

Formalising Filesystems in the ACL2 Theorem Prover

An Application To FAT32

Mihir Mehta

Department of Computer Science
University of Texas at Austin

`mihir@cs.utexas.edu`

05 November, 2018

Why filesystem verification matters

- ▶ Filesystems form the basis of our current computing paradigm.
- ▶ They help application programmers address data by pathnames (not numbers), and additionally address security/efficiency/redundancy concerns which the average dev doesn't want to be involved in.
- ▶ Thus, verifying the properties of filesystems in common use, and thereby making them reliable, is critically important.
 - ▶ FAT32 - once widely used on Windows, and still used by a large number of embedded systems - qualifies.

The plan

- ▶ Modelling a real, widely used filesystem in ACL2 - FAT32
- ▶ Verification through refinement
- ▶ Binary compatibility and execution efficiency
- ▶ Co-simulation tests for accuracy

Outline

FAT32

The models

Proofs and co-simulation

Related and future work

Our FAT32 model aims to have ...

- ▶ ... the same space constraints as a FAT32 volume of the same size.
- ▶ ... the same success and failure conditions for file operations, and the same error codes for the latter.
- ▶ ... a way to read a FAT32 disk image from a block device, and a way to write it back.
 - ▶ This is made easier by choosing to replicate the on-disk data structures of FAT32 in the model.

File operations in our model


- ▶ File operations categorised into *read operations*, which do not change the state of the filesystem, and *write operations* which do.
- ▶ Generic signature for read operations:
(read fs-inst) \mapsto (mv ret-val status errno)
- ▶ Generic signature for write operations:
(write fs-inst) \mapsto (mv fs-inst ret-val status errno)

The FAT32 specification

In a FAT32 volume, the unit of data storage is a *cluster* (also known as an *extent*). There are three on-disk data structures.

- ▶ *reserved area*, volume-level metadata such as the size of a cluster and the number of clusters.
- ▶ *file allocation table*, collection of *clusterchains* (linked lists of clusters), one for each regular file/directory file.
- ▶ *data region*, collection of clusters.

A FAT32 Directory Tree



	directory entry in /
0	"vmlinuz", 3
32	"initrd.img", 5
64	"tmp", 6
⋮	⋮

	directory entry in /tmp/
0	"ticket1", 7
32	"ticket2", 8
⋮	⋮

FAT index	FAT entry
0 (reserved)	
1 (reserved)	
2	<i>eoc</i>
3	4
4	<i>eoc</i>
5	<i>eoc</i>
6	<i>eoc</i>
7	<i>eoc</i>
8	<i>eoc</i>
9	0
⋮	⋮

Outline

FAT32

The models

Proofs and co-simulation

Related and future work

Abstract models

- ▶ Bootstrap - begin with abstract filesystem models, in order to explore the properties we require in a FAT32 model.
- ▶ Incrementally add the desired properties in a series of models
- ▶ Wherever possible, capture common features expected to exist in different filesystems.

Abstract models in brief

L1	Filesystem is a literal directory tree; contents of regular files are represented as strings stored in the nodes.
L2	A single element of metadata, <i>length</i> , is stored within each regular file.
L3	Regular files are divided into fixed-size blocks; blocks are stored in an external “disk” data structure; storage for these blocks remains unbounded as in L1 and L2.
L4	Disk size is now bounded; allocation vector data structure is introduced to help allocate and garbage-collect blocks.
L5	Additional metadata for file ownership and access permissions is stored within each regular file.
L6	Allocation vector is replaced by a file allocation table.

Beginning to model FAT32

Next, in models M1 and M2, we model FAT32 more concretely, providing the standard POSIX system calls.

- ▶ M1 - another tree model, but with directory entries exactly matching the FAT32 spec.
- ▶ M2 - stobj model with fields for all the metadata in the reserved area and arrays for the file allocation table and data region.

This way, we benefit from efficient stobj array operations in M2, and we can simplify our reasoning in M1 by continuing with directory trees.

Outline

FAT32

The models

Proofs and co-simulation

Related and future work

Read-after-write proofs

- ▶ *Read-over-write* properties show that write operations have their effects made available immediately for reads at the same location, and also that they do not affect reads at other locations.
- ▶ We've proved these properties for the abstract models L1-L6, and we've also proved them for our concrete models M1 and M2, with the caveat that the transformations between M1 and M2 are not yet verified.

Refinement proofs

- ▶ For the abstract models, we started by proving the read-over-write properties *ab initio* for L1.
- ▶ For each subsequent model in L2-L6, we proved a refinement relationship where possible, or an equivalence where a strict refinement did not hold, with a previous model and used it to prove read-over-write properties by analogy.
- ▶ An illustration of such a proof follows.

Proof example: first read-over-write in L2




Figure: l2-wrchs-correctness-1 (write is overloaded for L2 and L1)




Figure: l2-rdchs-correctness-1 (read is overloaded for L2 and L1)

Proof example: first read-over-write in L2


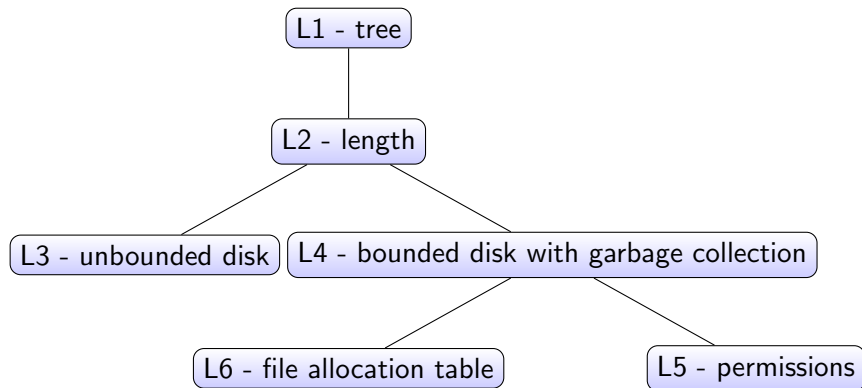


Figure: l2-read-over-write-1

Relationships between abstract models



Co-simulation

- ▶ Ensure that our implementation lines up with FAT32, the target filesystem.
- ▶ POSIX system calls supported - `lstat`, `open`, `pread`, `pwrite`, `close`, `mkdir` and `mknod`.
- ▶ Wherever `errno` is to be set, do what Linux does.
- ▶ Compare the output of our ACL2 programs (based on the FAT32 model) with the utilities (such as `cp` and `mkfs`) which they replicate.

Outline

FAT32

The models

Proofs and co-simulation

Related and future work

Related work - interactive theorem provers

- ▶ Bevier and Cohen - Synergy FS, executable model with processes and file descriptors, but no read-over-write theorems (ACL2).
- ▶ Klein et al. - COGENT, verifying compiler from a DSL to C code for a filesystem (Isabelle/HOL).
- ▶ Chen - FSCQ, high-performance filesystem with verified crash consistency properties (Coq).

Related work - non-interactive theorem provers

- ▶ Hyperkernel - microkernel with system calls simplified until the point where useful properties can be proved through SMT solving (Z3).
- ▶ Yggdrasil - filesystem verification through SMT solving, but constrained from modelling important features such as extents (Z3).

Future work

- ▶ Model the remaining POSIX system calls for FAT32 and use them to reason about sequences of file operations (i.e. do code proofs).
- ▶ Reuse FAT32 verification artefacts for a filesystem with crash consistency, for instance, ext4.
- ▶ Model concurrent file operations in a multiprogramming environment.

Recent progress

- ▶ Set of supported POSIX system calls expanded.
- ▶ Set of co-simulation tests, mostly based on coreutils programs, expanded based on these.
- ▶ Functions for converting M2 instances to FAT32 disk images and back proved to be inverses of each other.
- ▶ Equivalence relation developed to allow two FAT32 disk images to be compared modulo rearrangement of data and reordering of files within directories.
 - ▶ This gives us a means to co-simulate programs which modify filesystem state, such as `mv` and `rm`.

Conclusion

- ▶ FAT32 formalised, demonstrating the applicability of the refinement style to filesystem verification.
- ▶ Co-simulation infrastructure developed to compare filesystem models to a canonical implementation, such as that of Linux.
- ▶ FAT32's allocation and garbage collection algorithms certified.