

Contents

1. Preface/Introduction
2. Standardization and Implementation
3. File I/O
4. Standard I/O Library
5. Files and Directories
6. System Data Files and Information
7. Environment of a Unix Process
8. Process Control
9. Signals
10. Inter-process Communication

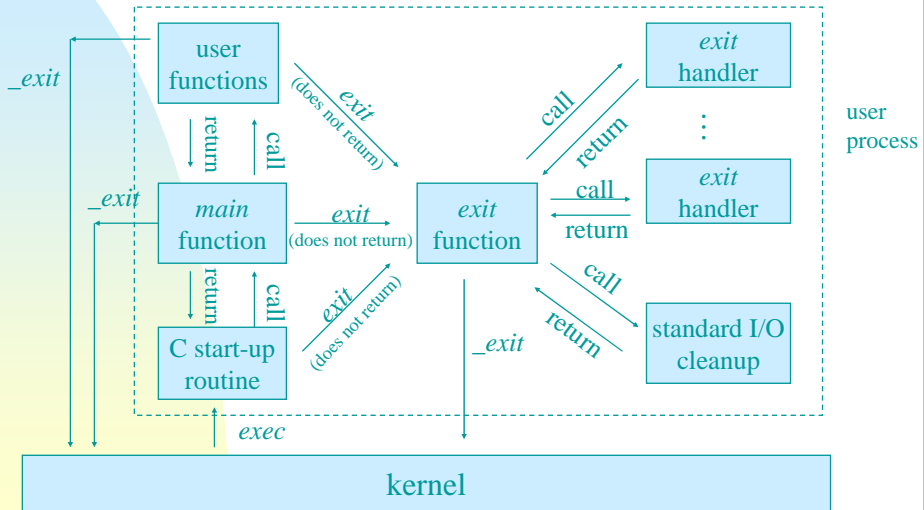
* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

The Environment of a Unix Process

- Objective:
 - How a process is executed and terminates?
 - What the typical memory layout is?
 - Related functions and resource limits.
- `int main(int argc, char *argv[])`
 - A special start-up routine is called to set things up first before call `main()`
 - Set up by the link editor (invoked by the compiler)

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

How a C program is started and terminated.



* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Process Termination

- Five ways to terminate:
 - Normal termination
 - Return from `main()`.
 - `exit(main(argc, argv));`
 - Call `exit()`.
 - Call `_exit()`
 - Abnormal termination
 - Call `abort()`
 - Be terminated by a signal.

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Process Termination

```
#include <stdlib.h>
```

```
void exit(int status);
```

- ANSI C
- Perform cleanup processing
 - Close and flush I/O streams (fclose)

```
#include <unistd.h>
```

```
void _exit(int status);
```

- POSIX.1
- Undefined exit status:
 - Exit/return without status.
 - Main falls off the end.

```
#include <stdio.h>
```

```
main() {  
    printf("hello world\n");  
}
```

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Process Termination

```
#include <stdlib.h>
```

```
int atexit(void (*func)(void));
```

- Up to 32 functions called by *exit* – ANSI C, supported by SVR4&4.3+BSD
- The same exit functions can be registered for several times.
- Exit functions will be called in reverse order of their registration.
- Program 7.1 – Page 165
 - Exit handlers

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Command-Line Arguments & Environment Variables

- Command-Line Arguments
 - `argv[argc]` is NULL under POSIX.1 & ANSI C
 - Program 7.2 – Page 166
- Environment Variables
 - `int main(int argc, char **argv, char **envp);`

`extern char **environ;`



environment list



environment strings

HOME=/home/stevens\0
PATH=:/bin:/usr/bin\0
SHELL=/bin/sh\0
USER=stevens\0
LOGNAME=stevens\0

- `getenv/putenv`

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Memory Layout

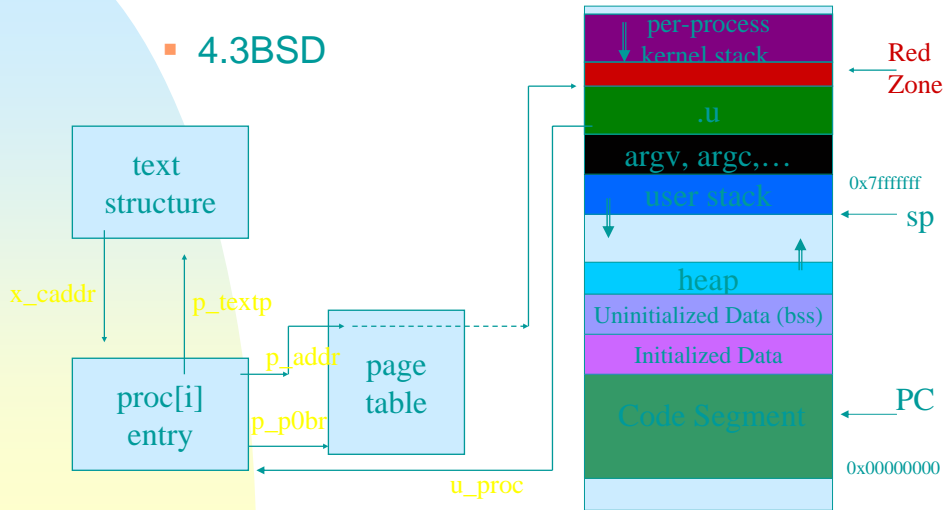
Read from
program file
by *exec*

- Pieces of a process
 - Text Segment
 - Read-only usually, sharable
 - Initialized Data Segment
 - `int maxcount = 10;`
 - Uninitialized Data Segment – bss (Block Started by Symbol)
 - Initialized to 0 by *exec*
 - `long sum[1000];`
 - Stack – return addr, automatic var, etc.
 - Heap – dynamic memory allocation

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Memory Layout

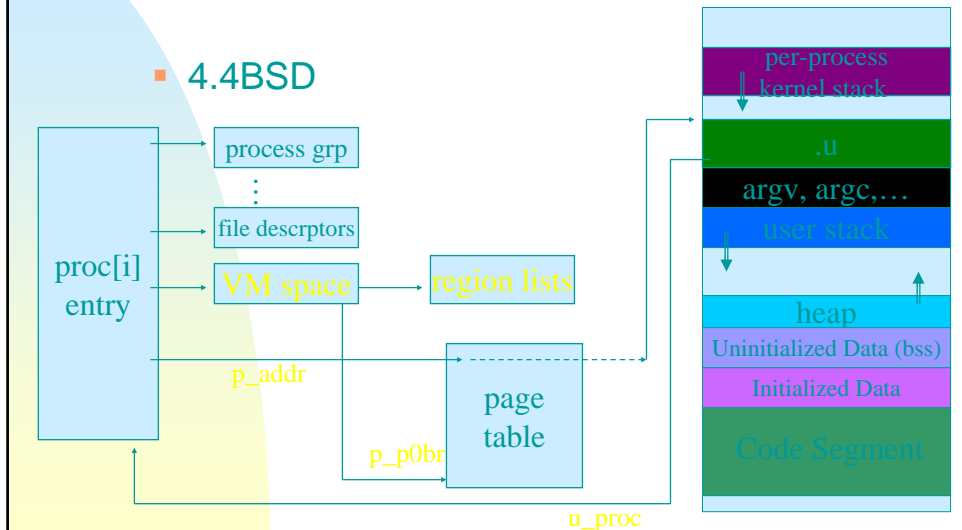
4.3BSD



* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Memory Layout

4.4BSD



* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Memory Layout

- > size ls1 hole

$$\text{ls1: } 6971 + 876 + 364 = 8211$$

$$\text{hole: } 6995 + 896 + 368 = 8259$$

↑ ↑ ↑
text data bss

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Shared Library

- Why a shared library?
 - Remove common library routines from executable files.
 - Have an easy way for upgrading
- Problems
 - More overheads
- Remark:
 - compiling options – gcc
 - Supported by many Unix systems

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Memory Allocation

- Three ANSI C Functions:
 - malloc – allocate a specified number of bytes of memory. Initial values are indeterminate.
 - calloc – allocate a specified number of objects of a specified size. The space is initialized to all 0 bits.
 - realloc – change the size of a previously allocated area. The initial value of increased space is indeterminate.

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

Memory Allocation

- ```
#include <stdlib.h>
void *malloc(size_t size);
void *calloc(size_t nobj, size_t size);
void *realloc(void *ptr, size_t newsize);
```
- Suitable alignments for any data obj
  - Generic void \* pointer
  - free(void \*ptr) to release space to a pool.
  - Leaking problem
  - Free already-freed blocks or blocks not from alloc().
    - malloc(M\_GRAINSet, value), mallinfo

```
#include <stdio.h>
#include <stdlib.h>
main() {
 char *ptr;

 ptr = malloc(100);
 free(ptr);
 free(ptr);
 ptr[0] = 'a'; ptr[1]=0;
 printf("%s - Done\n", ptr);
}
```

\* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

# Memory Allocation

- Remark
  - `realloc()` could trigger moving of data → avoid pointers to that area!
    - `prt == NULL` → `malloc()`
  - `sbrk()` is used to expand or contract the heap of a process – a `malloc` pool
  - Record-keeping info is also reserved for memory allocation – do not move data inside.
  - `alloca()` allocates space from the stack frame of the current function!
    - No needs for free with potential portability problems

\* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

# Environment Variables

- Name=value
  - Interpretation is up to applications.
    - Setup automatically or manually
    - E.g., HOME, USER, MAILPATH, etc.

```
setenv FONTPATH $X11R6HOME/lib/X11/fonts:$OPENWINHOME/lib/fonts
```

```
#include <stdlib.h>
```

```
char *getenv(const char *name);
```

- Figure 7.4 – environment variables
- ANSI C function, but no ANSI C environment variable.

\* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.



# Environment Variables

```
#include <stdlib.h>
```

```
int putenv(const char *name-value);
```

- Remove old definitions if they exist.

```
int setenv(const char *name, const char *value, int rewrite);
```

- `rewrite = 0` → no removing of existing names.

```
int unsetenv(const char *name);
```

- Remove any def of the *name*
- No error msg if no such def exists.

\* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

# Environment Variables

- Adding/Deleting/Modifying of Existing Strings

- Modifying

- The size of a new value  $\leq$  the size of the existing value → overwriting
    - Otherwise; `malloc()` & redirect the ptr

- Adding


- The first time we add a new name → `malloc` of pointer-list's room & update *envron*
    - Otherwise; copying. `realloc()` if needed.

- The heap segment could be involved.

\* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

# setjmp and longjmp

- Objective:
  - goto to escape from a deeply nested function call!
  - Program 7.3 – Program Skeleton
    - What if cmd\_add() suffers a fatal error? How to return to main() to get the next line?



Note: Automatic variables are allocated within the stack frames!

\* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

# setjmp and longjmp

- ```
#include <setjmp.h>
int setjmp(jmp_buf env);
int longjmp(jmp_buf env, int val);
```
- Return 0 if called directly; otherwise, it could return a value *val* from longjmp().
 - *env* tends to be a global variable.
 - longjmp() unwinds the stack and affect some variables.
 - Program 7.4 – Page 178
 - setjmp and longjmp

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

setjmp and longjmp

- Automatic, Register, and Volatile Variables
 - Compiler optimization
 - Register variables could be in memory.
 - Values are often indeterminate
 - Normally no roll back on automatic and register variables
 - Shown in Program 7.5 later
 - Global and static variables are left alone when longjmp is executed.
 - Portability Issues!

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

setjmp and longjmp

- Program 7.5 – Page 179
 - Effects of longjmp
 - Variables stored in memory have their values unchanged – no optimization...
- Potential Problems with Automatic Variables – Program 7.6 (Page 180)
 - Never be referenced after their stack frames are released.

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

getrlimit and setrlimit

```
#include <sys/time.h>
#include <sys/resource.h>

struct rlimit {
    rlim_t rlim_cur; /* soft limit */
    rlim_t rlim_max; /* hard limit */
}

int getrlimit(int resource, struct rlimit *rlp);
int setrlimit(int resource, const struct rlimit *rlp);
```

- Not POSIX.1, but supported by SVR4 and 4.3+BSD
- Rules:
 - Soft limit \leq hard limit, the lowering of hard limits is irreversible.
 - Only a superuser can raise a hard limit.

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

getrlimit and setrlimit

- Resources (SVR4&4.3BSD)
 - RLIMIT_CORE (both), RLIMIT_CPU (both, SIGXCPU), RLIMIT_DATA (both), RLIMIT_FSIZE (both, SIGXFSZ), RLIMIT_MEMLOCK (4.3+BSD, no implementation), RLIMIT_NOFILE (SVR4, _SC_OPEN_MAX), RLIMIT_NPROC (4.3+BSD, _SC_CHILD_MAX), RLIMIT_OFILE (4.3+BSD=RLIMIT_NOFILE), RLIMIT_RSS (4.3+BSD, max memory resident size), RLIMIT_STACK (both), RLIMIT_VMEM (SVR4)

* All rights reserved, Tei-Wei Kuo, National Taiwan University, 2003.

getrlimit and setrlimit

- Resource Limits → inheritance by processes
 - Built-in commands in shells
 - `umask`, `chdir`, `limit` (C shell), `ulimit -H` and `-S` (Bourne shell and KornShell)
- Program 7.7 – Page 183
 - Resource limits
 - `#define RLIM_NLIMITS 7`
 - `doit(RLIMIT_CORE) = pr_limits("RLIMIT_CORE", RLIMIT_CORE)`
 - `#define doit(name) pr_limits(#name, name)`