



The Triumph of the New Jersey Style: How Architectural Elegance, Components, and Reuse Shaped the Evolution of Unix

Diomidis Spinellis

Department of Management Science and Technology
Athens University of Economics and Business
&
Department of Software Technology
Delft University of Technology

Paris Avgeriou

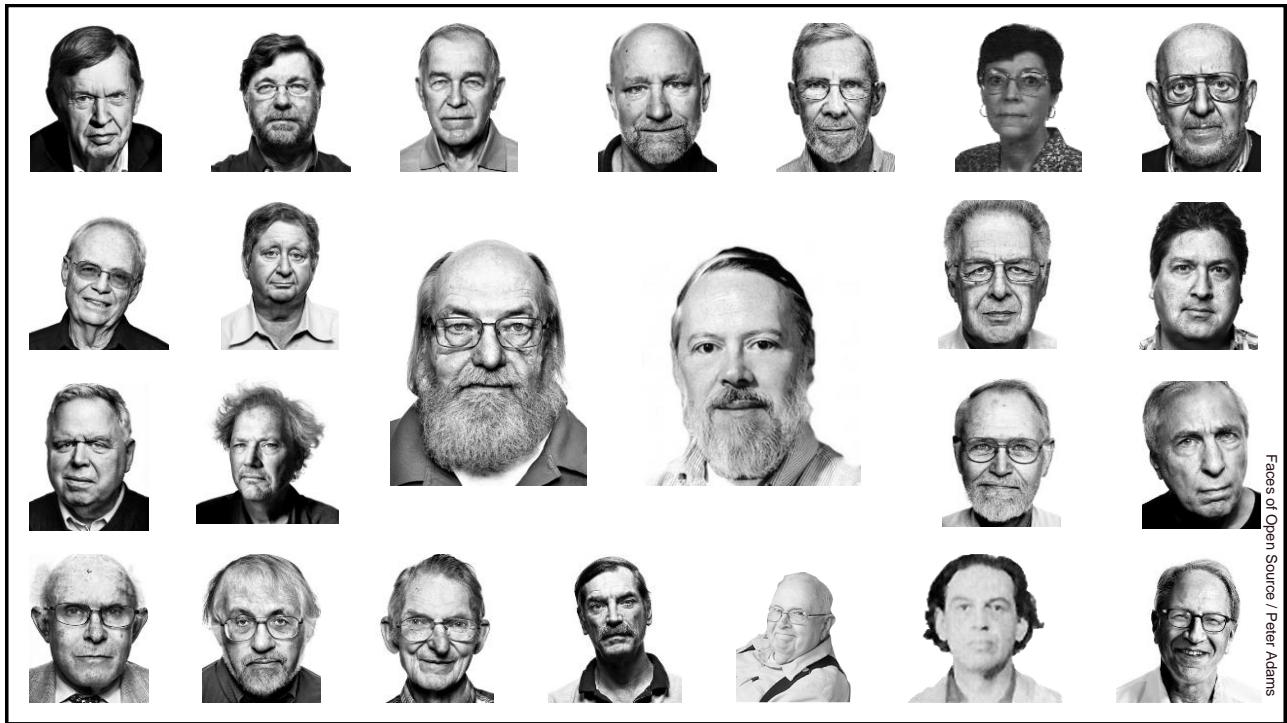
Institute for Mathematics and Computer Science
University of Groningen

www.spinellis.gr
@CoolSWEng

1

Unix 50th Anniversary 1969–2019

2



3



4



5



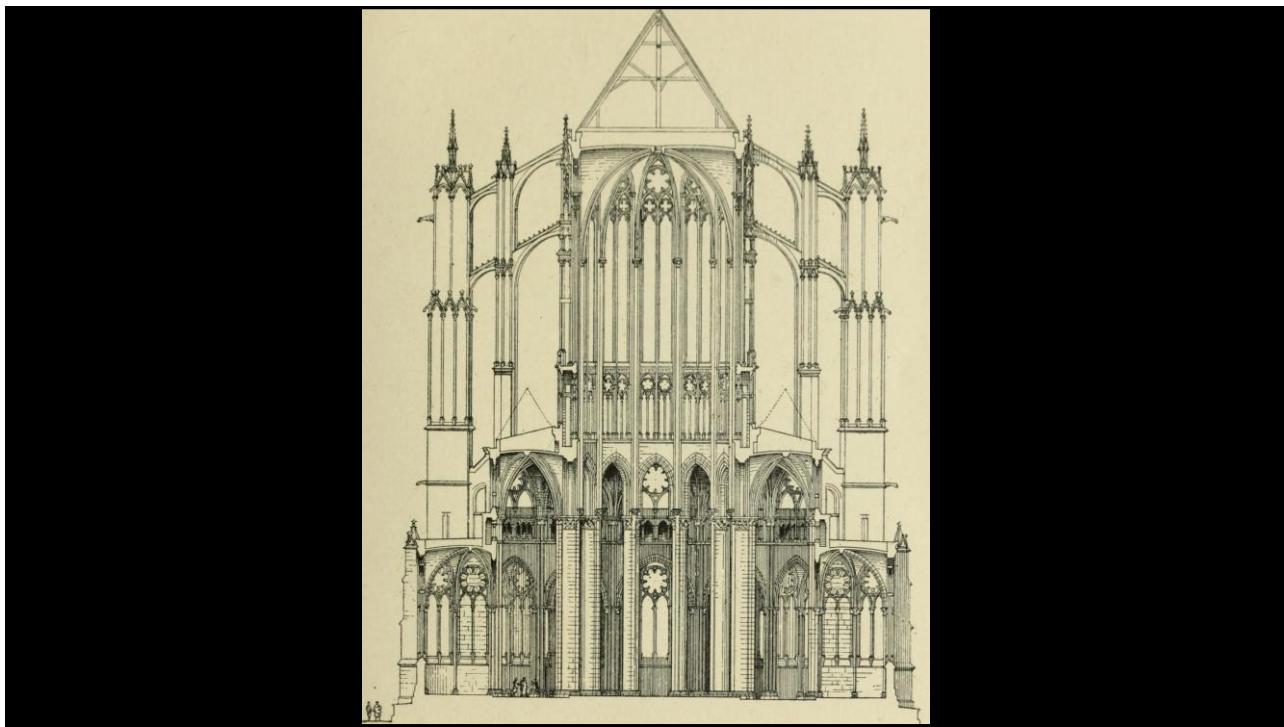
6



7



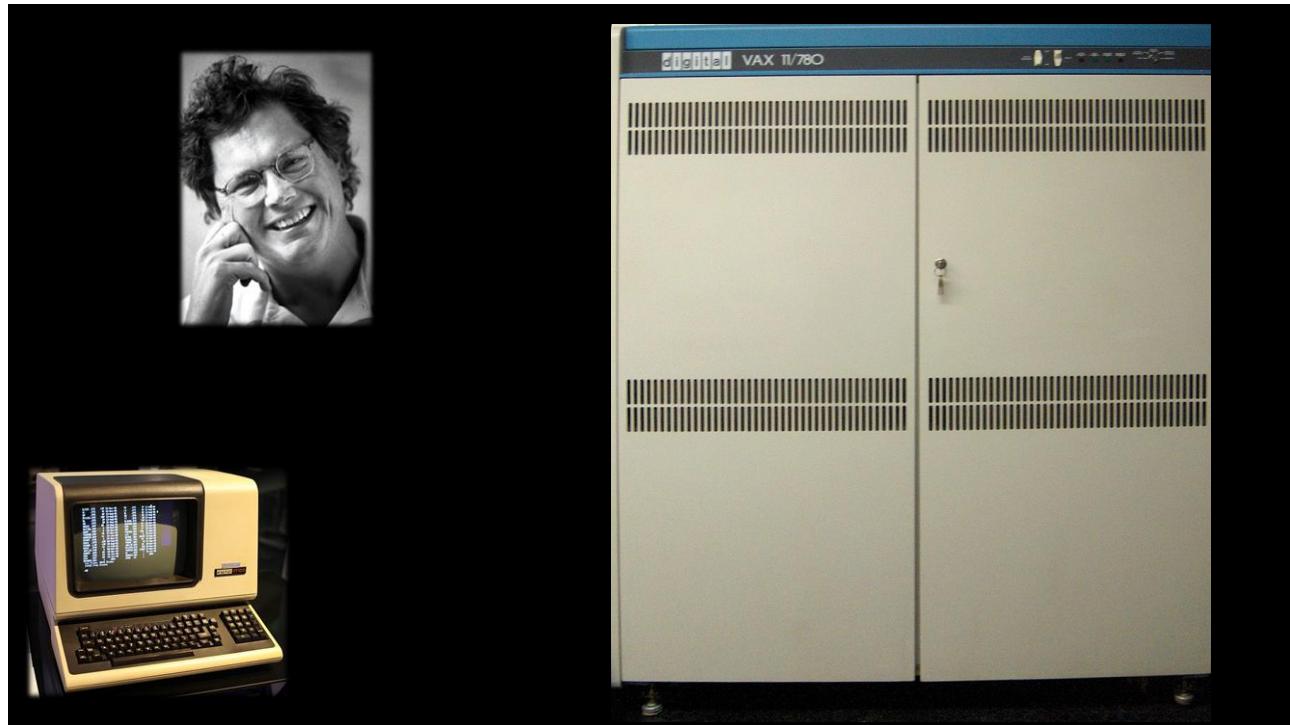
8



9



10



11



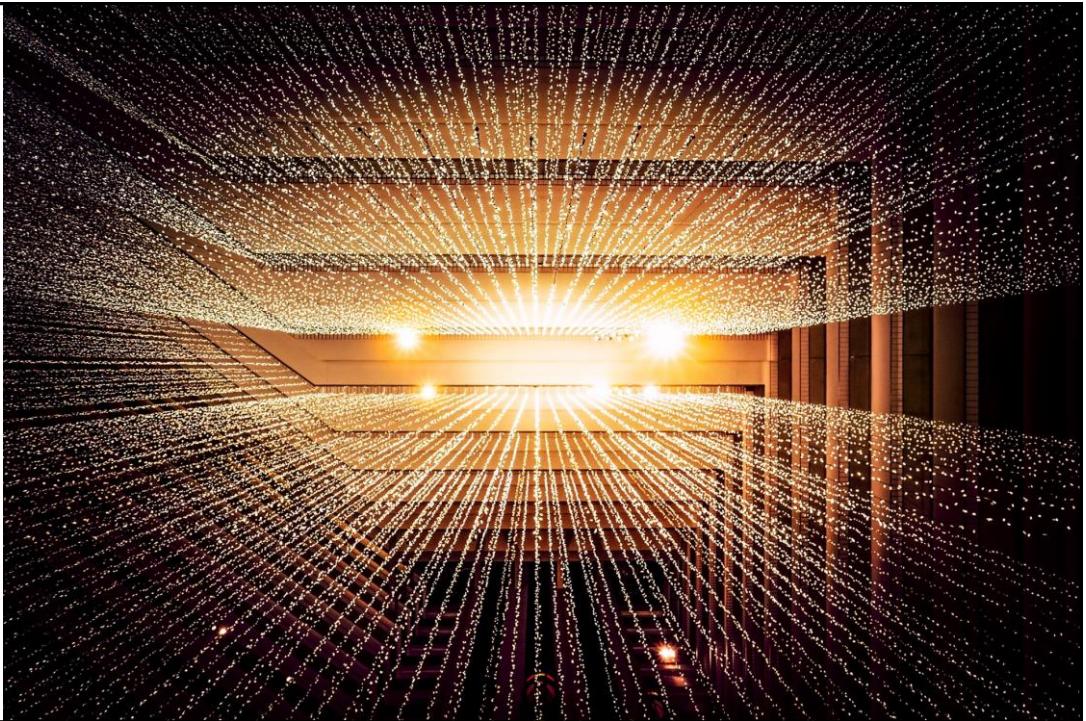
12

History



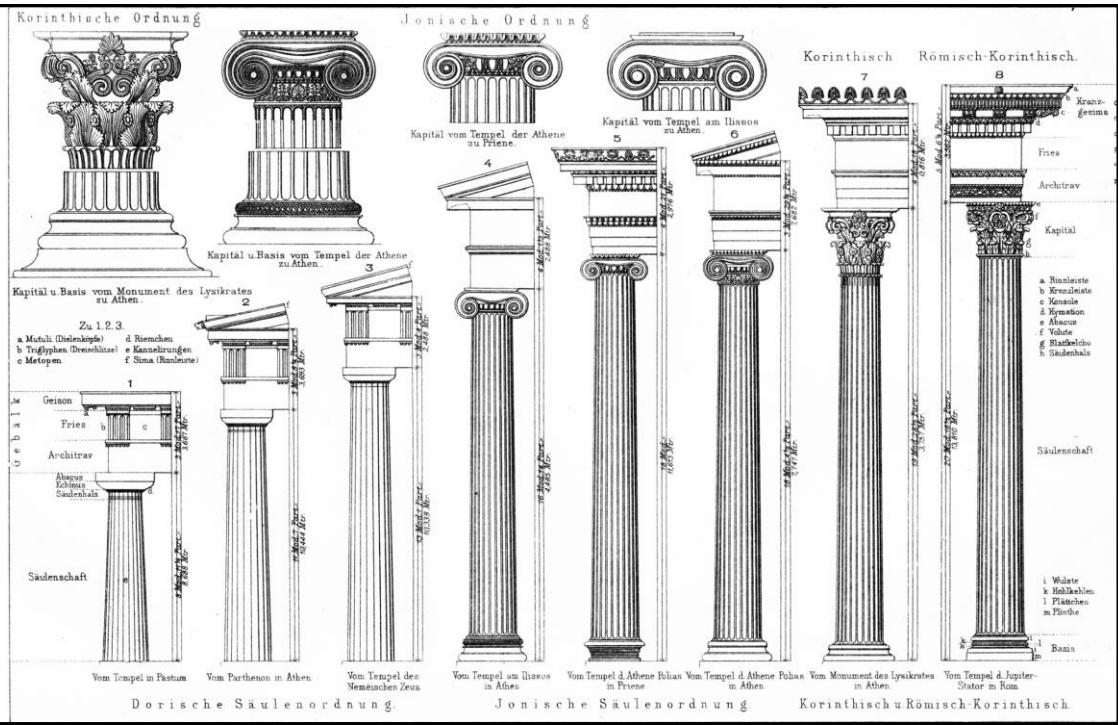
13

Data Sources



14

Architectural Design Decisions



15

Quantitative Results



16

Theory of OS Architectural Evolution



17

History



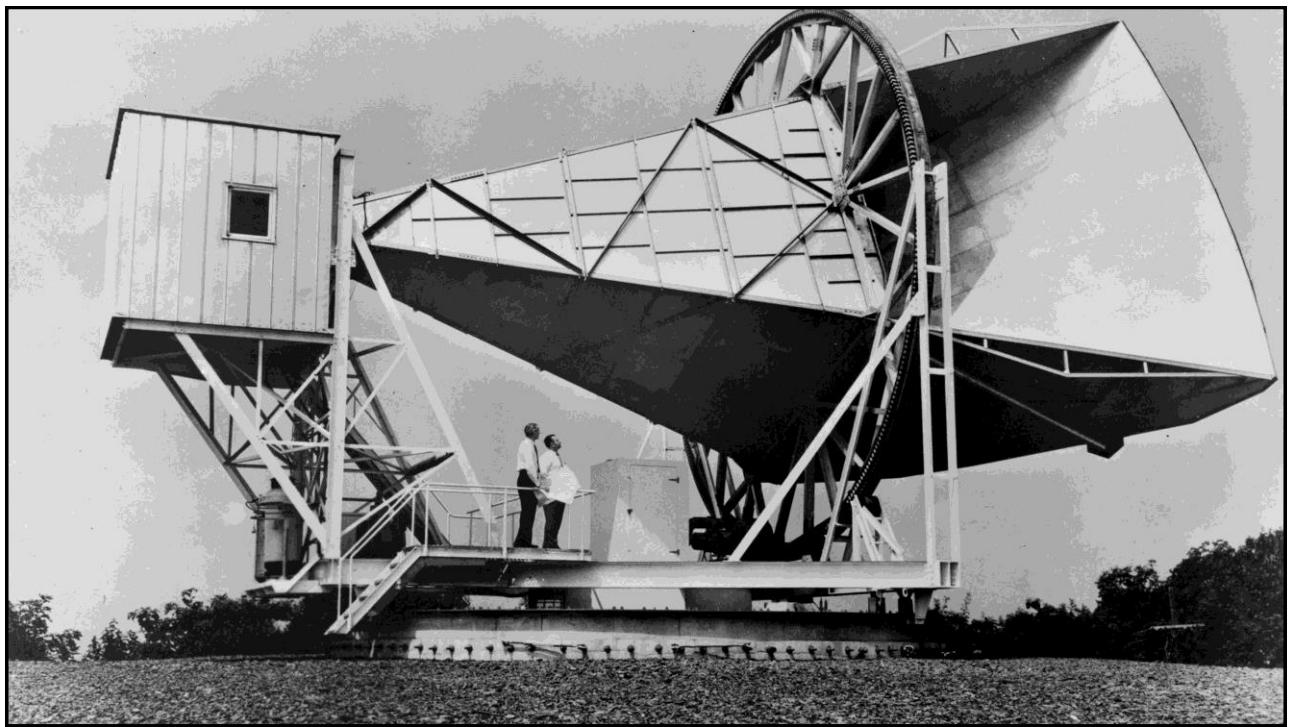
18



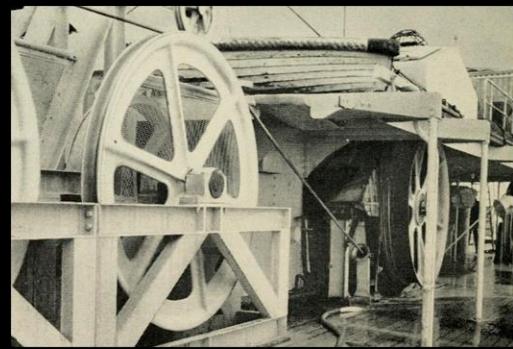
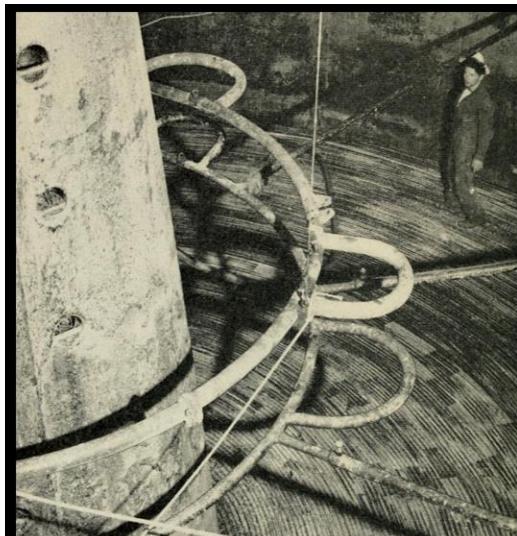
19



20



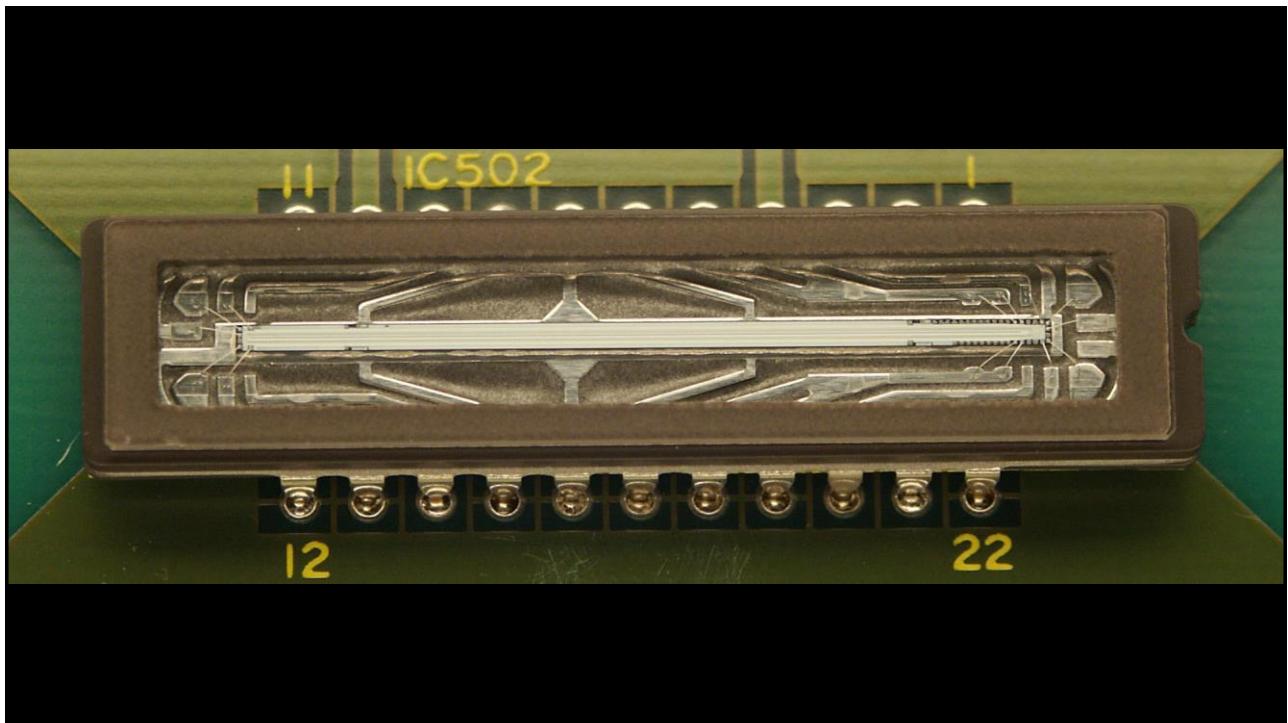
21



22



23



24

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 bandwidths 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 be concerned with 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 quantities such as the position of a Relays. 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 content of the set. 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:

1. 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 system which has N relays doubles the number of possible states of the relay. It adds 1 to the base 2 logarithm of this number. Doubling this number roughly squares the number of possible messages, or doubles the logarithm, etc.

2. It is closer to our intuitive feeling as to the proper measure. This is closely related to (1) since we intuitively measure 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.

3. 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.

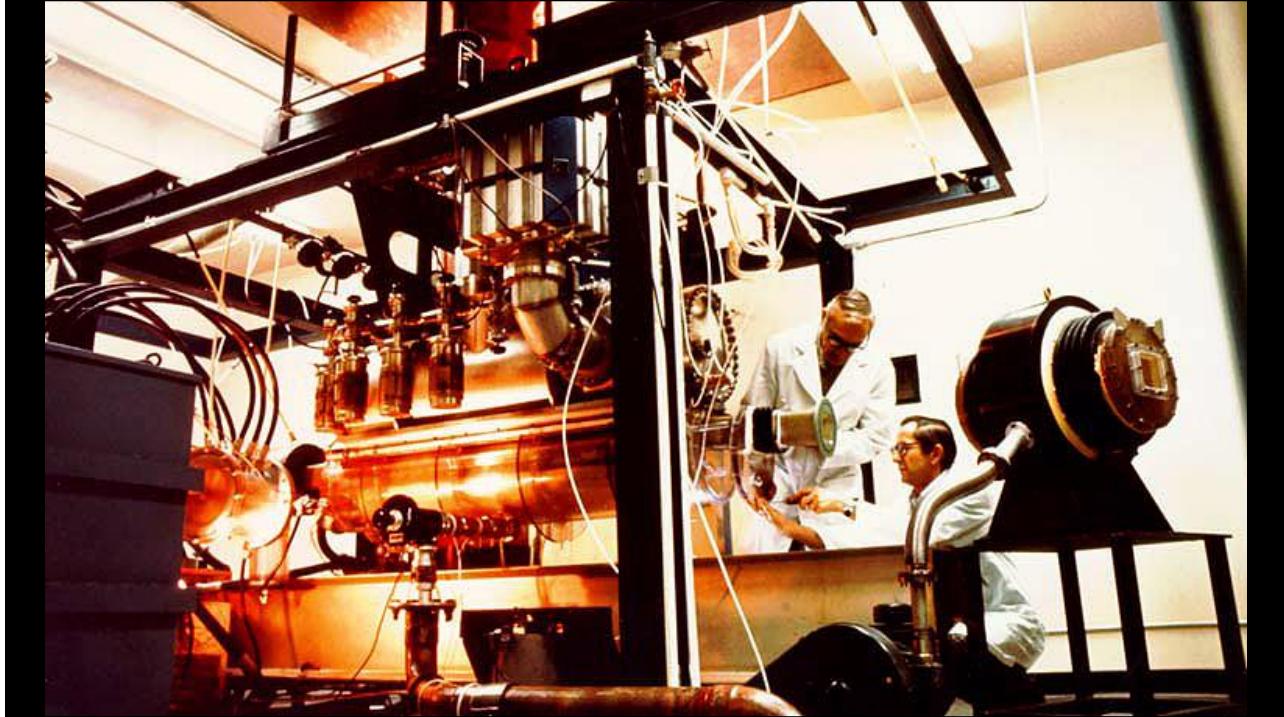
The choice of a logarithmic base corresponds to the choice of a unit for measuring information. If the base 2 is used the resulting units may be called binary digits, or more briefly *bites*, a word suggested by J. W. Tukey. A device with two stable positions, such as a relay or a flip-flop circuit, 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

$$\log_{10} M = \log_2 M / \log_2 10 \\ = 3.32 \log_2 M$$

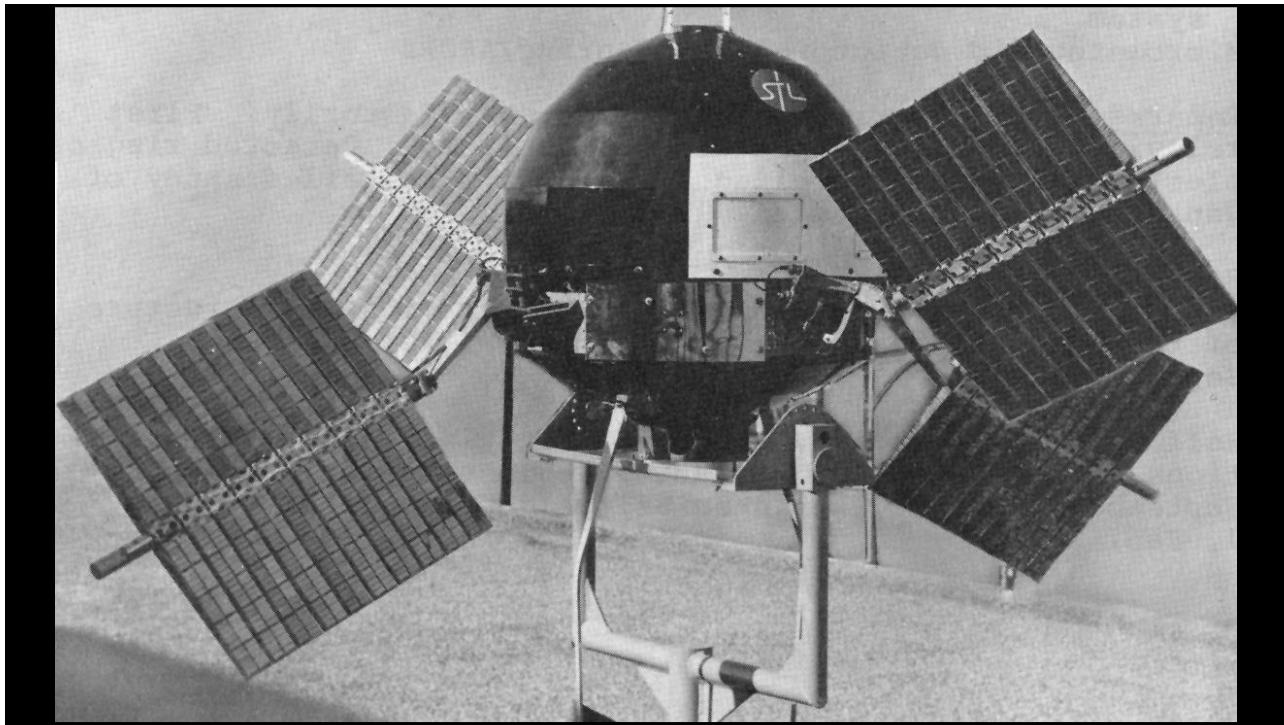
¹Nyquist, H., "Certain Factors Affecting Telegraph Speed," *Bell System Technical Journal*, April 1924, p. 324; "Certain Topics in Telegraph Transmission Theory," *A.I.E.E. Trans.*, v. 47, April 1928, p. 617.
²Hartley, R. V. L., "Transmission of Information," *Bell System Technical Journal*, July 1928,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

1

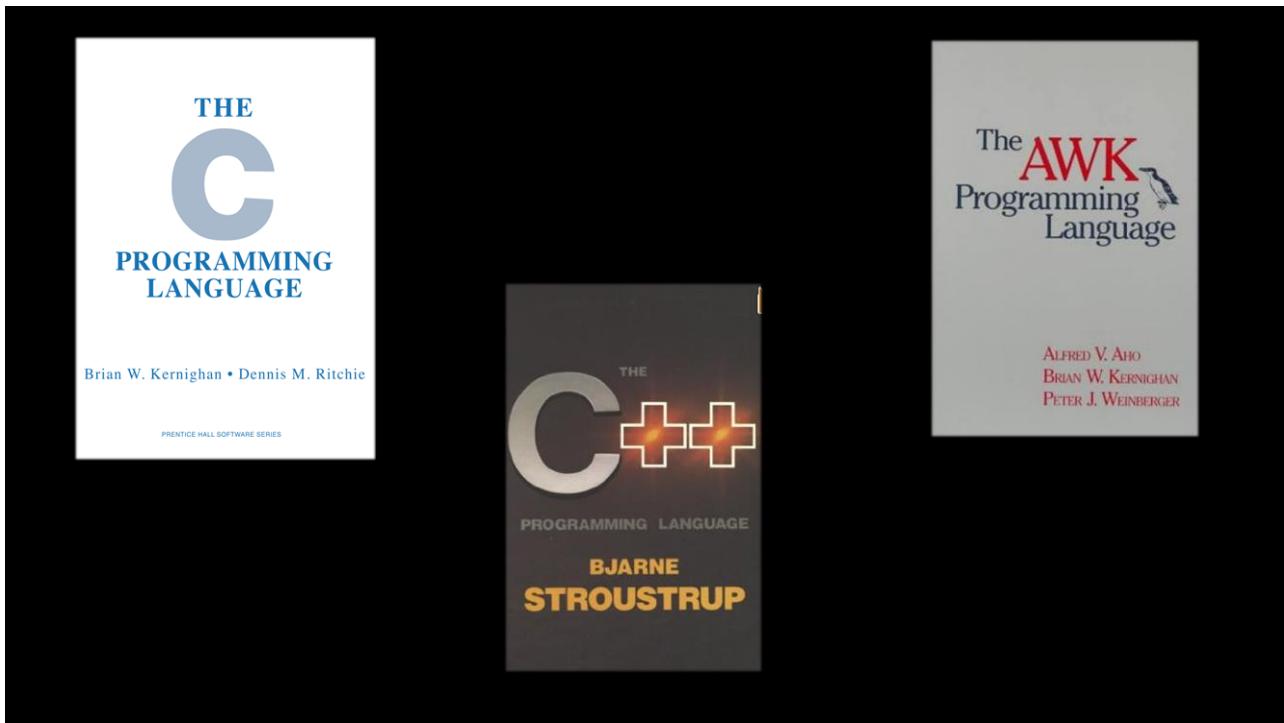
25



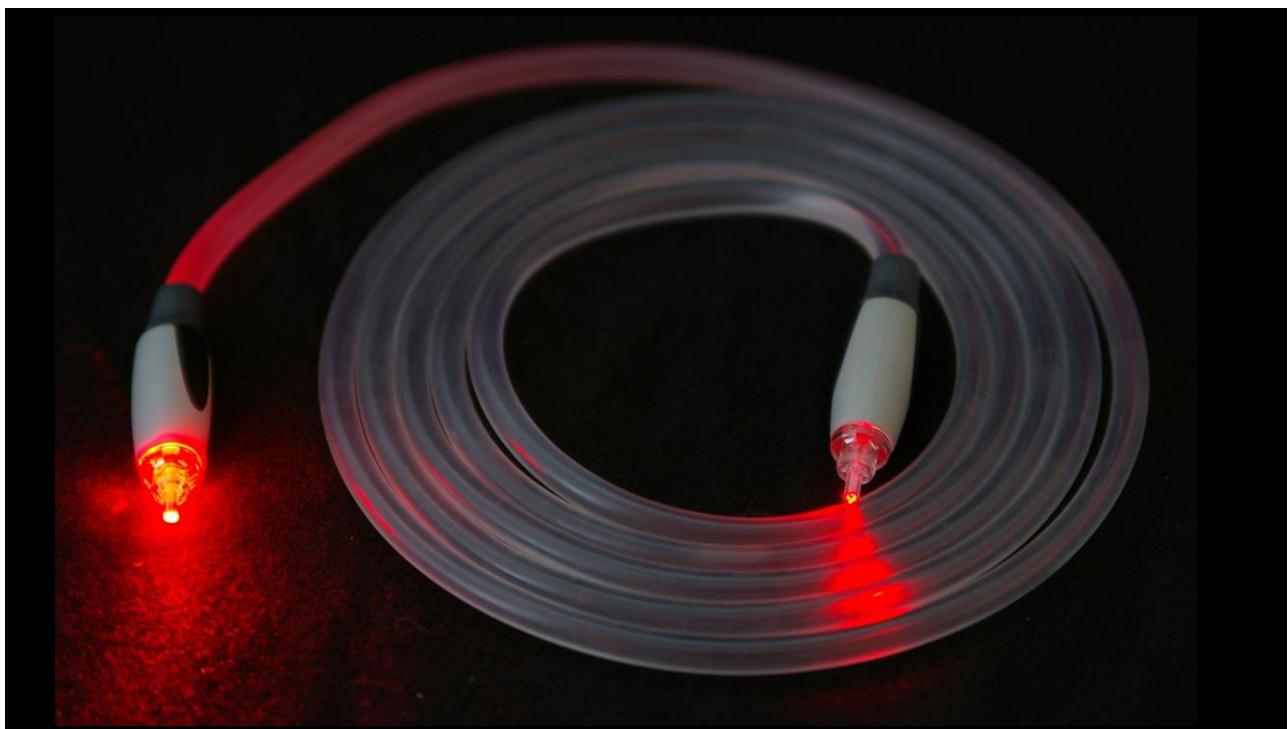
26



27



28



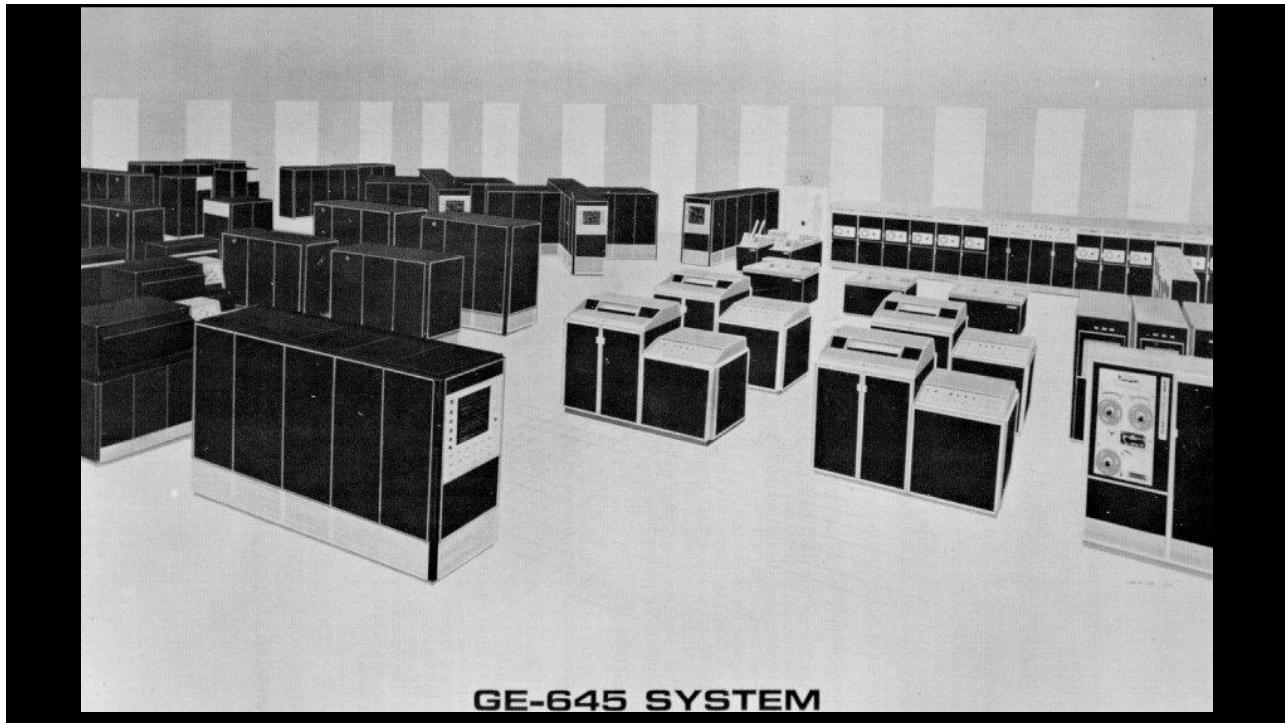
29



30



31



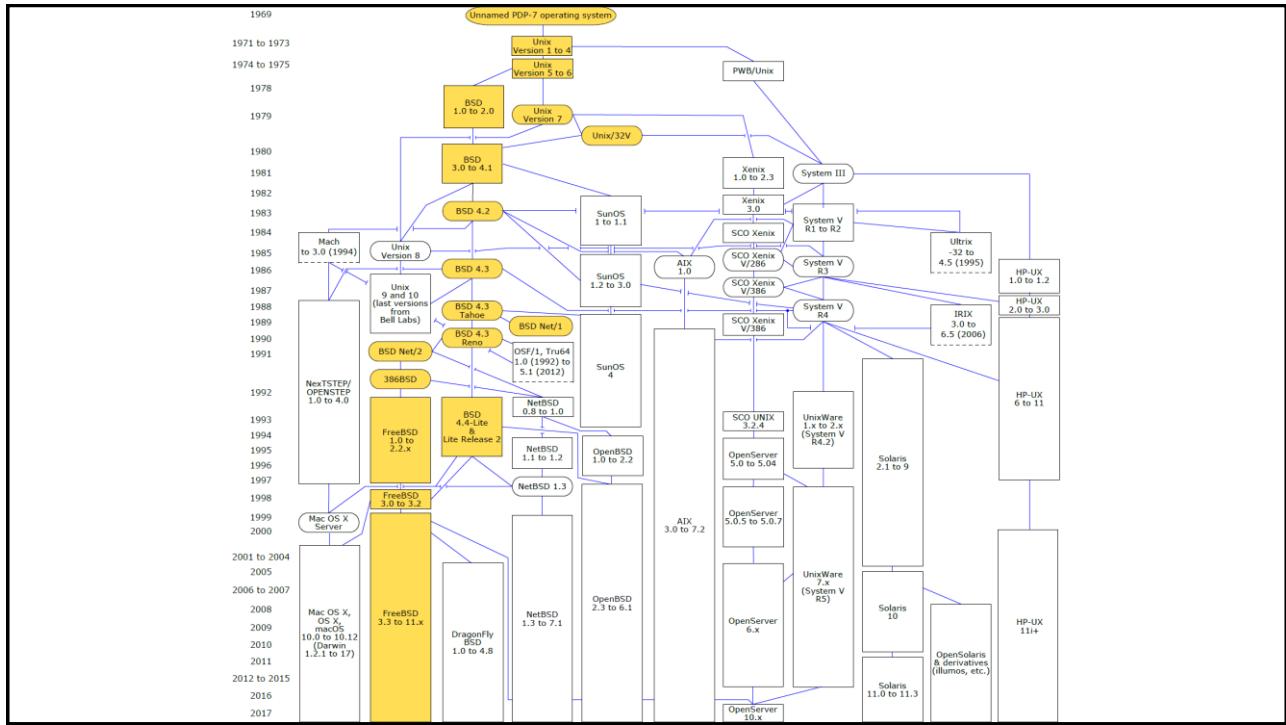
32



33

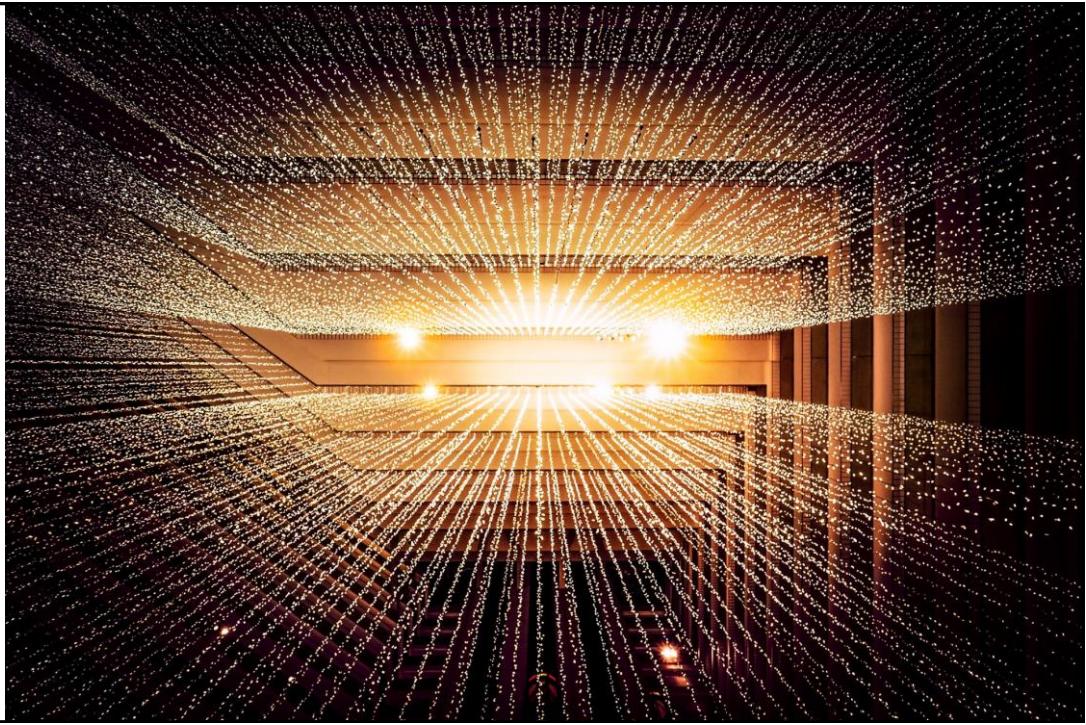


34



35

Data Sources



36



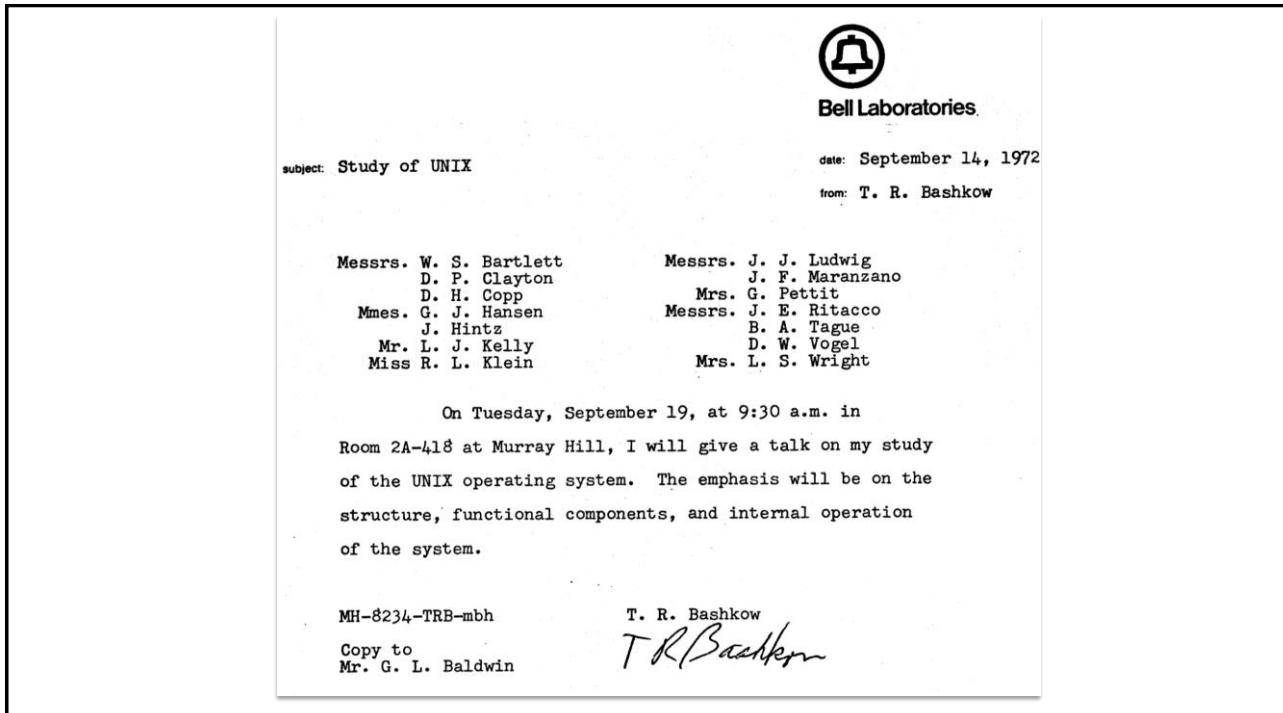
37



38



39



40

```

/ initialize inodes for special files (inodes 1 to 40.)
        mov    $40.,r1 / set r1=i-node-number 40.
1:      jsr    r0,iget / read i-node 'r1' from disk into inode area of
                  / core and write modified inode out (if any)
        mov    $100017,i.flgs / set flags in core image of inode to indi-
                  / cate allocated, read (owner, non-owner),
                  / write (owner, non-owner)
        movb   $1,i.nlks / set no. of links = 1
        movb   $1,i.uid / set user id of owner = 1
        jsr    r0,setwmod / set imod=1 to indicate i-node modified, also
                  / stuff time of modification into i-node
        dec    r1 / next i-node no. = present i-node no.-1
        bgt    1b / has i-node 1 been initialized; no, branch

/ initialize i-nodes r1,...,47. and write the root device, binary, etc..
/ directories onto fixed head disk. user temporary, initialization prog.

```

Issue D Date 3/17/72 ID IMO.1-1 Section E.O Page 4

Section E.0 Page 4

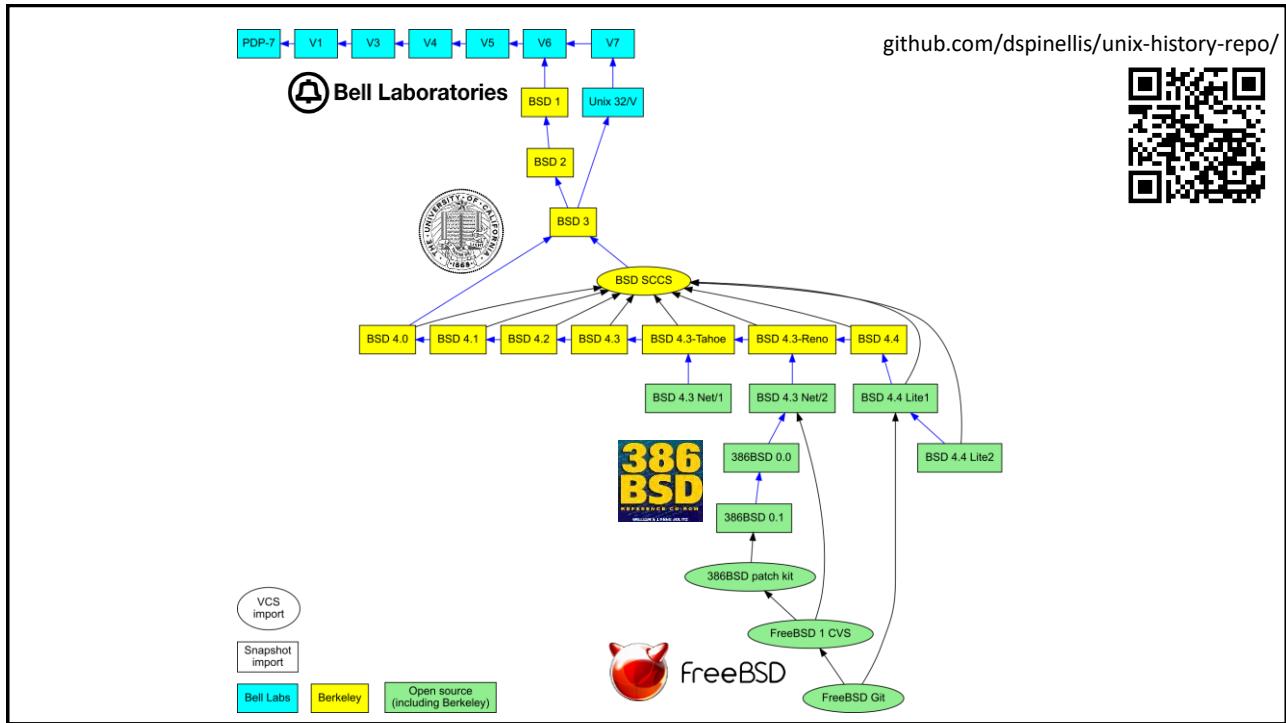
41

UNIX IMPLEMENTATION

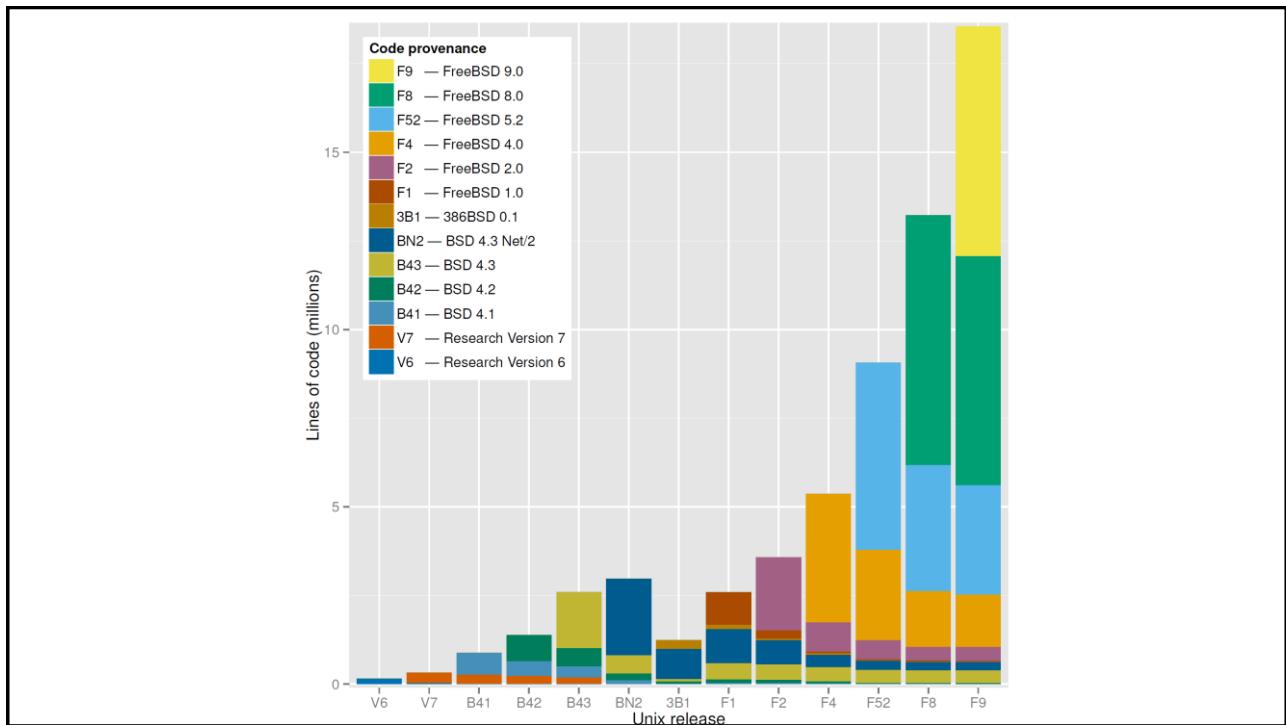
UNIX IMPLEMENTATION

Issue D Date 3/17/72 ID IMO.1-1 Section E.0

Issue D Date 3/17/72 ID IMO.1-1 Section E.0 Page 5



43

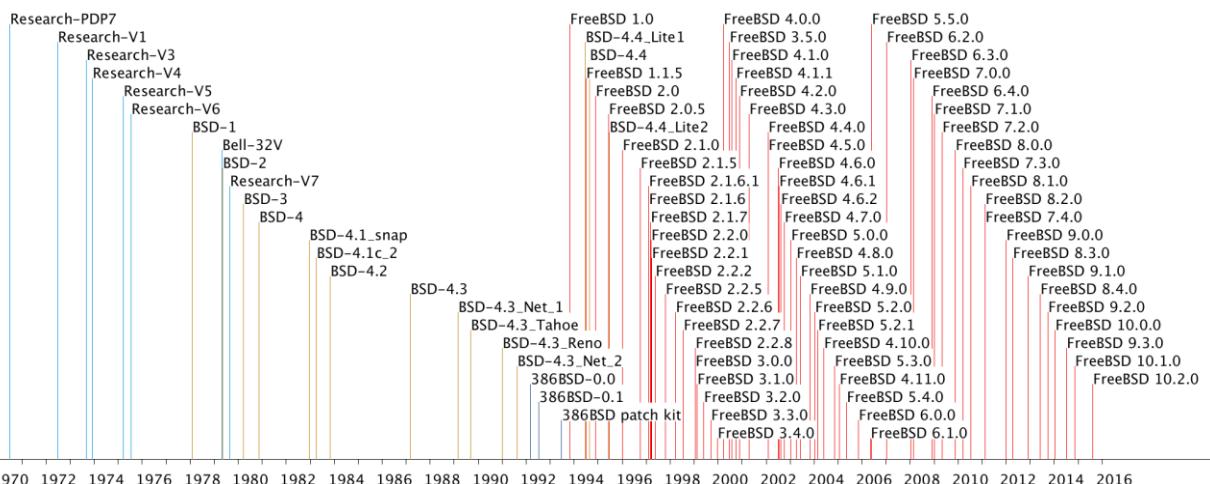


44

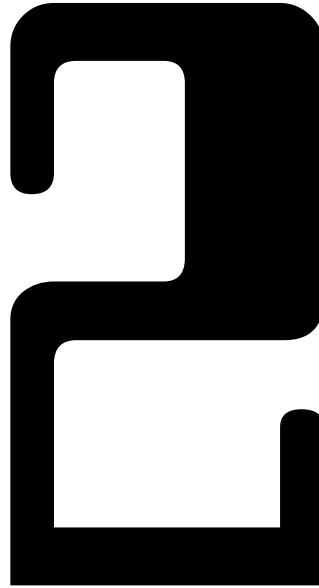
```
$ git checkout FreeBSD-release/10.0.0
$ git blame -M -M -C ./lib/libc/gen/timezone.c

usr/src/libc/gen/timezone.c  (Dennis Ritchie 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)     int offset;
usr/src/libc/gen/timezone.c  (Dennis Ritchie 1979-01-10 14:58:45 -0500 78)     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[] = {
lib/libc/gen/timezone.c    (Jordan K. Hubbard 1996-07-12 18:57:58 +0000 81)     {-1*60, "MET", "MET DST"}, ...
[...]
lib/libc/gen/timezone.c    (Jordan K. Hubbard 1996-07-12 18:57:58 +0000 96)     {-1}
usr/src/libc/gen/timezone.c (Bill Joy 1980-12-22 00:40:25 -0800 97) };
usr/src/libc/gen/timezone.c (Bill Joy 1980-12-22 00:40:25 -0800 98)
usr/src/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 2002-02-01 01:08:48 +0000 110)     char sign;
usr/src/libc/gen/timezone.c (Bill Joy 1980-12-22 00:40:25 -0800 111)
usr/src/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/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 113)
usr/src/libc/gen/timezone.c (Dennis Ritchie 1979-01-10 14:58:45 -0500 114)
usr/src/libc/gen/timezone.c (Dennis Ritchie 1979-01-10 14:58:45 -0500 115)
usr/src/libc/gen/timezone.c (Dennis Ritchie 1979-01-10 14:58:45 -0500 116)
usr/src/libc/gen/timezone.c (Dennis Ritchie 1979-01-10 14:58:45 -0500 117)
usr/src/libc/gen/timezone.c (Dennis Ritchie 1979-01-10 14:58:45 -0500 118)
usr/src/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 119)
usr/src/libc/gen/timezone.c (Keith Bostic 1987-03-28 19:27:07 -0800 120)     if (zone < 0) {
                                         zone = -zone;
                                         sign = '+';
                                         }
                                         else
                                         sign = '-';
                                         (void)sprintf(czone,
                                         "GMT%cd:%02d", sign, zone /
                                         sizeof(czone),
                                         lib/libc/gen/timezone.c (Warner Losh 1998-01-21 21:46:36 +0000 127)
                                         "GMT%cd:%02d", sign, zone /
                                         60, zone % 60);
                                         lib/libc/gen/timezone.c (Rodney Grimes 1994-05-27 05:00:24 +0000 128)
                                         lib/libc/gen/timezone.c (Rodney Grimes 1994-05-27 05:00:24 +0000 129) }
```

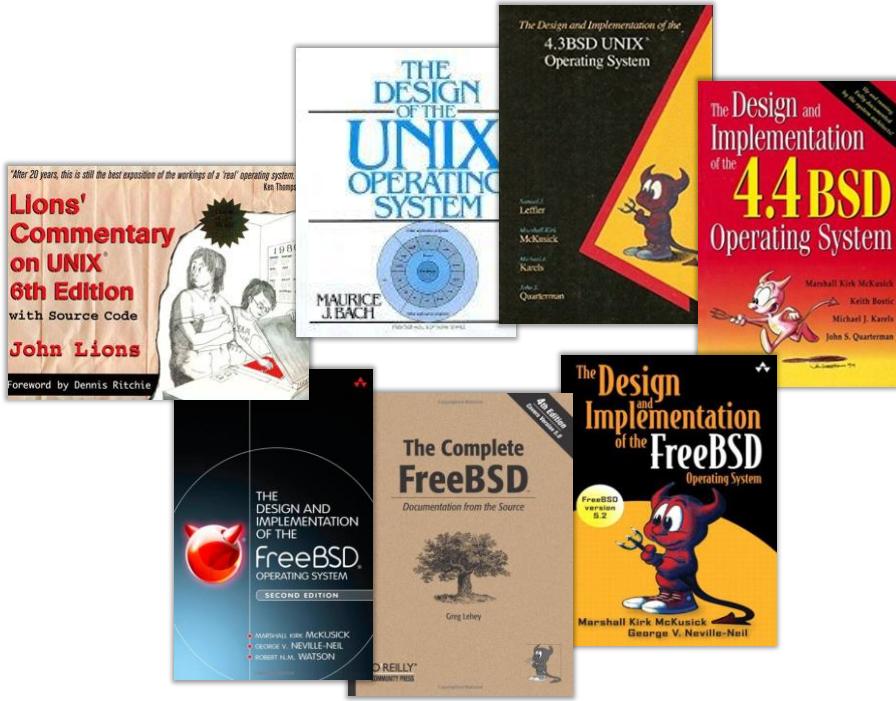
45



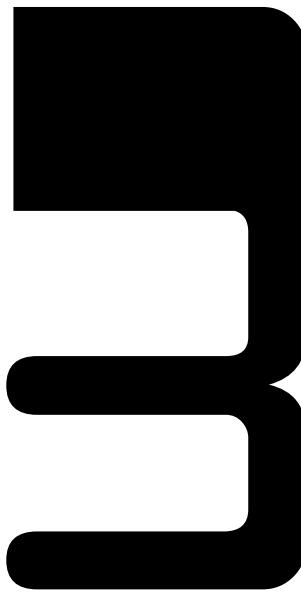
46



47



48



49

UNIX PROGRAMMER'S MANUAL
Third Edition

K. Thompson
D. M. Ritchie

February, 1973

Copyright © 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.

50



51

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

52

Evolution of Unix section 2: System calls

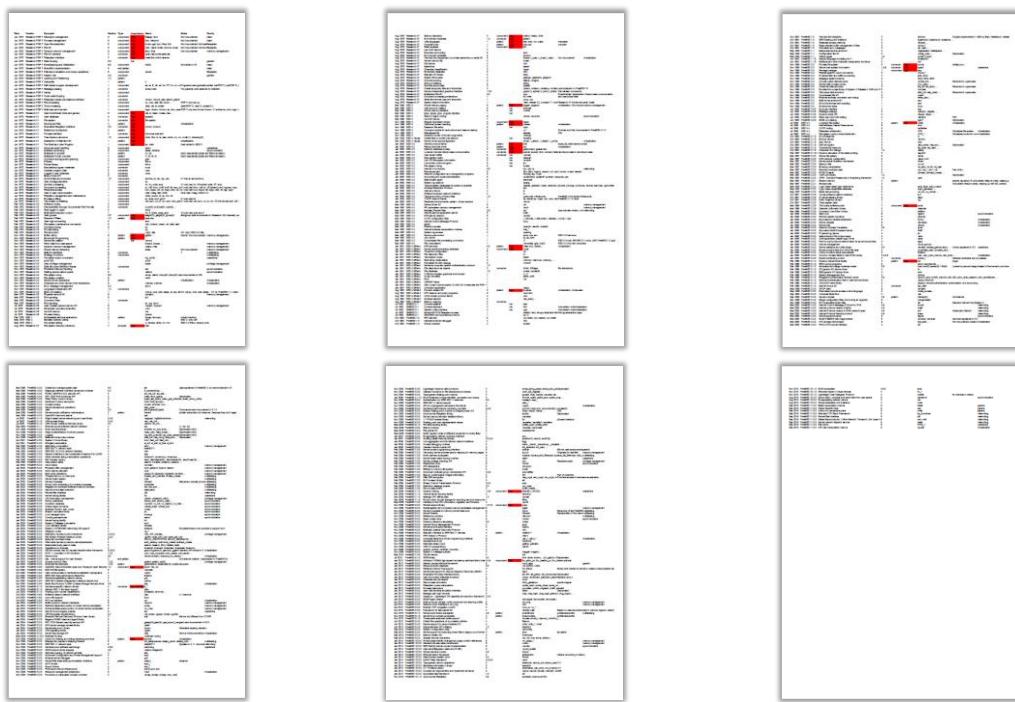
Facility	Appearance	Research V1	Research V2	Research V3	Research V4	Research V5	Research V6	BSD 1	BSD 2	Bell 32V	Research V7	BSD 3
time	Research V1											
umount	Research V1											
unlink	Research V1											
wait	Research V1											
write	Research V1											
chd	Research V2											
hog	Research V2											
kill	Research V2											
mkdir	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											
netmif	Research V4											

[Back to section index](#)

Disclaimers

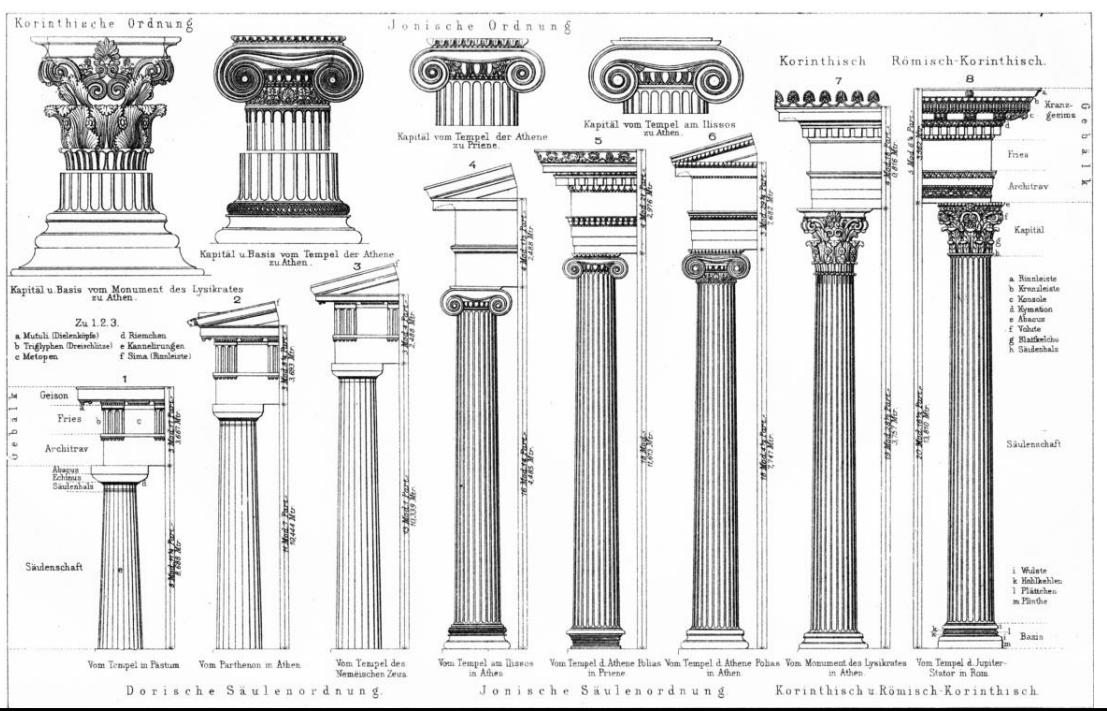
- The name of a facility may have been repurposed over time.
- Facilities in sections 1, 6, 8 moved across sections over time. To allow a continuous view of their evolution, all have been relocated to the section of the most recent FreeBSD release, if they still existed at the time.
- The evolution data of collapsed tree nodes depict the evolution of the tree's first child node.

53



54

Architectural Design Decisions



55



56

PDP-7 [Unix] (1970)

```

between: 0
    lme oma
    lac between i
    dac 9ft+
    lsz between
    lacq
    tad 9ft+ i
    sna
    dmp lf
    lac between i
    dac 9ft+
    lsz between
    lacq
    tad 9ft+ i
    sna
    spa sna
1:   isz between
    lacq
    cma
    dmp between i
copy: 0
    =1
    tad copy i
    dac 8
    lsz copy

```

dac 9ft+
 xct between i
 fad 9ft+ isz between
 spa
 jump lf
 xctbetween i
 ist between
 fad 9ft+
 spa sna
 isz between
 jump between i

2584 lines

57

... 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

58

Layering and Partitioning

```

adm.s  cat.s    dskio.s   init.s    s6.s
ald.s  check.s  dskres.s  lcase.b   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

```

59

Process Management (fork)

```

.fork!
jms lockfor 0 "not-used
    skp
    jms_error
    dae 9f*t
    lsz uniqpid
    iac uniqpid
    dae u,ac
    iav sysekit
    dae u,swapget
    iac 0200000
    tdd u,uclistp i
    dae u,uclistp i
    jms dsksvapl 07000
    iac 9f*t
    dae u,uclistp
    iac 0100000
    xor u,uclistp i
    dae u,uclistp i
    iac u,pid

```

60

Descriptor Management

```

fget: 0
    jms between d01 d9
    jmp fget i
    cli; muls 9
    ladd

    tad ofilesp
    das 9f+t
    das .+2
    jms copyi .,. fnodei 3
    lsz fget
    jms fget i

    fput: 0
    lac 9f+t
    das .+3
    jms copyi fnodei .,. 3
    jmp fput i
    t = t+1

```

61

Separation of File Metadata from File Naming

inodei i.flagst .E.+1 i.nlinki .E.+7 i.uidi .E.+1 i.sizei .E.+1 i.uniqi .E.+1 , = inode+12	namei: 0 jms iget -1 tad namei i das 9f+t+1 lsz namei lac i.flags and o20 sna jmp namei i -8 tad i.size cma irss 3 das 9f+t sna jmp namei i dzm di
---	--

62

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  
sh:  
    <sh>; 040040;040040;040040  
system:  
    <sy>;<st>;<em>; 040040  
-----
```

63

File I/O API

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

64

File System API

- creat
- rename
- link
- unlink

65

Interpreter

```

main $(
    extrn read, write;
    auto i, c, state, line 100;
)
loop:
    state = i = 0;
loop1:
    c = read();
    if(c==0) 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(' ');
noi:
    i = 0;
loop3:
    c = line[i];
    write(c);
    i = i+1;
    if(c!=012) goto loop3;
    goto loop;
)

```

End.6

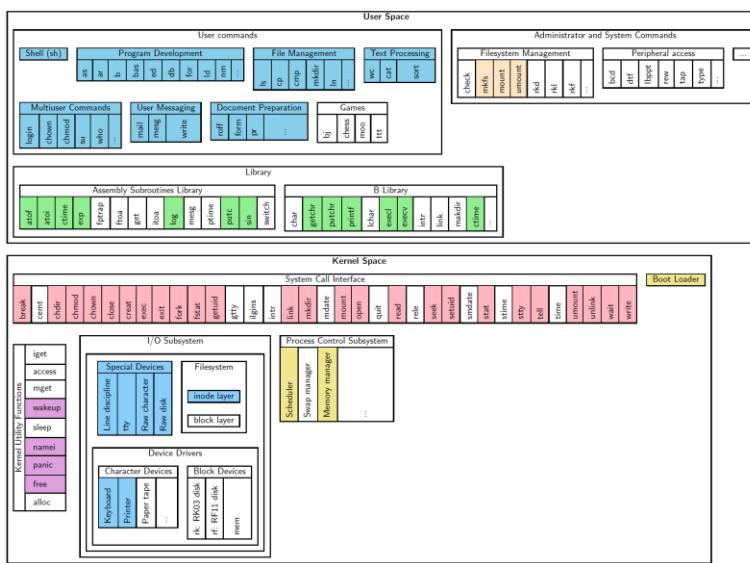
/x ind * /

66



67

First Research Edition (Nov 1971)



68

First Edition Unix 1972 FreeBSD 11.1 2018

sysrele / 0	0 { int nosys(void); } syscall nosys_args int
sysexit / 1	1 { void sys_exit(int rval); } exit \
sysfork / 2	sys_exit_args void
sysread / 3	2 { int fork(void); }
syswrite / 4	3 { ssize_t read(int fd, void *buf, \
sysopen / 5	size_t nbytes); }
sysclose / 6	4 { ssize_t write(int fd, const void *buf, \
syswait / 7	size_t nbytes); }
syscreat / 8	5 { int open(char *path, int flags, int mode); }
syslink / 9	6 { int close(int fd); }
sysunlink / 10	7 { int wait4(int pid, int *status, \
	int options, struct rusage *rusage); }
	8 { int creat(char *path, int mode); }
	9 { int link(char *path, char *link); }
	10 { int unlink(char *path); }

Issue D Date 3/17/72

ID IMO.1-1

Section E.1

Page 1

69

The Shell as a User Program

11/3/71

PASSWD (V)

NAME passwd -- password file
SYNOPSIS --
DESCRIPTION passwd contains for each user the following information:
 name (login name)
 password
 numerical user ID
 default working directory
program to use as Shell

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.

70

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 "**K**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.

71

Interoperability through Documented File Formats

V. FILE FORMATS

a.out	assembler and loader output
archive	archive file
bppt	binary paper tape format
core	core image file
directory	directory format
file system	file system format
passwd	password file
uids	map names to user ID's
utmp	logged-in user information

72

Format	Description	Clients
a.out	Assembler and linker output	as, ld, strip, nm, un
Archive	Object code libraries	ar, ld
Core	Crashed program image	Kernel, db
Directory	File system directories	du, find, ls, ln, mkdir, rmdir
File system	File system format	check, dump,* mkfs, restor*
Ident	GECOD ident card format	opr
Password	User accounts and passwords	chown, find, getpw,* login,* ls, passwd*
Tape*	DECtape file format	mt,* tap*
Uid	User identifier to name map	chown
utmp	Logged in users	init, login,* who,* write*
wtmp*	Users login history	acct, date, init, login, tacct, who

73

User-Contributed Tools and Games

VI. USER MAINTAINED PROGRAMS

```

basic ..... DEC supplied BASIC
bj ..... the game of black jack
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 ..... sort a file
ttt ..... the game of tic-tac-toe

```

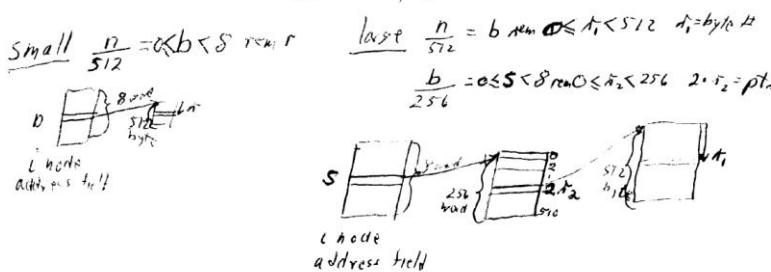
74

Tree Directory Structure

- mkdir(l)
- chdir(l)

--

$f_r/e = 8 \times 256 \times 512 \times 2^3 \times 2^3 \times 2^7 = 2^{26}$ bytes
but actually limited by offset in ftable = 2^{16} bytes



- chdir(l)
- find(l)
- ln(l)
- ls(l)
- stat(l)
- mkdir(l)
- mv(l)
- rm(l)
- rmdir(l)

75

Mountable File System Interface

- mount(l)
- umount(l)

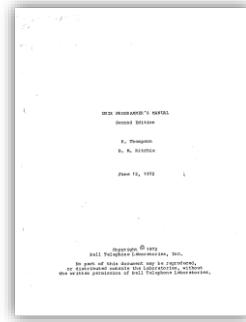
- mount(l)
- umount(l)

76

Second Research Edition (Jun 1972)

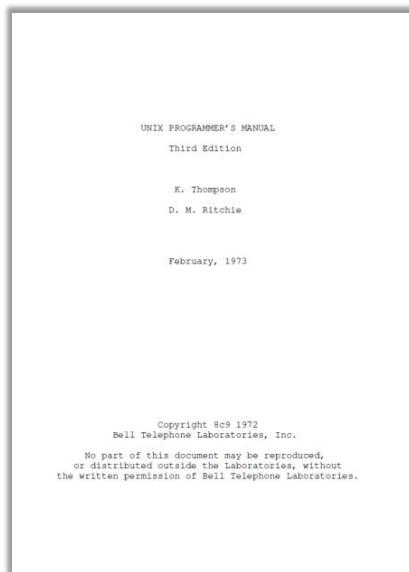
- Software Library

III. SUBROUTINES	
atan	arctangent
atof	convert ASCII to floating
atoi	convert ASCII to integer
const	floating-point constants
ctime	convert time to ASCII
exp	exponential function
ftrap	floating-point simulator
ftoa	convert floating to ASCII
gerts	communicate with GCOS
getc	get character
hypot	compute hypotenuse
itoa	convert integer to ASCII
log	logarithm base e
mesg	print string on typewriter
nlist	read name list
ptime	print time
putc	write character or word
qsort	quicker sort
salloc	storage allocator
sin	sine, cosine
sqrt	square root
switch	transfer depending on value



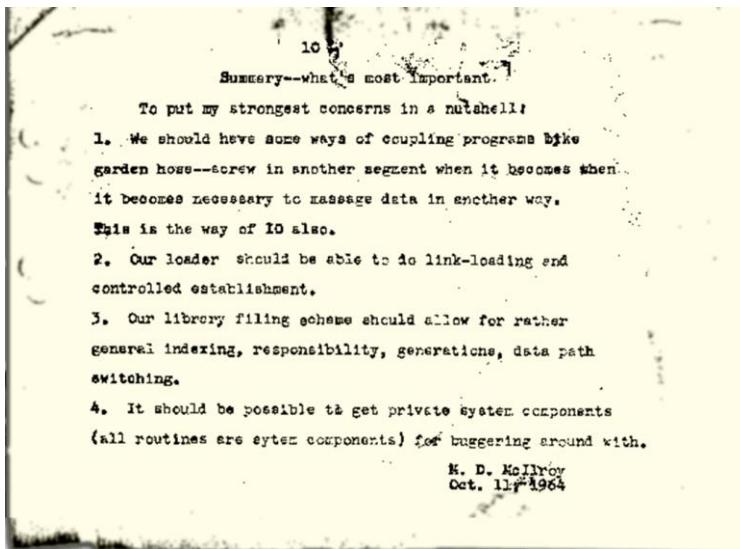
77

Third Research Edition (Feb 1973)



78

Pipes and Filters



79

Fourth Research Edition (Nov 1973)



80

Structured Programming

- Kernel implemented in “New B”
- 6373 lines New B
- 768 lines PDP-11 assembly
- 105 functions + 50 assembly symbols
- vs 248 global symbols in the First Edition

```

main()
{
    extern uchar;
    extern char end[], data[], etext[];
    int i, ii, *p;

    /* zero and free all of core
     */
    UIISA->r[0] = KISA->r[6] + USIZE;
    UIISO->r[0] = 6;
    for(; fubyte(0) >= 0; UIISA->r[0]++) {
        clearseg(UIISA->r[0]);
        mfree(coremap, 1, UIISA->r[0]);
    }
    mfree(swapmap, NSMAP, 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_proc = &proc[0];

    /* set up 'known' i-nodes
     */
    sureg();
    LKS->i_integ = 0115;
    cinit();
    binit();
    iinit();
    rootdir = igrand(ROOTDEV, ROOTINO);
    rootdir->i_flag = ~ILOCK;
    u.u_cdir = igrand(ROOTDEV, ROOTINO);
    u.u_cdir->i_flag = ~ILOCK;
}

```

81

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.

82

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
```

83

Dynamic Resource Management

```
int      coremap[CMAPSIZ];  
int      swapmap[SMAPSIZ];  
  
struct map {  
    char *m_size;  
    char *m_addr;  
};  
  
malloc(mp, size)  
struct map *mp;  
{  
    ...  
}
```

84

Device Driver Abstraction

IV. SPECIAL FILES

85

Driver Interface

```
struct {
    int      (*d_open)();
    int      (*d_close)();
    int      (*d_strategy)();
} bdevsw[];

struct {
    int      (*d_open)();
    int      (*d_close)();
    int      (*d_read)();
    int      (*d_write)();
    int      (*d_sgtty)();
} cdevsw[];
```

86

Buffer Cache

Fourth Edition

```
#define B_READ  01
#define B_DONE   02
#define B_ERROR  04
#define B_BUSY   010
#define B_XMEM   060
#define B_WANTED 0100
#define B_RELOC   0200
#define B_ASYNC   0400
#define B_DELWRI  01000
```

FreeBSD 11.1

```
#define B_ASYNC 0x00000004
/* Start I/O, do not wait.
 */
[...]
#define B_DELWRI 0x00000080
/* Delay I/O until buffer
reused. */
#define B_DONE 0x00000200
/* I/O completed. */
```

87

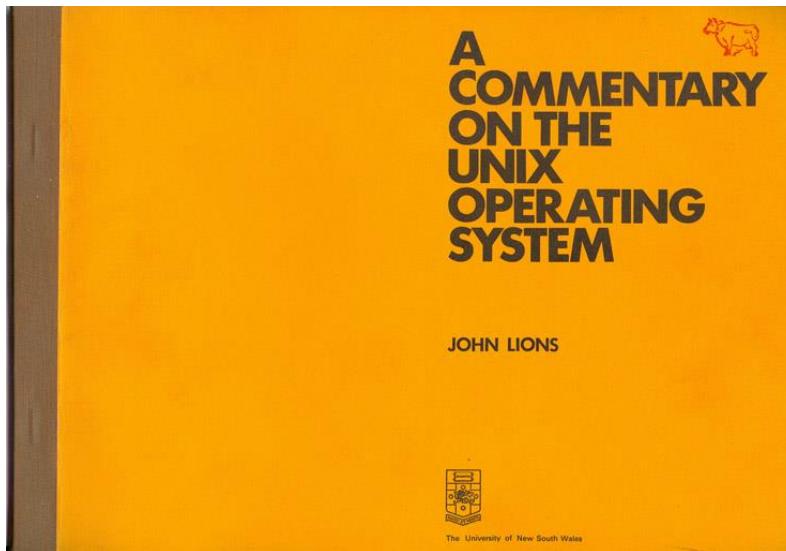
Fifth Research Edition (Jun 1974)

- 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
```

88

Sixth Research Edition (May 1975)



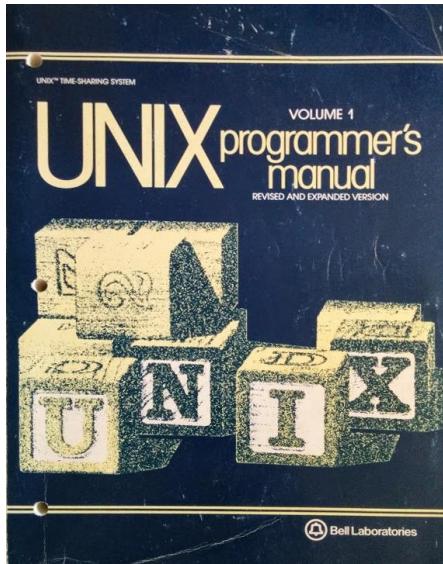
89

Portable C Library

<code>alloc.c</code>	<code>clenf.c</code>	<code>makbuf.c</code>	<code>scan1.c</code>
<code>calloc.c</code>	<code>copen.c</code>	<code>maktab.c</code>	<code>scan2.c</code>
<code>cclose.c</code>	<code>cputc.c</code>	<code>nexch.c</code>	<code>scan3.c</code>
<code>ceof.c</code>	<code>cwrd.c</code>	<code>nodig.c</code>	<code>system.c</code>
<code>cerror.c</code>	<code>dummy.s</code>	<code>printf.c</code>	<code>tmpnam.c</code>
<code>cexit.c</code>	<code>ftoa.c</code>	<code>putch.c</code>	<code>unget.c</code>
<code>cflush.c</code>	<code>getch.c</code>	<code>puts.c</code>	<code>unprnt.s</code>
<code>cfree.c</code>	<code>gets.c</code>	<code>relvec.c</code>	<code>wdleng.c</code>
<code>cgetc.c</code>	<code>getvec.c</code>	<code>revput.c</code>	
<code>ciodec.c</code>	<code>iehzap.c</code>	<code>run</code>	

90

Seventh Research Edition (Jan 1979)



91

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

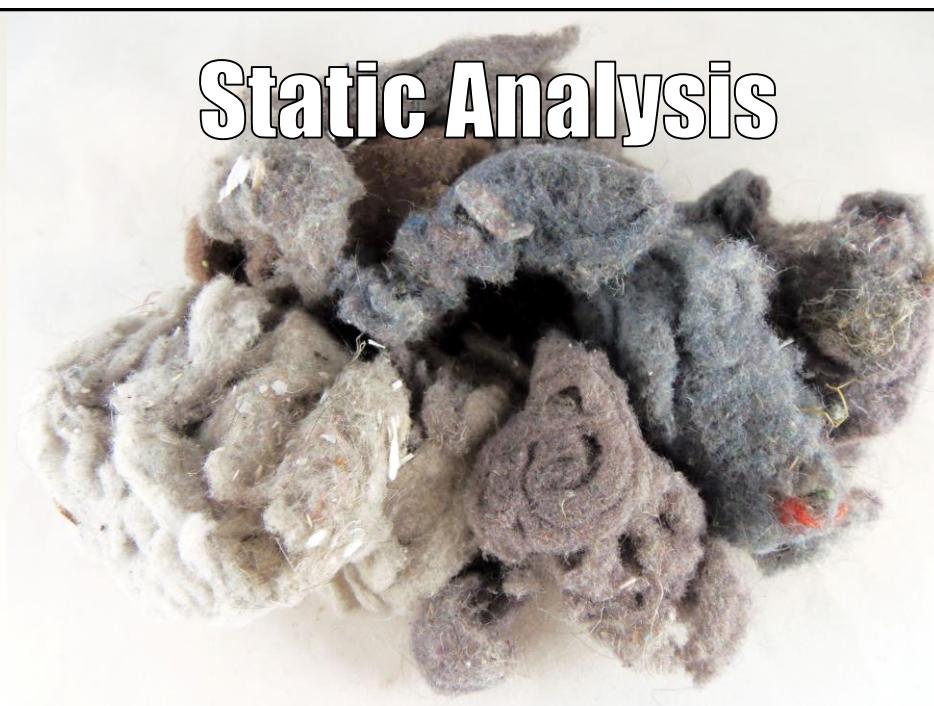
92

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)`

93

Static Analysis



94

Environment Variables

- KEY=value
- Kernel
- Shell
- C Library

ENVIRON(5) UNIX Programmer's Manual ENVIRON(5)

```

NAME
    environ - user environment

SYNOPSIS
    extern char **environ;

DESCRIPTION
    An array of strings called the 'environment' is made available by exec(2) when a process begins. By convention these strings have the form 'name=value'. The following names are used by various commands:
    PATH  The sequence of directory prefixes that sh, time, nice(1), etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by ':'. Login(1) sets PATH=/bin:/usr/bin.
    HOME  A user's login directory, set by login(1) from the password file passwd(5).
    TERM  The kind of terminal for which output is to be prepared. This information is used by commands, such as nroff or plot(1), which may exploit special terminal capabilities. See term(7) for a list of terminal types.
    Further names may be placed in the environment by the export command and 'name=value' arguments in sh(1), or by exec(2). It is unwise to conflict with certain Shell variables that are frequently exported by 'profile' files: MAIL, PS1, PS2, IFS.

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

```

95

Language Development Tools

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

Copyright © 1978 American Telephone and Telegraph Company
THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 57, No. 6, June/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.

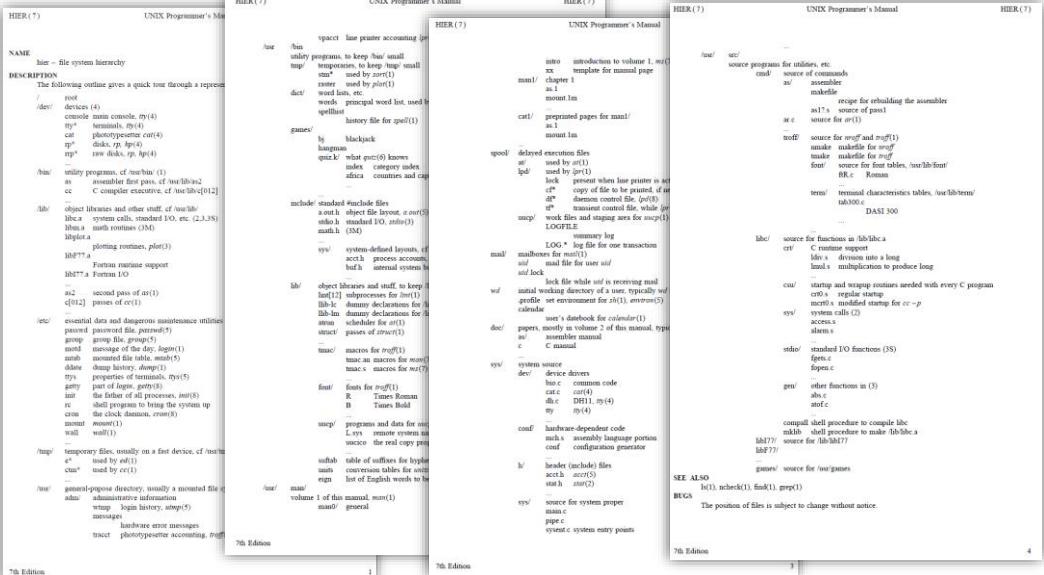
96

Domain-Specific Languages

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

97

Filesystem Directory Hierarchy



98



99

First and Second Berkeley Software Distributions (1978)

- Software Packages
 - csh
 - ex
 - Mail
 - Pascal
 - termlib

100

3BSD (1979)

- Virtual Memory Paging
 - `vm_*.`c
 - 2808 out of 16039 C source code
 - 17% of kernel source code
 - `vread(2)`, `vwrite(2)`, `vfork(2)`

101

4BSD (Oct 1980)



102

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

103

Optimized Screen Handling

- curses(3)
- termcap(5)



104

4.2BSD (Sep 1983)

- Internet Protocol Family
 - ARP, IP, TCP, UDP, ICMP
- Local and Remote Interprocess Communication
 - socket(2), etc.
- Network and User Database Access
 - getfsent(3x), getgrent(3), gethostent(3n),
getnetent(3n), getprotoent(3n), getpwent(3),
getservent(3n)
- Pseudo-Terminal Driver
 - pty(4)

105

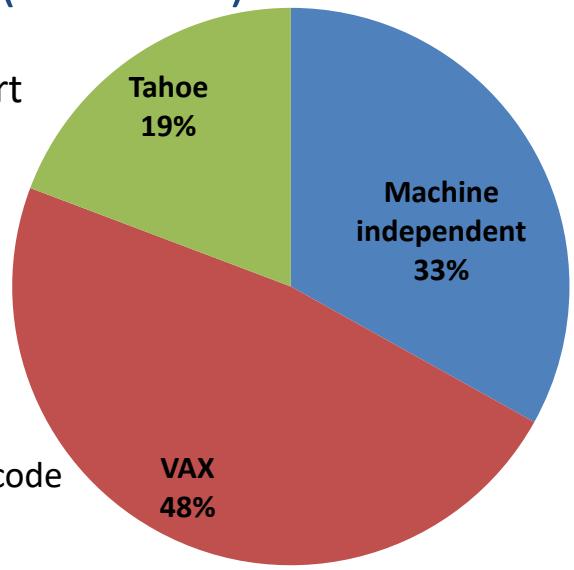
- 2 library functions
 - rcmd(3x)
 - rexec(3x)
- 11 system daemons
 - comsat(8c)
 - ftpd(8c)
 - gettable(8c)
 - implogd(8c)
 - rexecd(8c)
 - rlogind(8c)
 - af(8c)
 - rshd(8c)
 - rwhod(8c)
 - telnetd(8c)
 - tftpd(8c)
- 8 user-mode programs
 - ftp(1c)
 - rlogin(1c)
 - rsh(1c)
 - talk(1c)
 - telnet(1c)
 - tftp(1c)
 - whois(1c)
 - sendmail(1c)

System call	Uses
bind	23
connect	15
accept	13
select	12
listen	11
sendto	10
shutdown	9
recvfrom	8
getsockname	6
recv	2
send	2
sendmsg	1
getsockopt	0
recvmsg	0
socketpair	0

106

4.3BSD Tahoe (Jun 1988)

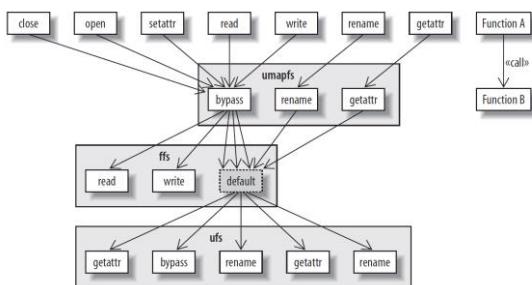
- Multiple CPU Architecture Support
 - VAX
 - CCI Power 6/32 (Tahoe)
- Timezone Handling
 - Community contribution
 - Separation of timezone rules from code



107

4.3BSD Reno (Jun 1990)

- Kernel Packet Forwarding Database
 - route(4)
 - routed(8), XNSrouted(8)
- Virtual Filesystem Interface



```
/*
 * Operations on vnodes.
 */
struct vnodeops {
    int (*vn_lookup)(...); /* ndp */ 
    int (*vn_create)(...); /* ndp, fflags, vap, cred */; 
    int (*vn_mknod)(...); /* ndp, vap, cred */; 
    int (*vn_open)(...); /* vp, fflags, cred */; 
    int (*vn_close)(...); /* vp, fflags, cred */; 
    int (*vn_access)(...); /* vp, fflags, cred */; 
    int (*vn_getattr)(...); /* vp, vap, cred */; 
    int (*vn_setattr)(...); /* vp, vap, cred */; 

    int (*vn_read)(...); /* vp, uiop, offp, ioflag, cred */; 
    int (*vn_write)(...); /* vp, uiop, offp, ioflag, cred */; 
    int (*vn_ioctl)(...); /* vp, com, data, fflag, cred */; 
    int (*vn_select)(...); /* vp, which, cred */; 
    int (*vn_mmap)(...); /* vp, ..., cred */; 
    int (*vn_fsync)(...); /* vp, fflags, cred */; 
    int (*vn_seek)(...); /* vp, (old)offp, off, whence */; 
};
```

...

108

4.3BSD Net/2 (Jun 1991)

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

109

4.4BSD (Jun 1994)

- Stackable Filesystems
 - `mount_null(8)`
 - `mount_union(8)`
- Generic System Control Interface (MIB)
 - `sysctl(1)`
 - `sysctl(3)`
 - `sysctl(9)`

110



111

386BSD Patch Kit (1992-1993)

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

112



FreeBSD®

113

FreeBSD 1.1 (May 1994)

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

114

FreeBSD 2.0 (Nov 1994)

- Process Filesystem
 - procfs(5)
- Dynamically Loadable Kernel Modules
 - lkm(4), then kld(4)
 - device drivers
 - file systems
 - emulators
 - system calls
 - **992** modules in FreeBSD 11.1

115

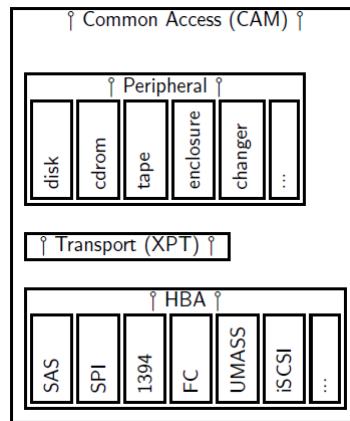
FreeBSD 2.1 (Nov 1995)

- Linux Emulation
- Packet Capture Library

116

FreeBSD 3.0 (Jan 1999)

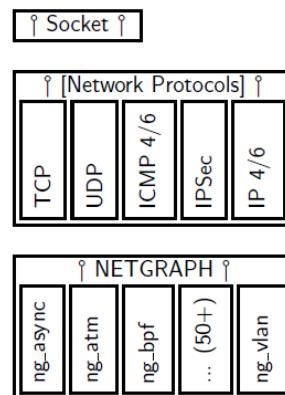
- Common Access Method I/O Subsystem (CAM)



117

FreeBSD 3.4 (Dec 1999)

- Graph-based Kernel Networking and User Library (NETGRAPH)



118

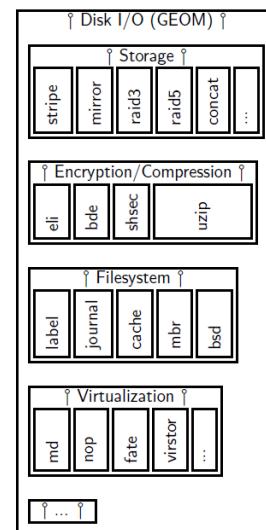
FreeBSD 4.0 (Mar 2000)

- OpenSSL Secure Sockets Layer and Transport Layer Security framework
 - Version 0.9.4
 - 1127 files, 227118 lines
 - libssl(3), libcrypto(3), openssl(1)
- Jail: Isolate a process and its descendants
- Access Control Lists

119

FreeBSD 5.0 (May 2006)

- Symmetric Multiprocessing
- Modular Disk I/O Request Transformation Framework (GEOM)
- Mandatory Access Control
- Pluggable Authentication Module



120

FreeBSD 5.3 (Nov 2004)

- Streaming Archive Access Library
- Miniport Driver Wrapper

121

FreeBSD 6.2 (Jan 2007)

- Basic Security Module Auditing

122

FreeBSD 7 (Feb 2008)

- ZFS File System and Storage Pools
- Dynamic Tracing

123

FreeBSD 9.0 (Jan 2012)

- Para-virtualized I/O
- Infiniband Support
- Application Compartmentalization

124

FreeBSD 10.0 (Jan 2014)

- Virtual Machine Monitor (bhyve)
- Fast User Space Raw Packet Processing

125

FreeBSD 11.0 (Sep 2016)

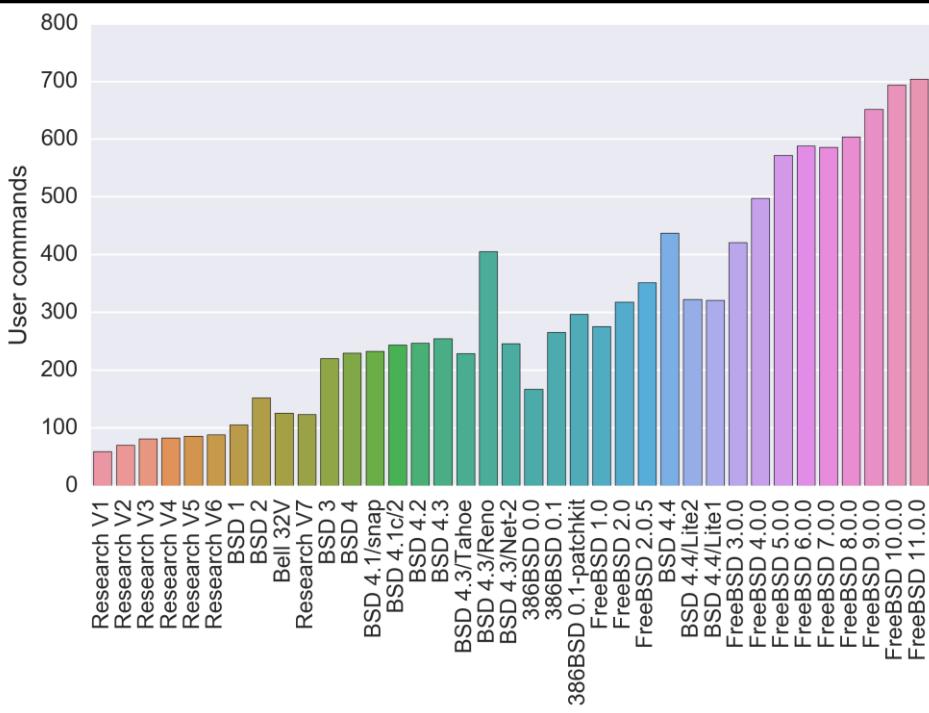
- Network Blacklisting

126

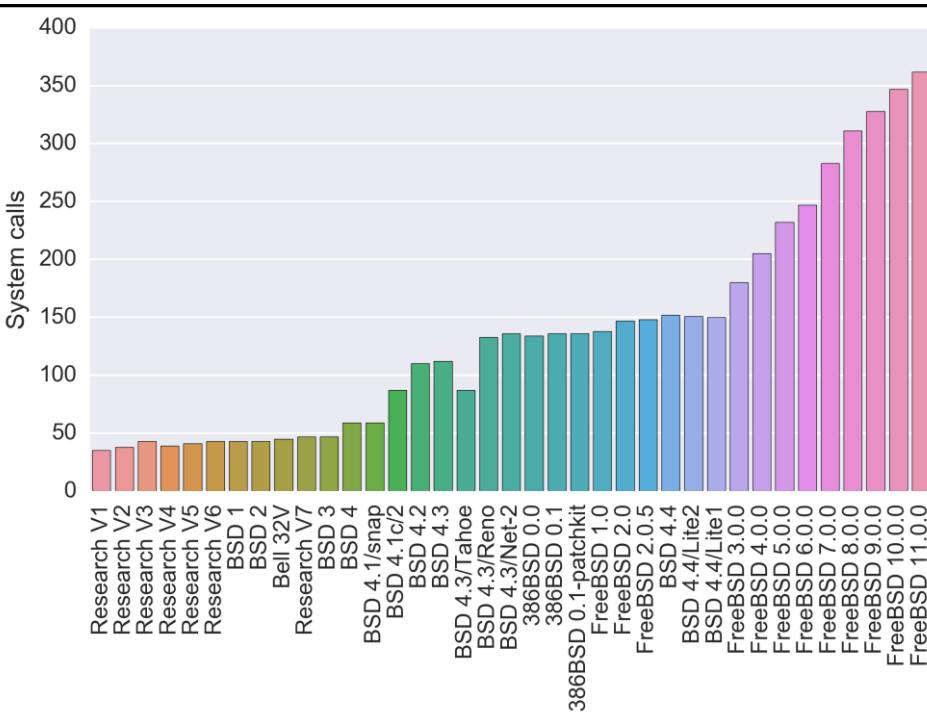
Quantitative Results



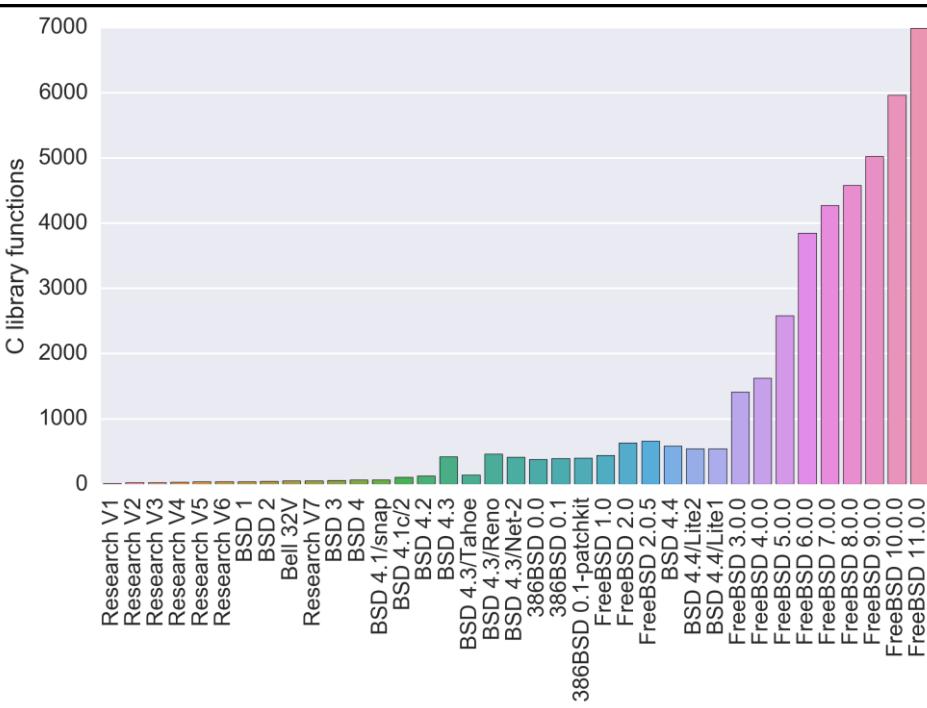
127



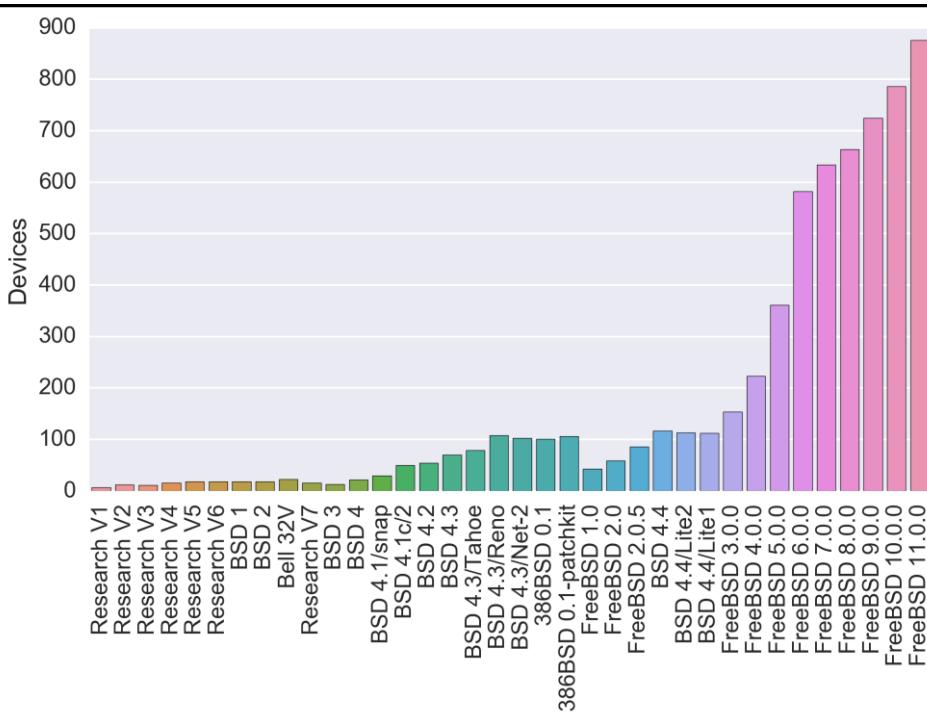
128



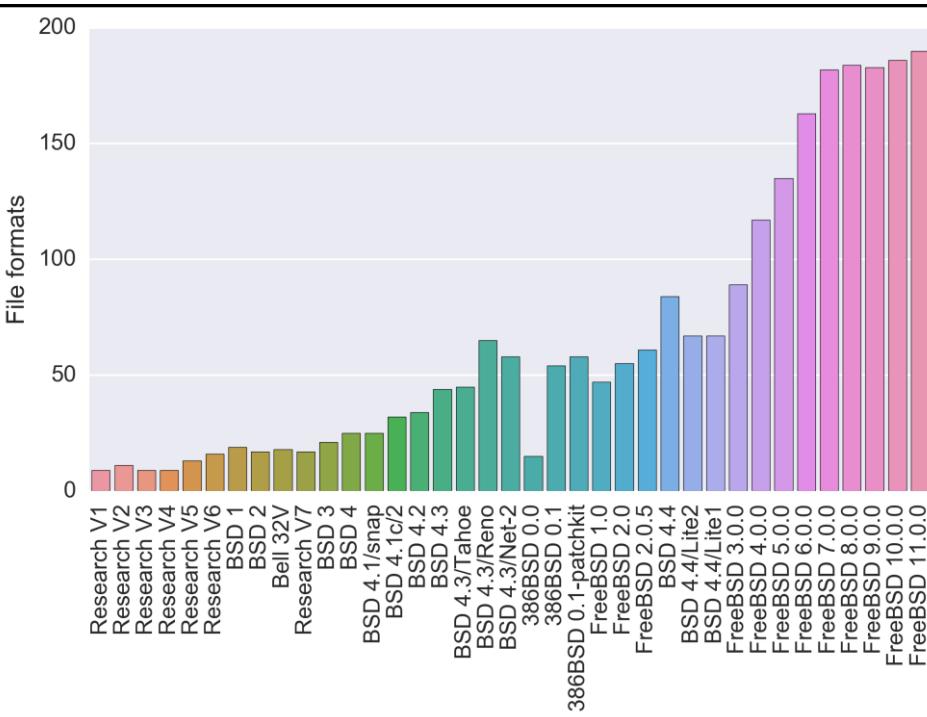
129



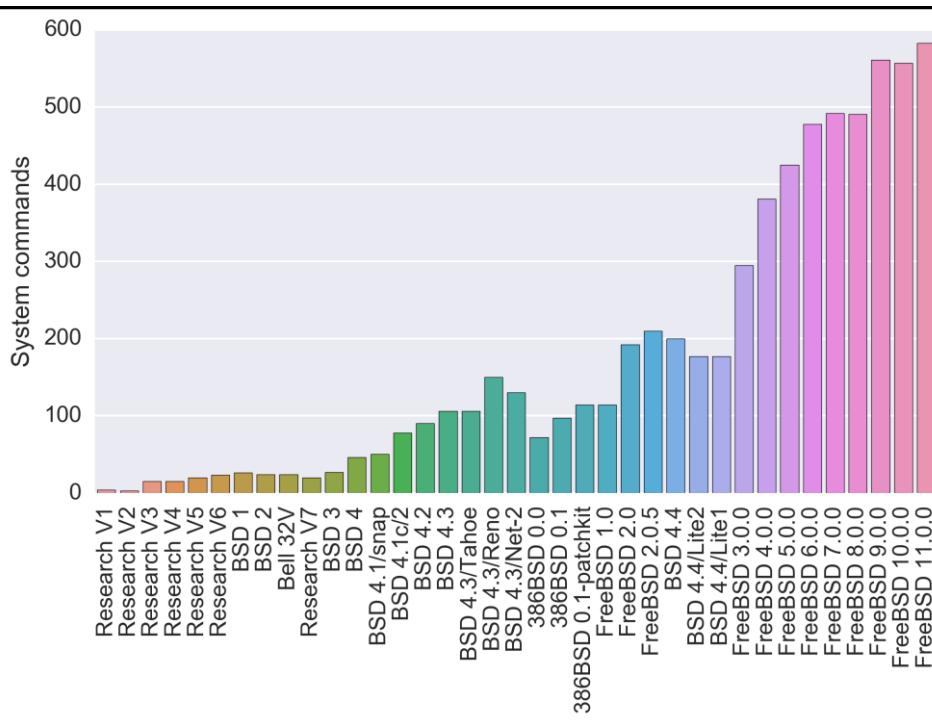
130



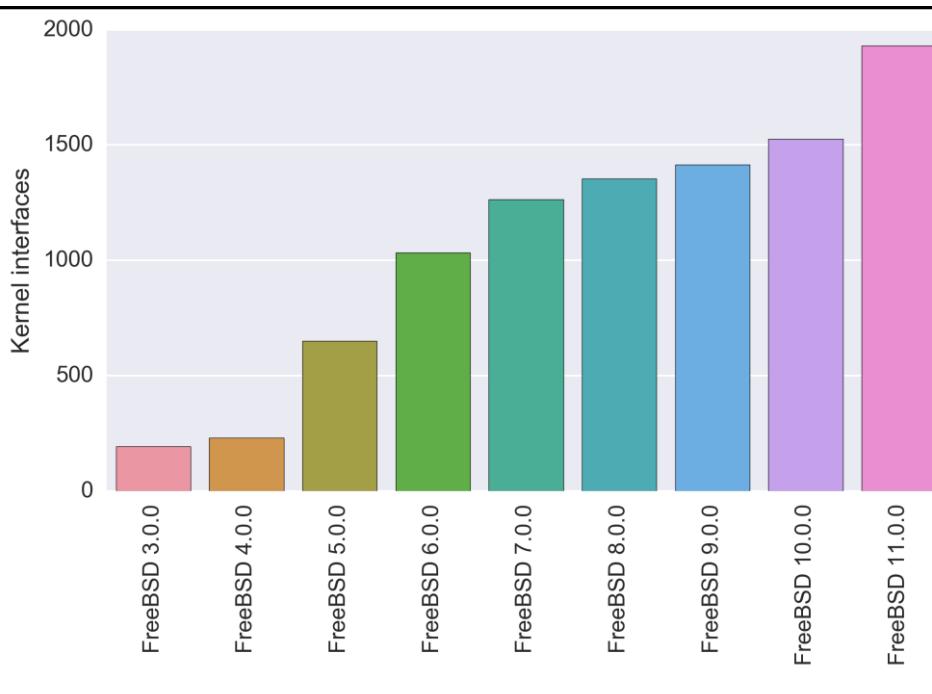
131



132

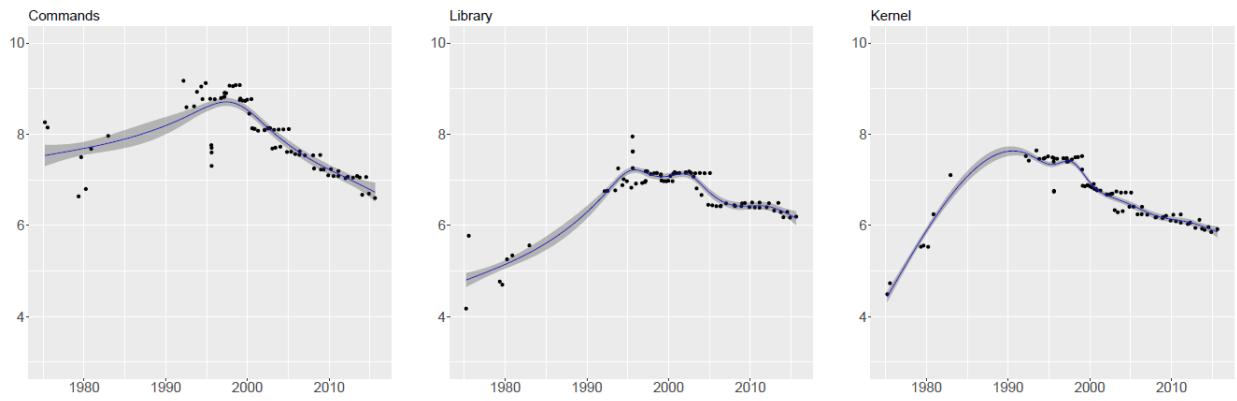


133



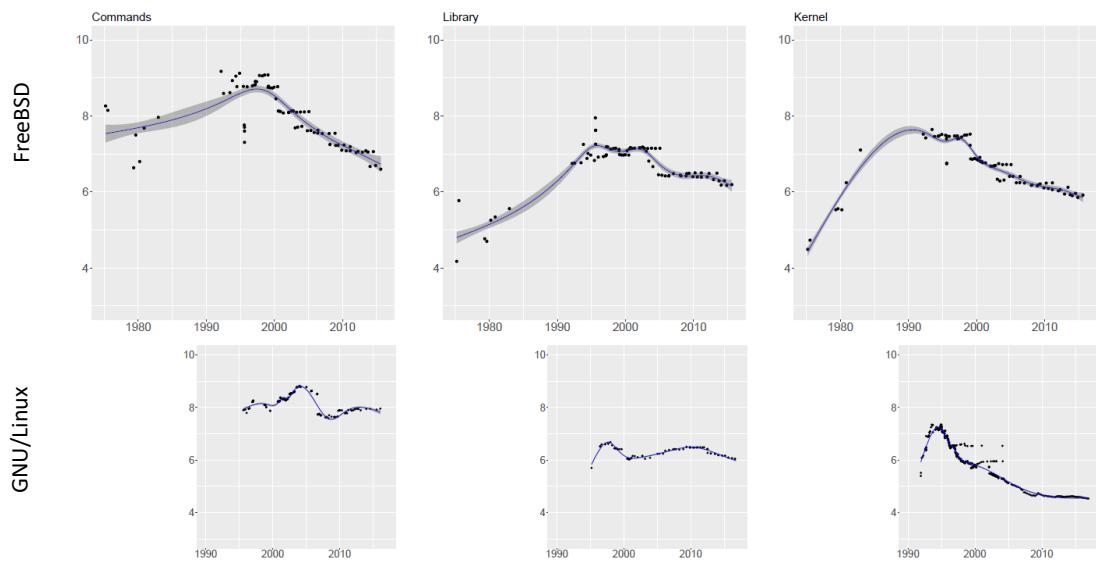
134

Mean Cyclomatic Complexity



135

Mean Cyclomatic Complexity



136

Theory of OS Architectural Evolution



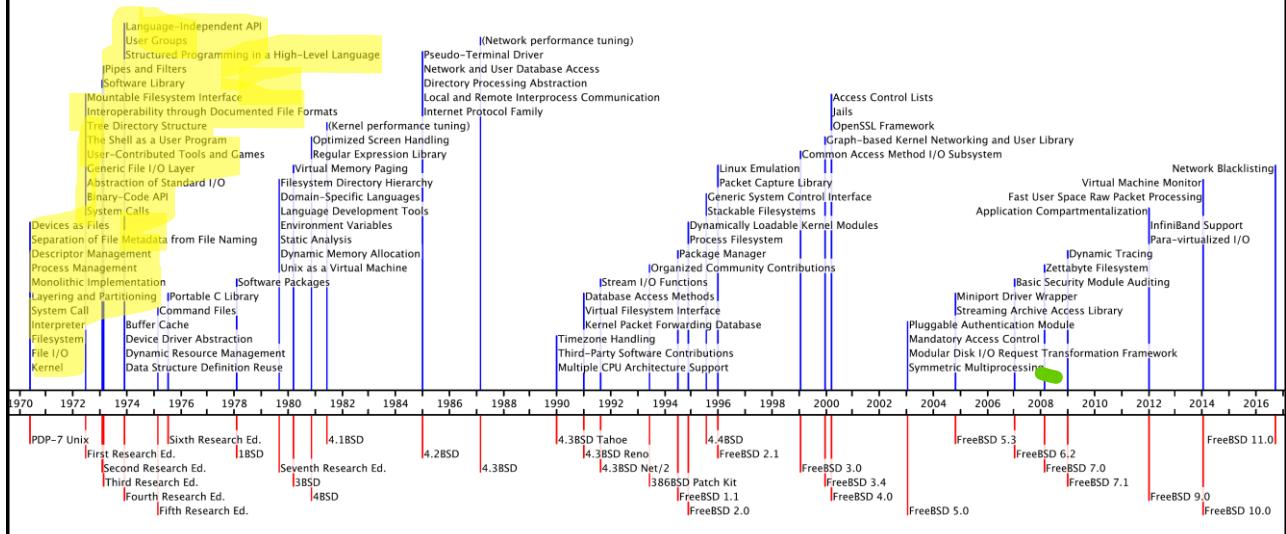
137

Form and Pace of Architectural Evolution



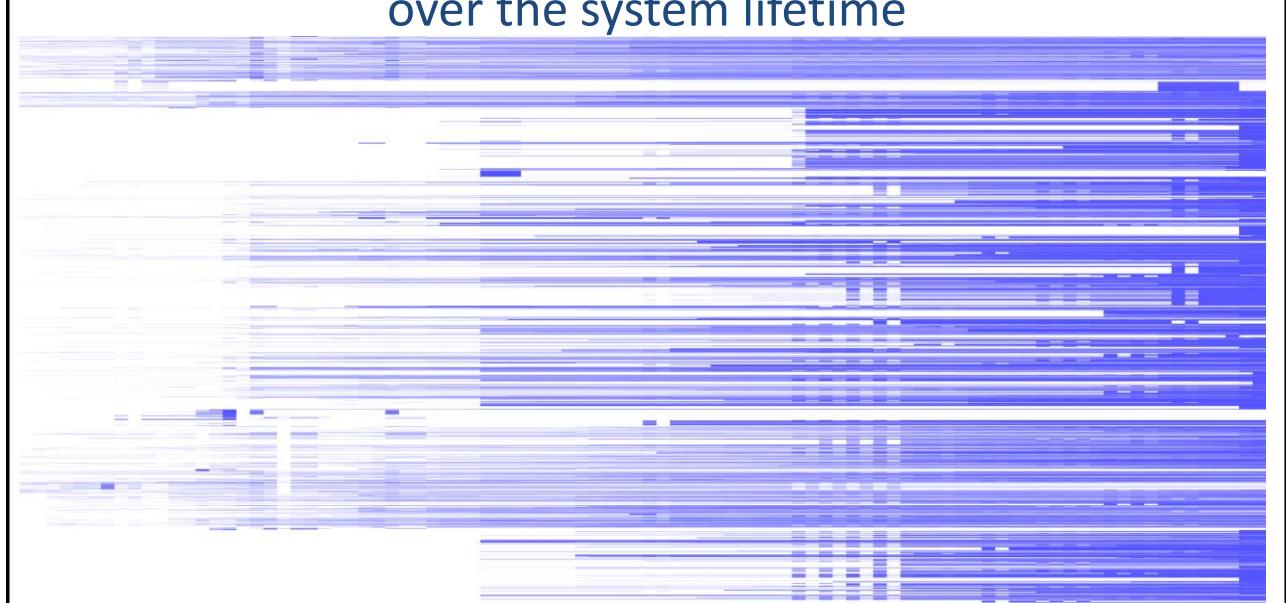
138

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

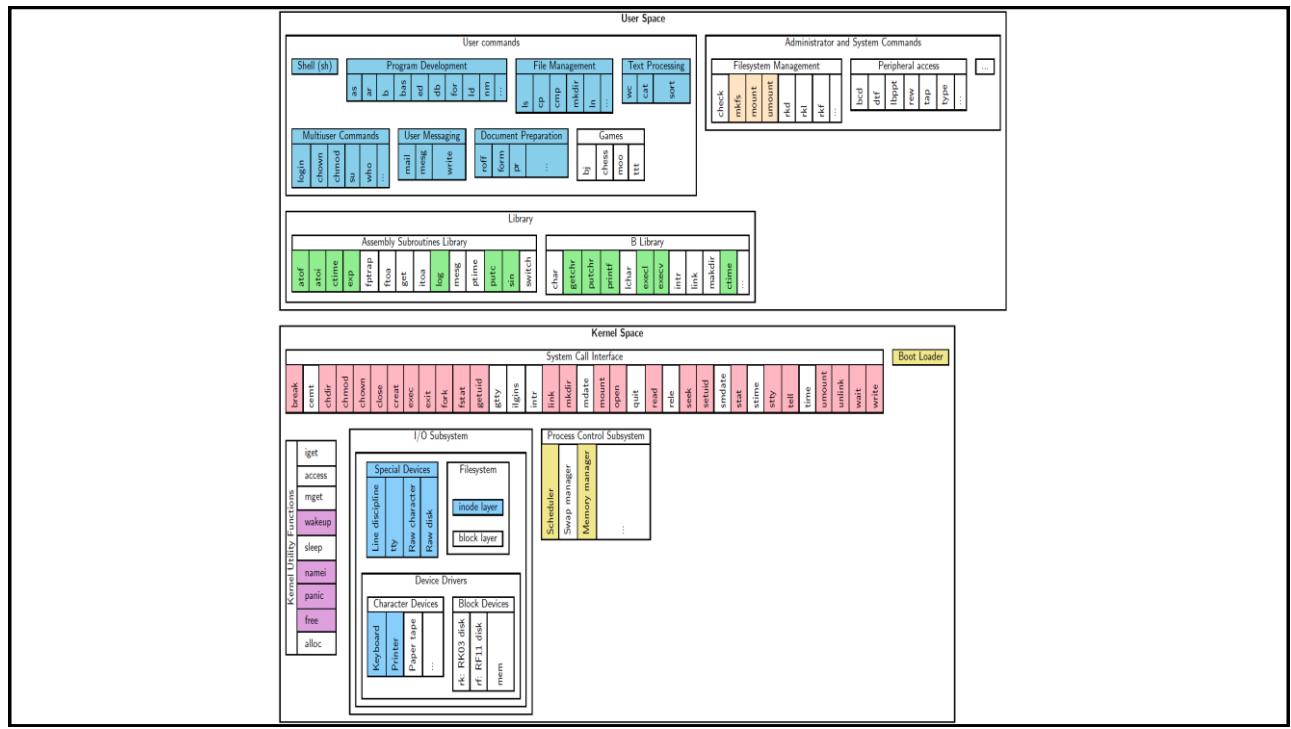


139

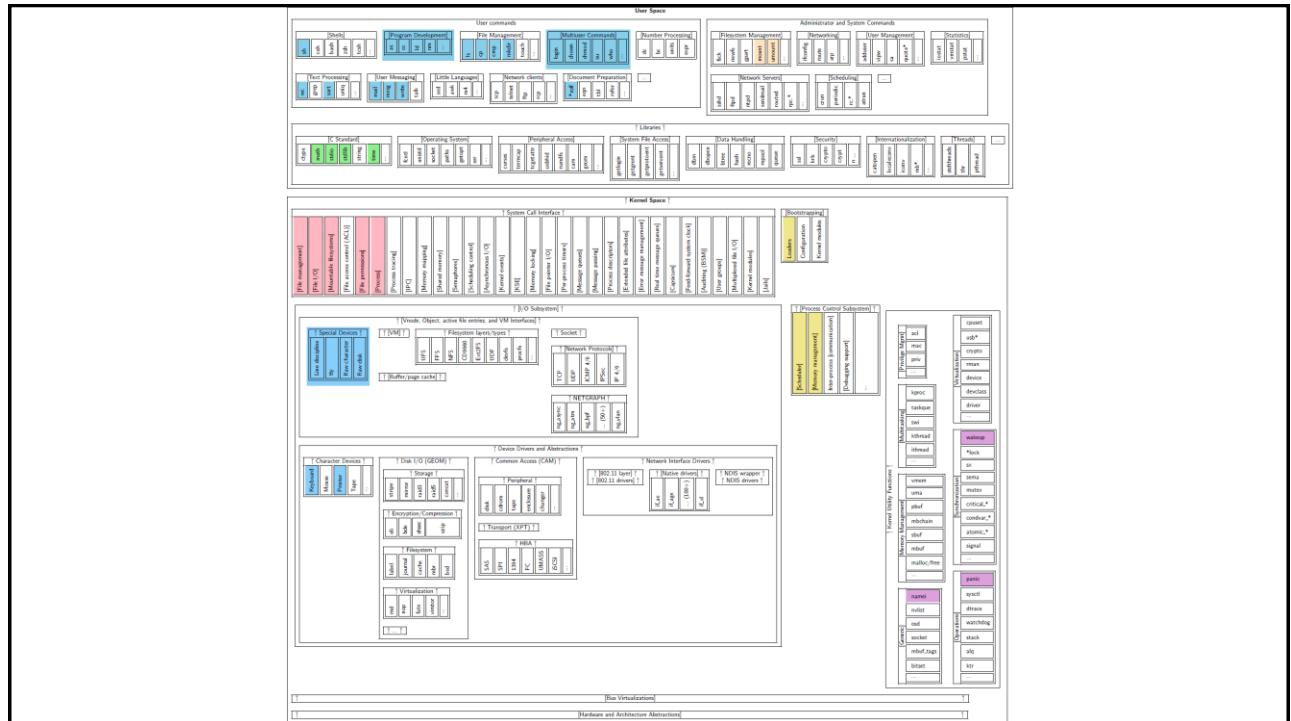
Most architecture decisions survive over the system lifetime



140

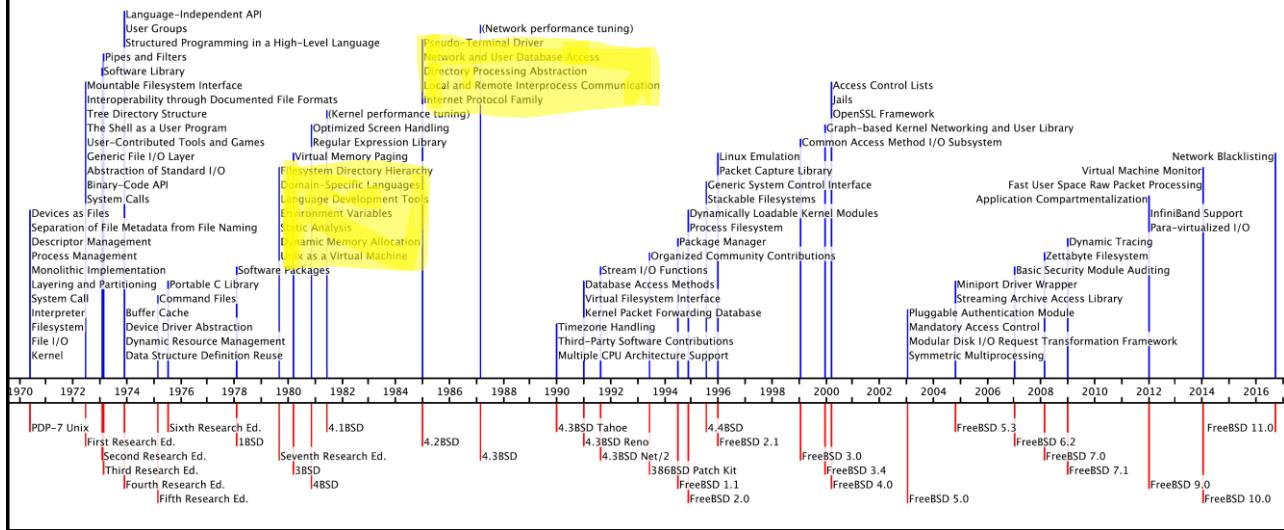


141



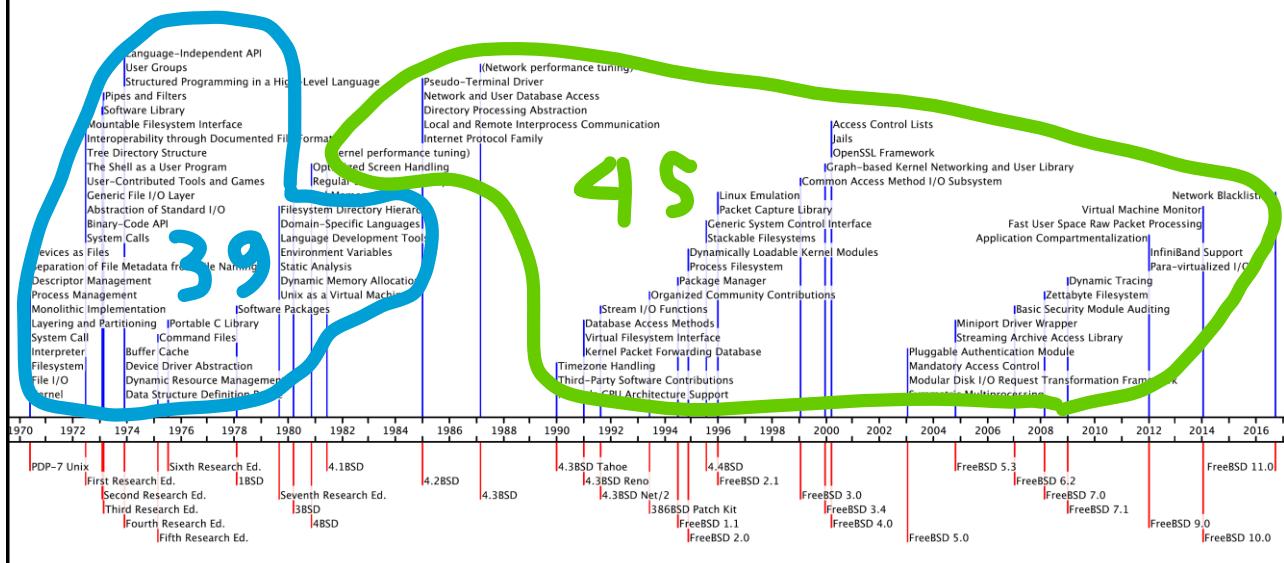
142

New architecture decisions are continuously made, further fueling architecture evolution



143

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



144

Accumulation of Architectural Technical Debt



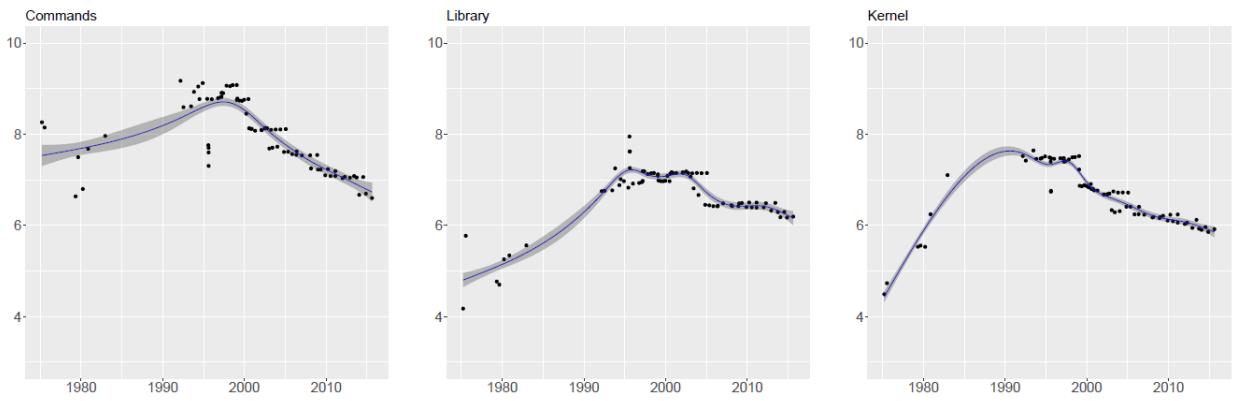
145

Architecture decisions offering features that are either similar to existing ones or remain under-used

- read, pread, readv, preadv, recv, recvfro, recvmsg
- /var/log, syslogd, acct, auditd
- UGO permissions, ACL, MAC
- Threads and processes for multitasking

146

Debt systematically paid back despite increasing system size and complexity



147

Forces of Architectural Evolution



148

Conventions instead of enforcement facilitate evolution by reducing effort and offering flexibility

- Text files
- Directory hierarchy
- Documented file formats
- Environment variables

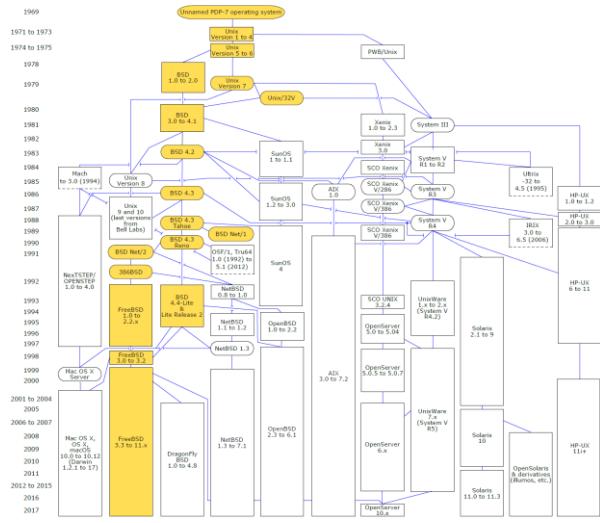
149

Portability, due to its inherent complexity,
is a key driver of architectural evolution



150

An ecosystem of other OSs and third parties constantly shapes the architecture evolution



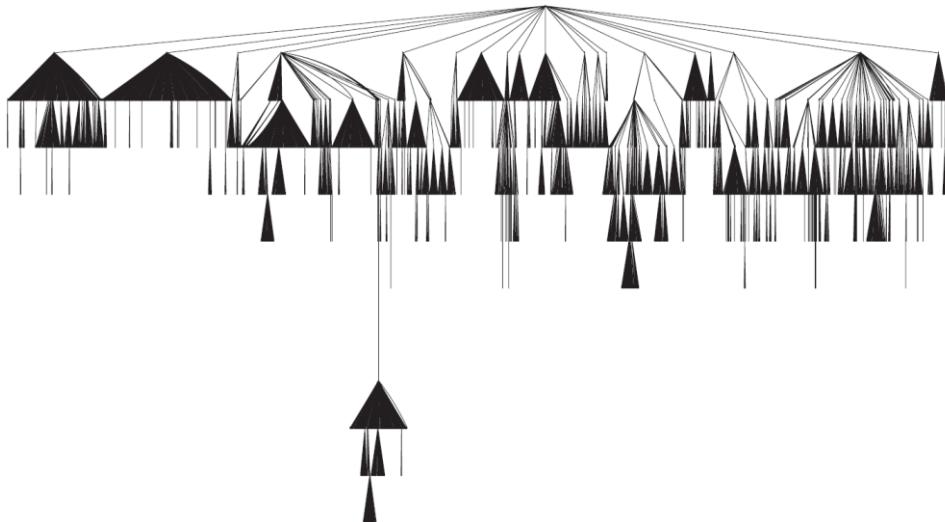
Major Third-Party Subsystems in FreeBSD 11.1

Major FreeBSD Third-Party Influences

Subsystem	kLoC	LoC %
TrustedBSD Project	1215	413 339
NetBSD	1166	2 665 223
OpenBSD	726	113 195
KAME	451	163 874
Semihalf sp.	330	214 289
DragonflyBSD	179	675 906
Linux	151	109 600
Qualcomm Atheros, Inc.	139	46 608
ABT Systems Ltd	133	8 704
Juniper Networks, Inc.	125	66 971
NetApp, Inc.	120	8 044
Illumos	97	56 618
OpenSolaris	95	125 503
Wheel Systems, Inc.	81	3 552
Yandex LLC	64	3 630
Apple, Inc.	58	13 378

151

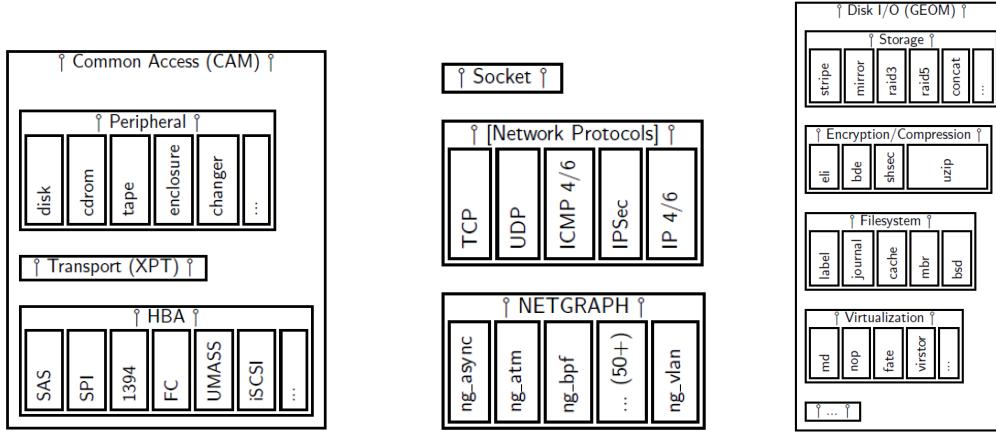
Third-party subsystems facilitate evolution through code reuse, but add technical debt



152

Large subsystems form their own independent architecture

- Netgraph, SSL, MAC, PAM, GEOM, BSM, ZFS, DTrace



153

Thank you!



Diomidis Spinellis and Paris Avgeriou. Evolution of the Unix system architecture: An exploratory case study. *IEEE Transactions on Software Engineering*, 2019.
doi:10.1109/TSE.2019.2892149

www.spinellis.gr
 @CoolSWEng

154

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?sbcars2020>



155

Image Credits

- Faces of Open Source / Peter Adams
 - Data: Joshua Sortino
 - Hackers at Junction 2015: [Vmuru](#)
 - [ASR-33 Teletype: Rama](#) & Musée Bolo
 - [VT100: Jason Scott](#)
 - [PDP 11/20](#): 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

156

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.

157

IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, VOL. X, NO. Y

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 systems. In this paper we study the evolution of the Unix system architecture over its lifetime, examining 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 the growth in size with some notable outliers, while system evolution is driven by a mix of forces, including technological evolution, market demand, and user requirements. 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 form of redundant code, legacy code, and unoptimized code, but the system's evolution is not driven by this debt, although it appears that 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 influence of the Bell Labs culture on the system's evolution. These findings have led us to formulate 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

Unix¹ 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 pioneers. As an operating system, it has had an immeasurable impact on the field of computing. It is also the root of the current state of the art in software, network, and hardware engineering.

Studying the evolution of operating system software is not just significant from a historical perspective; it can provide valuable insights into evolvability, best practices and anti-patterns, for large, complex, long-living systems. Unix is a good candidate for such a study, not only 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 maintain their quality over time, without succumbing to soaring technical debt or uncontrollable architectural decay.

In this paper we study the evolution of Unix along the four dimensions of the software architecture presented. 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 its four main architectural design decisions across the main

1 Unix is with the Athene University of Economics and Business, Greece.
E-mail: <http://research.athenaeconomics.gr/~spinelli/>
E-mail: paris.avgeriou@athenaeconomics.gr
Manuscript received December 12, 2016.

1. UNIX® is a registered trademark of The Open Group. For the sake of clarity, we refer to the original version of "Unix" to take both the UNIX systems developed at Bell Labs and to Unix-like systems, such as FreeBSD, that descended from them.

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 more than on the ground.

2.1 Software Evolution

There have been several studies on the longitudinal evolution of large systems. The seminal work of Lehman [1] and its numerous 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*, 2019.
doi:10.1109/TSE.2019.2892149

158

Data Sources

Tag	Data source(s)
Research-V1	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v1/svntrree-20081216.tar.gz
Research-V3	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v3/nsys.tar.gz
Research-V4	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v4/v4man.tar.gz
Research-V5	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v5/v5root.tar.gz
Research-V6	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v6/v6root.tar.gz
	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v6/v6src.tar.gz
	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Dennis_v6/v6doc.tar.gz
BSD-1	http://www.tuhs.org/Archive/PPD-11/Distributions/ucb/1bsd.tar.gz
BSD-2	http://www.tuhs.org/Archive/PPD-11/Distributions/ucb/2bsd.tar.gz
Research-V7	http://www.tuhs.org/Archive/PPD-11/Distributions/research/Henry_Spencer_v7/v7.tar.gz
	http://www.tuhs.org/Archive/PPD-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:///OSRG-CD-ROMs/cd1/4.0
BSD-4_1_snap	file:///OSRG-CD-ROMs/cd1/4.1.snap
BSD-4_1c_2	file:///OSRG-CD-ROMs/cd1/4.1c.2
BSD-4_2	file:///OSRG-CD-ROMs/cd1/4.2
BSD-4_3	file:///OSRG-CD-ROMs/cd1/4.3
BSD-4_3.Tahoe	file:///OSRG-CD-ROMs/cd2/4.3tahoe
BSD-4_3_Net_1	file:///OSRG-CD-ROMs/cd2/net.1
BSD-4_3_Reno	file:///OSRG-CD-ROMs/cd2/4.3reno
BSD-4_3_Net_2	file:///OSRG-CD-ROMs/cd2/net.2
BSD-4_4	file:///OSRG-CD-ROMs/cd3/4.4
BSD-4_4_Lite1	file:///OSRG-CD-ROMs/cd2/4.4BSD-Lite1
BSD-4_4_Lite2	file:///OSRG-CD-ROMs/cd3/4.4BSD-Lite2
BSD-SCCS	file:///OSRG-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
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
FreeBSD-release/2...	https://github.com/freebsd/freebsd

159

Empirical Software Engineering
10.1007/s10664-016-9445-5

A Repository of Unix History and Evolution

Dionisios 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 to store the history of the system. The data set was collected and curated by synthesizing custom software 24 snapshots of systems developed at Bell Labs, the University of California at Berkeley, and the 386ist team, two legacy repositories, and the modern repository of the open source FreeBSD 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 · Unix · configuration management · 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

The work has been partially funded by the Research Centre of the Athens University of Economics and Business, under the Original Scientific Publication framework [project code EP-2379] and by the Ministry of Education and Religious Affairs of the Greek Republic and the Research Network (GREENET) in the National eRC facility – ARIS – under project ID PA000005-CIOLPOT.

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

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

This is a muchabbreviated version of a working paper draft that led to a publication. The full version is available and cited in reference 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>.



160

Documented Unix Facilities Over 48 Years

Diomidis Spinellis
Department of Management Science and Technology
Athens University of Economics and Business
Athens, Greece
dds@hua.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 page descriptions of 93 Unix releases, on the analysis of the source code 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, system headers, system files, system utilities, file formats, games et al., miscellany, system maintenance procedures and commands, and system kernel interfaces. A timeline view allows the visualization of the evolution of the facilities. 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 → Software evolution;

ACM Reference Format:
Diomidis Spinellis. 2018. Documented Unix Facilities Over 48 Years. In *ASDI '18: Proceedings of the 2018 ACM SIGART Workshop on Mining Software Repositories*, May 28–29, 2018, Göteborg, Sweden. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3193939.319476>

1 INTRODUCTION
The Unix operating system is being continuously developed from the initial code base for more than a century now, and as a major engineering effort due to its early 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 language, its facilities, and its associated tools and libraries has been characterized as offering unusual simplicity, power, and elegance [5, 7]. Consequently, empirical data regarding how Unix facilities have evolved over time can be profitably used for empirical research on API evolution, system design, as well as technology adoption and trends.

Although the Unix facilities data set evolution through its source code [1, 10], the very large size of most systems can hinder the recognition of the relevant parts. Fortunately, another avenue is

Permissions to make digital or hard copies of all or part of this work for personal or private 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 to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to publish from <http://www.acm.org> or <http://diverse.csail.mit.edu>.
ASDI '18, May 28–29, 2018, Göteborg, Sweden
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5700-0 /18/05, Article 10, © 2018, ACM, Inc.
<https://doi.org/10.1145/3193939.319476>

available for studying Unix systems, namely their documentation. From the first version of the Unix system until today, all provided facilities (commands, APIs, file formats, and device drivers) are annotated with detailed documentation (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 Laboratories included a printed manual page instance. Some early editions of the manuals have not survived in a machine-readable format, but most are available in text markup that can be processed by any simple script.

The data set presented here is based on the printed and machine-readable Unix reference manuals released over a period of 48 years. It documents the evolution of 15 596 facilities through 93 versions of Unix releases, from the Sixth Edition to the Seventh Edition, showing how the data were produced, and Section 4 sketches two examples of how data can be used for quantitative and qualitative empirical studies.

2.1. UNRELEASED AND THEIR FACILITIES

The primary data source is available in the form of 93 Unix files containing 897 736 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 the following fields: the name of the facility (the name of the Unix manual section associated with a facility (1–8 see Table 1) followed by the facility's name, optionally followed by a cut identifier), the facility's documentation in troff markup [6]. In total, the set contains data for 15 596 facilities across 93 unique Unix releases, identifying 48 250 distinct manual page instances.

As an example, the following lines show the documentation (first section) associated with the assembly language compiler (*cc*), and the assembler (*as*), as documented in Section 1 of the 1973 Third Edition Unix manual.

```
1       Research->/3/as/as(1).1
1       as      Research->/3/as/as(1).1
1       as      Research->/3/as/as(1).1
```

By sending an email to the Unix History Repository GitHub pullrequest base URL (<https://github.com/dspinellis/uxc-history-repo/blob/>) to an entry's URL, one can obtain a URL for viewing the documentation manually or for the corresponding entry.

A separate text file (*releases.txt*) summarizes each of the data files with the year, month, and day of the corresponding release. For each release, the file also contains the timestamp list for the data associated with the Sixth and Seventh Research Editions and the first Berkeley Software Distribution.

Research->/3/as/as(1).1 1973-05-07 19
BSD-1 1978 02 01
Research->/3/as/as(1).1 1973-08 26
Research->/3/as/as(1).1 1973-08 26

161

81