

Lessons from Unix History

Diomidis Spinellis

Department of Management Science and Technology Athens University of Economics and Business

Department of Software Technology Delft University of Technology

www.spinellis.gr

X @CoolSWEng OCoolSWEng@mastodon.acm.org





Unix and Linux are powering modern IT

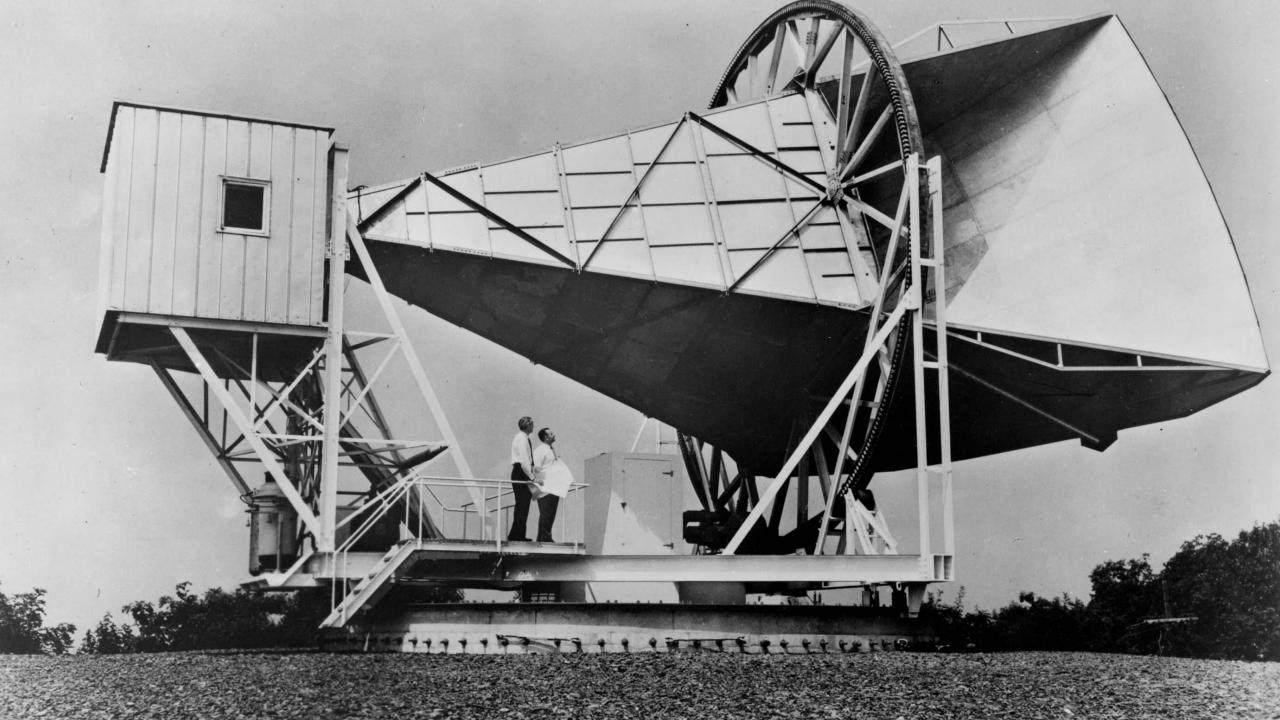
- Android (1 bn shipments / year)
- Google, Facebook, Amazon, X technology stacks
- macOS
- 60% of the top one million web servers,
- 75% of major cloud providers' instances,
- 97% of embedded systems,
- All of the world's top-500 supercomputers



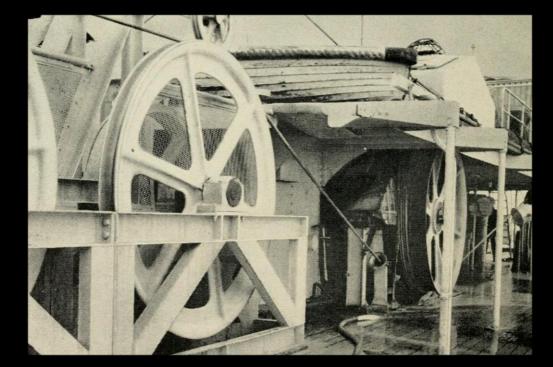


















Reprinted with corrections from *The Bell System Technical Journal*, Vol. 27, pp. 379–423, 623–656, July, October, 1948.

A Mathematical Theory of Communication

By C. E. SHANNON

INTRODUCTION

THE recent development of various methods of modulation such as PCM and PPM which exchange bandwidth for signal-to-noise ratio has intensified the interest in a general theory of communication. A basis for such a theory is contained in the important papers of Nyquist¹ and Hartley² on this subject. In the present paper we will extend the theory to include a number of new factors, in particular the effect of noise in the channel, and the savings possible due to the statistical structure of the original message and due to the nature of the final destination of the information.

The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have *meaning*; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one *selected from a set* of possible messages. The system must be designed to operate for each possible selection, not just the one which will actually be chosen since this is unknown at the time of design.

If the number of messages in the set is finite then this number or any monotonic function of this number can be regarded as a measure of the information produced when one message is chosen from the set, all choices being equally likely. As was pointed out by Hartley the most natural choice is the logarithmic function. Although this definition must be generalized considerably when we consider the influence of the statistics of the message and when we have a continuous range of messages, we will in all cases use an essentially logarithmic measure.

The logarithmic measure is more convenient for various reasons:

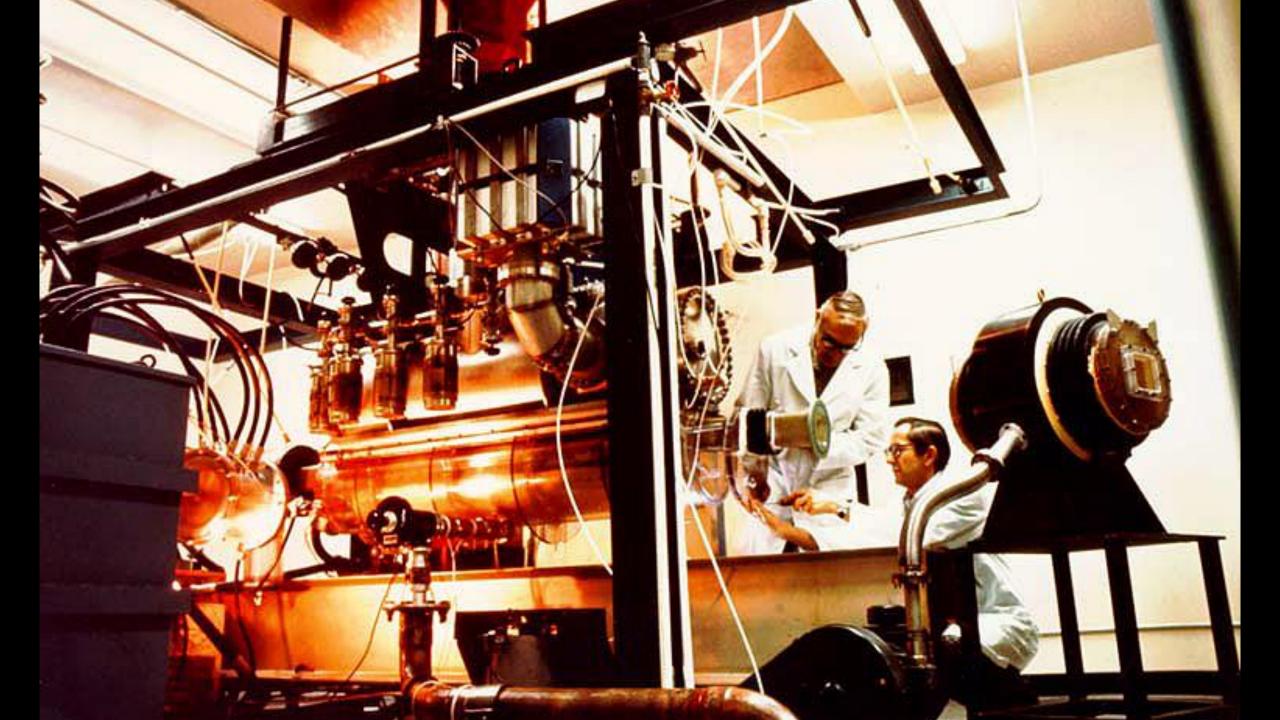
- It is practically more useful. Parameters of engineering importance such as time, bandwidth, number of relays, etc., tend to vary linearly with the logarithm of the number of possibilities. For example, adding one relay to a group doubles the number of possible states of the relays. It adds 1 to the base 2 logarithm of this number. Doubling the time roughly squares the number of possible messages, or doubles the logarithm, etc.
- 2. It is nearer to our intuitive feeling as to the proper measure. This is closely related to (1) since we intuitively measures entities by linear comparison with common standards. One feels, for example, that two punched cards should have twice the capacity of one for information storage, and two identical channels twice the capacity of one for transmitting information.
- It is mathematically more suitable. Many of the limiting operations are simple in terms of the logarithm but would require clumsy restatement in terms of the number of possibilities.

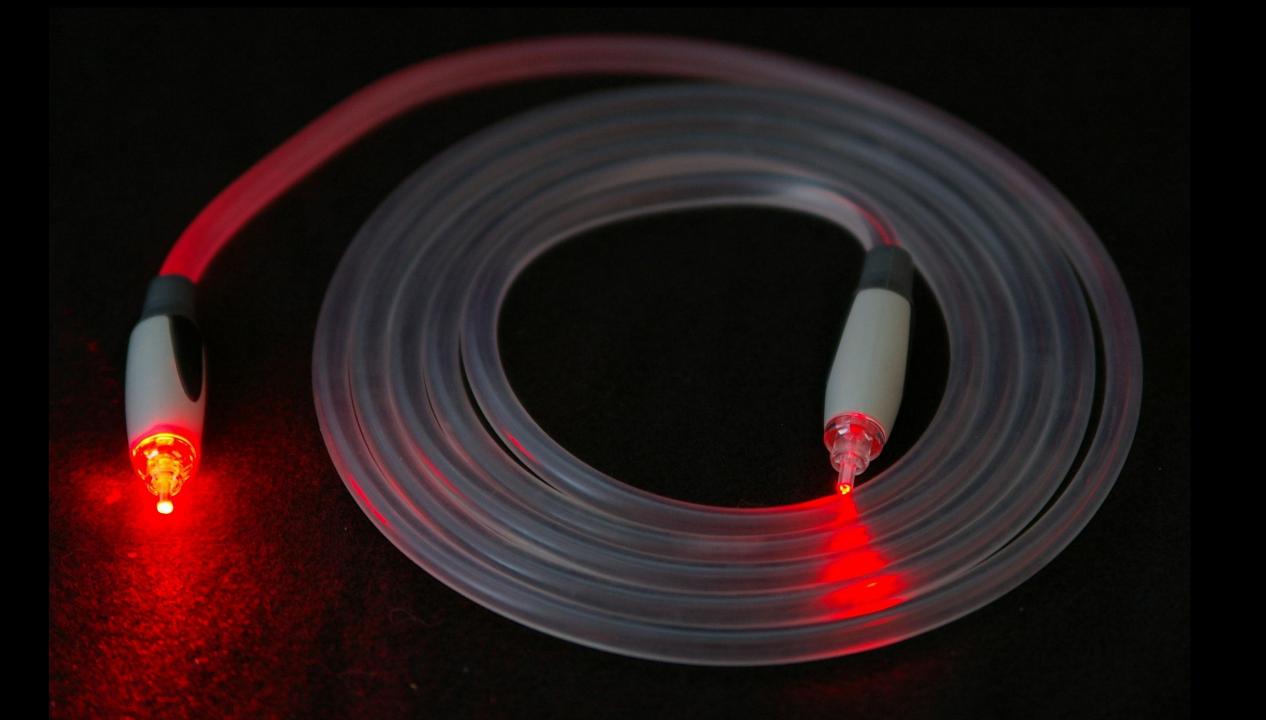
The choice of a logarithmic base corresponds to the choice of a unit for measuring information. If the base 2 is used the resulting units may be called binary digits, or more briefly *bits*, a word suggested by J. W. Tukey. A device with two stable positions, such as a relay or a flip-flop circuit, can store one bit of information. N such devices can store N bits, since the total number of possible states is 2^N and $\log_2 2^N = N$. If the base 10 is used the units may be called decimal digits. Since

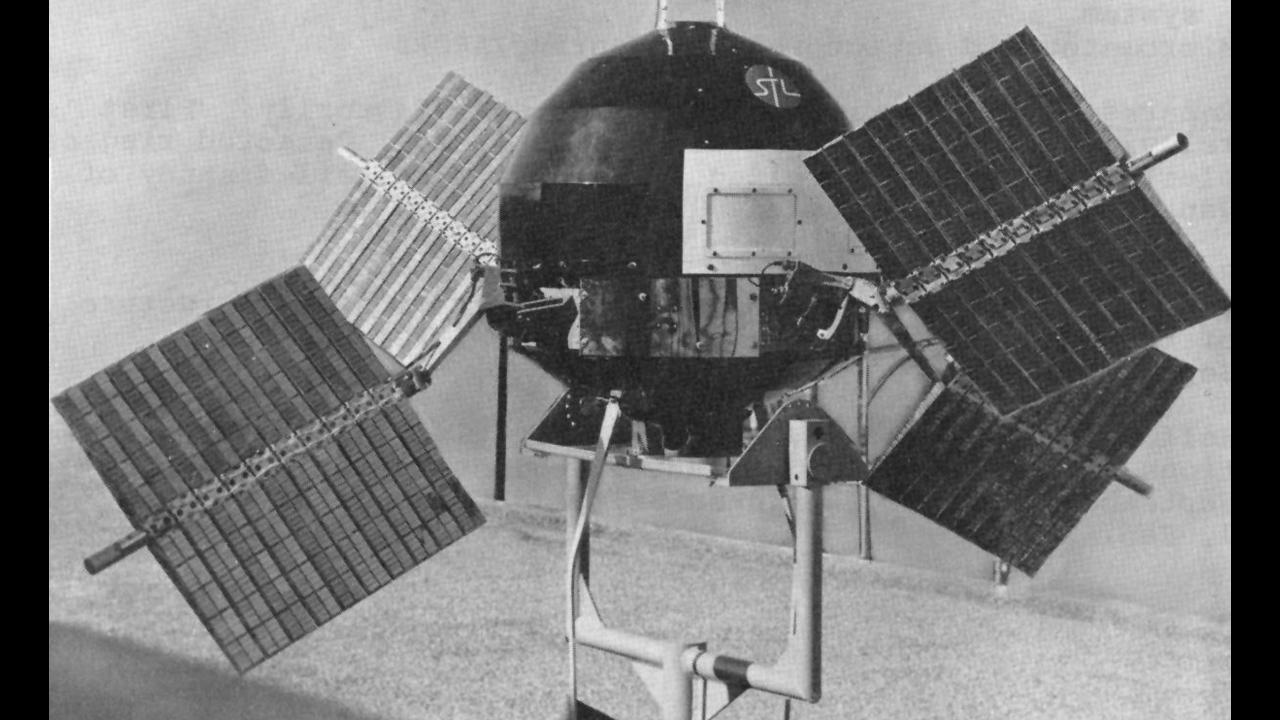
$\log_2 M = \log_{10} M / \log_{10} 2$

 $= 3.32 \log_{10} M$

¹Nyquist, H., "Certain Factors Affecting Telegraph Speed," *Bell System Technical Journal*, April 1924, p. 324; "Certain Topics in Telegraph Transmission Theory," *ALE E. Trans.*, v. 47, April 1928, p. 617. ²Hartley, R. V. L., "Transmission of Information," *Bell System Technical Journal*, July 1928, p. 535.





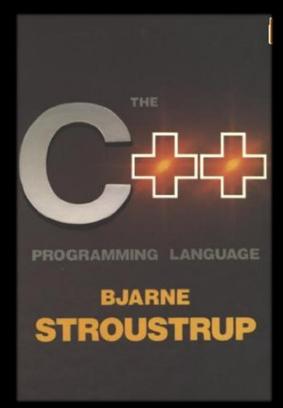


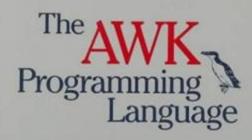


THE C PROGRAMMING LANGUAGE

Brian W. Kernighan • Dennis M. Ritchie

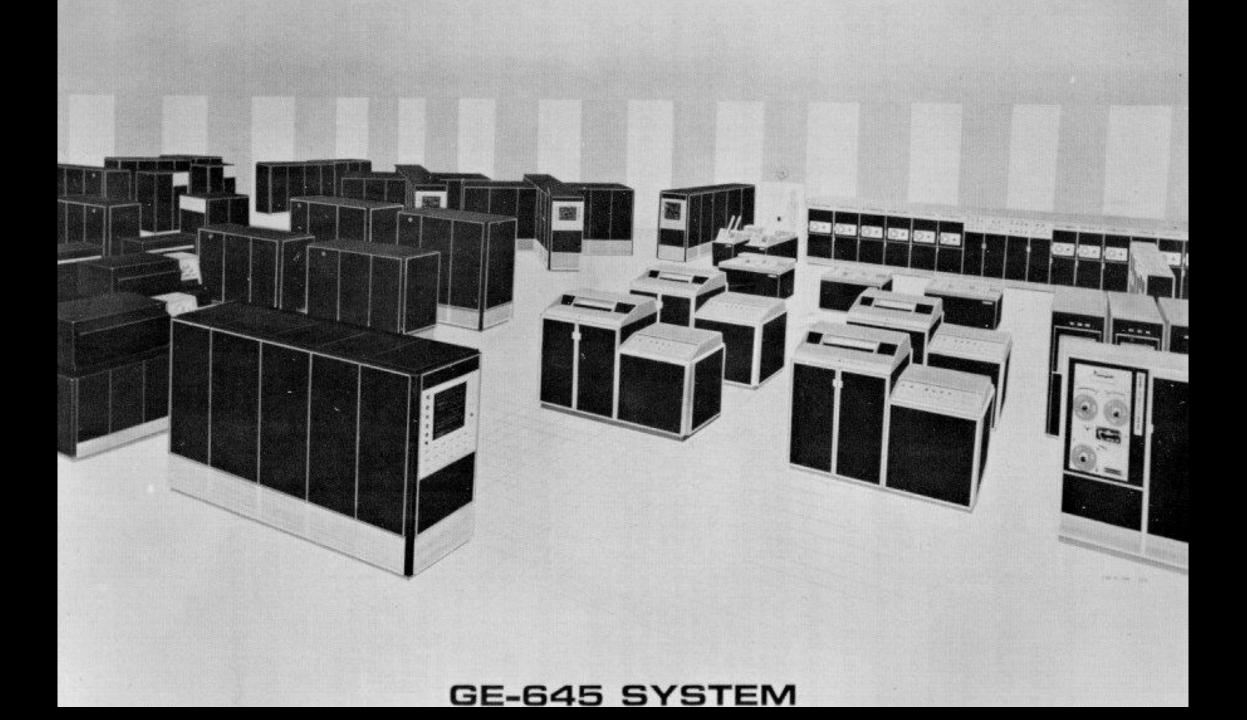
PRENTICE HALL SOFTWARE SERIES





Alfred V. Aho Brian W. Kernighan Peter J. Weinberger



























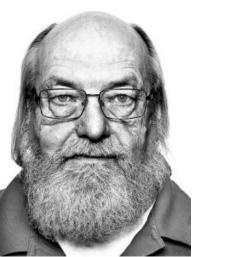
























Faces of Open Source / Peter Adams

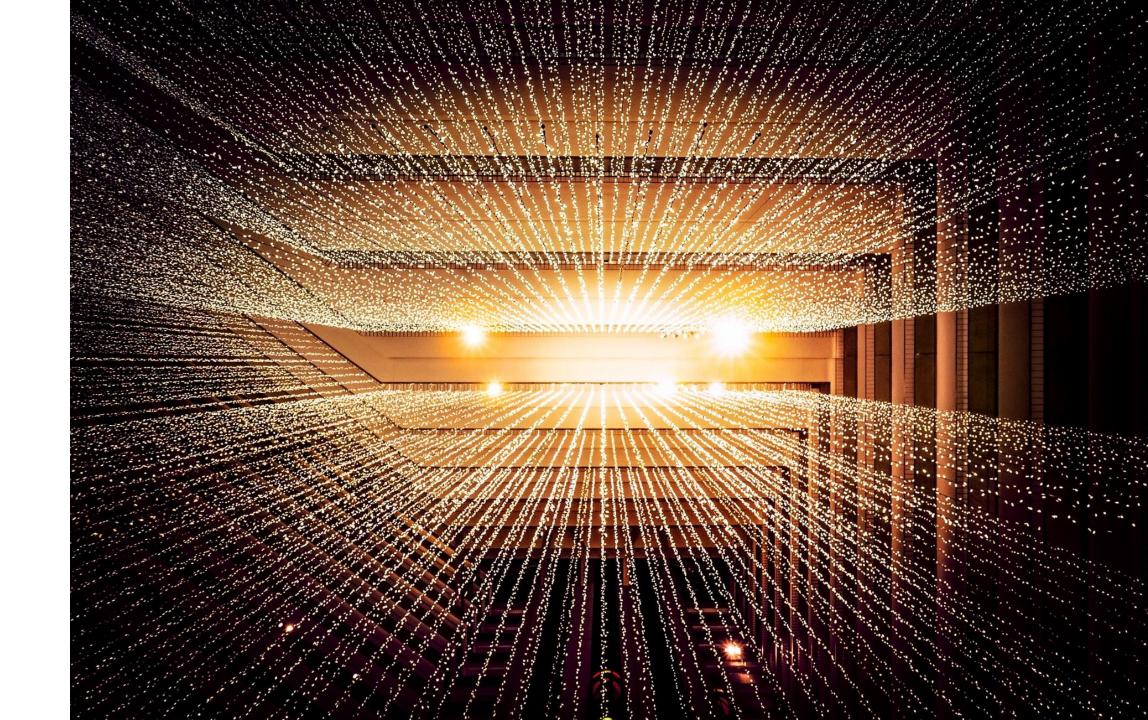


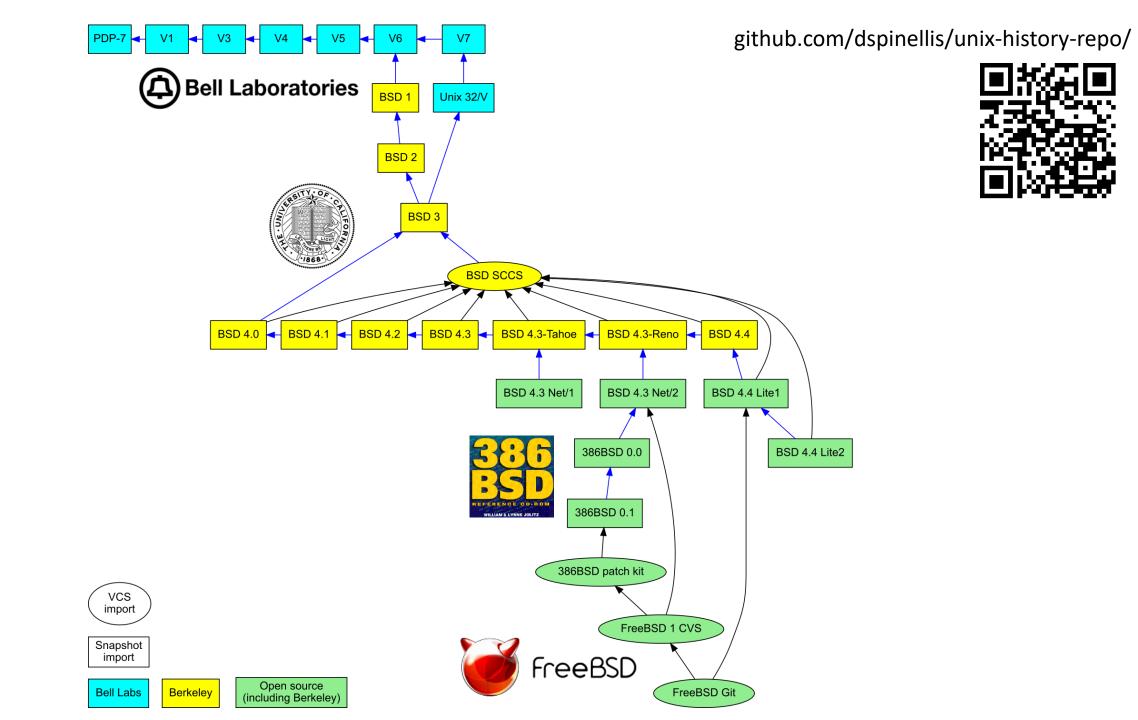






Data Sources







January 23, 2002

Dear UNIX® enthusiasts,

Caldera International, Inc. hereby grants a fee free license that includes the rights use, modify and distribute this named source code, including creating derived binary products created from the source code. The source code for which Caldera International, Inc. grants rights are limited to the following UNIX Operating Systems that operate on the 16-Bit PDP-11 CPU and early versions of the 32-Bit UNIX Operating System, with specific exclusion of UNIX System III and UNIX System V and successor operating systems:

32-bit 32V UNIX 16 bit UNIX Versions 1, 2, 3, 4, 5, 6, 7

Caldera International, Inc. makes no guarantees or commitments that any source code is available from Caldera International, Inc.

The following copyright notice applies to the source code files for which this license is granted.

Copyright(C) Caldera International Inc. 2001-2002. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code and documentation must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed or owned by Caldera International, Inc.

Neither the name of Caldera International, Inc. nor the names of other contributors may be used to endorse or promote products derived from this software without specific prior written permission.

USE OF THE SOFTWARE PROVIDED FOR UNDER THIS LICENSE BY CALDERA INTERNATIONAL, INC. AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL CALDERA INTERNATIONAL, INC. BE LIABLE FOR ANY DIRECT, INDIRECT INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Very truly yours,

/signed/ Bill Broderick

Bill Broderick Director, Licensing Services

* UNIX is a registered trademark of The Open Group in the US and other countries

Bell Laboratories

subject: Study of UNIX

date: September 14, 1972

from: T. R. Bashkow

Messrs. W. S. BartlettMessrs. J. J. LudwigD. P. ClaytonJ. F. MaranzanoD. H. CoppMrs. G. PettitMmes. G. J. HansenMessrs. J. E. RitaccoJ. HintzB. A. TagueMr. L. J. KellyD. W. VogelMiss R. L. KleinMrs. L. S. Wright

On Tuesday, September 19, at 9:30 a.m. in Room 2A-418 at Murray Hill, I will give a talk on my study of the UNIX operating system. The emphasis will be on the structure, functional components, and internal operation of the system.

MH-8234-TRB-mbh

Copy to Mr. G. L. Baldwin T. R. Bashkow

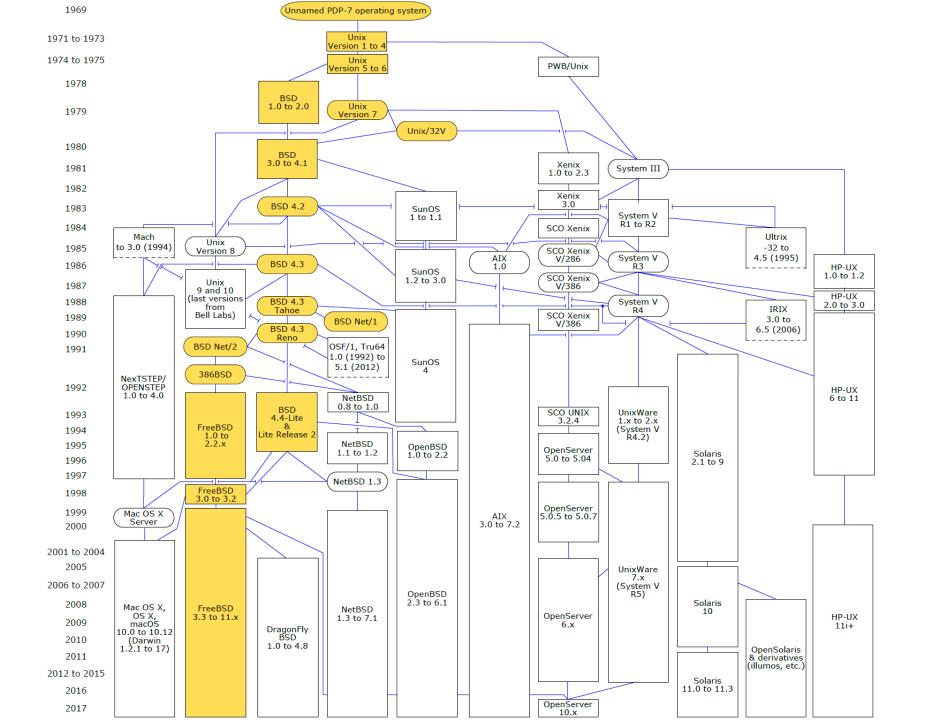
/ initialize inodes for special files (inodes 1 to 40.)

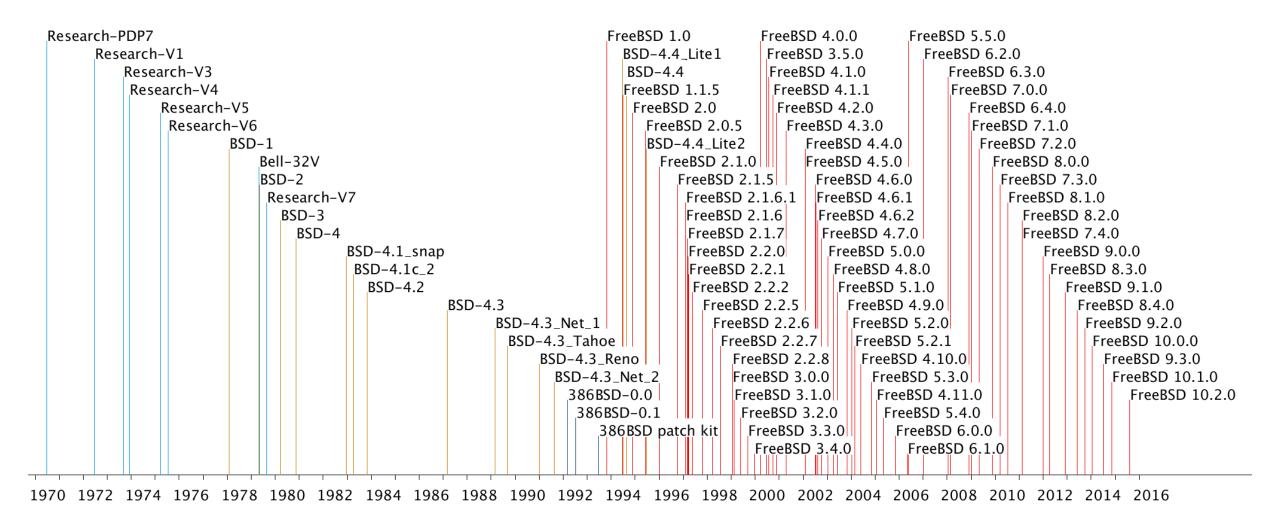
```
$40..r1 / set r1=i-node-number 40.
       mov
             r0, iget / read i-node 'r1' from disk into inode area of
       isr
                    / core and write modified inode out (if any)
             $100017, i.flgs / set flags in core image of inode to indi-
       mov
                          / cate allocated, read (owner, non-owner),
                          / write (owner. non-owner)
             $1.i.nlks / set no. of links = 1
       movb
             $1.i.uid / set user id of owner = 1
       movb
             r0.setimod / set imod=1 to indicate i-node modified, also
       isr
                       / stuff time of modification into i-node
             r1 / next i-node no. = present i-node no.-1
       dec
             1b / has i-node 1 been initialized; no. branch
       bat
/ directories onto fixed head disk. user temporary, initialization prog.
Issue D Date 3/17/72 ID IMO.1-1
                                       Section E.O Page 4
```

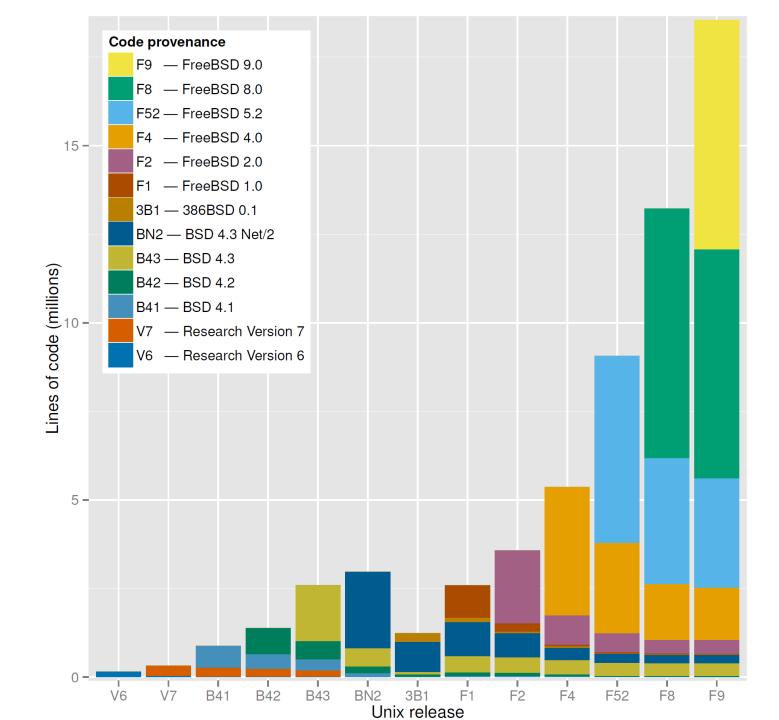
1:

/ u0 unix		mov	<pre>\$idata,r0 / r0=base addr. of assembled directories.</pre>
		mov	su.off, u.fofp / pointer to u.off in u.fofp (holds file
cold = 0			/ offset)
orig = 0 . / orig = 0 . relocatable	1:		··
	••	mov	(r0)+,r1/r1=41.,,47; "0" in the assembled directory
rkda = 177412 / disk address reg rk03/rk11		mov	/ header signals last
rkds = 177400 / driv status reg rk03/rk11		•	
rkcs = 177404 / control status reg rk03/rk11		beq	1f / assembled directory has been written onto drum
rcsr = 174000 / receiver status reg dc-11		jsr	r0, imap / locate the inode map bit for i-node 'r1'
rcbr = 174002 / receiver buffer reg dc-11		bisb	$mq_{i}(r2)$ / set the bit to indicate the i-node is not
tcsr = 174002 / receiver status reg dc-11			/ available
LCSI = 1/4004 / Amer Beacab reg		jsr	r0,iget / read inode 'r1' from disk into inode area of
		2	/ core and write modified i-node on drum (if any)
		mov	(r0)+,i.flgs / set flags in core image of inode from
		alov	/ assembled directories header
		movb	(r0)+.i.nlks / set no. of links from header
			(r0)+,i.uid / set user id of owner from header
		movb	(i), i ulu / Bet user iu of owner from a medified, plac
dcs = 177460 / drum control status rf11/rs11		jsr	r0, setimod / set imod=1 to indicate inode modified: also,
dae = 177470 / drum address extension rf11/rs11			/ stuff time of modification into i-node
1ks = 177546 / clock status reg kw11-1		mov	(r0)+,u.count / set byte count for write call equal to
prs = 177550 / papertape reader status pc11			/ size of directory
prb = 177552 / buffer pc11		mov	r0,u.base / set buffer address for write to top of directory
pps = 177554 / punch status pc11		clr	u.off / clear file offset used in 'seek' and 'tell'
ppb = 177556 / punch buffer pc11		add	u.count,r0 / r0 points to the header of the next directory
/lps = 177514 line printer status (future)		isr	r0, writei / write the directory and i-node onto drum
/lpb = 177516 line printer buffer (future)		br	1b / do next directory
tks = 177560 / console read status asr-33			
tkb = 177562 / read buffer asr-33		<pre>.endif</pre>	
tps = 177564 / punch status asr-33			
tpb = 177566 / punch buffer asr-33	/ next		uctions not executed during cold boot.
ps = 177776 / processor status		bis	\$2000,sb0 / sb0 I/O queue entry for superblock on drum;
			/ set bit 10 to 1
halt = 0		jsr	r0,ppoke / read drum superblock
wait = 1	1:	-	
rti = 2		tstb	sb0+1 / has I/O request been honored (for drum)?
		bne	1b / no. continue to idle.
and the surface of processes			
nproc = 16. / number of processes	1:		
nfiles = 50.	1.	decb	sysflg / mormally sysflag=0, indicates executing in system
ntty = 8+1			
nbuf = 6		sys	exec; 2f; 1f / generates trap interrupt; trap vector =
.if cold / ignored if cold = 0			/ sysent; 0
nbuf = 2		br	panic / execute file/etc/init
.endif			This file#117 lists on E0, d See E0,10
	1:		I I A HUM A TO TAD Son FO.10
core = orig+40000 / specifies beginning of user's core		2f;0	This is file # 11 leads on EU, 4
ecore = core+20000 / specifies end of user's core (4096 wo	2:		
loce and that he const		<td>(init) > / UNIX looks for strings term, noted by nul()</td>	(init) > / UNIX looks for strings term, noted by nul()
		(,,	
A'(unkni;0 bus error	panic:		
/ 4' (unkni;0 " bus error / 10/11 fpsym;0 " illg in tr	pante.	clr	22
14:16 unkni;0 / trace and trap (see Sec. B.1 page)	4.	CII	ps
10,11 unkni;0 / trap	1:		10
14:2 panic:0 / pwr		dec	\$0
30,71 rtssym:0 / emt		bne	1b
34)34 sysent;0 / sys		dec	\$5
		bne	16
Issue D Date 3/17/72 ID IMO.1-1 Section E.0		jmp	*\$173700 / rom loader address

Issue D Date 3/17/72 ID IMO.1-1







\$ git checkout FreeBSD-release/10.0.0

\$ git blame -M -M -C -C ./lib/libc/gen/timezone.c

(Dennis Ritchie usr/src/libc/gen/timezone.c 1979-01-10 14:58:45 -0500 76) static struct zone { usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 77) offset; int 78) usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 char *stdzone; usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 79) char *dlzone: usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 80) $\}$ zonetab[] = { "MET", lib/libc/gen/timezone.c (Jordan К. Hubbard 1996-07-12 18:57:58 +0000 81) $\{-1*60,$ "MET DST"}, [...] lib/libc/gen/timezone.c (Jordan K. Hubbard 1996-07-12 18:57:58 +0000 96) $\{-1\}$ usr/src/lib/libc/gen/timezone.c (Bill Joy 1980-12-22 00:40:25 -0800 97) }; usr/src/lib/libc/gen/timezone.c (Bill Joy 98) 1980-12-22 00:40:25 -0800 usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 106) char * lib/libc/gen/timezone.c (Ed Schouten 2009-12-05 19:31:38 +0000 107) _tztab(int zone, int dst) lib/libc/gen/timezone.c (Rodney Grimes 1994-05-27 05:00:24 +0000 108) { lib/libc/gen/timezone.c (David E. O'Brien 2002 - 02 - 01 01:08:48 + 0000 109)struct zone *zp; lib/libc/gen/timezone.c (David E. O'Brien sign; $2002 - 02 - 01 \ 01:08:48 + 0000 \ 110)$ char usr/src/lib/libc/gen/timezone.c (Bill Joy $1980 - 12 - 22 \quad 00:40:25 \quad -0800 \quad 111)$ usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 112) for (zp = zonetab; zp->offset != -1:++zp) /* static tables */ usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 113) if (zp->offset == zone) { usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 114) if (dst && zp->dlzone) usr/src/libc/gen/timezone.c (Dennis Ritchie $1979 - 01 - 10 \quad 14:58:45 \quad -0500 \quad 115)$ return(zp->dlzone); usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 116) if (!dst && zp->stdzone) usr/src/libc/gen/timezone.c (Dennis Ritchie **1979-01-10 14:58:45** -0500 117) return(zp->stdzone); usr/src/libc/gen/timezone.c (Dennis Ritchie 1979 - 01 - 10 14:58:45 -0500 118)} usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 119) usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 120) if (zone < 0) { /* create one */ usr/src/lib/libc/gen/timezone.c (Bill Joy $1980 - 12 - 22 \quad 00:40:25 \quad -0800 \quad 121)$ zone = -zone;usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 122) sign = '+': usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 123) } usr/src/lib/libc/gen/timezone.c (Keith Bostic else 1987-03-28 19:27:07 -0800 124) sian = '-': usr/src/lib/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 125) lib/libc/gen/timezone.c (void)snprintf(czone, (Warner Losh 1998-01-21 21:46:36 +0000 126) sizeof(czone), "GMT%c%d:%02d",sign,zone / lib/libc/gen/timezone.c (Warner Losh 1998-01-21 21:46:36 +0000 127) 60, zone % 60); lib/libc/gen/timezone.c (Rodney Grimes 1994 - 05 - 27 05:00:24 +0000 128) return(czone); lib/libc/gen/timezone.c (Rodney Grimes 1994-05-27 05:00:24 +0000 129) }

UNIX PROGRAMMER'S MANUAL

Third Edition

K. Thompson

D. M. Ritchie

February, 1973

Copyright 8c9 1972 Bell Telephone Laboratories, Inc.

No part of this document may be reproduced, or distributed outside the Laboratories, without the written permission of Bell Telephone Laboratories.

UNIX PROGRAMMER'S MANUAL

-

Fourth Edition

K. Thompson D. M. Ritchie

November, 1973

Copyright © 1972, 1973 Bell Telephone Laboratories, Inc.

No part of this document may be reproduced, or distributed outside the Laboratories, without the written permission of Bell Telephone Laboratories. dspinellis.github.io/unix-history-man

Evolution of Unix Facilities

- 1. User commands
- 2. System calls
- 3. C library functions
- 4. Devices and special files
- 5. File formats and conventions
- 6. Games et. al.
- 7. Miscellanea
- 8. System maintenance procedures and commands
- 9. System kernel interfaces



Evolution of Unix section 2: System calls

Facility	Appearance	Research V1	Research V2	Research V3	Research V4	Research V5	Research V6	BSD 1	BSD 2	Bell 32V	Research V7	BSD
time	Research V1											
umount	Research V1											
unlink	Research V1											_
wait	Research V1											_
write	Research V1											_
chd	Research V2											
hog	Research V2											
kill	Research V2											_
makdir	Research V2											
sleep	Research V2											
sync	Research V2											_
boot	Research V3											
csw	Research V3											
dup	Research V3											_
fpe	Research V3											
nice	Research V3											_
pipe	Research V3											
times	Research V3											_
aetaid	Research V4							İ				

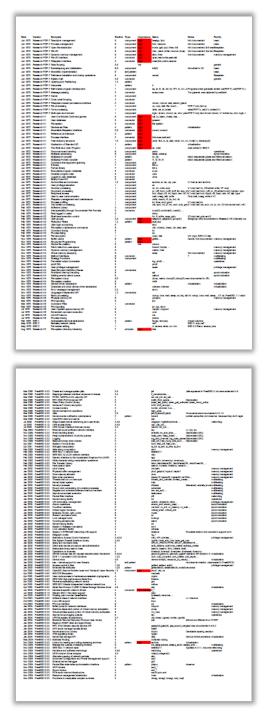
Back to section index

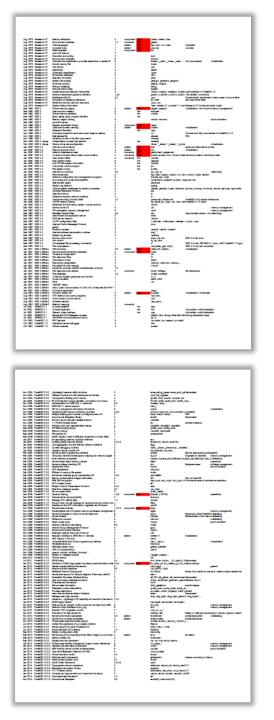
Disclaimers

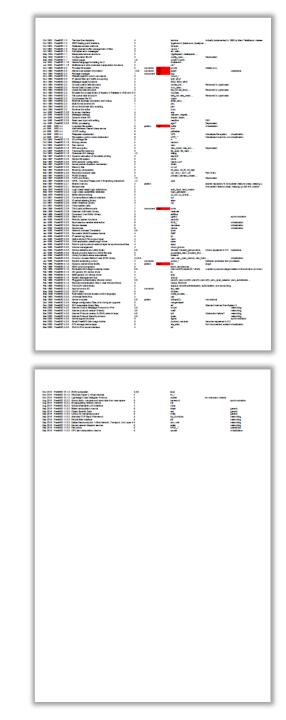
> The name of a facility may have been repurposed over time.

Facilities in sections 1, 6, 8 moved across sections over time. To allow a continuous view of their evolution, all have been relocated to the section of the most recent FreeBSD release, if they still existed at the time.

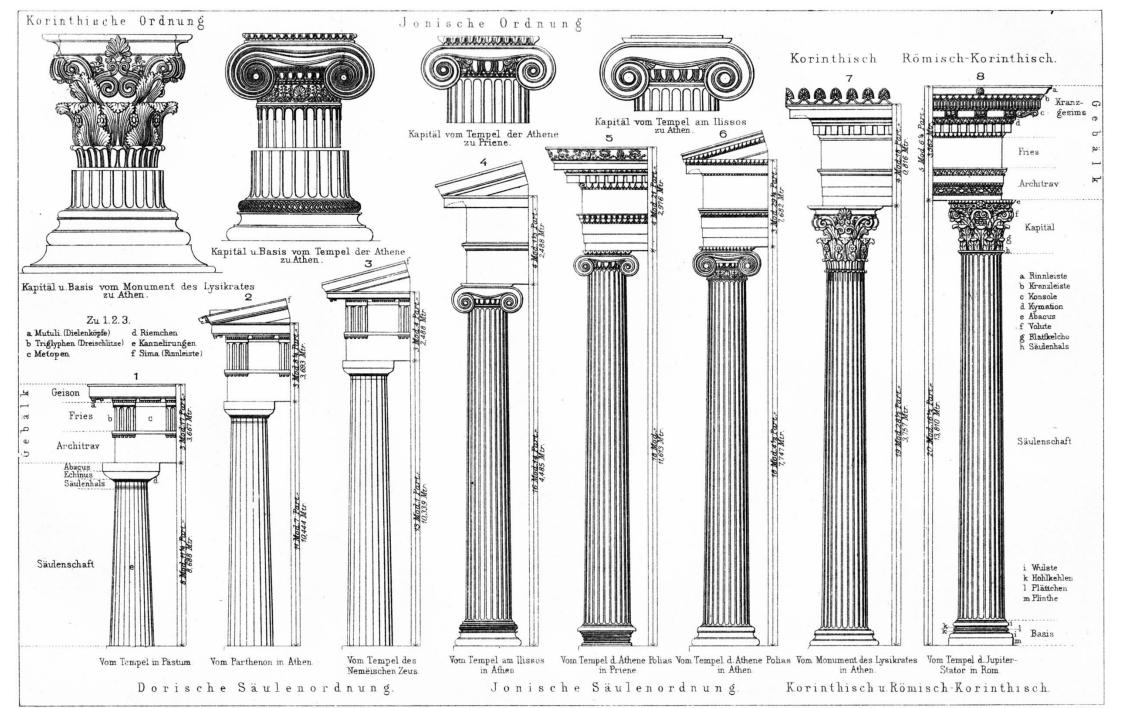
> The evolution data of collapsed tree nodes depict the evolution of the tree's first child node.







00 voluti Ú Unix Ð Archite



Lessons from the Unix Architecture Evolution

Write small programs.



PDP-7 [Unix] (1970)

betveni 0	dae 9fit.	····
lmg omg	oac stri	
las betwen i	Xct between a heliver	BAIMes
686 9£+t	-isz between	
1sz betwen	tral OFAIL 182 between	
1869	CODA.	V //,
486 9F++ 1	Spar 66-	
sma	jmp If	L'N
dmp 1£		$\mathbf{\Theta}$
las betwen 1	Ketbetwer L	
440 9f+t		
ist betwen	187 between	
1869	pro 9Ft	
tad 9fit i	The TIFE	
C 11 B	Spo sna	
spa sna	Spor Stiel	
11	1: 182 between core gftt	
isz betwen		
1869	Trup between i	
C ma		
dme betwen 1		
COPY: 0		
••• 1		
tad copy i		
ã 3 e - 8		
LSZ COPY		

... and I once heard an old-timer growl at a young programmer:

"I've written boot loaders that were shorter than your variable names!"

— Stephen C. Johnson

Build modular code through partitioning, composition, and layering.

Layering and Partitioning

adm.s	cat.s	dskio.s	init.s	s6.s
ald.s	check.s	dskres.s	<pre>lcase.b</pre>	s7.s
apr.s	chmod.s	dsksav.s	maksys.s	s8.s
as.s	chown.s	ds.s	s1.s	s9.s
bc.s	chrm.s	dsw.s	s2.s	scope.v
bi.s	cp.s	ed1.s	s3.s	sop.s
bl.s	db.s	ed2.s	s4.s	trysys.s
cas.s	dmabs.s	ind.b	s5.s	

Value developer time over machine time.

Separation of File Metadata from File Naming

inodel	namei: 0
i.flagst .# + 1	ims iget
1. ASKOSI .# .*7	
i.uid: "F.+1	tad namel i
i.nlkst #1	dae 9f+t+1
	isz namel
i,sizet .E.+1	lac i,flags
i.uniqi .#.#1	and 020
. # inode+12	sha
• ••••••••••••••••••••••••••••••••••••	jmp narel 1
	- B
	tad i size
	CDA
	1r65 3
	dae 9f+t
	SA
	jmp narel i

Devices as Files

```
ttyin:
   <tt>;<yi>;<n 040;040040
ttyout:
   <tt>;<yo>;<ut>; 040040
keybd:
   <ke>;<yb>;<oa>;<rd>
displ:
   <di>;<sp>;<la>;<y 040</pre>
sh:
   <sh>; 040040;040040;040040
system:
   <sy>;<st>;<em>; 040040
```

File I/O API

- open
- read
- write
- seek
- tell
- close

File System API

- creat
- rename
- link
- unlink

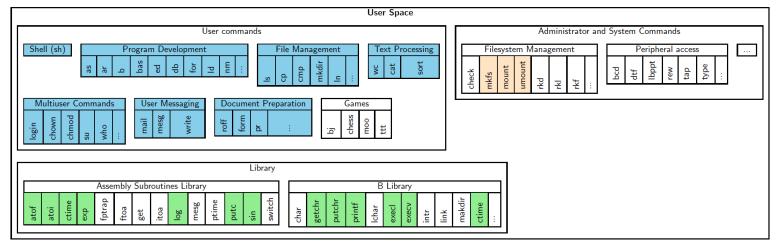
Interpreter

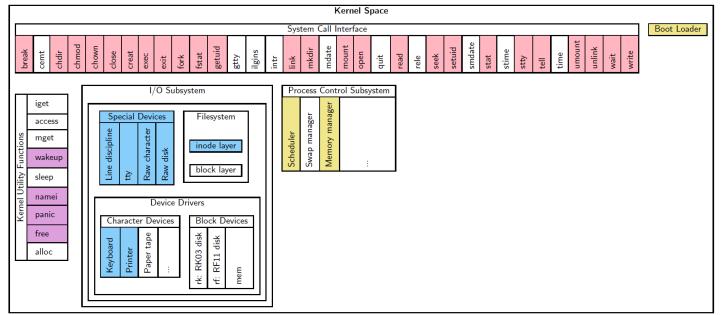
```
ind.b
nain $(
  extrn read, vrite;
  auto i, c. state, line 100;
                                                     1x ind X1
LOOPI
  state = i = 0;
LOOP
   c = read()
   if(c==4) return;
   if(c=='i' \& state==0) state = 2;
  if((c<'0' * c>'9'&c<'a' * c>'z') & state==0} state = 1;
   line[i] = c;
   i = i+1;
   if(c!=012) goto loop1;
   if(state==2 " i==1) goto noi;
   write(' ');
   write(' ');
1011
   i = 0;
.00031
                         1
   c = line[i];
   write(c);
   i = i + 1i
   if(c!=012) goto loop3;
   goto loop;
S)
```

Prototype software before polishing it.



First Research Edition (Nov 1971)





First Edition — 1972 FreeBSD 11.1 — 2018 Linux 6.10 — 2024

<pre>sysrele / 0 0 { int nosys(void); } syscall nosys_args int sysexit / 1 1 { void sys_exit(int rval); } exit \ sys_exit_args void sysread / 3 3 { size_t read(int fd, void *buf, \ size_t nbyte); } sysclose / 6 5 { int open(char *path, int flags, int mode); } sysscreat / 8 sysslink / 9 8 { int creat(char *path, int mode); } sysunlink / 10 10 { int unlink(char *path); }</pre>	0 1 2 3 4 5 6 7 8 9 10	i386 i386 i386 i386 i386 i386 i386 i386	restart_syscall exit fork read write open close waitpid creat link unlink
---	--	--	---

Issue D Date 3/17/72 ID IMO.1-1 Section E.1 Page 1

Make each program do one thing well.

15t edition
151 2000 000
TABLE OF CONTENTS
TABLE OF CONTENTS Nov 3 197/
I. COMMANDS
ar archive (combine) files
asassembler
bb program basbas
basbasic dialect bcdto BCD
boot reboot system
cat concatenate (or print) files
chdir change working directory
check of file system
chmod change access mode of files
chown change owner of files
<pre>cmpfile contents</pre>
cpopy file dateof day
date det date and time of day db
dbppt write binary paper tape
dc desk calculator
df
dswdelete files interactively
dtf format DECtape
dufind disk usage ed
find given name for compile Fortran program
form generate form letter
hup hang up typewriter
lbppttape
ld link editor (loader)
lnlink to file
mail user user mail to another user mess mess description and the messages
mkdir create directory
mkfs
mount mount detachable file system
mv move or rename file
nm print namelist
odoctal dump of file prprint file with headings
pr headings rew
rkd dump disk to tape
rkf format RK disk
rkltape
rm remove (delete) file
rmdir remove (delete) directory
roff text sdate run off (format) text
sdate
stat get file status
strip its strip st
susu become super-user

	<i>C(1)</i>
sum	sum file
tap	manipulate DECtape
tm	get time information
	find name of terminal
tty	
type	print file on IBM 2741
umount	dismount removable file system
un	find undefined symbols
	get (English) word count
WC	who is on the system
who	who is on the system
write	write to another user
II. SYSTEM CALLS	
break	set program break
cemt	catch EMT traps
	change working directory
chdir	
chmod	change mode of file
chown	change owner of file
close	close open file
	create file
creat	
exec	execute program file
exit	terminate execution
fork	create new process
	status of open file
fstat	
getuid	get user ID
gtty	get typewriter mode
ilgins	catch illegal instruction trap
	catch or inhibit interrupts
intr	
link	link to file
mkdir	create directory
mount	mount file system
open	
guit	catch or inhibit quits
read	read file
rele	
seek	TD -
setuid	
smdate	set date modified of file
stat	
stime	
stty	
tell	find read or write pointer
timo	get time of year
umount	dismount file system
unlink	remove (delete) file
unlink	mit for process
wait	walt for process
write	write file
III. SUBROUTINES	
atof	convert ASCII to floating
	convert ASCII to integer
atoi	convert time to ACCTT
ctime	COnvert time to ASCII
exp	exponential function
- vi	_
- 11	

11/3/71	CAT (I)
NAME	cat concatenate and print
SYNOPSIS	<pre>cat file1</pre>
DESCRIPTION	<u>cat</u> reads each file in sequence and writes it on the standard output stream. Thus:
	cat file
	is about the easiest way to print a file. Also:
	cat filel file2 >file3
	is about the easiest way to concatenate files.
	If no input file is given \underline{cat} reads from the standard input file.
FILES	
SEE ALSO	pr, cp
DIAGNOSTICS	none; if a file cannot be found it is ignored.
BUGS	
OWNER	ken, dmr

CAT(1)		User Commands	CAT(1)						
NAME		oncatenate files and print on the standard output							
SYNO									
51110		PTION] [FILE]							
DESCI	RIPTIO								
	Concate	enate FILE(s) to standard output.							
	With no	FILE, or when FILE is –, read standard input.							
	-A,	show-all equivalent to -vET							
	-b,1	number–nonblank number nonempty output lines, overrides – n							
	-е	equivalent to -vE							
	- E ,	show—ends display \$ at end of each line							
	-n,1	number number all output lines							
	-s,s	queeze-blank suppress repeated empty output lines							
	-t	equivalent to -vT							
	- T ,	show-tabs display TAB characters as ~I							
	-u	(ignored)							
	-v,s	how-nonprinting							
		use ^ and M- notation, except for LFD and TAB							
	-	help display this help and exit							
	versi	output version information and exit							
EXAM	PLES	oupla version mornation and exa							
LININ	cat f - g								
		Output f's contents, then standard input, then g's contents.							
	cat	Copy standard input to standard output.							
AUTH		her Taskian Country day d Dishard M. Chillener							
DEDO		by Torbjorn Granlund and Richard M. Stallman.							
KEFU		ary translation bugs to ">https://www.gnu.org/software/coreutils/>							
COPY	<https: <br="">This is</https:>	, ght © 2022 Free Software Foundation, Inc. License GPLv3+: GNU GPL version 3 //gnu.org/licenses/gpl.html>. free software: you are free to change and redistribute it. There is NO WARRANTY, to the d by law.							
SEE A	•								
	Full do	cumentation <https: cat="" coreutils="" software="" www.gnu.org=""> able locally via: info '(coreutils) cat invocation'</https:>							
GNU co	oreutils 9	.1 September 2022	1						

Avoid captive interfaces.

The Shell as a User Program

11/3/71

PASSWD (V)

NAME	passwd password file
SYNOPSIS	
DESCRIPTION	<u>passwd</u> contains for each user the following information:
	name (login name) password numerical user ID default working directory <mark>program to use as Shell</mark>
	This is an ASCII file. Each field within e

This is an ASCII file. Each field within each user's entry is separated from the next by a colon. Each user is separated from the next by a new-line. If the password field is null, no password is demanded; if the Shell field is null, the Shell itself is used.

Write extensible programs and protocols.

Abstraction of Standard I/O

11/3/71

SH (I)

Two characters cause the immediately following string to be interpreted as a special argument to the shell itself, not passed to the command. An argument of the form "<arg" causes the file arg to be used as the standard input file of the given command; an argument of the form ">arg" causes file "arg" to be used as the standard output file for the given command.

User-Contributed Tools and Games

VI. USER MAINTAINED PROGRAMS

basic	DEC supplied BASIC
bj	
cal	print calendar
chess	the game of chess
das	disassembler
dli	load DEC binary paper tapes
dpt	read DEC ASCII paper tapes
moo	the game of MOO
sort	
ttt	the game of tic-tac-toe





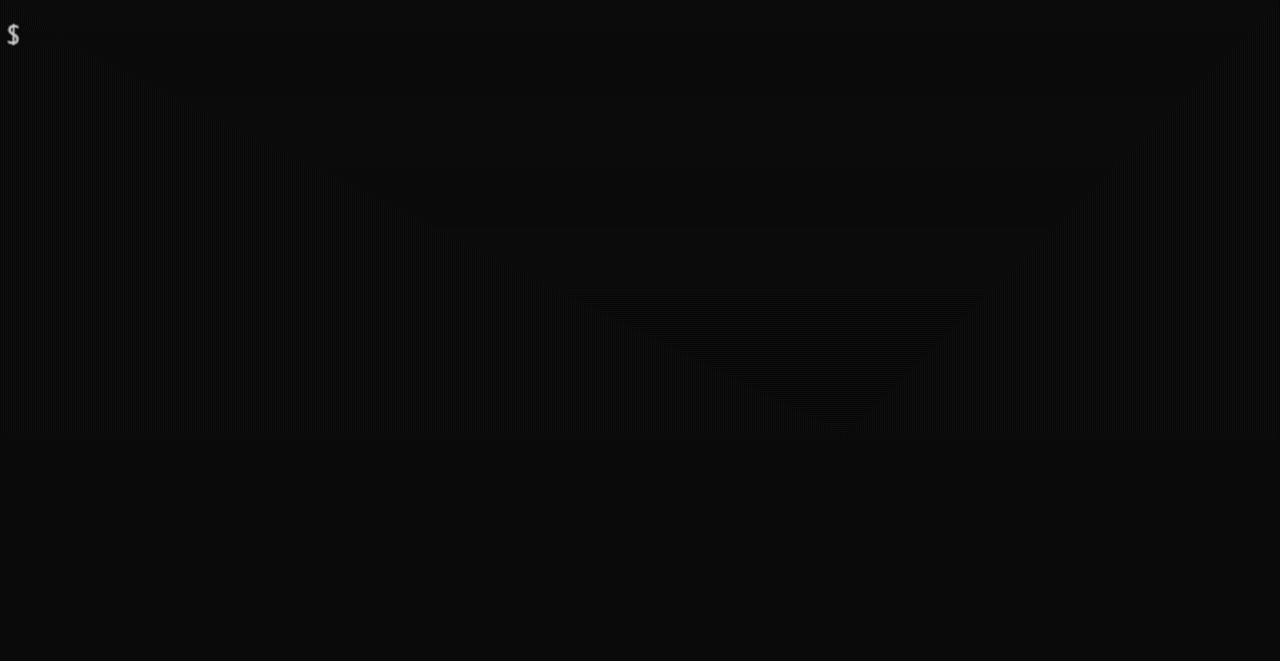
• • • • • • • •

•••••

.

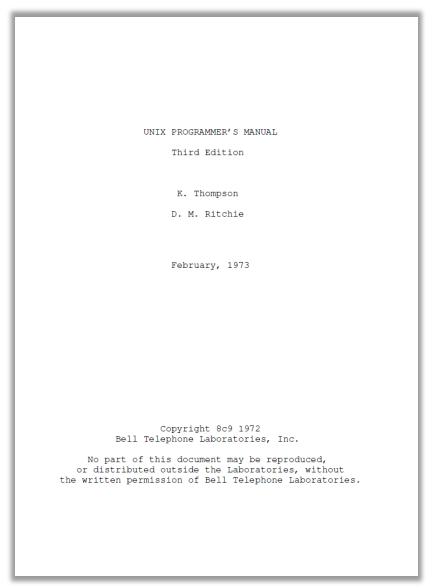
. . .

• •



Write programs that work together as filters that process text streams.

Third Research Edition (Feb 1973)



Pipes and Filters

10

Summery--what S most important. To put my strongest concerns in a nutshell: 1. We should have some ways of coupling programs bike garden hose--acrew in another segment when it becomes then it becomes necessary to massage data in another way. This is the way of 10 also.

Oct

. .

PIPE (II)

NAME pipe -- create a pipe

SYNOPSIS sys pipe / pipe = 42.; not in assembler (file descriptor in r0)

DESCRIPTION The <u>pipe</u> system call creates an I/O mechanism called a pipe. The file descriptor returned can be used in both read and write operations. When the pipe is written, the data is buffered up to 504 bytes at which time the writing process is suspended. A read on the pipe will pick up the buffered data.

> It is assumed that after the <u>pipe</u> has been set up, two (or more) cooperating processes (created by subsequent <u>fork</u> calls) will pass data through the pipe with <u>read</u> and <u>write</u> calls.

The shell has a syntax to set up a linear array of processes connected by pipes.

Write maintainable programs.

Fourth Research Edition (Nov 1973)

UNIX PROGRAMMER'S MANUAL Fourth Edition K. Thompson D. M. Ritchie November, 1973 Copyright © 1972, 1973 Bell Telephone Laboratories, Inc. No part of this document may be reproduced, or distributed outside the Laboratories, without the written permission of Bell Telephone Laboratories.

Structured Programming

- Kernel implemented in "New B"
 - 6373 lines New B
 - 768 lines PDP-11 assembly
- Improvement:
 - First Ed.: 248 global symbols
 - Fourth Ed.: 105 functions, 50 assembly symbols

main() extern schar; extern char end[], data[], etext[]; int i, i1, *p; /* * zero and free all of core */ $UISA \rightarrow r[0] = KISA \rightarrow r[6] + USIZE;$ $UISD \rightarrow r[0] = 6;$ for(; fubyte(0) >= 0; UISA->r[0]++) { clearseg(UISA->r[0]); mfree(coremap, 1, UISA->r[0]); mfree(swapmap, NSWAP, SWPLO); /* * set up system process */

```
proc[0].p_addr = KISA->r[6];
proc[0].p_size = USIZE;
proc[0].p_stat = SRUN;
proc[0].p_flag =| SLOAD|SSYS;
u.u_procp = &proc[0];
```

```
/*
* set up 'known' i-nodes
*/
```

sureg(); LKS->integ = 0115; cinit(); binit(); iinit(); rootdir = iget(ROOTDEV, ROOTINO); rootdir->i_flag =& ~ILOCK; u.u_cdir = iget(ROOTDEV, ROOTINO); u.u_cdir->i_flag =& ~ILOCK;

Language-Independent API

PIPE(II)

8/5/73

PIPE(II)

NAME

pipe – create a pipe

SYNOPSIS

(pipe = 42.) **sys pipe** (read file descriptor in r0) (write file descriptor in r1)

pipe(fildes) int fildes[2];

DESCRIPTION

The *pipe* system call creates an I/O mechanism called a pipe. The file descriptors returned can be used in read and write operations. When the pipe is written using the descriptor returned in r1 (resp. fildes[1]), up to 4096 bytes of data are buffered before the writing process is suspended. A read using the descriptor returned in r0 (resp. fildes[0]) will pick up the data.

It is assumed that after the pipe has been set up, two (or more) cooperating processes (created by subsequent *fork* calls) will pass data through the pipe with *read* and *write* calls.

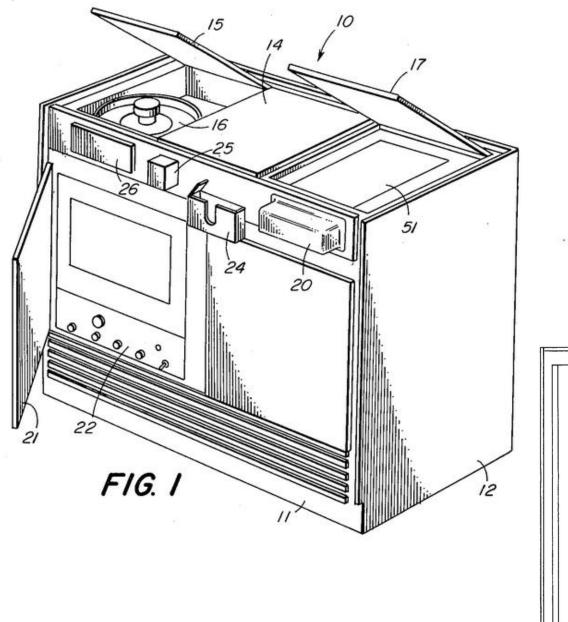
The shell has a syntax to set up a linear array of processes connected by pipes.

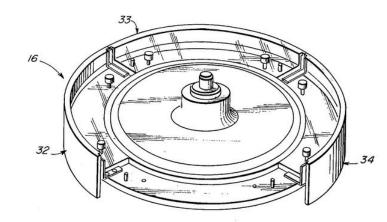
Read calls on an empty pipe (no buffered data) with only one end (all write file descriptors closed) return an end-of-file. Write calls under similar conditions are ignored.

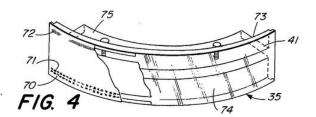
Data Structure Definition & Reuse

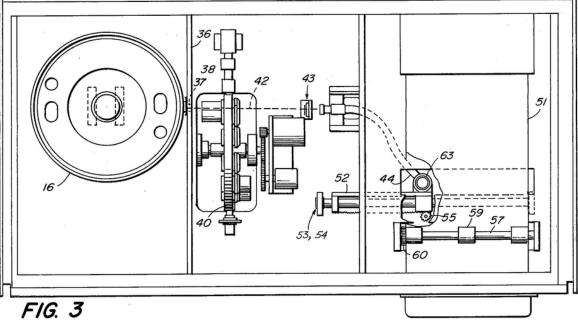
- buf.h filsys.h proc.h text.h
- conf.h inode.h reg.h tty.h
- file.h param.h systm.h user.h

Avoid unnecessary output and make failures easy to diagnose.









ASTRICT LOW 3 Œ / 08 11 2 Q k q 0 ĸ å Ø ø А х B x n Ν b 5 . 9 1 Æ 6 20 W w J h H G g 11 £ 2 2 5 11 O ö D d 5 č Ü ü v S v 1 1 . = * -. Ä ä F U U i 0 1 P M 0 2 m p a 0 ß % 4 Z C. C. C Z -12 Berthold AG Berlin an and the second second

NAME

cat - phototypesetter interface

DESCRIPTION

Cat provides the interface to a Graphic Systems C/A/T phototypesetter. Bytes written on the file specify font, size, and other control information as well as the characters to be flashed. The coding will not be described here.

Only one process may have this file open at a time. It is write-only.

FILES

/dev/cat

SEE ALSO

troff (I), Graphic Systems specification (available on request)

BUGS

⁻aces of Open Source / Peter Adams

"After phototypesetting, you had to take a long wide strip of paper and feed it carefully into a smelly, icky machine which eventually (several minutes later) spat out the paper with the printing visible."

"One afternoon several of us had the same experience typesetting something, feeding the paper through the developer, only to find a single, beautifully typeset line: "cannot open file foobar" The grumbles were loud enough and in the presence of the right people, and a couple of days later the standard error file was born..."

— Stephen C. Johnson

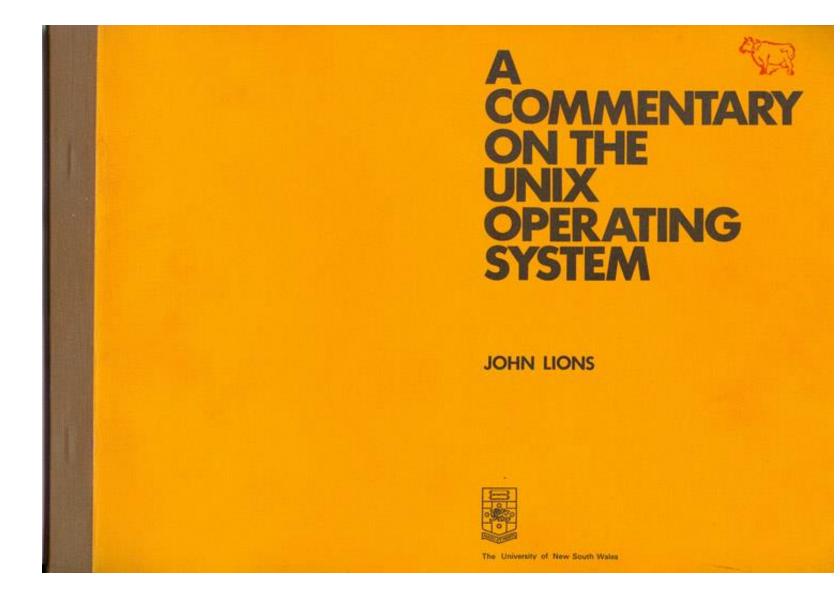
Use shell scripts to increase leverage and portability.

Fifth Research Edition (Jun 1974)

 Command Files chdir /usr/source/s3 cc -c ctime.c ar r /lib/liba.a ctime.o rm ctime.o chdir /usr/source/s1 cc -s -n date.c cp a.out /bin/date cc -s -n dump.c cp a.out /bin/dump cc -s -n ls.c cp a.out /bin/ls rm a.out

Choose appropriately powerful abstractions.

Sixth Research Edition (May 1975)

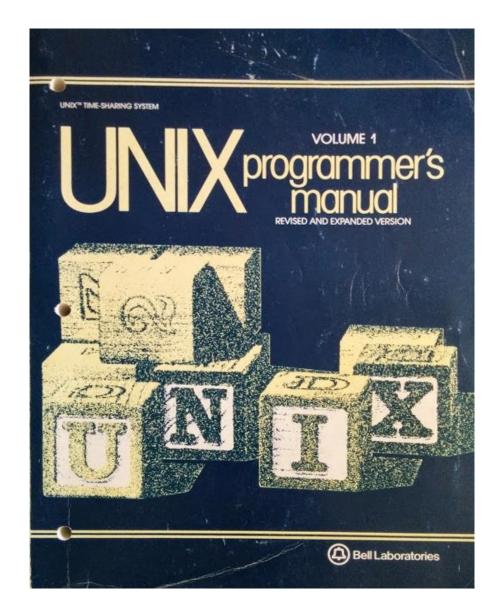


Portable C Library

alloc.c clenf.c calloc.c copen.c cclose.c cputc.c ceof.c cwrd.c dummy.s cerror.c ftoa.c cexit.c cflush.c getch.c cfree.c gets.c getvec.c cgetc.c ciodec.c iehzap.c

makbuf.c scan1.c maktab.c scan2.c nexch.c scan3.c nodig.c system.c printf.c tmpnam.c putch.c unget.c puts.c unprnt.s relvec.c wdleng.c revput.c run

Seventh Research Edition (Jan 1979)



Unix as a Virtual Machine

Also, about this time [1973] I had a fateful discussion with Dennis, in which he said:

"I think it may be easier to port Unix to a new piece of hardware than to port a complex application from Unix to a new OS"

- Steve Johnson

Separate mechanisms from policy.

Dynamic User Memory Allocation

- malloc(3), free(3)
 - Used by 26 programs: awk cc col cron dc dcheck diff ed eqn expr graph icheck learn ls m4 neqn nm quot ratfor spline struct tar tsort uucp xsend quiz
- stdio(3), mp(3)

Static Analysis

Environment Variables

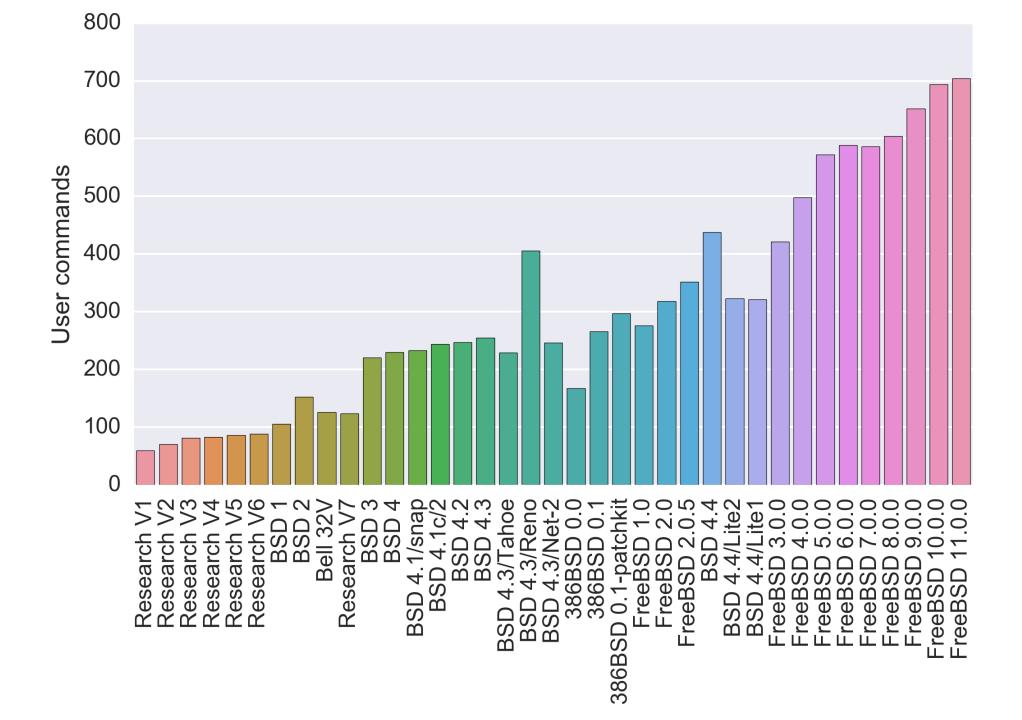
- KEY=value
- Kernel
- Shell
- C Library

ENVIRON(5)	UNIX Programmer's Manual	ENVIRON(5)		
NAME environ	– user environment			
SYNOPSIS extern	char **environ;			
	y of strings called the 'environment' is made available by <i>exec</i> (2) when ion these strings have the form 'name=value'. The following names ar			
PATH	The sequence of directory prefixes that <i>sh</i> , <i>time</i> , <i>nice</i> (1), etc., apply known by an incomplete path name. The prefixes are separated PATH=:/bin:/usr/bin.			
HOME	A user's login directory, set by <i>login</i> (1) from the password file <i>passwd</i> (5).		
TERM	The kind of terminal for which output is to be prepared. This inform mands, such as $mroff$ or $plot(1)$, which may exploit special terminal can for a list of terminal types.			
in $sh(1)$	names may be placed in the environment by the <i>export</i> command and ', or by <i>exec</i> (2). It is unwise to conflict with certain Shell variables that file' files: MAIL, PS1, PS2, IFS.			
SEE ALSO				

exec(2), sh(1), term(7), login(1)

Filesystem Directory Hierarchy

		HIER(7)	UNIX Programmer's Manual			HIER(7)					
ER(7)	UNIX Programmer's Mar				_		HIER(7)			UNIX Programmer's Manual	HIE
	· ·		HIER(7)			UNIX Programmer's Manual					
			vpacct line printer accounting lpr								
		/usr	/bin								
ME	a sector data and a		utility programs, to keep /bin/ small			intro introduction to volume 1, n	/usr/	src/			
	e system hierarchy		tmp/ temporaries, to keep /tmp/ small			xx template for manual page	5()	source program			
SCRIPTION			stm* used by sort(1)		man1/	chapter 1		cmd/		of commands	
The follo	owing outline gives a quick tour through a represer		raster used by plot(1)			as.1			as/	assembler makefile	
/ 1	root		dict/ word lists, etc.			mount.1m				recipe for rebuilding the assembler	
/dev/ o	devices (4)		words principal word list, used b spellhist							as1?.s source of pass1	
	console main console, $tty(4)$		history file for spell(1)		cat1/	preprinted pages for man1/			ar.c	source for ar(1)	
t	tty* terminals, tty(4)		games/			as.1			ar.c	source for an (1)	
6	cat phototypesetter cat(4)		bj blackjack			mount.1m			troff/	source for <i>nroff</i> and <i>troff</i> (1)	
	rp* disks, rp, hp(4)		hangman							nmake makefile for nroff	
1	rrp* raw disks, rp, hp(4)		quiz.k/ what quiz(6) knows	spool/		execution files				tmake makefile for troff	
			index category index		at/ lpd/	used by at(1) used by lpr(1)				font/ source for font tables, /usr/lib/font/	
	utility programs, cf /usr/bin/ (1)		africa countries and cap		ipa/	lock present when line printer is	aat			ftR.c Roman	
	as assembler first pass, cf /usr/lib/as2					cf* copy of file to be printed in					
	cc C compiler executive, cf /usr/lib/c[012]					df* daemon control file, <i>lpd</i> (8)	ne			term/ terminal characteristics tables, /usr/lib/term/	
/lib/ 0	object libraries and other stuff, cf /usr/lib/		include/ standard #include files			tf* transient control file, while	'nr			tab300.c	
	libc.a system calls, standard I/O, etc. (2,3,3S)		a.out.h object file layout, a.out(5)		uucp/	work files and staging area for uucp				DASI 300	
	libm.a math routines (3M)		stdio.h standard I/O, stdio(3)		and b	LOGFILE					
	libplot.a		math.h (3M)			summary log		libc/	source	for functions in /lib/libc.a	
	plotting routines, plot(3)		sys/ system-defined layouts, cf			LOG.* log file for one transaction		11007	crt/	C runtime support	
1	libF77.a		acct.h process accounts,	mail/		xes for mail(1)			cit	ldiv.s division into a long	
	Fortran runtime support		buf.h internal system by		uid	mail file for user uid				Imul.s multiplication to produce long	
1	libI77.a Fortran I/O		out in internal system of		uid.loc						
			lib/ object libraries and stuff, to keep /l			lock file while uid is receiving mail			csu/	startup and wrapup routines needed with every C program	1
	as2 second pass of as(1)		lint[12] subprocesses for lint(1)	wd		working directory of a user, typically				crt0.s regular startup	
6	c[012] passes of cc(1)		llib-lc dummy declarations for /li		.profile calenda	set environment for sh(1), environ(5)			mcrt0.s modified startup for $cc - p$	
			llib-lm dummy declarations for /li		calenda	user's datebook for calendar(1)			sys/	system calls (2)	
	essential data and dangerous maintenance utilities		atrun scheduler for $at(1)$	doc/		mostly in volume 2 of this manual, t				access.s	
-	passwd password file, passwd(5)		struct/ passes of struct(1)	doc/	as/	assembler manual	pic			alarm.s	
	group group file, group(5)				c alar	C manual					
	motd message of the day, <i>login</i> (1)		tmac/ macros for troff(1)			C mandar			stdio/	standard I/O functions (3S)	
	mtab mounted file table, $mtab(5)$		tmac.an macros for man(7	svs/	system	source				fgets.c	
	ddate dump history, dump(1) ttys properties of terminals, ttys(5)		tmac.s macros for $ms(7)$	- / -	dev/	device drivers				fopen.c	
	ttys properties of terminals, ttys(5) getty part of login, getty(8)		Suntil Sunta Sun tra (1/1)			bio.c common code			gen/	other functions in (3)	
	init the father of all processes, <i>init</i> (8)		font/ fonts for troff(1) R Times Roman			cat.c cat(4)			gen/	abs.c	
	rc shell program to bring the system up		B Times Bold			dh.c DH11, tty(4)				atof.c	
	cron the clock daemon, cron(8)		D Thics Dold			tty tty(4)					
	mount mount(1)		uucp/ programs and data for uuc						compal	ll shell procedure to compile libc	
1	wall wall(1)		L.sys remote system na		conf/	hardware-dependent code				shell procedure to make /lib/libc.a	
			uucico the real copy prog			mch.s assembly language portion conf configuration generator			source	for /lib/libI77	
/tmp/ t	temporary files, usually on a fast device, cf /usr/tm					conf configuration generator		libF77/			
e	e* used by ed(1)		suftab table of suffixes for hyphe		h/	 header (include) files					
	ctm* used by cc(1)		units conversion tables for units		10	acct.h acct(5)		games/	source	for /usr/games	
			eign list of English words to be			stat.h stat(2)	SEE ALSO	1 1/12 0 1/23	(**		
	general-pupose directory, usually a mounted file sy	/usr/	man/			Start Star(2)		ncheck(1), find(1),	grep(1)		
e	adm/ administrative information		volume 1 of this manual, man(1)		sys/	source for system proper	BUGS		dia e e	-to	
	wtmp login history, utmp(5)		man0/ general			main.c	The p	ostuon of files is s	ubject to	change without notice.	
	messages					pipe.c					
	hardware error messages					sysent.c system entry points					
	tracct phototypesetter accounting, troff(7th Edition									
		/ III EUIUOII					7th Edition				
							. ur zonioli				
			7th Edition						3		



Write abstract programs that generate code instead of writing code by hand.

Language Development Tools

- lex(1)
- yacc(1)
- 12 clients:
 - awk
 - bc
 - срр
 - egrep
 - eqn
 - lex
 - m4
 - make
 - pcc
 - neqn
 - struct

Copyright © 1978 American Telephone and Telegraph Company THE BELL SYSTEM TECHNICAL JOURNAL Vol. 57, No. 6, July-August 1978 Printed in U. S. A.

UNIX Time-Sharing System:

Language Development Tools

By S. C. JOHNSON and M. E. LESK (Manuscript received December 27, 1977)

The development of new programs on the UNIX^{*} system is facilitated by tools for language design and implementation. These are frequently program generators, compiling into C, which provide advanced algorithms in a convenient form, while not restraining the user to a preconceived set of jobs. Two of the most important such tools are Yacc, a generator of LALR(1) parsers, and Lex, a generator of regular expression recognizers using deterministic finite automata. They have been used in a wide variety of applications, including compilers, desk calculators, typesetting languages, and pattern processors.

Raise abstraction through DSLs.

Domain-Specific Languages

- sh
- awk
- sed
- find
- expr
- egrep
- m4
- make

Architectural innovations are sticky and face increasing resistance.



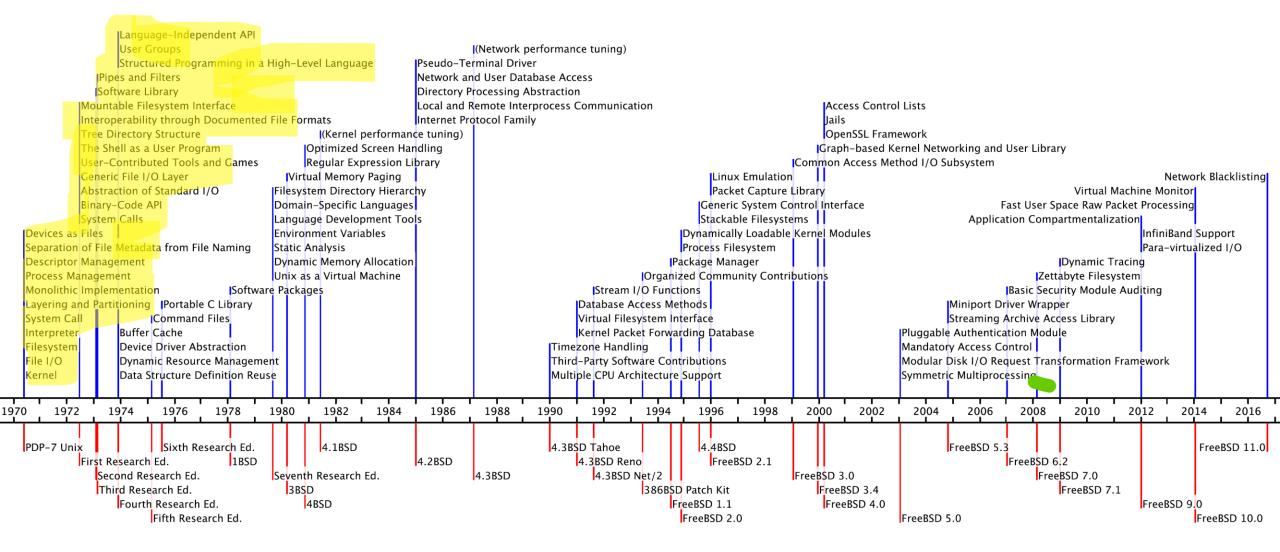
Regular Expression Library: regex(3)

- 5 implementations: awk, sed, ed, grep, expr
- 1 client: more(1)
- 2 more by 4.3: dbx(1), rdist(1)
- 4 replacements in FreeBSD: ed, grep, sed, expr

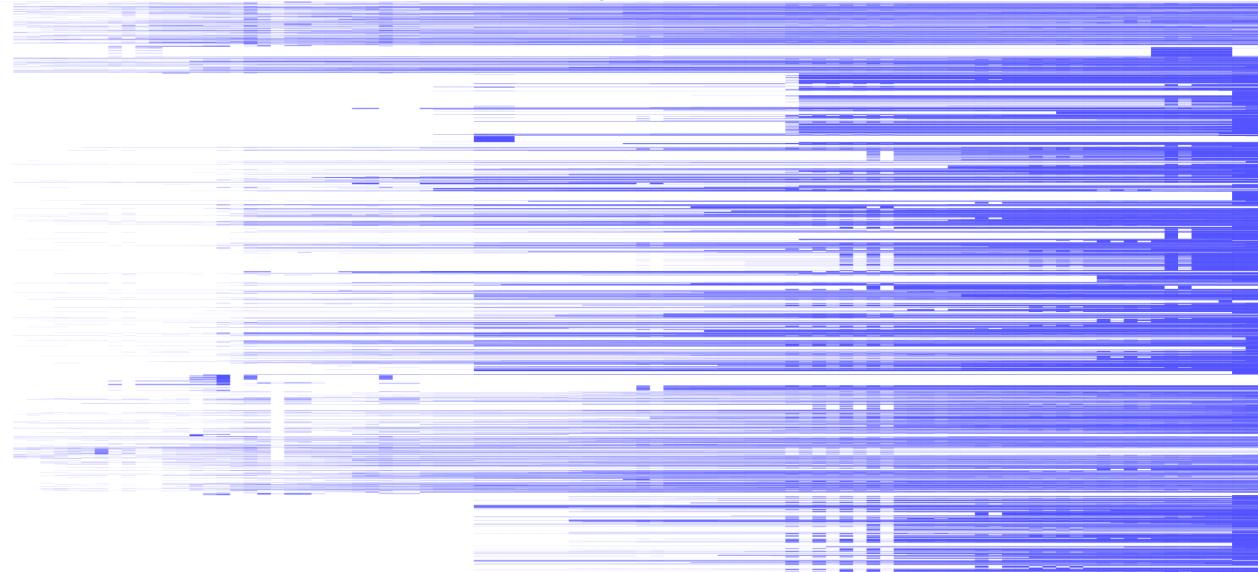
4.3BSD Net/2 (Jun 1991)

- Stream I/O Functions
 - funopen(3)
 - GNU funopencookie(3) added in FreeBSD 11

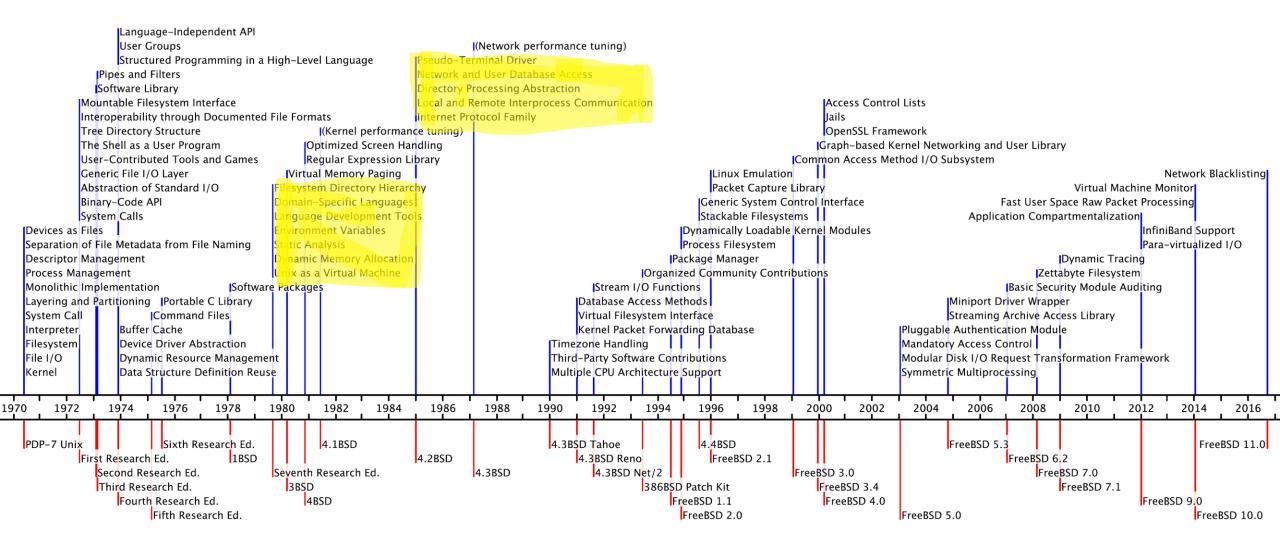
Many core architecture decisions are taken at the beginning of the system's lifetime



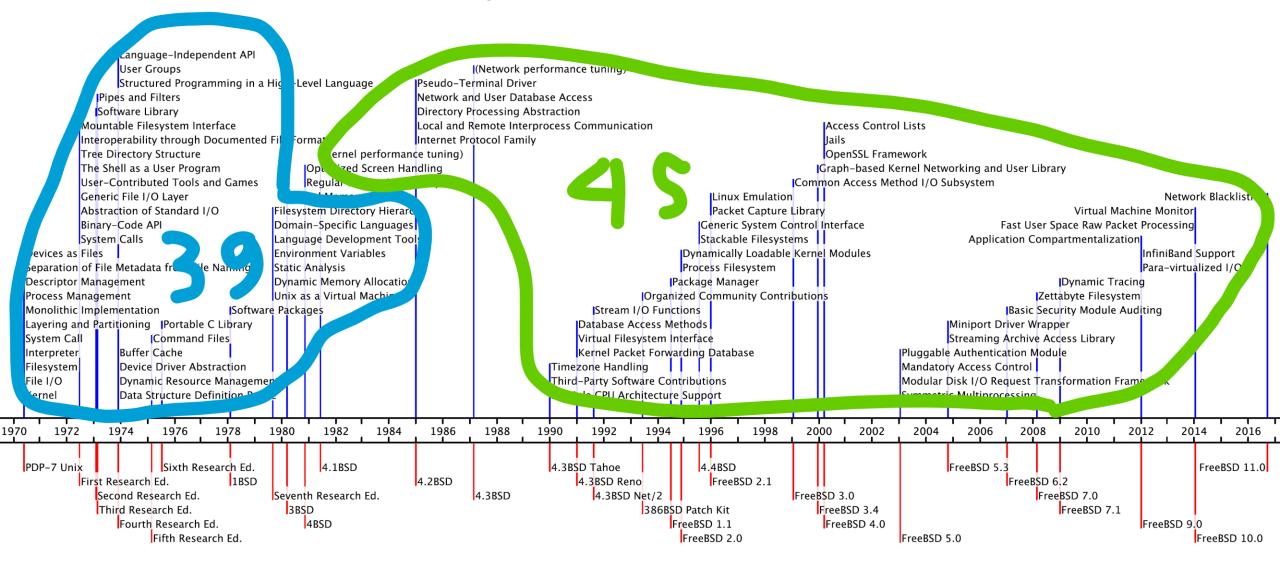
Most architecture decisions survive over the system lifetime

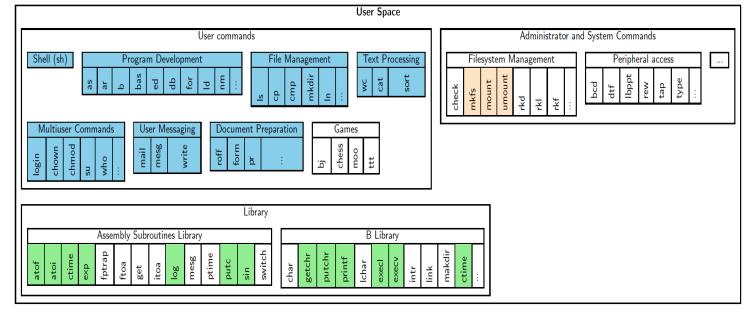


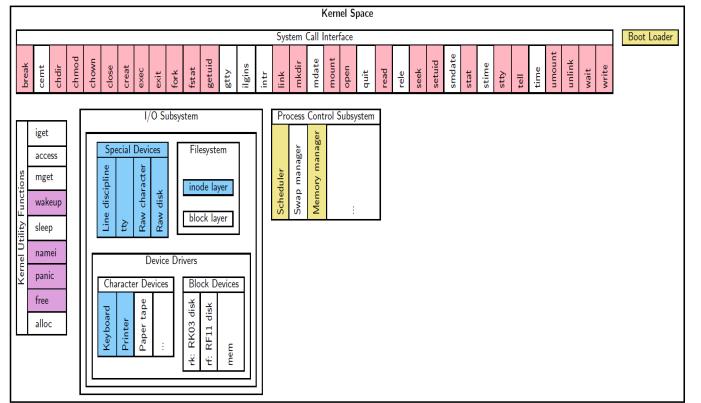
New architecture decisions are continuously made, further fueling architecture evolution

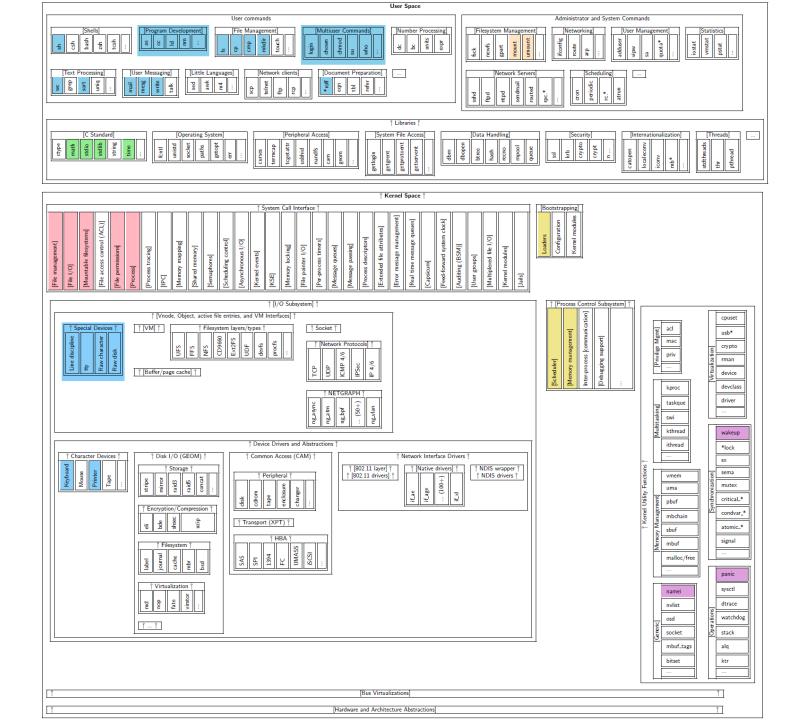


The rate of architecture decisions declines over the system's lifetime









Package managers grow ecosystems and communities.



386BSD Patch Kit (1992-1993)

- Patch metadata
 - title
 - author
 - description
 - prerequisites
- Organized Community Contributions
 - From open-source software ...
 - ... to an open-source project



FreeBSD 1.1 (May 1994)

- Package Manager
 - Patch
 - Compile
 - Install
 - Uninstall
 - Handling of dependencies

Package ecosystems

Operating system

• GNU/Linux: Debian, Fedora, Ubuntu, ...

لسك ليا

• FreeBSD, NetBSD, OpenBSD

Package repository / manager

- Maven / mvn, Gradle
- PyPI / pip
- NPM / npm, yarn
- CRAN
- RubyGems / gem

Activity

• Data science (Python Anaconda)



ndex

Maven^{**}



Thank you!

- www.spinellis.gr
- X @CoolSWEng
- @CoolSWEng@mastodon.acm.org
- ⊠ dds@aueb.gr





Free open edX course (MOOC): Unix Tools: Data, Software and Production Engineering

Grow from being a Unix novice to Unix wizard status! Process big data, analyze software code, run DevOps tasks and excel in your everyday job through the amazing power of the Unix shell and command-line tools.

https://www.spinellis.gr/unix





Image Credits

- Faces of Open Source / Peter Adams
- Data: Joshua Sortino
- Hackers at Junction 2015: <u>Vmuru</u>
- <u>ASR-33 Teletype</u>: <u>Rama</u> & Musée Bolo
- VT100: Jason Scott
- <u>PDP 11/20</u>: Image courtesy of Computer History Museum

(Creative commons licenses)

- PDP11/40: Stefan_Kögl, CC BY-SA 3.0
- Digital VAX 11/780: Emiliano Russo, PD
- Numbers: Nick Hillier
- Building construction: chuttersnap
- Technical Debt: Jacob Duck Die Pfändung
- Twisted skyscraper: Mitya Ivanov
- SPARCstation, Mike Chapman, PD

Funding Credit

The research described has been partially carried out as part of the CROSSMINER Project, which has received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 732223.

Evolution of the Unix System Architecture: An Exploratory Case Study

Diomidis Spinellis, Senior Member, IEEE, and Paris Avgeriou, Senior Member, IEEE

Abstract—Unix has evolved for almost five decades, shaping modern operating systems, key software technologies, and development practices. Studying the evolution of this remarkable system from an architectural perspective can provide insights on how to manage the growth of large, complex, and long-lived software systems. Along main Unix releases leading to the FreeBSD lineage we examine core architectural design decisions, the number of features, and code complexity, based on the analysis of source code, reference documentation, and related publications. We report that the growth in size has been uniform, with some notable outliers, while cyclomatic complexity has been religiously safeguarded. A large number of Unix-defining design decisions were implemented right from the very early beginning, with most of them still playing a major role. Unix continues to evolve from an architectural perspective. but the rate of architectural innovation has slowed down over the system's lifetime. Architectural technical debt has accrued in the forms of functionality duplication and unused facilities, but in terms of cyclomatic complexity it is systematically being paid back through what appears to be a self-correcting process. Some unsung architectural forces that shaped Unix are the emphasis on conventions over rigid enforcement, the drive for portability, a sophisticated ecosystem of other operating systems and development organizations and the emergence of a federated architecture, often through the adoption of third-party subsystems. These findings have led us to form an initial theory on the architecture evolution of large, complex operating system software.

Index Terms—Unix, Software Architecture, Software Evolution, Architecture Design Decisions, Operating Systems.

1 INTRODUCTION

T NIX¹ has a long and celebrated history. Its evolution spans five decades and is a result of the work by thousands of developers, including several distinguished work, and hardware engineering.

Studying the evolution of operating system software is accumulation of architectural technical debt. not just significant from a historical perspective; it can provide valuable insights into evolvability best practices and anti-patterns, for large, complex, and long-lived systems. Unix is a unique case among all operating systems, both due to its longevity, and its impact on the operating systems that followed. The evolution of a system of this size, complexity and age can shed light on how similar systems can sustainably grow without the perils of software aging like soaring technical debt or uncontrolled architectural decay.

In this paper we study the evolution of Unix along the FreeBSD lineage from a software architecture perspective. While there have been studies on how Unix evolved (see Section 2), these have mostly focused at the source code level and were limited to the kernel. On the contrary, we turn our attention to the system architecture and study a) the core architectural design decisions across the main more thin on the ground.

- D. Spinellis is with the Athens University of Economics and Business, Greece E-mail: see http://www.dmst.aueb.gr/dds/
- P. Avgeriou is with the University of Groningen
- Manuscript received December 19, 2016.

1. UNIX[®] is a registered trademark of The Open Group. For the sake of simplicity, in this paper we use the word "Unix" to refer both to UNIX systems developed at Bell Labs and to Unix-like systems, such as FreeBSD, that descended from them.

releases, and b) the evolution in the number of the system's features (obtained from the Unix reference documentation) and in the code's complexity. The former entails qualitapioneers. As an operating system, it has left an undeniable tive analysis, while the latter quantitative. These analyses mark on the history of computing, while it has influenced subsequently lead to forming an initial theory on the architremendously the current state of the art in software, net-tecture evolution of large and complex operating systems, regarding their form, pace, driving forces, as well as the

> The rest of the paper is structured as follows: In Section 2 we present related work, whereas in Section 3 we elaborate on the case study design. In sections 4 and 5, we present the qualitative results (main architectural design decisions), and the quantitative results (evolution of size and complexity) respectively. Next, in Section 6 we discuss the main findings, and in Section 7 the threats to this study's validity. Finally, in Section 8 we conclude the paper with a summary and discussion of our findings.

2 RELATED WORK

The work reported here covers mainly two areas: a) software evolution in general, which has been intensely studied, and b) the evolution of Unix in particular, where related work is

2.1 Software Evolution

There have been several studies on the longitudinal evolution of large systems. The seminal work of Lehman [1] and its subsequent refinements attempted to establish laws of software evolution, not unlike those of biological evolution. Those laws have been the subject of much discussion and research work [2]: their validity has been long debated, their



Diomidis Spinellis and Paris Avgeriou. Evolution of the Unix system architecture: An exploratory case study. IEEE Transactions on Software Engineering, 47:1134–1163, June 2021. doi:10.1109/TSE.2019.2892149

 $\begin{array}{l} \textbf{Empirical Software Engineering} \\ 10.1007/s10664\text{-}016\text{-}9445\text{-}5 \end{array}$

A Repository of Unix History and Evolution

Diomidis Spinellis

Abstract The history and evolution of the Unix operating system is made available as a revision management repository, covering the period from its inception in 1972 as a five thousand line kernel, to 2016 as a widely-used 27 million line system. The 1.1GB repository contains 496 thousand commits and 2,523 branch merges. The repository employs the commonly used Git version control system for its storage, and is hosted on the popular GitHub archive. It has been created by synthesizing with custom software 24 snapshots of systems developed at Bell Labs, the University of California at Berkeley, and the 386BBD team, two legacy repositories, and the modern repository of the open source FreeBBD system. In total, 973 individual contributors are identified, the early ones through primary research. The data set can be used for empirical research in software engineering, information systems, and software archaeology.

Keywords Software archeology \cdot Unix \cdot configuration management \cdot Git

1 Introduction

The Unix operating system stands out as a major engineering breakthrough due to its exemplary design, its numerous technical contributions, its impact, its development model, and its widespread use (Gehani 2003, pp. 27–29). The design of the Unix programming environment has been characterized as one offering unusual

This is a machine-readable rendering of a working paper draft that led to a publication. The publication should always be cited in preference to this draft using the reference in the previous footnote. The final publication is available at Springer via http://dx.doi.org/10.1007/s10664-016-9445-5.



Diomidis Spinellis. A repository of Unix History and evolution. *Empirical Software Engineering*, 22(3):1372–1404, 2017. doi:10.1007/s10664-016-9445-5

The work has been partially funded by the Research Centre of the Athens University of Economics and Business, under the Original Scientific Publications framework (project code EP-2279-01) and supported by computational time granted from the Greek Research & Technology Network (GRNET) in the National HPC facility — ARIS — under project ID PA003005-CDOLPOT.

D. Spinellis, Department of Management Science and Technology, Athens University of Economics and Business. E-mail: dds@aueb.gr

Diomidis Spinellis. A Repository of Unix History and Evolution. *Empirical Software Engineering*, 2017 (available online; to appear in print).

Documented Unix Facilities Over 48 Years

Diomidis Spinellis Department of Management Science and Technology Athens University of Economics and Business Athens, Greece dds@aueb.gr

ABSTRACT

The documented Unix facilities data set provides the details regarding the evolution of 15 596 unique facilities through 93 versions of Unix over a period of 48 years. It is based on the manual transcription of early scanned documents, on the curation of text obtained through optical character recognition, and on the automatic extraction of data from code available on the Unix History Repository. The data are categorized into user commands, system calls, C library functions, devices and special files, file formats and conventions, games et. al., miscellanea, system maintenance procedures and commands, and system kernel interfaces. A timeline view allows the visualization of the evolution across releases. The data can be used for empirical research regarding API evolution, system design, as well as technology adoption and trends.

CCS CONCEPTS

• Software and its engineering \rightarrow Software evolution;

ACM Reference Format:

Diomidis Spinellis. 2018. Documented Unix Facilities Over 48 Years. In MSR '18: 15th International Conference on Mining Software Repositories, May 28– 29, 2018, Gothenburg, Sweden, ACM, New York, NY, USA, 4 pages. https: //doi.org/10.1145/3196398.3196476

1 INTRODUCTION

The Unix operating system is being continuously developed from the same code base for almost over half a century. It stands out as a major engineering artefact due to its exemplary design, its numerous technical contributions, its impact, its development model, and its widespread use [2, pp. 27–29], [9]. The design of the Unix programming environment, which nowadays offers thousands of tools and libraries, has been characterized as offering unusual simplicity, power, and elegance [5, 7]. Consequently, empirical data regarding how the facilities Unix provides grew and changed over time can be used for empirical research on API evolution, system design, as well as technology adoption and trends.

Although one can study a system's evolution through its source code [1, 10], the very large size of modern systems can hinder the recognition of the relevant parts. Fortunately, another avenue is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. MKR '18, May 2–29, 2018, Contenburg, Sweden @ 2018 Copyright held by the owner?author(s). Publication rights licensed to the Association for Computing Machinery. ACM ISBN 978-1-4503-5716-6/18(05...515.00 https://doi.org/10.1145/310583.3196476 available for studying Unix systems, namely their documentation. From the first version of the Unix system until today, every release is accompanied by a complete reference manual, where all provided facilities (commands, APIS, file formats, and device drivers) are neatly organized into several corresponding sections (see Table 1). The central role of the reference manual in the Unix system is evidenced by the fact that early Unix versions coming out of AT&T Bell Labs were named after the edition of the accompanying manual. Some early editions of the manuals have not survived in a machinereadable format, but most are available in text markup that can be processed through scripts to extract relevant data.

The data set presented here is based on the printed and machinereadable Unix reference manuals released over a period of 48 years. It documents the evolution of 15 596 facilities through 93 versions of Unix. Section 2 outlines the provided data, Section 3 describes how the data were produced, and Section 4 sketches two examples of how the data can be used for quantitative and qualitative empirical studies.

2 UNIX RELEASES AND THEIR FACILITIES

The primary data are made available in the form of 93 text files containing 405726 records. The files are named after the associated Unix release, following the tags and branches nomenclature established in the Unix History Repository [9]. A record is a text line with tab-separated fields. Each record contains the number of the Unix manual section associated with a facility (1–9; see Table 1) followed by the facility's name, optionally followed by a URI identifying the facility's documentation in *troff* markup [6]. In total, the set contains data about 15 596 facilities pointing to 193 781 unique URIs, identifying 48 250 distinct manual page instances.

As an example, the following lines show the documentation files associated with the label command (:), the archiver (*ar*), and the assembler (*as*), as documented in Section I of the 1973 Third Edition Unix manual.

- : Research-V3/man/man1/:.1
- ar Research-V3/man/man1/ar.1
- as Research-V3/man/man1/as.1

By prepending the Unix History Repository GitHub permalink base URL "https://github.com/dspinellis/unix-history-repo/blob" to an entry's URI, one can obtain a URL for viewing the documentation markup source code for the corresponding entry.

A separate text file, named *timeline* associates each of the data files with the year, month, and day of the corresponding release. For instance, the following entries of the *timeline* file list the dates associated with the Sixth and Seventh Research Editions and the first Berkeley Software Distribution.

Research-V6 1975 07 18 BSD-1 1978 02 01 Research-V7 1979 08 26

Diomidis Spinellis. Documented Unix facilities over 48 years. In *MSR '18: Proceedings of the 15th Conference on Mining Software Repositories*, pages 58–61, New York, NY, USA, May 2018. Association for Computing Machinery. <u>doi:10.1145/3196398.3196476</u>

Data Sources

Tag	Data source(s)
Research-V1	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v1/svntree-20081216.tar.gz
Research-V3	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v3/nsys.tar.gz
Research-V4	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v4/v4man.tar.gz
Research-V5	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v5/v5root.tar.gz
Research-V6	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v6/v6root.tar.gz
	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v6/v6src.tar.gz
	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Dennis_v6/v6doc.tar.gz
BSD-1	http://www.tuhs.org/Archive/PDP-11/Distributions/ucb/1bsd.tar.gz
BSD-2	http://www.tuhs.org/Archive/PDP-11/Distributions/ucb/2bsd.tar.gz
Research-V7	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Henry_Spencer_v7/v7.tar.gz
	http://www.tuhs.org/Archive/PDP-11/Distributions/research/Henry_Spencer_v7/v7.patches.tar.gz
Bell-32V	http://www.tuhs.org/Archive/VAX/Distributions/32V/32v_usr.tar.gz
BSD-3	http://www.tuhs.org/Archive/4BSD/Distributions/3bsd.tar.gz
BSD-4	file://CSRG-CD-ROMs/cd1/4.0
BSD-4_1_snap	file://CSRG-CD-ROMs/cd1/4.1.snap
$BSD-4_1c_2$	file://CSRG-CD-ROMs/cd1/4.1c.2
BSD-4_2	file://CSRG-CD-ROMs/cd1/4.2
BSD-4_3	file://CSRG-CD-ROMs/cd1/4.3
BSD-4_3_Tahoe	file://CSRG-CD-ROMs/cd2/4.3tahoe
$BSD-4_3_Net_1$	file://CSRG-CD-ROMs/cd2/net.1
BSD-4_3_Reno	file://CSRG-CD-ROMs/cd2/4.3reno
$BSD-4_3_Net_2$	file://CSRG-CD-ROMs/cd2/net.2
$BSD-4_4$	file://CSRG-CD-ROMs/cd3/4.4
$BSD-4_4_Lite1$	file://CSRG-CD-ROMs/cd2/4.4BSD-Lite1
$BSD-4_4_Lite2$	file://CSRG-CD-ROMs/cd3/4.4BSD-Lite2
BSD-SCCS	file://CSRG-CD-ROMs/cd4
386BSD-0.0	http://www.oldlinux.org/Linux.old/distributions/386BSD/386bsd-0.0/floppies/3in/src/
386BSD-0.1	http://www.oldlinux.org/Linux.old/distributions/386BSD/0.1/386BSD/
FreeBSD-release/1.0	http://ftp-archive.freebsd.org/pub/FreeBSD-Archive/old-releases/i386/ISO-IMAGES/1.0/1.0-disc1.
	iso
$\rm FreeBSD$ -release/1.1	http://ftp-archive.freebsd.org/pub/FreeBSD-Archive/old-releases/i386/ISO-IMAGES/FreeBSD-1.
	1-RELEASE/cd1.iso
FreeBSD-release/1.1.5	http://ftp-archive.freebsd.org/pub/FreeBSD-Archive/old-releases/i386/ISO-IMAGES/FreeBSD-1.1.5.
	1/cd1.iso
${ m FreeBSD}$ -release/2	https://github.com/freebsd/freebsd