIX without the *make* utility [Feld1979a]. *Make* is so taken for granted these days that the distribution of software in source form without a *makefile* is an event worthy of attention and inquiry. But there was a time when the name of the file with the instructions to build a system were chosen at random from the names "build", "rc", "run", "runfile" and others. And these were shell files which built the entire component.

Make builds a program or component from individual pieces, and recompiles only the minimum needed to rebuild it as changes are made. The edit-make-debug cycle is well known to all UNIX programmers. Since the topic has been treated in detail in *The UNIX File* by one of us, we will not expand on it here. Suffice to say that *make* was written by Stuart Feldman at Research, and first appeared outside Bell in the PWB release of the system. The sketchiness of the Source Code Control System user interface, its slowness, and its incompatibility with most of UNIX including *make* led not to the correction of SCCS, but to an "enhanced" *make* which appears publicly in System III and System V, and to superior revision control systems such as RCS and CVS. (Variants of *make* proliferated, culminating much later in *mk* by Andrew Hume [Hume1987a] and *mash* for Inferno by Bruce Ellis.)

4.5. ratfor, efl

Brian Kernighan at Research realized that Fortran would not go away, so he did something about it. He fixed it. He fixed it by adding the control structures of C and the definition and inclusion capabilities of the C preprocessor. The converter which takes in "rationalised Fortran" and produces ugly conventional Fortran he called "Ratfor" [Kern1977a]. Several versions were written; one in Ratfor (translated by hand or machine into Fortran) for bootstrapping onto other systems, another with *yacc* and *lex* as mentioned above. The meaning of Ratfor is best told in the book *Software Tools* co-written with P. J. Plauger. The source code for the programs in the book was made available on magnetic tape by Addison-Wesley, in a move that was very far-sighted for 1976. This lead to the formation of the Software Tools User Group at Lawrence Berkeley Labs in Berkeley; this group was quite active and co-sponsored meetings with USENIX.

The *Software Tools* book would later be re-done in Pascal (see 'Pascal' section). There are no plans announced for doing a 'Software Tools in C' book; most of what you need is in Kernighan's & Pike's book *The UNIX Programming Environment*.

After the Fortran-77 compiler, Stuart Feldman turned his attention to Fortran extensions, and produced the *efl* language [Feld1979b]. This combines the control structures of *ratfor* (which in turn derive from C) with the data structuring capabilities of C including the *struct* capability to group related data items, analogous to Pascal's *record* capability. *efl* is included in some System V and some 4.?BSD systems. Some microcomputer ports (i.e., UniSoft) include *efl* even though they don't have a Fortran compiler.

4.6. awk

Aho, Weinberger, and Kernighan. The names of three authors put together in the most pronounceable way. That's what they did when they couldn't think of a more imaginative name for a wonderful program they'd devised. A is Al Aho, of compiler book fame. W is Peter Weinberger of Research. And K is Brian Kernighan, just described for his work on Ratfor. (Kernighan and Pike's book remarks that "Naming a language after its authors also shows a certain poverty of imagination." [page 131]) *Awk* is not at all awk-ward, it is a great simplification. You can think of it as the combination of most of the best ideas of the other tools all rolled into one [Aho1978a]. We use it all the time. You can enter the *awk* commands from the command line if they are simple enough, so that awk -F: '{print \$1}' /etc/passwd

is a complete program to print the names of all the accounts shown in */etc/passwd*, the standard place for the names of all accounts on the system. By all means learn about *awk* and use it. It will make life far less awk-ward.

Awk was first described in *Software Practice and Experience* in July, 1978, and first distributed with Seventh Edition UNIX.

4.7. Pascal

The Pascal language never did catch on at Research. In 1981, Brian Kernighan wrote a paper published as a CSTR entitled "Why Pascal is not my Favorite Programming Language." [Kern1981b] The note was not based solely on introspection, for he and P. J. Plauger had just converted their book *Software Tools* into *Software Tools in Pascal*, including re-coding all the programs in Pascal. In the process they came to regard Pascal as their not favorite language.

Berkeley UNIX has included Pascal for a long time. Ken Thompson wrote the first version of Berkeley Pascal while working at Berkeley as Visiting Mackay Lecturer in Computer Science in 1975/76. He spent the academic year at UC Berkeley, and taught several courses in Computer Science. He recalls:

"When I arrived, the CS department shared an 11/45 with Statistics. It was 50-50 UNIX and RSTS. The first advance was an 11/70 dedicated to teaching. I put my first 155 [operating systems] course on it. Between the first and second quarters I wrote the Berkeley Pascal and talked Bob Fabry into using it on his 153 [data structures] class. It has been used for that ever since. By the time I left, there were several (2 or 3) 11/70s in the computing center providing UNIX service. CS had the 11/70 for teaching, they had almost completely taken over the stat 11/45 and there was a research 11/40" in an AI lab [Thompson, personal correspondence].

Pascal compilers can be had for most 68000-based UNIX boxes. These are available from commercial software firms and OEMs – see the annual UNIX software directory in the April Microsystems.

4.8. S

Finally, we cannot overlook an interesting "application" language from Research. *S: An Interactive Environment for Data Analysis and Graphics.* The title of the 1984 book by Richard A. Becker and John M. Chambers puts it succinctly. And we put it as follows: The S package is to conventional mainframe statistics packages as the UNIX shell is to batch Job Control languages. It provides interactive, exploratory statistics, interactive and offline graphics, and data modelling and time series manipulation.

The S language was developed at Research. The first public release was in 1981, which was for V7 and 32V. There were several interim releases; the next major release was in early 1984. This release was accompanied by a change in licensing and an order-of-magnitude cost increase for non-educational users, as part of the swing to the commercialization of UNIX by AT&T Technologies.

The only remotely similar products that we know of in all of computerdom are APL and Speakeasy. APL was first implemented on IBM systems; at least one version for UNIX was developed (Ken Thompson's for the PDP-11). Speakeasy similarly arose on IBM hardware; a subset version called SpeakeC was developed at Purdue. Speakeasy was developed some time before S, but in quite different circles of influence. There appears to be no cross-pollination between the two although many of the ideas are similar. S uses *yacc* to interpret its grammar; the *yacc* specification appeared in a *CACM* paper in 1984.

5. Interlude

The Computing Science Research Center, Center 1127, at Bell Labs has had impact on computerdom far out of proportion to the number of people working there. A small group of talented people, started in motion by Ken Thompson and Dennis Ritchie with the original design of UNIX, aided and abetted by those mentioned here and others, developed Research UNIX and its related tools. Many reasons are given for the success of UNIX, but one we'd like to add is the consistency of the system in all its facets. As a single example, the syntax used for pattern matching (a notation for the abstract concept of Regular Expressions)

allows you easily to develop tremendous skill in pattern matching. This skill once learned can be applied in the editor to find text, in *awk* to find records to be acted on, in *lex* to specify partitioning of an input, (with simple modification) in shell commands to match filenames, and in a dozen or so contexts. This kind of consistency is a rare treat in any computer system; the extent to which it permeates UNIX is exemplary.

This concludes the first installment of our history of the natural creation of the UNIX timesharing system. The next installment will cover the relationships among many, many different versions of UNIX, most of them never released or publicised outside of Bell Labs and the telephone companies. There will be a "family tree" diagram illustrating the descent of UNIX. We hope you enjoyed this edition, and that you'll be looking forward to the next (assuming that somebody contributes it).

6. Acknowledgements

A myriad of UNIX experts helped with this paper. While it's not possible to thank everyone who offered to help, we should single out for special thanks Dennis Ritchie, who was there from the start and helped us to find it. We also got special help from Ken Thompson, who answered many questions about the past. Henry Spencer and Laura Creighton made useful comments on the work in progress. Lorinda Cherry and Nina Macdonald helped by describing the evolution of Style, Diction and the Writer's Workbench. While it's not really possible for people outside "the Labs" to get an exact picture of UNIX history, we have done our best, and assume responsibility for the accuracy of the material presented here.

7. References

1978a. Bell Sys. Tech. J., AT&T (July-August 1978). part 2

- 1984a. AT&T Bell Laboratories Tech. J., AT&T (October 1984).
- Aho1978a. A. V. Aho, P. J. Weinberger, and B. W. Kernighan, "AWK A Pattern Scanning and Processing Language," *Software—Practice and Experience* 9, pp. 267-280 (April 1979).
- Amer1989a. American National Standards Institute, X3J11 committee, *American National Standards Institute X3.159-1989 -- Programming Language C, = ISO/IEC 9899:1990*, ANSI, New York (1989).
- Bent1984a. Jon L. Bentley and Brian W. Kernighan, "GRAP A Language for Typesetting Graphs (Tutorial and User Manual)," Comp. Sci. Tech. Rep. No. 114 (August 1984).
- Carg1984a. T.A. Cargill, "Debugging C programs with the Blit," *Bell Laboratories Technical Journal* **63**(8 (part 2)), pp. 1633-1647 (October 1984).
- Cher1981a. L. L. Cherry, "Writing Tools The STYLE & Diction Programs," Comp. Sci. Tech. Rep. No. 91 (February 1981).
- Comp1990a. Computing Science Research Center, AT&T Bell Laboratories, Murray Hill, New Jersey, UNIX Research System Programmer's Manual, Tenth Edition, Saunders College Publishing (1990).
- Darw1984a. Ian Darwin and Geoff Collyer, "The Evolution of UNIX 1974 to the Present," *Microsystems* **5**(11), pp. 44-56 (November 1984).
- Darw1985a. Ian Darwin and Geoff Collyer, "Can't Happen or /* NOTREACHED */ or Real Programs Dump Core," *Proc. Winter Usenix Conf. Dallas 1985*, pp. 136-151 (January 1985).
- Deut1967a. L. P. Deutsch and B. W. Lampson, "An online editor," *Comm. Assoc. Comp. Mach.* 10(12), pp. 793-799, 803 (December 1967).
- Dolo1978a. T. A. Dolotta, R. C. Haight, and J. R. Mashey, "UNIX Time-Sharing System: The Programmer's Workbench," *Bell Sys. Tech. J.* **57**(6), pp. 2177-2200 (1978).
- Feld1976a. S. I. Feldman, "Fortlex A General Purpose Lexical Analyzer for Fortran," Comp. Sci. Tech. Rep. No. 51 (October 1976).
- Feld1979a. S. I. Feldman, "MAKE A Program for Maintaining Computer Programs," *Software Practice and Experience* 9, pp. 255-265 (April 1979).
- Feld1979b. Stuart I. Feldman, "The Programming Language EFL," Comp. Sci. Tech. Rep. No. 78 (June 1979).

- Hume1987a. A. Hume, "Mk: A Successor to MAKE," Comp. Sci. Tech. Rep. No. 141 (November 5, 1987).
- ISO1990a. ISO, Programming Language C, International Standards Organization (ISO), revision and redesignation of American National Standards Institute (ANSI) X3.159-1989, Amendment 1, ISO (1990).
- John1975a. S. C. Johnson, "Yacc Yet Another Compiler-Compiler," Comp. Sci. Tech. Rep. No. 32, Bell Laboratories, Murray Hill, New Jersey (July 1975).
- John1978a. S. C. Johnson, "A Portable Compiler: Theory and Practice," Proc. 5th ACM Symp. on Principles of Programming Languages, pp. 97-104 (January 1978).
- Kern1975a. B. W. Kernighan and L. L. Cherry, "A System for Typesetting Mathematics," Comm. Assoc. Comp. Mach. 18, pp. 151-157 (March 1975).
- Kern1976a. Brian W. Kernighan and P. J. Plauger, Software Tools, Addison-Wesley (1976).
- Kern1977a. B. W. Kernighan, "Ratfor A Preprocessor for a Rational Fortran," Software—Practice and Experience 5, pp. 395-406 (October 1975).
- Kern1978a. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, Englewood Cliffs, New Jersey (1978).
- Kern1979a. Brian W. Kernighan and Dennis M. Ritchie, "UNIX Programming Second Edition," UNIX Programmer's Manual, Seventh Edition (January, 1979).
- Kern1981a. B. W. Kernighan, "PIC A Crude Graphics Language for Typesetting," Comp. Sci. Tech. Rep. No. 85 (January 1981).
- Kern1981b. B. W. Kernighan, "Why Pascal is Not My Favorite Programming Language," Comp. Sci. Tech. Rep. No. 100 (July 1981).
- Kern1982a. B. W. Kernighan, "A Typesetter-Independent TROFF," Comp. Sci. Tech. Rep. No. 97 (1981, revised March 1982).
- Kern1984a. Brian W. Kernighan and Rob Pike, *The UNIX Programming Environment*, Prentice-Hall (1984).
- Kris1986a. David M. Kristol, "Four Generations of Portable C Compiler," pp. 335-343 in USENIX Conference Proceedings, USENIX, Atlanta, GA (Summer 1986).
- Labo1986a. AT&T Bell Laboratories, in *UNIX Programmer's Manual, Ninth Edition, Volume One*, ed. M. D. McIlroy, Murray Hill, New Jersey (September 1986).
- Lesk1975a. M. E. Lesk, "The Portable C Library (on UNIX)," UNIX Programmer's Manual, Sixth Edition (1975).
- Lesk1975b. M. E. Lesk, "Lex A Lexical Analyzer Generator," Comp. Sci. Tech. Rep. No. 39, Bell Laboratories, Murray Hill, New Jersey (October 1975).
- Lesk1976a. M. E. Lesk, "Tbl A Program to Format Tables," UNIX Programmer's Manual 2, Section 10 (January 1979).
- Lesk1977a. M. E. Lesk, Typing Documents on UNIX and GCOS: The -ms Macros for Troff, 1977.
- Lesk1978a. M. E. Lesk, "Some Applications of Inverted Indexes on the UNIX System," UNIX Programmer's Manual 2, Section 11 (January 1979).
- Lion1977a. John Lions, *Commentary on UNIX Sixth Edition*, Peer-to-peer Communications (1977). The famed underground classic exegesis of the workings of the kernel.
- McMa1978a. L. E. McMahon, L. L. Cherry, and R. Morris, "UNIX Time-Sharing System: Statistical Text Processing," *Bell Sys. Tech. J.* 57(6), pp. 2137-2154 (1978).
- Pike1984a. R. Pike, "The Blit: a multiplexed graphics terminal," *Bell Laboratories Technical Journal* **63**(8 (part 2)), pp. 1607-1631 (October 1984).
- Pike1985a. Rob Pike, "The Blit Programmer's Manual (TTY Edition)," Comp. Sci. Tech. Rep. No. 121 (October 1985).

- Pike1990a. Rob Pike, Dave Presotto, Ken Thompson, and Howard Trickey, "Plan 9 from Bell Labs," *Proc. of the Summer 1990 UKUUG Conf., London*, pp. 1-9 (July, 1990).
- Pike1995a. R. Pike, D. L. Presotto, S. M. Dorward, B. Flandrena, K. Thompson, H. W. Trickey, and P. Winterbottom, "Plan 9 from Bell Labs," *Computing Systems* 8(3), pp. 221-254 (Summer 1995).
- Pres1985a. D.L. Presotto and D.M. Ritchie, "Interprocess Communication in the Eighth Edition Unix System," pp. 309-316 in USENIX Conference Proceedings, USENIX, Portland, OR (Summer 1985).
- Pres1986a. D. L. Presotto, "The Eighth Edition Unix Connection Service," *European UNIX Systems User Group Autumn Conference*, Florence, Italy, pp. 1-9 (April 1986).
- Stro1981a. B. Stroustrup, "Classes: An Abstract Data Type Facility for the C Language," ACM Sigplan Notices Jan. 1982 (Revised August 1981).
- Stro1984a. Bjarne Stroustrup, "The C++ Programming Language Reference Manual," Comp. Sci. Tech. Rep. No. 108 (January 1984).
- Stro1984b. Bjarne Stroustrup, "A C++ Tutorial," Comp. Sci. Tech. Rep. No. 113 (September 1984).
- Stro1989a. B. Stroustrup, "The Evolution of C++: 1985 to 1989," Comp. Sci. Tech. Rep. No. 144 (July 1989).
- Thom1975a. K. Thompson and D. M. Ritchie, UNIX Programmer's Manual, Bell Laboratories (May 1975). Sixth Edition
- Thom1978a. K. Thompson and D. M. Ritchie, *UNIX Programmer's Manual*, Bell Laboratories (1978). Seventh Edition.
- Tils1983a. Michael Tilson, "A Tutorial on C Portability," pp. 315-323 in USENIX Conference Proceedings, USENIX, San Diego, CA (Winter 1983).
- Wals1982a. Daniel Walsh, "UTS: UNIX on the Amdahl 470," pp. 210 in USENIX Conference Proceedings, USENIX, Boston, MA (Summer 1982). Abstract only