

Operating Systems

Tutorial 2 & 16

Michael Tänzer

os-tut@nhng.de
<http://os-tut.nhng.de>

Calendar Week 3

Outline

- 1 Review
- 2 File System Basics
- 3 Implementing Random Access to Files
- 4 Files in Linux
- 5 Virtual File System

True or False

- When executing a TestAndSet instruction interrupts are disabled on all processors
- It's possible to execute the TestAndSet instruction on multiple processors simultaneously
- Many caches use write-through in combination with write-to-memory

True or False

- When executing a TestAndSet instruction interrupts are disabled on all processors
- It's possible to execute the TestAndSet instruction on multiple processors simultaneously
- Many caches use write-through in combination with write-to-memory

True or False

- When executing a TestAndSet instruction interrupts are disabled on all processors
- It's possible to execute the TestAndSet instruction on multiple processors simultaneously
- Many caches use write-through in combination with write-to-memory

True or False

- When executing a TestAndSet instruction interrupts are disabled on all processors
- It's possible to execute the TestAndSet instruction on multiple processors simultaneously
- Many caches use write-through in combination with write-to-memory

What are the basic methods for accessing a file?

What are the basic methods for accessing a file?

Sequential Access

- File is accessed one record/byte after the other
- Read operation reads the next n records and advances the current position in the file accordingly
- Writes are appended to the end of the file

Direct Access

- Allows to read and write the records/bytes in any order
- Programmer has to specify explicitly which record to read/write

List two ways how the OS can determine which programme to run when clicking on a file icon

List two ways how the OS can determine which programme to run when clicking on a file icon

- ‘Remember’ which programme created the file
 - Can be stored in the file header or some other attached meta data store
 - Breaks when switching to another programme (e. g. Photoshop to GIMP)
- File type indicates the programme to run
 - File type can be determined by using file extensions (Windows) or magic numbers (Unix)
 - Look up suitable applications for this type in a table
 - If two applications use the same extension but a different file format almost certainly always the wrong one is started (Murphy’s law)

Relative vs. absolute path name

Absolute Pathname starts at the root of the directory structure

Relative Pathname starts at the current directory, the current directory is usually process specific

Symbolic vs. hard link

Symbolic vs. hard link

Symbolic Link

- A special kind of file which contains the path to the file it links to
- Path can be absolute or relative
- Can easily be identified as a link

Symbolic vs. hard link

Symbolic Link

- A special kind of file which contains the path to the file it links to
- Path can be absolute or relative
- Can easily be identified as a link

Hard Link

- Directory entry that refers to the file data of the original file
- Indistinguishable from the original file
- Only possible within the same file system

What is an access control list (ACL)?

What problems can occur within ACLs, and how can they be solved?

What is an access control list (ACL)?

What problems can occur within ACLs, and how can they be solved?

- A list associated with a file which contains entries in the form `username: access rights`

What is an access control list (ACL)?

What problems can occur within ACLs, and how can they be solved?

- A list associated with a file which contains entries in the form `username: access rights`
 - + Allows for fine-grained control which user is allowed to perform which operations on the file
 - ACLs can become quite large especially on systems with many users
 - High management overhead, when a new user is added to the system he has to be added to the ACL of *every* file he should have access to
- ⇒ Allow to groups in ACLs

Discuss kernel data structures required for a Unix-like handling of open files

Discuss kernel data structures required for a Unix-like handling of open files

- View on open files local to each process
- ⇒ On `open` a new entry is added to the process-local open file table in the PCB
- A new entry in a global open file table is allocated and assigned to the entry in the local open file table
- An entry in the global open file table contains mapping to a `vNode` structure (virtual representation of the actual file) and meta data like seek position and access rights
- Every `open` crates a new entry in the global open file table even if they point to the same `vNode` (might for example have a different seek position)
- Multiple local open file table entries might point to the same global open file table entry (e. g. after `fork` or `dup`)

What's the name of the system calls that move the current file position in Windows and Unix

What's the name of the system calls that move the current file position in Windows and Unix

Unix: `lseek`

Windows: `SetFilePointer` **and** `SetFilePointerEx`

How do these syscalls differ?

How do these syscalls differ?

`off_t lseek(int fd, off_t offset, int whence)`

- Sets seek position in the file referenced by `fd` to `offset` and returns the absolute seek position (or -1 on error)
- `whence` specifies how `offset` is interpreted
 - `SEEK_SET` Absolute addressing
 - `SEEK_CUR` `offset` is added to the current seek position
 - `SEEK_END` `offset` is added to the end of the file
- It's possible to set the seek position after the end of the file
- If data is written to that point there will be a gap in between (no disk space occupied and reads will return zeros) until data is actually written to it

How do these syscalls differ? I

```
DWORD SetFilePointer(HANDLE hFile,  
LONG lDistanceToMove,  
PLONG lpDistanceToMoveHigh,  
DWORD dwMoveMethod)
```

- Sets file pointer of the file referenced by `hFile` is to `lDistanceToMove` and returns the new value of the file pointer
- To specify bigger values than fit in a `LONG` `lpDistanceToMoveHigh` can contain a pointer to the high order 32 bit (the new file pointer high order bits are also returned via this pointer). Pass `NULL` if not used

How do these syscalls differ? II

- `dwMoveMethod` specifies the starting point for the file pointer move

`FILE_BEGIN` Start of the file (like `SEEK_SET`)

`FILE_CURRENT` Current file pointer (like `SEEK_CUR`)

`FILE_END` End of the file (like `SEEK_END`)

- On error the return value is

`INVALID_SET_FILE_POINTER` but this is a valid value

⇒ Need to call `GetLastError()` to determine whether an error occurred

Discuss alternative implementations without such an additional system call

Discuss alternative implementations without such an additional system call

- Keep the state in the application instead of the kernel
- ⇒ Add a parameter to the access functions (`read()`, `write()`, etc.) in which the application passes the seek position for each call
- + Kernel needs to store less information
- + One syscall less ⇒ less code and errors inside the kernel, less overhead on random access
- One more parameter passed to often called functions even though most files are accessed sequentially ⇒ need to validate that parameter
- Each application needs to implement handling of that state ⇒ could be done in a library

Is the syscall `open` absolutely necessary?

Consequences for not having it

Is the `syscall open` absolutely necessary?

Consequences for not having it

- Again we could add additional parameters to the access methods
- ⇒ Same arguments apply
 - Access methods have to do all the error checking which would normally be done by `open()`
- ⇒ On every access the kernel would have to look up the file in the directory structure and check the access rights
 - If a file is renamed between the operations it can't be found
- ⇒ As the state can't be extracted as easily as with the seek position it only pays off when accessing many files and only doing few operations on the same file

How can you easily create a new, empty file?

How can you easily create a new, empty file?

```
touch filename
```

How can you add execute-rights for the owner of a file to that file?

How can you add execute-rights for the owner of a file to that file?

```
chmod u+x filename
```

chmod u+x filename **VS.** chmod u+X filename

`chmod u+x filename` VS. `chmod u+X filename`

- Execute rights have different semantics on files and directories

File: File can be executed as application/script

Directory: Directory can be traversed (one can access subdirectories and files but without read permission one can't list their contents)

- `chmod u+X filename` only changes the execute permission if the operand is a directory
- This is especially useful in the combination with the `-R` option which traverses directories recursively (and thus makes all subdirectories traversable but doesn't make files executable which weren't before)

Write a C programme that creates a new file with a hole in it

Write a C programme that creates a new file with a hole in it

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <unistd.h>
#include <errno.h>
int main(int argc, char *argv[] ){
    char *txt = "Hallo Welt";

    int fd = open( "test.txt", O_CREAT | O_RDWR, S_IRUSR | S_IWUSR );

    write( fd, txt, 10 );
    lseek( fd, 10 * 1024 * 1024, SEEK_CUR );
    write( fd, txt, 10 );

    close( fd );
    return 0;
}
```

How can you find out whether a file contains a hole or not?

How can you find out whether a file contains a hole or not?

- The `du` command (Disk Usage) normally gives the size the file occupies in the file system
 - When given the `--apparent-size` option `du` gives the size as seen by applications
- ⇒ If the two sizes differ the file contains a hole

Write a C programme that copies `srcfile` to `dstfile` given as input parameters

Write a C programme that copies `srcfile` to `dstfile` given as input parameters

```
/* include many things */
#define BUF_SIZE    4096
int main( int argc, char *argv[] ){
    int src_fd = open( argv[1], O_RDONLY );
    int dst_fd = open( argv[2], O_CREAT | O_WRONLY, S_IRUSR | S_IWUSR );
    char buf[BUF_SIZE];
    int readsize, writesize;

    while ((readsize = read(src_fd, buf, BUF_SIZE)) > 0){
        char *write_buf = buf;
        while ((writesize = write(dst_fd, write_buf, readsize)){
            < readsize)
            readsize -= writesize;
            write_buf+= writesize;
        }
    }
    close( src_fd );
    close( dst_fd );
    exit( EXIT_SUCCESS );
}
```

What is the purpose of the VFS layer in an OS?

What is the purpose of the VFS layer in an OS?

- Abstraction from specific file system implementations
- ⇒ Kernel subcomponents (and possibly applications) which use file services can use any of the VFS file systems without modification
- Doesn't matter whether the file system is a normal file system (e. g. ExtX, ReiserFS) on a disk, ramdisk or a network file system (e. g. NFS, SMB)

Discuss potential drawbacks of using a VFS

Discuss potential drawbacks of using a VFS

- Lower performance through indirection. Instead of a direct call to a function of which the address can be determined at compile time the functions for the specific file system are called by following a function pointer in the VFS data structure (e. g. `vnode->vn_ops->vop_read()`)
 - Special features and optimisations on some file systems (e. g. meta data channels) can't be used
- ⇒ To use them one could call these special functions directly
⇒ destroys portability, but at least the standard tasks don't have to be ported for every file system

Which layer (VFS or underlying FS) resolves hard respectively symbolic links?

