

M[a]fs – Plan 9 userspace file systems

Mfs and Mafs wants you to be able to understand it, so you can be self-sufficient and fix a crash at two in the morning or satisfy your need for speed or a feature. This empowerment is priceless for those with skin in the game.


Mfs and Mafs are user space file systems to provide system stability and security. They are based on kfs.

As this document aims to also provide working knowledge, it gratuitously uses the actual commands and the relevant C data structure definitions to convey information.

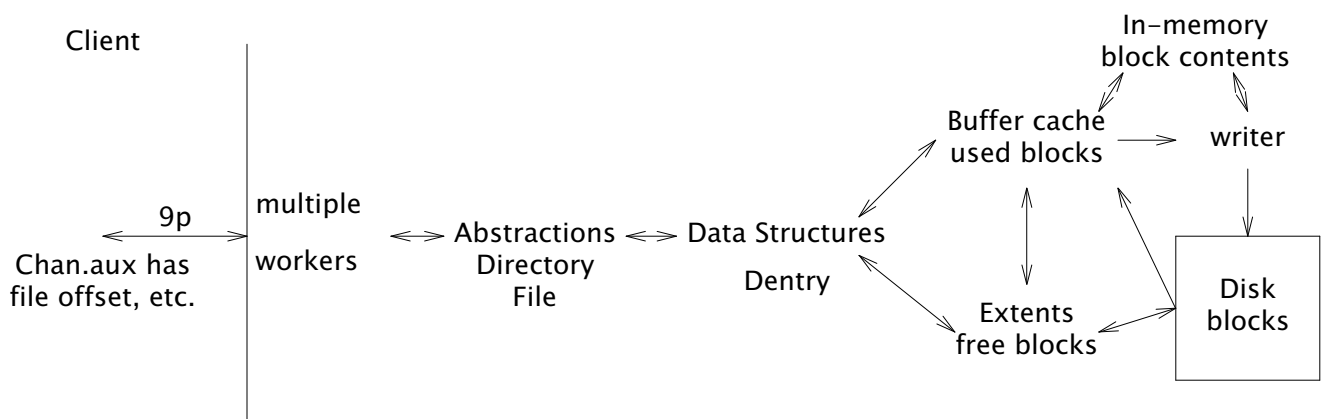
Mfs writes synchronously to the disk. It ensures that what has been said to be written has been passed along to the disk driver. Mafs writes asynchronously to the disk. It stores the writes in memory for a writer process to write to the disk at leisure.

This document uses the word M[a]fs to refer to both mfs and mafs.

Mfs Workflow



Mafs Workflow




Disk Contents

M[a]fs organizes and saves content on a disk as directories and files, just like any other filesystem.

The unit of storage is a logical block (not physical sector) of data. Disk space is split into 512 byte units. A directory entry block uses 1 unit (512 bytes) and a data block uses 16 units (8192 bytes).

A sample disk of 2048 bytes with 4 units.



A block is stored to the disk with a tag in the first byte and the Qid.path in the last 8 bytes. The different types of blocks on a disk are:

```
enum
{
    Tblank = 0,
    Tfree = 0,    /* free block */
    Tnone = 0,
    Tdentry,      /* directory entry, size = Dentrysize */
                  /* Tdata & indirect blocks are last, to allow for greater depth */
    Tdata,        /* actual file contents */
    Tind0,         /* contains a list of Tdata block numbers for files
                  and Tdentry block numbers for directories.*/
    Tind1,         /* contains a list of Tind0 block numbers */
    Tind2,         /* contains a list of Tind1 block numbers */
                  /* gap for more indirect block depth in future.
                  It can be put upto Tind5 without needing any code changes */
    Maxtind,      /* should be Tind0+Niblock */
    MAXTAG,

    Tmaxind = Maxtind - 1,
};
```

Every file or directory is represented on the disk by a directory entry (Dentry). A directory entry uses a unit sized block (tag = Tdentry) and is uniquely identifiable by a Qid.

A file's contents are stored in the directory entry itself if they are 320 bytes or lesser. A file stores its contents in blocks with a tag of Tdata if the file size is more than that. A directory holds the directory entries of its children in blocks each with a tag of Tdentry.

The blocks used by a file or directory entry are listed in their directory entry. As it is not possible to represent big files using the list of blocks available in the directory entry, the blocks are structured to use multiple levels of indirection as the file size increases.

A file's data blocks are identified by a tag of Tdata and that file's Qid.path. A directory's data blocks are identified by a tag of Tdentry and Qid.path of the child directory entry. (Is this quirky? Should the child's directory entry have a tag with the parent's Qid.path?)

A block number of zero represents the end of the file's contents. If a file is truncated, the data and indirect blocks are given up and the dentry.dblocks[0] = 0.

M[a]fs does not store the last access time of a file or directory.

A directory entry is defined as:

```
enum {
    Unit = 512ULL,          /* minimum data unit size */
    Dentryunits = 1,        /* number of units per Tdentry */
    Rawblockunits = 16,      /* number of units for others */

    Dentrysize = Dentryunits*Unit, /* real block size of a Tdentry */
    /* real block size of others. a multiple of Unit to keep the extents code uniform */
    Rawblocksize = Rawblockunits*Unit,
    Blocksize = Rawblocksize - sizeof(u8 /* tag */) - sizeof(u64 /* path */),
    Nindperblock = (Blocksize-7)/sizeof(u64), /* number of pointers per block. -7 for pad */

    Namelen = 129, /* maximum length of a file name, calculated manually */
    Ndblock = 32, /* number of direct blocks in a Dentry */
    Niblock = 3, /* maximum depth of indirect blocks, can increase it to 6 without issues */
};
struct Qid9p1
{
    u64 path; /* unique identifier */
    u32 version;
};
struct Dentryhdr
{
    u8 tag;
    s8 name[Namelen]; /* 130, Namelen = 129 to align to a u64 after the ids */
    s16 uid; /* 132 */
    s16 gid; /* 134 */
    s16 muid; /* 136 */
    u64 size; /* 0 for directories. For files, size in bytes of the content - 8 */
    u64 pdblkn; /* block number of the parent directory entry. Will be 0 for root. - 16 */
    u64 pppath; /* parent qid.path - 24 */
    u64 mtime; /* modified time in nano seconds from epoch - 32 */
    Qid9p1 qid; /* 44 bytes */
    u32 mode; /* same bits as defined in /sys/include/libc.h:/Dir.mode/ - 48 */
};
enum {
    /* max possible size of data that can be stuffed into a Dentry */
    Ddatasize = Dentrysize - sizeof(u64 /* path */) - sizeof(Dentryhdr),
};
```

```
struct Dentry
{
    Dentryhdr;
    union
    {
        struct
        {
            u64 dblocks[Ndblock];    /* direct blocks. */
                                   /* List of Tdata block numbers for files and
                                   Tdentry block numbers for directories */
            u64 iblocks[Niblock];    /* indirect blocks */
        };
        /* when size <= Dentrysize-184-sizeof(Tag), store the data here itself */
        s8 buf[Ddatasize];
    };
    u64 path;
};
```

A directory entry once assigned is not given up until the parent directory is removed. It is zero'ed if the directory entry is removed. It is reused by the next directory entry created under that parent directory. This removes the need for garbage collection of directory entries on removals and also avoids zero block numbers in the middle of a directory entry's list of blocks. A zero block number while traversing a directory entry's dblocks or iblocks represents the end of directory or file contents. When a directory is removed, the parent will have a directory entry with a tag of Tdentry and Qpnone and the rest of the contents set to zero.

A directory's size is always zero.

```
; tests/6.sizes
Unit 512 Dentrysize 512 Rawblocksize 8192
Dentryhdr size 184 Ddatasize 320
Dentry size 512 Namelen 129
Namelen 129 Ndblock 32 Niblock 3
Blocksize 8183 Nindperblock 1022
A Tind0 unit points to 1 data blocks (8183 bytes)
    block points to 1022 data blocks
    reli start 32    max 1053
    max size 1054*Blocksize = 8624882 bytes    = 8 MiB
A Tind1 unit points to 1022 data blocks (8363026 bytes)
    block points to 1044484 data blocks
    reli start 1054 max 1045537
    max size 1045538*Blocksize = 8555637454 bytes = 7 GiB
A Tind2 unit points to 1044484 data blocks (8547012572 bytes)
    block points to 1067462648 data blocks
    reli start 1045538 max 1068508185
    max size 1068508186*Blocksize = 8743602486038 bytes    = 7 TiB
```

On an empty m[a]fs filesystem mounted at /n/mafs, the disk contents added by the below commands are:

```
mkdir /n/mafs/dir1
```

```
echo test > /n/mafs/dir1/file1
```

Representation of a file in a directory: /dir1/file1

Block 11 contents: /dir1 Dentry

```
Tdentry 64
name dir1
qid.version 0
qid.path 64
size 0
pdblkn0 10
pqpath 10
mtime 1653302180819962729
mode 20000000777
uid 10006
gid -1
muid 10006
direct blocks
0 12
1 0
2 0
.
.
.
30 0
31 0
indirect blocks
0 0
1 0
2 0
```

Block 12 contents: file1 Dentry

```
Tdentry 65
name file1
qid.version 0
qid.path 65
size 5
pdblkn0
pqpath 64
mtime 1653302180823455071
mode 666
uid 10006
gid -1
muid 10006
test
```

Representation of two files in a directory (/dir2/file1 and /dir2/file2)

Block 13 contents: /dir2 directory entry

```
Tdentry 66
name dir2
qid.version 0
qid.path 66
size 0
pdblkn0 10
pqpath 10
mtime 1653302180819962729
mode 20000000777
uid 10006
gid -1
muid 10006
direct blocks
  0 14
  1 15
  .
  .
  .
 31 0
indirect blocks
  0 0
  1 0
  2 0
```

Block 14 contents: file1 directory entry

```
Tdentry 67
name file1
qid.version 0
qid.path 67
size 5
pdblkn0 13
pqpath 66
mtime 1653302180823455071
mode 666
uid 10006
gid -1
muid 10006
test
```

Block 15 contents: file2 directory entry

```
Tdentry 68
name file2
qid.version 0
qid.path 68
size 5
pdblkn0 13
pqpath 66
mtime 1653302180823455071
mode 666
uid 10006
gid -1
muid 10006
test
```

iblocks[0] holds the block number of a Tind0 block. A Tind0 block contains a list of Tdata block numbers for files or a list of Tdentry block numbers for directories.

iblocks[1] has the block number of a Tind1 block. A Tind1 block contains a list of Tind0 block numbers.

Similarly, for other iblocks[n] entries, iblocks[n] has the block number of a Tind n block. A Tind n block contains a list of Tind $(n-1)$ block numbers.

Relative index

The zero'th relative index in a directory entry is the first data block. The next relative index is the second data block of the directory entry, and so on.

tests/6.reli shows how a relative index (reli) is translated into an actual disk block number.

To find the actual block number where the first block (zero'th as zero indexed) of a file is stored:

```
tests/6.reli 0 # command, below is the output of this command
reli 0
dblock[0]
```

To find the actual block number where the second block of a file is stored:

```
tests/6.reli 1
reli 1
dblock[1]
```

And so on, for the 32nd and 33rd blocks of a file:

```
tests/6.reli 31
reli 31
dblock[31]
```

```
tests/6.reli 32
reli 32
iblock[0]
Tind0 reli 0 is at [0]
```

Representation of a 2 MB file (/dir3/2MB.file)

Block 17 contents

```
Tdentry 70
name 2MB.file
qid.version 0
qid.path 70
size 2056192
pdblkn0 16
pqpath 69
mtime 1653302180819962729
mode 20000000777
uid 10006
gid -1
muid 10006
direct blocks
  0 18
  1 34
  2 50
.
.
indirect blocks
  0 546
  1 0
  2 0
```

Block 18 contents

```
Tdata 70
0 0123456789
```

contents of 2MB.file

Representation of a 25MB file (/dir4/25MB.file)

Block 4067 contents

```
Tdentry 72
name 25MB.file
qid.version 0
qid.path 72
size 26214400
pdblknno 4066
pqpath 71
mtime 1653302180819962729
mode 664
uid 10006
gid -1
muid 10006
direct blocks
  0 4068
  1 4084
  2 4100
  .
  .
  31 4564
indirect blocks
  0 4596
  1 20964
  2 0
```

Block 4068 contents

```
Tdata 72
0 0123456789
.
.
.
```

starting contents
of 25MB.file

Block 4596 contents

```
Tind0 72
0 4580
1 4612
2 4628
.
.
```

Block 4580 contents

```
Tdata 72
89
.
.
```

more content
of 25MB.file

System Files	
Block	Description
0	magic
1	/adm/config
2	/adm/super
3	/adm
4	/adm/users
5	/adm/bkp/
6	/adm/users/inuse
7	/adm/frees
8	/adm/ctl (virtual file, empty contents)
9	/adm/users/staging
10	/

The /adm/ctl file is used to halt or sync the file system. /adm/users is a r/w file that will reload users when written to it. The owner of the /adm/ctl file or any user belonging to the sys group can read the disk.

There is no /adm/magic directory entry as the block number of the magic block is zero and zero block in a directory entry signifies the end of the directory contents.

Backup blocks

Three copies of Config, Super and Root blocks are maintained. This ensures two backups of config, Super and root blocks.

The backup block numbers on the disk are calculated during read based on the disk size.

Block	Description	Backup Blocks	
		1	2
1	/adm/config	last block number -1	middle block number -1
2	/adm/super	last block number -2	middle block number -2
10	/	last block number -3	middle block number -3

M[a]fs needs at least $N_{minblocks}=16$ blocks (8 KiB). The middle block number is $N_{minblocks} + ((nblocks - N_{minblocks})/2)$, where $nblocks$ = total number of blocks.

kfs and cwfs use 8192 byte blocks. Hence, they store multiple directory entries (Dentry) per block. They use slot numbers to identify a particular directory entry in a block of directory entries. M[a]fs avoids that by using 512 byte blocks thus having only one directory entry per block. This avoids locking up other sibling directory entries on access.

Users

Users are defined in /adm/users/inuse file. Any changes to it are made through the /adm/users/staging file. All changes are written to the staging file and then inuse file is updated by writing the command users to the /adm/ctl file.

Either all changes to /adm/users/inuse are installed or nothing is installed from the /adm/users/staging file.

The format of /adm/users/inuse is described in users(6).

Buffer cache – Hash buckets with a circular linked list of lobuf's for collisions.

An lobuf is used to represent a block in memory. An lobuf is unique to a block. All disk interaction, except for free block management, happens through an lobuf. We read a block from the disk into an lobuf. To update a block on the disk, we write to an lobuf, which, in-turn gets written to the disk.

An lobuf is protected by a read-write lock (RWlock). This ensures synchronization across multiple processes updating the same file.

getbuf(), putbuf(), putbufs() and putbuffree() are used to manage lobuf's. The contents of an lobuf is not touched unless it is locked by getbuf(). It is unlocked by putbuf(), putbufs() or putbuffree() calls. The lobuf.dirties Ref is decremented by the mafs writer's dowrite() without a lock(). This is to avoid deadlocks between putbuf() and the writer especially when the writer queue is full.

allocblock() allocates a free block into an lobuf. allocblocks() allocates a bunch of free blocks with their own lobuf's.

freeblock() erases the lobuf and returns the block to the free block management routines.

lobuf's are organized into a list of hash buckets to speed up access.

```
Hiob *hiob = nil;    /* array of nbuckets */
struct Hiob          /* Hash bucket */
{
    lobuf* link;      /* least recently used lobuf in the circular linked list */
    QLock;            /* controls access to this hash bucket */
};
struct Content /* used to unmarshall the disk contents */
{
    u8 tag;
    union
    {
        u8 buf[Blocksize];
        struct
        {
            u8 pad[7]; /* unused, to align to a multiple of 8 */
            u64 bufa[Nindperblock];
        };
    };
    u64 path; /* same as qid.path */
};
struct lobuf
{
    Ref;
    RWLock;          /* controls access to this lobuf */
};
```

```
u64 blkno;          /* block number on the disk, primary key */
lobuf  *fore;        /* for lru */
lobuf  *back;        /* for lru */
union{
    u8  *xiobuf;     /* "real" buffer pointer */
    Content *io;     /* cast'able to contents */
    Dentry *d;
};
/*
```

This field is used by mafs to ensure that lobufs are not reused while there are pending writes.

dowrite() uses a Ref instead of a wlock() to mark lobuf's with pending writes.

Using a wlock() in dowrite() causes a deadlock with putwrite() especially when the writer queue is full.

getbuf() guarantees that even a free'd block cannot be stolen until the dirties == 0. This avoids dirty blocks being stolen for other block numbers.


incrcf(dirties) only happens while holding a wlock() in putwrite().

```
*/
Ref  dirties;        /* number of versions of this block yet to be written by the writer */
};
```

The lobuf's are arranged into a list of hash buckets. Each bucket points a circular linked list of lobuf's to handle collisions. If all the lobuf's in the circular linked list are locked, new lobuf's are added to this linked list. This circular list is ordered on a least recently used basis. lobuf's once added to this list are not removed. When an lobuf is not in the list, the oldest unlocked lobuf is reused.

Hiob hiob[nbuckets] is a valid representation of the list of hash buckets. The block number is hashed to arrive at the relevant hash bucket index.

hiob[hash(block number)].link = Address of lobuf0, where lobuf0 is the least recently used lobuf.



The size of the buffer cache is: number of hash buckets * collisions per hash bucket * block size. The approximate size of the buffer cache = Nbuckets * Ncollisions * Raw-blocksize = 256 * 10 * 512 bytes = 1.28GiB. The -h parameter can be used to change the number of hash buckets.

If you have RAM to spare, increase Nbuckets instead of Ncollisions as the hash index lookup is faster than searching through a linked list.

lobuf.Ref is used to avoid locking up the hash bucket when a process is waiting for a lock on an lobuf in that hash bucket.

lobuf.Ref ensures that an lobuf is not stolen before another process can get to wlock()'ing it after letting go of the lock on the hash bucket. We cannot hold the lock on the hash bucket until we wlock() the lobuf as that blocks other processes from using the hash bucket. This could also result in a deadlock. For example, the directory entry is block 18, which hashes to a hash index of 7. A writer() locked the directory entry lobuf and wants to add a data block 84 to the directory entry. Block 84 hashes to the same hash index of 7. Another process wanting to access the directory entry is waiting for a lock on that io buffer. While doing so, it has locked the hash bucket. Now, this has caused a deadlock between both these processes. The first process cannot proceed until it can lock the hash bucket holding block 84 and is still holding the lock on the directory entry in block 18. The second process cannot lock block 18 and is holding the lock on the hash bucket.

for locking a buffer:

```
qlock(hash bucket); incref(buffer); qunlock(hash bucket);
wlock(buffer); decref(buffer);
```

for stealing an unused buffer:

```
qlock(hash bucket);
find a buffer with ref == 0 and wlock()'able.
qunlock(hash bucket);
```

for unlocking a buffer:

```
wunlock(buffer);
```

Asynchronous writes of Mafs

The blocks to be written to a disk are stored in a linked list represented by:

struct Dirties

```
{
    QLock lck;           /* controls access to this writer queue */
    Wbuf *head, *tail;   /* linked list of dirty blocks yet to be written to the disk */
    s32 n;                /* number of dirty blocks in this linked list */
    Rendez isfull;        /* write throttling */
    Rendez isempty;       /* writer does not have to keep polling to find work */
} drts = {0};
```

struct Wbuf

```
{
    u64 blkno;           /* block number on the disk, primary key */
    Wbuf *prev, *next;   /* writer queue */
    lobuf *lobuf;         /* pointer to the used lobuf in the buffer cache */
    union{
        u8 payload;      /* "real" contents */
        Content io;       /* cast'able to contents */
    };
};
```

A single writer process takes the blocks from the Dirties linked list on a FIFO (first-in-first-out) basis and writes them to the disk. putbuf() and putbufs() add blocks to the end of this linked list, the writer queue.

The dirty blocks not yet written to the disk remain in the buffer cache and cannot be stolen when a need for a new lobuf arises.

Free'd blocks are not written to the disk to avoid writing blanks to a disk.

The writer throttles input when there are more than `Npendingwrites` waiting to be written. This can be adjusted with the `-w` parameter.

The alternative to having a single writer process is to have each worker process write to the disk, as mfs does. Synchronous writes throttle writes to disk write speed. With asynchronous writes, memory is used to hold the data until written to the disk. This shows increased write throughput until we fill up memory. After filling up memory, writes happen at disk speed. Asynchronous writes have the side effect of a single disk write queue.

The ideal `npendingwrites` = $((\text{ups time in seconds})/2) * (\text{disk speed in bytes/second}) / \text{Rawblocksize}$.

Free blocks

Free blocks are managed using Extents. The list of free blocks is stored to the disk when shutting down. If this state is not written, then the file system needs to be checked and the list of free blocks should be updated.

When shutting down, the Extents are written to free blocks. This information can be accessed from `/adm/frees`. Also, `fsok` in the super block is set to 1. M[a]fs does not start until `fsok` is 1. When `fsok` = 0, run the sanity check that the unused blocks and the free blocks in `/adm/frees` match up. `disk/reconcile` identifies any missing blocks or blocks that are marked as both used and free.

This process of fixing issues and setting `fsok` to 1 is manual. There is no automatic file system checker as in other file systems. This document aims to empower you with the knowledge to fix your file system issues instead of entrusting your precious data to an arbitrary decision maker such as the file system checker.

A tag of `Tfree` and `Qpnone` represent a free block. If a directory entry is removed, the parent will have a zero'ed out child directory entry (`Qid.path` = 0) and a tag of `Tdentry` and `Qpnone`.

Extents

Free blocks and memory are managed using Extents, an abstraction used to manage a continuous list of items.

An Extent represents a continuous list of items. An Extents is a list of such Extent's.

```
struct Extent {
    struct Extent *low, *high; /* sorted by start */
    u64 start;                /* where this extent starts from */
    u64 len;                  /* how many units in this extent */

    /* circular least recently used linked list limited to Nlru items */
    struct Extent *prev, *next;
```

```
};
struct Extents {
    Extent *head; /* find the first block in a jiffy */
    QLock lck;
    u32 n; /* number of extents */
    Rendez isempty; /* fully used, nothing available */

    u8 nlru; /* number of items in the lru linked list */
    Extent *lru; /* least recently used extent in the circular lru linked list */
};
```

To allocate n items from Extents, we find the lowest (by block number or memory address) extent that can satisfy our request. If a bigger Extent is available, slice it and take the portion we need.

If there is no available Extent to satisfy our request, panic().

allocblock() and freeblock() use balloc() and bfree() respectively. balloc() assigns blocks from an extent and bfree() adds them to an extent for next allocation.

Extents at memory location 1

lru	100
el	0
n	3

assuming that the Extent at 100 was used last
unlocked

Extent at 100

blkno	10
len	1

low	0
high	200

Extent at 200

blkno	20
len	3

low	100
high	300

Extent at 300

blkno	30
len	2

low	200
high	0

+ freed block numbers =
11,12,13,14

Extent at 100

blkno	10
len	5

low	0
high	200

Extent at 200

blkno	20
len	3

low	100
high	300

Extent at 300

blkno	30
len	2

low	200
high	0

Extents before				Extents after		
↓	blkno	len		↓	blkno	len
	20	3	+		20	3
					40	4

Block number 40 followed by 3 free blocks =

Extents before				Extents after		
↓	blkno	len		↓	blkno	len
	100	5	+		100	13
	110	3				

Block number 105 followed by 4 free blocks =

Extents before				Extents after		
↓	blkno	len		↓	blkno	len
	105	4	+		101	8

Block number 101 followed by 3 free blocks =

Extents before				Extents after		
↓	blkno	len		↓	blkno	len
	101	4	+		100	8

Block number 105 followed by 3 free blocks =

Extents before				Extents after		
↓	blkno	len		↓	blkno	len
	180	4	+		180	4
					250	4

Block number 250 followed by 3 free blocks =

Extents before				Extents after		
↓	blkno	len		↓	blkno	len
	250	4	+		180	4
					250	4

Block number 180 followed by 3 free blocks =

Kfs stores the list of free blocks in a Tfree block and the Superblock. Instead we use block management routines, similar to pool.h, to allocate and monitor free blocks. On shutdown(), the block management routines (extents.[ch]) store state into the free blocks. This can be read from /adm/frees. On startup, this is read back by the block management routines. On a crash, the fsck can walk the directory structure to identify the free blocks and recreate /adm/frees.

Code details

Program	Description
disk/mfs	Start mfs on a disk.
disk/mafs	Start mafs on a disk.
disk/free	List the free blocks. It reads the contents of /adm/frees.
disk/used	List the used blocks by traversing all directory entries.
disk/block	Show the contents of a block.
disk/unused	Given a list of used blocks, lists the unused blocks.
disk/updatefrees	Update the contents of /adm/frees.

File	Description	chatty9p
9p.c	9p transactions	2
blk.c	routines to show blocks.	
console.c	obsolete. /adm/ctl is the console.	
ctl.c	/adm/ctl operations.	
dentry.c	encode/decode the file system abstraction into block operations.	3
extents.[ch]	routines to manage the free blocks.	6
iobuf.c	routines on lobuf's. The bkp() routines operate on lobuf's.	5
sub.c	initialization and super block related routines.	2
tag.c	routines to manage a relative index (reli) in a directory entry.	
user.c	user management routines.	
writer.c	disk writer routines.	

A Chan's state could get out of sync with the contents if another process changes the on-disk state. Ephase error occurs when that happens.

For throughput, multiple processes are used to service 9p i/o requests when the -s flag is not used.

Useful commands:

disk/mfs and disk/mafs have the same arguments. The following commands use disk/mafs to avoid duplicating them for disk/mfs.

Ream and start single process M[a]fs on a disk and also mount it for use.

```
mount -c <{disk/mafs -s -r mafs_mysevice mydisk <[0=1]} /n/mafs_mysevice
-s: use stdin and stdout for communication
```

-r mafs_mysevice: ream the disk using mafs_mysevice as the service name
mydisk: running mafs on the disk, mydisk

Ream and start multiple-process mafs on a disk.

```
disk/mafs -r mafs_mysevice mydisk
mount -c /srv/mafs_mysevice /n/mafs_mysevice
```

Ream and start mafs on a file. Also, mount the filesystem at /n/mafs_mysevice.

```
dd -if /dev/zero -of myfile -bs 512 -count 128 # 64KB file
mount -c <{disk/mafs -s -r mafs_service myfile <[0=1]} /n/mafs_mysevice

# to reuse the contents of myfile later, remove -r (ream) from the above command.
mount -c <{disk/mafs -s myfile <[0=1]} /n/mafs_mysevice
```

Prepare and use a disk (/dev/sdF1) for mafs.

```
disk/fdisk -bawp /dev/sdF1/data # partition the disk
echo '
a fs 9 $-7
w
p
q' | disk/prep -b /dev/sdF1/plan9 # add an fs plan 9 partition to the disk
disk/mafs -r mafs_sdF1 /dev/sdF1/fs # -r to ream the disk
mount -c /srv/mafs_sdF1 /n/mafs_sdF1

# for using the mafs file system on the disk later on
disk/mafs /dev/sdF1/fs # no -r
mount -c /srv/mafs_sdF1 /n/mafs_sdF1
```

Starting mafs on a 2MB byte file. The below commands create a disk.file to use as a disk. Mount /n/mafs_disk.file for the file system.

```
dd -if /dev/zero -of disk.file -bs 512 -count 4096;
mount -c <{disk/mafs -s -r mafs_disk.file \
    <[0=1]} /n/mafs_disk.file
```

Starting mafs on a RAM file. The below commands create a ramfs filesystem to use as a disk.

```
ramfs -m /n/mafs_ramfs
touch /n/mafs_ramfs/file
dd -if /dev/zero -of /n/mafs_ramfs/file -count 700 -bs 1m
disk/mafs -r mafs_ramfs_file /n/mafs_ramfs/file
mount -c /srv/mafs_ramfs_file /n/mafs_ramfs_file
```

Sync M[a]fs. This command does not return until all the writes are written to the disk. So, could take a long time if you have a long writer queue.

```
echo sync >> /n/mafs_mysevice/adm/ctl
```

Stop M[a]fs. This command does not return until all the writes are written to the disk. So, could take a long time if you have a long writer queue.

```
echo halt >> /n/mafs_myservice/adm/ctl
```

Interpret the contents of a block based on the tag and write out a single formatted block based on the tag

```
disk/block tests/test.0/disk 22
```

Traverse the directory heirarchy and write out all the used block numbers. disk/reconcile uses the output of this to reconcile the list of used blocks with the list of free blocks. Also, writes the invalid blocks to stderr. Starting from root, walk down each directory entry printing out the linked blocks with invalid tags. (Why not just write out the list of dirty blocks too? instead of using a different command for it?)

```
disk/used tests/test.0/disk
```

From the contents of /adm/frees show the list of free blocks. disk/reconcile uses the output of this to reconcile the list of used blocks with the list of free blocks.

```
disk/free tests/test.0/disk
```

Read two lists of block numbers and flag the common and missing blocks.

```
disk/reconcile -u <{disk/used tests/test.0/disk} \  
-F <{disk/free tests/test.0/disk} 32
```

Find traverses the directory heirarchy and identifies the file that a block number belongs to.

```
disk/find tests/test.0/disk 17
```

Find the total number of blocks on a disk.

```
dd -if /dev/sdF1/fs -bs 512 -iseek 1 -count 1 -quiet 1 | awk '$1 == "nblocks" { print $2 }'
```

```
disk/block /dev/sdF1/fs 1 | awk '$1 == "nblocks" { print $2 }'
```

Build the list of free blocks. This should match the contents of /adm/frees.

```
disk/unused <{disk/used /dev/sdF1/fs} 11721040049 # 11721040049 = total number of disk blocks  
disk/unused <{disk/used test.0/disk} 32 # 32 = total number of disk blocks
```

Change the contents of /adm/frees.

```
disk/updatefrees tests/test.0/disk <{disk/unused <{disk/used tests/test.0/disk} 32}  
disk/updatefrees /dev/sdF1/fs <{disk/unused <{disk/used /dev/sdF1/fs} 11721040049}
```

A sanity check that the file system is not corrupt by comparing that the unused blocks and free blocks match up. \$nblocks is the total number of disk blocks. \$disk is the disk.

```
diff <{disk/unused -l <{disk/used tests/test.0/disk} 32}} <{disk/free tests/test.0/disk}
```

Changing the service name without a ream.

```
disk/block /dev/sdF1/fs 1 | wc
```

```
Tdata 2
size 6001172505088
nblocks 11721040049
backup config 1 to 11721040048 5860520032
backup super 2 to 11721040047 5860520031
backup root 3 to 11721040046 5860520030
service mafs_ddf_1
```

```
dd -if /dev/sdF1/fs -count 10 -skip 682 -bs 1
mafs_ddf_110+0 records in
10+0 records out
```

```
dd -if <{echo m_ddf_1; cat /dev/zero} -of /dev/sdF1/fs -count 11 -oseek 682 -bs 1
7+0 records in
7+0 records out
```

```
disk/block /dev/sdF1/fs 1
Tdata 2
size 6001172505088
nblocks 11721040049
backup config 1 to 11721040048 5860520032
backup super 2 to 11721040047 5860520031
backup root 3 to 11721040046 5860520030
service m_ddf_1
```

Changing the magic phrase in the magic block.

```
disk/block /dev/sdF1/fs 0
Tmagic 1
mafs device
512
```

```
dd -if /dev/sdF1/fs -count 16 -iseek 256 -bs 1
mafs device
512
20+0 records in
20+0 records out
```

```
dd -if <{echo m[a]fs device; echo 512; cat /dev/zero} -of /dev/sdF1/fs -count 18 -oseek 256 -bs 1
18+0 records in
18+0 records out
```

```
dd -if /dev/sdF1/fs -count 18 -iseek 256 -bs 1
m[a]fs device
512
18+0 records in
18+0 records out
```

```
disk/block /dev/sdF1/fs 0
Tmagic 1
m[a]fs device
512
```

Tests

Program	Description
tests/regress.rc	All regression tests
tests/chkextents.rc	Unit tests on extents
tests/chkreli.rc	Unit tests on relative index lookups
tests/6.offsets	Write file using different offsets to test mafswrite()
tests/6.sizes	Show the effects of the different parameters
tests/6.testextents	Test extents.[ch] state changes
tests/6.reli	Translate relative index to block number on a disk

The below disk state tests:

1. Initialize a disk for mafs.
2. Run mfs or mafs on that disk.
3. Stop mfs or mafs.
4. Compare the contents with the expected contents (tests/test.0/blocks/*).

Disk State	
Test	Description
tests/test.0	empty disk
tests/test.1	create a file /dir1/file1 and echo test into it
tests/test.2	writes at different offsets to a file and then removes the file
tests/test.3	write, read and delete files with sizes upto 16384 blocks
tests/test.4	directory copy
tests/test.5	fcg gzipped files
tests/test.6	df
tests/test.7	multiple processes working on the filesystem simultaneously
tests/test.8	check backup blocks locations
tests/test.9	examples used by this document
tests/test.a	write, read and delete a 100MB file
tests/test.b	duplicate of test.2 but seeded with random data
tests/test.d	seed with random data and do mkdir -p a/b/c/d/e/f/g/h
tests/test.e	seed with random data and test directory and file deletions

Extents behaviour	
Test	Description
tests/extents/addabove	Figure 1 of the Extents section
tests/extents/addabove1	Figure 2 of the Extents section
tests/extents/addbelow	Figure 3 of the Extents section
tests/extents/mergeabove	Figure 4 of the Extents section
tests/extents/mergenext	Figure 5 of the Extents section
tests/extents/mergeprevious	Figure 6 of the Extents section

To run all the regression tests:

```
cd tests/; ./regress.rc
```

To loop through all the blocks of a test:

```
for(t in tests/test.2/blocks/^{seq 0 39}*){ echo $t; echo '-----'; cat $t; echo }
```

Performance metrics

```
ramfs -m /n/ramfs
touch /n/ramfs/file
cat /dev/zero | tput -p > /n/ramfs/file
```

```
196.00 MB/s
198.76 MB/s
187.58 MB/s
176.96 MB/s
175.87 MB/s
180.42 MB/s
183.52 MB/s
185.99 MB/s
187.96 MB/s
189.54 MB/s
190.83 MB/s
191.89 MB/s
192.80 MB/s
```

```
dd -if /dev/zero -of /n/ramfs/file -count 700 -bs 1m
```

```
disk/mfs -r mfs_ramfs_file /n/ramfs/file
mount -c /srv/mfs_ramfs_file /n/mfs_ramfs_file
cat /dev/zero | tput -p > /n/mfs_ramfs_file/zeros.file
```

```
31.86 MB/s
31.53 MB/s
33.68 MB/s
34.98 MB/s
35.69 MB/s
36.13 MB/s
36.65 MB/s
37.17 MB/s
37.47 MB/s
36.58 MB/s
36.67 MB/s
36.93 MB/s
37.20 MB/s
37.55 MB/s
37.26 MB/s
36.91 MB/s
36.54 MB/s
36.25 MB/s
36.07 MB/s
```

```
echo halt >> /n/mfs_ramfs_file/adm/ctl; lc /srv
umount /n/mfs_ramfs
```

```
disk/mafs -r mafs_ramfs_file /n/ramfs/file
mount -c /srv/mafs_ramfs_file /n/mafs_ramfs_file
cat /dev/zero | tput -p > /n/mafs_ramfs_file/zeros.file # increase memunits for speed
57.94 MB/s
55.27 MB/s
```

```
46.67 MB/s
41.87 MB/s
39.00 MB/s
36.72 MB/s
35.18 MB/s
33.89 MB/s
32.81 MB/s
echo halt >> /n/mafs_ramfs_file/adm/ctl; lc /srv

dd -if /dev/zero -of /n/ramfs/file -count 700 -bs 1m
hjfs -f /n/ramfs/file -r
echo allow >> /srv/hjfs.cmd
mount -c /srv/hjfs /n/hjfs/
cat /dev/zero | tput -p > /n/hjfs/zeros.file
70.85 MB/s
71.02 MB/s
70.88 MB/s
70.62 MB/s
70.46 MB/s
70.54 MB/s
69.74 MB/s
68.63 MB/s
67.74 MB/s
67.02 MB/s
63.81 MB/s
echo halt >> /srv/hjfs.cmd
umount /n/ramfs
```

Limitations

As we use packed structs to store data to the disk, a disk with m[a]fs is not portable to a machine using a different endian system.

Design considerations

For exclusive use (mode has DMEXCL bit set) files, there is no timeout.

Use an fs(3) device for RAID or other configuration involving multiple disks.

Why are you not using a checksum to verify the contents?

Checksums are probabilistic and can be implemented as a bespoke application instead of complicating the file system implementation.

Source

<http://git.9front.org/plan9front/mafs/HEAD/info.html>

References

[1] Sean Quinlan, "A Cached WORM File System," Software--Practice and Experience, Vol

21., No 12., December 1991, pp. 1289-1299

[2] Ken Thompson, Geoff Collyer, "The 64-bit Standalone Plan 9 File Server"