

init { setup data structures

runtime {

- tracking names (fds, etc.)
- mem allocation
- counter update
- wrappers





Shutdown {

- identify shared
- aggregation (reduction)
- compression
- write

Compiler script:

- optionally calls out to files or scripts provided by modules to get link args

prepare_for_shutdown()



- no file names
- instead give bits of some len

Core lib at shutdown:

- spit out filename -> ID mapping
- its own aggregation for shared files

- job info
- tracks file names (full paths)
-> ID

- permits memory allocation (to track hard bounds)

Common/Utility lib

- ~~is~~ kicking off shutdown
- generic (zlib) compression
- utility fns for common patterns

iler script:
tionally calls out to
les or scripts provided
by modules to get
link args

core lib API

A: -register (char * mod-name, int * ^{runtime-}mem-limit, struct mod-fns * mfuncs)
-lookup-id (void * name, int len, int64 * ID, int printable-flag)

↑
path string
or
obj ID
or
?

-IDs don't go
away once generated

B' module API

-prep-for-shutdown() → triggers aggregation
-get-output (char **buffer, int size)

→ -no file name
-instead give buffer
of some len

Common API

-wtime
-detect shared IDs
-...

Posix

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

MPA-10

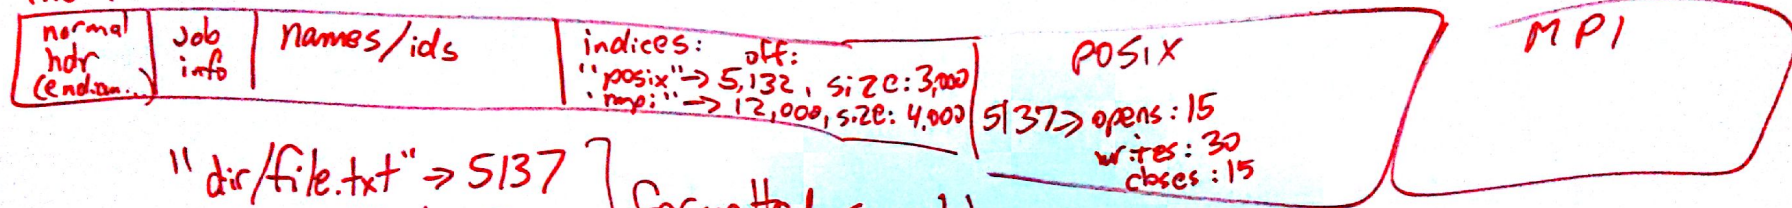
MPA-10

MPA-10

MPA-10

lib util (post processing)

file fmt:



"dir/file.txt" → 5137
↳ variable length } formatted sensibly
+ compressed

- each module can
define its own parser, grapher, whatever

devel notes

- git branch
- ignore compatibility
- strip down examples
- just one or 2 counters to start