Model Checking Software Systems: A Case Study

Jeannette M. Wing*and Mandana Vaziri-Farahani†
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213

Abstract

Model checking is a proven successful technology for verifying hardware. It works, however, on only finite state machines, and most software systems have infinitely many states. Our approach to applying model checking to software hinges on identifying appropriate abstractions that exploit the nature of both the system, S, and the property, ϕ, to be verified. We check ϕ on an abstracted, but finite, model of S.

Following this approach we verified three cache coherence protocols used in distributed file systems. These protocols have to satisfy this property: "If a client believes that a cached file is valid then the authorized server believes that the client's copy is valid." In our finite model of the system, we need only represent the "beliefs" that a client and a server have about a cached file; we can abstract from the caches, the files' contents, and even the files themselves. Moreover, by successive application of the generalization rule from predicate logic, we need only consider a model with at most two clients, one server, and one file. We used McMillan's SMV model checker; on our most complicated protocol, SMV took less than 1 second to check over 43,600 reachable states.

Keywords: model checking, verification, finite state machines, abstraction mappings, distributed systems, cache coherence protocols

*Computer Science Department, wing@cs.cmu.edu.
†Electrical and Computer Engineering Department, mv1k+@cs.cmu.edu.

0

1 Motivation: Theorem Proving and Model Checking

Software systems keep growing in size and complexity. Many large, complex software systems must guarantee certain critical functional, real-time, fault-tolerant, and performance properties. Proving that such a system satisfies these kinds of properties can increase our confidence that it will operate correctly and reliably. Proofs based on formal, rather than informal, techniques make our reasoning precise; moreover, they are amenable to mechanical aids such as syntax and semantics checkers.

Formal reasoning entails comparing two formal objects, e.g., establishing the correctness of a program with respect to a specification or showing that one concurrent process simulates another. The starting point is having two formal objects. There are two general approaches to showing the correspondence between these two objects: theorem proving and model checking. We argue that model checking should and will play a larger role in reasoning about software systems than it does today.

The traditional approach to formal reasoning about software is program verification where one formal object is the *program text* and the other is a *specification* written in some mathematical logic. The formal technique used to show a correspondence between the two objects is based on *theorem proving*. Over time this approach has been shown to work on increasingly larger and larger programs, especially as the tool support like theorem provers and proof checkers has become more and more sophisticated. Yet, it still has drawbacks:

- The size of a program we can prove correct is on the order of only a couple thousand lines of code [13].
- To do such a proof requires highly-skilled people, such as theorem-proving experts, domain experts, or both.

- The human time to do such a proof is on the order of months or even years; the machine time, on the order of hours [13].
- In the course of such a proof, we are often forced to prove "obvious" or "uninteresting" theorems; the amount of effort to prove them is often the same as that for proving the "essential" property of interest.

We believe that there is a time and place for program verification, e.g., for key components of a safety-critical system. In this paper, however, we directly address the concerns of the vast majority of the software engineering community, which questions whether the cost in time, effort, and resources for program verification is worth the eventual benefit gained.

We suggest a radically different tack: *model checking*. The two formal objects compared are a *finite state machine model* of the software system, and as before, a specification written in some mathematical logic.

Model checking is a proven successful verification technology in the hardware community. For example, it has been used to find bugs in published circuits [4, 14], the cache coherence protocol for the Encore Gigamax multiprocessor [23], the IEEE Futurebus+ Standard [9], and telephone switching systems [15]. It has been used to prove safety and liveness properties of the T9000 virtual channel processor [3].

Fundamental to model checking is its reliance on *finite* state machines. Model checking exploits this finiteness property by performing an exhaustive case analysis of the machine's set of states. Recent technological advances have greatly improved the ability to apply this technique to real systems: model checkers can now check systems on the order of 10^{120} states, and for some systems this number can be as large as 10^{1300} [10].

The rationale behind why theorem proving has been the primary approach for reasoning about software is that software systems are, in general, *infinite* state machines. We thus rely on induction to prove in a finite number of steps a property over infinite domains. Model checking, at first glance, seems inappropriate.

There are three methodological reasons for why model checking *is* appropriate. First, the inductive arguments used in proof work well for highly structured components (e.g., generic data types like sets and mappings), but fail at the system level, because of discontinuities in value spaces and irregularities in software structure. We are forced to resort to huge case analyses anyway, perhaps with inductive proofs performed only locally. Thus, though it may seem restrictive to use model checking because we cannot prove something about all possible values drawn from an infinite domain, it is exactly the kind of technology needed to handle the huge case analyses at the system level.

Second, checking a model of the system rather than the system itself raises the level of abstraction at which we do our formal reasoning. Though, we may fall short of doing "exact" reasoning about the original system, we can more quickly, with less effort, and completely automatically do "approximate" reasoning. (An argument must be made, of course, that the model checked is not so abstract that it trivially satisfies the property of interest.)

Finally, as the hardware community has discovered, model checking has been tremendously successful at finding bugs in hardware designs. It is more common to find that a system has errors than that it is correct. The same is true, if not more so, for software. Thus, model checking can help software designers find bugs in their designs, where a design is a natural abstraction of the actual working system. Moreover, if the goal of formal reasoning is to find bugs, then it matters less that we do only "approximate," rather than "exact," reasoning.

Thus, though theorem proving has its place, e.g., for doing inductive arguments and low-level program verification, model checking can complement theorem proving efforts. It is worth exploring all avenues as to how.

We present the gist of our approach in Section 2. We give details of a case study in Section 3; we used model checking to verify three cache coherence protocols for distributed file systems: two implemented for the Andrew File System (AFS) and one for the Coda File System. In Section 4, we use the case study to illustrate how we followed our approach; we explain different kinds of specific abstractions that software engineers can in general apply to their systems to produce finite state machine models. We close with related work and conclusions.

2 One Approach: Finding Good Abstractions

The successes in reasoning about hardware systems raise the obvious question: *How can model checking be applied to reasoning about software systems?* In this paper, we elaborate on this answer:

Approach: Model check a finite state machine *abstraction* of a software system.

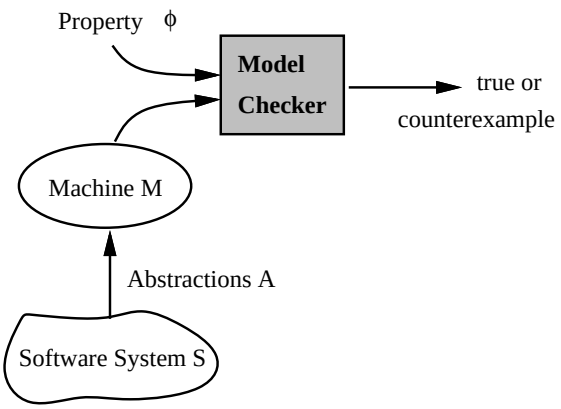

Figure 1: Model Checking Software

This approach (see Figure 1) relies on finding mappings, A, to apply to a possibly infinite state machine, S, and then subjecting the abstract model, M, to a model checker. We use a model checker as a black box to check M against the specification, ϕ. The model checker outputs either true, if M satisfies ϕ, or a counterexample, if it does not.

Key to our approach is to exploit the nature of ϕ to determine what abstractions are reasonable to define and apply. Based on our case study, we broadly classify the ones we identified as follows (elaborated on in Section 4):

- Exploiting the *form* of ϕ. In our example, because of the form of our correctness condition, we use the generalization rule of first-order logic to justify that checking a finite case suffices to show that ϕ holds for the infinite case.
- Exploiting *domain-specific knowledge*. In our distributed systems domain, we collapse distinct failures like crashed nodes and links into a single type of failure. Also, we place a bound on transmission delay, and furthermore model it as a single step taken by the system's environment. Both abstractions are reasonable since implementors of distributed systems make similar simplifications.
- Exploiting *problem-specific knowledge*. ϕ might express an abstract property about an object, not its value. For example, suppose ϕ is a property about the size of a bounded integer set; we do not care about the value of the set itself. If the bound is small, we can model each possible size.

The first kind of abstraction can be applied to any problem or domain. The second can be applied when considering any problem for distributed systems since failures and timing behavior cannot be completely ignored in those systems. The third class of abstractions, though particularly appropriate for our case study, are just examples of abstractions that are generally applicable to other problems. All can be viewed as software *design rules* ("-of-thumb") that when applied raise the level at which we can think about the essence of a software system.

3 Case Study: Cache Coherence Protocols for Distributed File Systems

In a distributed file system, *servers* store files; *clients* store local copies of these files in their caches. Caching increases performance at the client when connectivity and bandwidth are low, and increases availability and reliability when temporary failures occur. Clients communicate with servers by exchanging messages and data (e.g., files). Clients do not communicate with each other. Each file is associated with exactly one designated "home" or *authorized* server.

A problem arises when there are several copies of a file in a system. A file is *valid* if it is the most recent copy in the system. Recency is typically determined by a timestamp associated with the file. If a client updates its copy and it is the most recent update in the system, then all other copies of that file become invalid. The goal of a cache coherence protocol is to make sure a client performs work on only files it believes are valid.

There are two ways to ensure cache coherence. Either the client asks the server whether its copy is valid (*validation-based*) or the server tells the client when the client's copy is no longer valid (*invalidation-based*). A cache coherence protocol defines a set of possible *runs*, i.e., message exchanges, between clients and servers about files. Before a run a client considers all files in its cache suspect, which means that it does not have a belief on their validity. During a run clients and servers gain beliefs about cached files. At the end of the run, they all discard all their beliefs. If failures occur, beliefs are discarded but clients do not discard their cached files. Failures are detected by message timeouts.

Cache coherence in a distributed system is more difficult to achieve than on a multiprocessor because of the presence of failures and transmission delay. Thus, since global knowledge is impossible to achieve in a distributed system [18], we settle for pairwise

knowledge between clients and servers. Our notion of belief captures this pairwise knowledge [25].

An invariant property to prove of all cache coherence protocols is that if a client believes that a cached file is valid, then the server that is the authority on the file believes the client's copy is valid. More formally,

$$CC: \ \forall\, C : client\ .\ \forall\, S : server\ .\ \forall\, f : file\ .$$
$$C\ believes\ valid(f_C)\ \Rightarrow\ S\ believes\ valid(f_C)$$

where f_C stands for the copy of f at C [25].

A validation-based technique, AFS-1, was used in the Andrew File System from 1984-1985 [27]; for performance reasons, an invalidation-based technique, AFS-2, replaced AFS-1 and has been in use in Andrew since 1985 [20]. In 1993 Mummert, as part of her Ph.D. thesis work, started implementing a more complicated invalidation-based cache coherence protocol [24], similar to AFS-2, as a variation for the Coda Distributed File System [28]. It was this work that inspired our initial investigation of this case study since Mummert wanted a way to prove formally that her protocol design was correct. Inspired by the logic of authentication [6], Mummert et al. [25] formalized the notions of belief and validity, as used above, and applied the extended logic to reason about cache coherence for AFS-1, AFS-2, and Mummert's variation of Coda's protocol. The actual proofs were done by hand. We observed, however, that the state machine models for all protocols described in [25] are finite and small—trivial for a model checker. So to complete the formal analysis, we subjected all three protocols to model checking.

In this section, we present the details of how we used McMillan's SMV model checker to verify the AFS-1 and AFS-2 protocols. They are small enough to present in their entirety but "big" enough to let us illustrate common abstractions others can apply to their own systems. We only briefly present the results of the more complicated Coda example; see [29] for details. We give SMV input and output for each example. SMV expects input specifications (ϕ of Figure 1) in the form of Computational Tree Logic (CTL), a subset of branching time temporal logic. We explain SMV and CTL as needed in the examples.

3.1 AFS-1

In AFS-1, a client has two initial states: either it has no files or it has one or more files but no beliefs about their validity. If the protocol starts with the client having no file in its cache, then the client may request a copy from the server and the protocol terminates when the file is received. If the protocol starts with the client having suspect files, then the client may request a validation for a file from the server. If the file is invalid then the client requests a new copy and the run terminates. If the file is valid, the protocol simply terminates.

3.1.1 State Machine Model

Let's consider a simplified model with just one client, one server, and one file. The top graph in Figure 2 shows the state transition graphs for the client, and the bottom, for the server. The nodes are labeled by the value for the state variable, `belief`; the arcs, by the name of the message received that causes the state transition. A run of the protocol begins in an initial state (one of the leftmost nodes) and ends in a final state (one of the rightmost nodes).

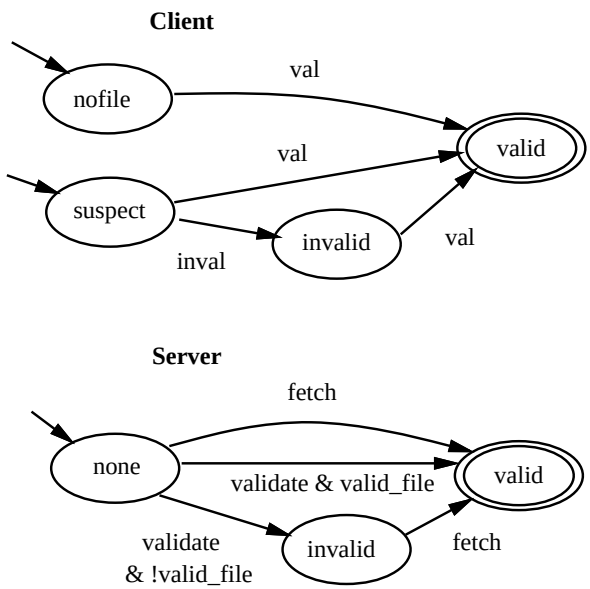

Figure 2: State Transition Graphs for AFS-1

The client's belief about a file ranges over { `nofile`, `valid`, `invalid`, `suspect` }. The client's belief is `nofile` if the client cache is empty; `valid`, if the client believes its cached file is valid; `invalid` if it believes its cached file is not valid; `suspect`, if it has no belief about the validity of the file (it could be valid or invalid).

The server's belief about the file cached by the client ranges over { `valid`, `invalid`, `none` }. The server's belief is `valid` if the server believes that the file cached at the client is valid; `invalid`, if the server believes it is not valid; `none`, if the server has no belief about the existence of the file in the client's cache or its validity.

The set of messages that the client may send to the server is { `fetch`, `validate` }. The message `fetch` stands for a request for the file. The `validate` message is used by the client to determine the validity of the file in its cache.

The set of messages that the server may send to the client is { `val`, `inval` }. The server sends the `val` (`inval`) message to indicate to the client that its cached file is valid (invalid).

3.1.2 Specification of Cache Coherence in CTL

Given the state variables for a client state machine, Client, and a server state machine, Server, we can express *CC* as the following CTL formula:

```
AG ((Client.belief = valid ) ->
   (Server.belief = valid))
```

CTL formulas are interpreted with respect to an infinite computation tree derived from the state transition graphs. `AG` is the operator that stands for "for all paths (A) for all states (G)." Thus `AG P` says "invariably, everywhere P."

3.1.3 SMV Input for AFS-1

The SMV input program for AFS-1, shown in Figure 3, is a textual representation of the state transition graphs of Figure 2. We show it for completeness. The fifth line gives the CTL specification against which the state machine is checked.

To show how SMV finds counterexamples, we add in the sixth line another CTL specification for SMV to check, the converse of our correctness condition:

```
AG ((Server.belief = valid) ->
   (Client.belief = valid))
```

The SMV input program is composed of the modules `main`, `server`, and `client`. The third and fourth lines declare instances of the server and client modules.

The module `server` takes a parameter `input` that can be any message coming from the client, indicated by the instantiation of the parameter by `Client.out` in the fourth line. The server module has two other state variables besides `belief`. The `out` variable ranges over the messages that the server may send to the client; the `0` message stands for no message and is needed to model the initial state. The `valid-file` boolean variable models the effect of the environment as perceived by the server. It is used when the client

```
MODULE main  -- afs1
VAR
Client : client (Server.out);
Server : server (Client.out);
SPEC AG ((Client.belief = valid) ->
         (Server.belief = valid))
SPEC AG ((Server.belief = valid) ->
         (Client.belief = valid))
MODULE server(input)
VAR
out          : { 0, val, inval };
belief       : { none, valid,invalid };
valid-file   : boolean;
ASSIGN
valid-file := { 0,1 };
init(belief) :=none;
next(belief) :=
 case
  (belief = none) & (input = fetch) : valid;
  (belief = none) & (input = validate) &
                       valid-file : valid;
  (belief = none) & (input = validate) &
                       !valid-file : invalid;
  (belief = invalid) & (input = fetch) : valid;
  1 : belief;
 esac;
init(out) := 0;
next(out) :=
 case
   (belief = none) & (input = fetch) : val;
   (belief = none) & (input = validate) &
                        valid-file : val;
   (belief = none) & (input = validate) &
                       !valid-file : inval;
   (belief = invalid) & (input = fetch) : val;
   1 : 0;
 esac;
MODULE client(input)
VAR
out : {0, fetch, validate};
belief : {valid, invalid, suspect, nofile};
ASSIGN
init(belief) := nofile, suspect;
next(belief) :=
 case
  (belief = nofile) & (input = val) : valid;
  (belief = suspect) & (input = val) : valid;
  (belief = suspect) & (input = inval) : invalid;
  (belief = invalid) & (input = val) : valid;
  1: belief;
  esac;
init(out) := 0;
next(out) :=
 case
  (belief = nofile) : fetch;
  (belief = invalid) : fetch;
  (belief = suspect) : validate;
  1 : 0;
 esac;
```

Figure 3: SMV Input Program for AFS-1

```
-- specification AG (Client.belief = valid ->
                     Server.beli... is true
-- specification AG (Server.belief = valid ->
                     Client.beli... is false
-- as demonstrated by the following execution ...
state 1.1:
   Client.out = 0
   Client.belief = nofile
   Server.out = 0
   Server.belief = none
   Server.valid-file = 0
state 1.2:
   Client.out = fetch
state 1.3:
   Server.out = val
   Server.belief = valid
resources used:
user time: 0.133333 s, system time: 0.116667 s
BDD nodes allocated: 1048
Bytes allocated: 917504
BDD nodes representing transition relation:112+1
reachable states: 26 (2^ 4.70044)
                  out of 216 (2^ 7.75489)
```

Figure 4: SMV Output for AFS-1

has a suspect file in its cache and requests a validation from the server. We need to model both the possibility that the server has received an update by some other client and the possibility that it has not. If an update by some other client has occurred then the server reflects that by nondeterministically setting the value of `valid-file` to 0; otherwise, the server sets the value to 1 (the file cached at the client is still valid). The initial belief of the server is `none`; the final belief is `valid`.

In the module `client`, in addition to the state variable `belief`, as for the server, we use the `out` variable to range over the messages that the client may send to the server. The client's initial belief is `nofile` or `suspect`. If the initial belief is `suspect` and a failed validation message is received, then the client believes its file is `invalid`. It then sends a `fetch` message to the server, as indicated in the definition of the transitions for `out`. The client's final belief is `valid`.

3.1.4 SMV Output for AFS-1

Figure 4 shows the output of SMV for AFS-1.

SMV indicates that (1) the first property, the cache coherence invariant, is true and (2) the second property, the converse, is false. The counterexample SMV finds corresponds to the following scenario. Initially, the client has no file and the server has no beliefs. The client then requests a copy from the server. Then, the server receives this fetch request and sends a copy to the client, believing, of course, that the file sent and thus cached by the client is valid. Thus, the server believes the file is valid, but in this last state, the client has not received the server's message and still believes that it has no file. The line after `resources used:` says that SMV takes fractions of a second to check both properties and the last line says that the number of reachable states for AFS-1 is 26.

3.2 AFS-2

In AFS-2, clients can update their files and failures must be handled explicitly. The protocol is based on callbacks. When a client caches a valid file, the server promises to notify that client if the file is updated. This promise is called a *callback* [20].

AFS-2 works as follows. Initially, a client may have one of two beliefs about a file. It either believes it has no copy of the file or it has a suspect copy. If the client's initial belief is that it has no file, the client may request a copy from the server. The server then has a callback on that file. If the file is ever updated, the server notifies the client and the client discards its copy. If client's initial belief is that there is a suspect file in its cache, the client may request a validation from the server. If the file is valid, then the server has a callback on that file. If the file is invalid, the client discards its copy. If at any time during a run a failure occurs in the system, the clients then consider their copies of the file suspect and the server discards its beliefs about the validity of the files cached by the clients.

3.2.1 State Machine Models

For AFS-2, we also consider a simplified model with just one server, two identical clients (Client1 and Client2), and one file. Figure 5 gives the state transition graphs for each client and the server.

The client's belief about a file ranges over { `valid`, `suspect`, `nofile` }. The belief `valid` indicates that the file is in cache and it is valid; `suspect`, that the file is in cache but the client has no belief about its validity; `nofile`, that there is no file in cache. Since a file that is believed to be invalid is immediately discarded by the clients, we have chosen not to represent the belief `invalid` to simplify the system.

For each client, the server has a belief about the validity of the file cached by that client. Each belief ranges over { `valid`, `nocall` }, where `valid` indicates that there is a file in the client's cache and it is valid; `nocall`, that the server has no callback on the file cached by a client.

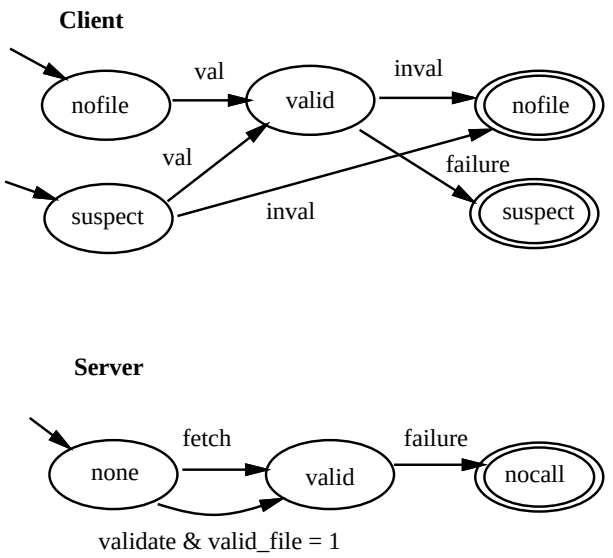

Figure 5: State Transition Graphs for AFS-2

The messages that the clients may send to the server are { `fetch`, `validate`, `update` }. An `update` message to the server indicates that the file cached by the client has been updated. The server's messages are the same as for AFS-1.

3.2.2 SMV Input and Output for AFS-2

We omit the actual SMV input and output for AFS-2 because they are similar to that for AFS-1 (see [29] for the details). We point out the major differences in this section.

If we use the same correctness criterion for AFS-2 as we did for AFS-1, SMV gives a counterexample. Our "invariant" is not an invariant! The reason is that the cache coherence invariant holds for AFS-2 only within certain timing constraints due to transmission delay. Consider the following scenario. Client1 has a valid file in its cache and the server has a callback on that file. Client2 suddenly updates its copy of the file. Then the server immediately believes that the file cached by Client1 is not valid and sends a message to Client1 to notify it. In this state, Client1 has not yet received the server's message and it still believes that its file is valid. So there is a period of time due to transmission delay during which the invariant does not hold. Let τ represent the upper bound on this time interval, then the following property is true:

> CC': If a client believes its file is valid at the present, then before some past interval of time of length τ, the server believed the file cached by that client is valid.

In CTL, there are operators about the future, not the past, so we reformulate this property using its contrapositive. We also model the transmission delay by the amount of time it takes to go from one state to another, i.e., a single step. This leads us to the following CTL formula, which is the cache coherence condition (CC' above) we use for AFS-2:

```
AG((Server.  belief1 = nocall) ->
  AX ((Client1.belief = nofile) |
      (Client1.belief = suspect)))
```

where `AX` means "invariably, in the next state."

The SMV input for AFS-2 consists of the instances Client1, Client2, Server, and Env. We introduce the `env` module, representing the environment, so that we can explicitly model failures that occur between Client1 and the Server and between Client2 and the Server. (The module has two state variables `failure1` and `failure2`. Each one of them can be independently set to 1. Once a variable is set to 1, it remains at that value for the rest of the run.)

The module `server` has two beliefs, `belief1` and `belief2`. The `server` module in AFS-2 is otherwise similar to the `server` module in AFS-1.

The module `client` is also similar as that for AFS-1. The only difference is that in AFS-2 a client may also send an update message to the server when it believes its file is valid. This difference is captured in the third to last line in the definition of `out` in the `client` module.

SMV's output for AFS-2 indicates that the cache coherence invariant (CC') is true. The total time used is less than .5 seconds and the number of reachable states is 7776.

3.3 Coda

Mummert's extension to the Coda protocol is complicated by the addition of a new type of cached data, called a *volume* (Section 4.3). This addition leads to a richer state space, more state transitions, and more cases in which failures can arise (see the figure in the Appendix). For example, whereas AFS-2 has only three classes of runs, Mummert's Coda protocol has fifteen. SMV takes less than one second to check 43,684 reachable states.

4 Discussion: Different Kinds of Abstractions

What made it possible for us to use model checking in our case study is that we chose different abstraction mappings to apply to the real system. We

certainly did not check the actual C code that implements AFS-1 or AFS-2, but then the level at which we would want to verify a protocol like cache coherence is much higher than the code level; if the design is incorrect or incomplete (misses some cases) then the code is apt to reflect those mistakes and omissions. Designs, especially for distributed systems protocols, are thus good subjects for model checking.

In this section we explain in more detail some of the abstraction mappings that we applied, which can be used more generally in other settings. We exploit (1) the form of the property, ϕ to be satisfied, (2) domain-specific knowledge, and (3) problem-specific knowledge.

4.1 Exploiting the Form of ϕ

In a real instance of a distributed file system like Andrew or Coda, there are an unbounded number of files, clients, and servers. To analyze the state space of a real instance would be beyond the capability of any model checker today.

If ϕ is in the form of $\forall x.P(x)$ then if for any instance, a, proving $P(a)$ will let us deduce ϕ by the generalization rule from logic:

$$\frac{P(a)}{\forall x.P(x)}$$

Our cache coherence invariant (CC) is stated as "for all clients, for all servers, for all files ...." Thus, for example, for AFS-1, we simply successively apply the generalization rule three times upon proving the property for a model with just one client, one server, and one file.[1]

4.2 Exploiting Domain-Specific Knowledge

In our domain of distributed systems, we need to worry about failures and transmission delays. First, we abstract from different types of failures like crashed nodes and downed communication links since our protocols act the same regardless of the type of failure. We do, however, need to model that a failure may occur; we use an environment machine (`env`) to model this possibility. We also need to model that clients and servers can detect a failure; we use the `failure1` and `failure2` parameters for this purpose.

[1] Roscoe and MacCarthy make a similar point in their work on model checking (using FDR) *data-independent* properties of concurrent processes [26]. We agree with their comment that further work on a formal justification is needed.

Second, transmission delay in conjunction with failures complicate reasoning about correctness in distributed systems. During the interval in which a message is traveling, the sender may have made some state change which will render the correctness condition false until the receiver processes the message and changes its state. Transmission delay is loosely bounded by the timeout period used by the underlying communication protocol, denoted by β. It also takes time for clients and servers to detect failures. The interval between the occurrence of a failure and its detection, denoted by τ, defines a window of vulnerability. To bound the failure detection interval, clients and servers probe each other periodically, and declare failures if messages time out. Let ρ be the probe interval and assume that clients and servers use the same probe interval, but do not necessarily probe each other at the same time. Then the failure detection interval is at most $\tau = \rho + \beta$. [2]

We abstract from the exact times that contribute to τ and model this interval as a single step in a state machine. This interpretation lets us characterize the correctness condition, CC', in terms of the next-state temporal logic operator (`AX`) rather than the original, simpler CC. If transmission were instantaneous and there were no failures, there would be no need for modeling τ since client and server beliefs could be updated simultaneously.

4.3 Exploiting Problem-Specific Knowledge

We use ϕ to drive the choice of "appropriate" abstractions. For example, if ϕ is a property about an integer, x, we may care only that x is negative, positive, or zero. We can define an abstraction function that maps an infinite set of values to three values. Or, if ϕ is a property about an integer set, s, we may care only about whether the set is empty or not; we may not care at all about what its elements are. We can define an abstraction function that maps an infinite set of set values to two values.

In our problem of cache coherence between clients and servers, the property ϕ that we want to verify is that the client's beliefs about the validity of cached files are consistent with a server's beliefs. We let ϕ guide us in determining what details of the real system we can safely ignore.

First, we abstract from the clients' caches, i.e., sets of files. Since we have simplified our model to handle just one file, we need only model whether the cache,

[2] In Coda τ is composed of a probe interval of 10 minutes and a message timeout of 15 seconds.

C, is empty or not, i.e, $C = \emptyset$ or $C = \{f\}$. This abstraction gives us two cases to consider instead of an exponential number of cases (2^n) for an unbounded number, n, of files.

Second, we need only represent the belief a client has about the cached file. We do not need to model sets of beliefs (given we have only one file about which to have a belief). For example, for AFS-2, in reality, we might have a cache, C, and a belief set B (for file f), that ranges over $\emptyset$ and $\{valid(f)\}$. Rather than two state variables, we need only one (`belief`), and rather than four possibilities (2×2), we need represent only three. We define a (partial) abstraction function, A, that maps a pair $\langle C, B \rangle$ to a belief in {`valid`,`nofile`,`suspect`}:

$\langle \{f\}, \{valid(f)\} \rangle \quad \mapsto$ `valid`
$\langle \emptyset, \{valid(f)\} \rangle \quad \mapsto \perp$ (unreachable case)
$\langle \{f\}, \emptyset \rangle \quad \mapsto$ `suspect`
$\langle \emptyset, \emptyset \rangle \quad \mapsto$ `nofile`

Third, we do not even need to model the file itself since we do not care at all about the contents (value) of the file. We would if we needed to compare the values of two different files or extract information about the file based on its value. Thus, we identify the transmission of the file with the transmission of the message about the file; for example, in AFS-2, the `val` message can be thought of as abstractly containing the file itself as well as the status about its validity.

Fourth, we abstract from the type of data cached. In the AFS protocols, the types include files and directories. For our analysis, however, that there are different types of data is completely irrelevant to the correctness of the protocol. We use the generic term "file" to stand for any kind of data. (In the Coda file system example, a third type, volumes, is treated specially, and thus must be modeled explicitly; as mentioned earlier, it is this new data type that complicates the Coda protocol.)

Fifth, we abstract from the notion of validity, which is determined by the recency of a file. Suppose as in the implementation, recency is determined by comparing the totally ordered timestamps associated with files. For any pair of timestamped files, f_{t1} and f_{t2}, we can determine whether one is more recent by comparing their timestamps, $t1 < t2$; but since this always returns true or false, we can model recency, and hence validity, as a boolean variable representing whether a file is valid or not. This abstraction appears explicitly in the way we use the `valid-file` boolean variable in the server module, letting the server nondeterministically choose between the two possibilities.

Finally, we abstract from individual runs of protocol. We consider *classes* of runs, categorized by a protocol's sets of initial and final states. For our most complex example, this reduces the number of cases to consider from forty-four to fifteen [25].

5 Related Work

Model checking originated with Clarke and Emerson's work in 1981 [8]. As mentioned in the introduction, it has already proven to be extremely successful in debugging hardware [4, 14, 23, 10, 9, 3]. Tool support for model checking includes SMV [22], FDR [16], COSPAN [19], and the Concurrency Workbench [11]. There are more and more documented case studies; for example, the proceedings of the 1995 Workshop of Industrial-Strength Formal Techniques contains four model checking case study papers [17].

We are not the first to explore the use of model checking in the software domain. Three other approaches complement ours and each other. Since they are all recent (dated 1993-94), we expect that over time results from one approach will carry over to the others.

- Atlee and Gannon follow a specification-language based approach [2]. They verify safety properties for event-driven systems described by the SCR tabular requirements language. Their case studies include an automobile cruise control system and a water-level monitoring system. They show how to represent any specification written in a subset of SCR as a finite state machine. They use an extended version of SMV for their model checker.

- Jackson explores the richness of types in software systems. State variables for hardware (and SCR) are of simple types like boolean, but in software they range over more complex type like sets, graphs, and relations. He exploits symmetry in mathematical relations to reduce the state space; he shows how to model check Z specifications, which is essentially based on his relational calculus [21].

- Allen and Garlan's use of model checking focuses at the level of software architecture, a level of abstraction far above the real system, but again where many design flaws can be detected. They use FDR to detect deadlocks in software architectures described in the Wright architectural description language [1]. Wright is based on a sub-

set of CSP, and thus it leaves states completely uninterpreted.

Our approach complements all three of the above since in each case, the researchers first build some finite model of the real system and express it in terms of SCR, Z, or Wright. In doing so, they implicitly apply the kinds of abstractions we used in our examples. By restricting themselves to a specific language, they have the advantage of avoiding having to define different abstractions per problem, since they do this mapping once and for all. In our approach, we let ϕ drive the choice of abstractions and simply express our models directly in terms of SMV input. We have the advantages of bypassing the "intermediate" specification language translation step, and of not being restricted to the domain of systems that a given specification language is most suited for describing. Thus, our focus is on finding "appropriate" abstractions that work across different domains and different problems, not on checking models expressed in a particular specification language.

Finally, Cheung and Kramer's two-step analysis approach applied to reasoning about large-scale distributed systems is similar in spirit to ours [7]. They use dataflow analysis as a way to approximate a system's behavior and then contextual analysis to do an exhaustive search of the resulting state space.

6 Conclusions and Future Directions

Model checking has the significant advantage over more traditional forms of software verification in that much of the hard work is done automatically by the machine. Moreover, both the inputs and the results of model checking tools are straightforward to understand by non-experts: since a model checker's interface is well-defined and it is straightforward to provide its expected inputs, we can readily use it as a "black box." This suggests that for gaining assurance about software, model checking can become a technology that is much more broadly accessible to practitioners than other techniques.

The choice of what abstractions to apply takes some good judgment. After all, we could define a model of a system that is so abstract that any property would be trivially satisfied or that would allow any possible concrete realization. Further research is needed to characterize more formally what makes an abstraction "good."

In our own work, toward making progress both in demonstrating feasibility and in understanding characteristics of good abstractions, we plan to push on more examples. We are currently exploring how to use model checking to validate recovery protocols [12] for redundant disk arrays, in particular for the RAID Level-5 architecture and its successors. Our primary goal is to provide more convincing evidence to systems designers and builders that formal reasoning tools are ready for day-to-day use. As a useful by-product, we expect to identify other kinds of abstractions appropriate to apply to real software systems.

We are optimistic about the future of model checking software systems. One way to measure the practicality of model checking is by how easy it is to teach and learn. The second author did the model checking case study in this paper as part of her senior honor's thesis. In CMU's Master's of Software Engineering core course on Analysis of Software Artifacts, we have students do a two-week project using SMV. Professor Clarke regularly offers a graduate-level course on model checking at CMU, and some CMU faculty have even proposed the idea of teaching model checking in undergraduate core courses [5].

Acknowledgments

Jeannette Wing would like to credit Daniel Jackson for his insight as to why model checking is especially appropriate for proofs about software at the system level. Discussions with Daniel and David Garlan have clarified our optimistic views toward the use of model checking for software.

We would also like to thank Ed Clarke and indirectly Randy Bryant for providing the model checking technology that made it possible to do our case study. Mandana Vaziri-Farahani especially thanks Xudong Zhao and Sergio Campos for their help in learning how to use SMV.

Finally, we like to thank the fellow "systems" faculty at CMU (like Satyanarayanan) who have continually challenged their formal specification and verification colleagues to demonstrate the utility of their research results on systems that "matter." That we were able to use model checking for showing Coda's cache coherence protocol correct is a step towards meeting that challenge.

This research is sponsored by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, and the Advanced Research Projects Agency (ARPA) under grant number F33615-93-1-1330. Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of Wright Laboratory or the United States Government.

References

[1] Allen, R., and Garlan, D. Formalizing architectural connection. In *Proceedings of the Sixteenth International Conference on Software Engineering* (May 1994).

[2] Atlee, J., and Gannon, J. State-based model checking of event driven systems requirements. *IEEE Trans. Soft. Eng.* (Jan. 1993).

[3] Barrett, G. Model checking in practice: The t9000 virtual channel processor. *IEEE Trans. on Soft. Eng.*, 2 (Feb. 1995), 69–78.

[4] Browne, M. C., Clarke, E. M., Dill, D. L., and Mishra, B. Automatic verification of sequential circuits using temporal logic. *IEEE Trans. Comput. C-35*, 12 (1986), 1035–1044.

[5] Bryant, R., 1994. private communication.

[6] Burrows, M., Abadi, M., and Needham, R. A logic of authentication. *ACM Transactions on Computer Systems 8*, 1 (Feb. 1990), 18–36.

[7] Cheung, S., and Kramer, J. An integrated method for effective behavior analysis of distributed systems. In *Proc. 16th Int'l Conf. on Soft. Eng.* (Sorrento, Italy, May 1994).

[8] Clarke, E. M., and Emerson, E. A. Synthesis of synchronization skeletons for branching time temporal logic. In *Logic of Programs: Workshop, Yorktown Heights, NY, May 1981* (1981), vol. 131 of *Lecture Notes in Computer Science*, Springer-Verlag.

[9] Clarke, E. M., Grumberg, O., Hiraishi, H., Jha, S., Long, D., McMillan, K. L., and Ness, L. A. Verification of the futurebus+ cache coherence protocol. In *Proc. of CHDL '93* (1993).

[10] Clarke, E. M., Grumberg, O., and Long, D. E. Model checking and abstraction. In *Proc. of POPL '92* (1992).

[11] Cleaveland, R., Parrow, J., and Steffen, B. The concurrency workbench: A semantics-based tool for the verification of concurrent systems. Tech. Rep. ECS-LFCS-89-83, Edinburgh Univeristy, 1983.

[12] Courtright, W., and Gibson, G. Backward error recovery in redundant disk arrays. In *Proceedings of the 1994 Computer Measurement Group Conference (CMG)* (Orlando, FL, Dec. 1994), pp. 63–74.

[13] Craigen, D., Gerhart, S., and Ralston, T. Formal methods reality check: Industrial usage. *IEEE Trans. on Soft. Eng. 21*, 2 (Feb. 1995), 90–98.

[14] Dill, D. L., and Clarke, E. M. Automatic verification of asynchronous circuits using temporal logic. *IEE Proceedings Part E 133*, 5 (1986).

[15] Flora-Holmquist, A., and Staskauskas, M. Formal validation of virtual finite state machines. In *Proceedings of the Workshop on Industrial-Strength Formal Specification Techniques* (Apr. 1995). to appear.

[16] Formal Systems Europe, L. *FDR: User Manual and Tutorial*. Oxford, England, Oct. 1992.

[17] France, R., and Gerhart, S. *Proceedings of the Workshop on Industrial-Strength Formal Specification Techniques*. IEEE, 1995.

[18] Halpern, J., and Moses, Y. Knowledge and common knowledge in a distributed environment. In *Proceedings of the Third ACM Symp. on Principles of Distributed Computing* (Aug. 1984), pp. 50–61.

[19] Har'El, Z., and Kurshan, R. P. Software for analytical development of communications protocols. *AT&T Technical Journal 69*, 1 (Jan.–Feb. 1990), 45–59.

[20] Howard, J., Kazar, M., Menees, S., Nichols, D., Satyanarayanan, M., Sidebotham, R., and West, M. Scale and Performance in a Distributed File System. *ACM Transactions of Computer Systems 6*, 1 (Feb. 1988), 51–81.

[21] Jackson, D. Abstract model checking of infinite specifications. In *Proc. of FME '94* (Oct. 1994).

[22] McMillan, K. L. *Symbolic Model Checking: An Approach to the State Explosion Problem*. PhD thesis, Carnegie Mellon University, 1992. CMU-CS-92-131.

[23] McMillan, K. L., and Schwalbe, J. Formal verification of the gigamax cache consistency protocol. In *Shared Memory Multiprocessing*, N. Suzuki, Ed. MIT Press, 1992.

[24] Mummert, L., and Satyanarayanan, M. Large Granularity Cache Coherence for Intermittent Connectivity. In *USENIX Summer Conference Proceedings* (June 1994), USENIX Association, pp. 279 – 289.

[25] Mummert, L., Wing, J., and Satyanarayanan, M. Using belief to reason about cache coherence. In *Proceedings of the Symposium on Principles of Distributed Computing* (Aug. 1994). Also CMU-CS-94-151, May 1994.

[26] Roscoe, A., and MacCarthy, H. A case study in model-checking CSP. submitted for publication, Oct. 1994.

[27] Satyanarayanan, M., Howard, J., Nicols, D., Sidebotham, R., Spector, A., and West, M. The ITC Distributed File System: Principles and Design. In *The Tenth ACM Symposium on Operating Systems Principles* (Dec. 1985), ACM, pp. 35–50.

[28] Satyanarayanan, M., Kistler, J., Kumar, P., Okasaki, M., Siegel, E., and Steere, D. Coda: A highly available file system for a distributed workstation environment. *IEEE Transactions on Computers 39*, 4 (Apr. 1990).

[29] Vaziri-Farahani, M. Using symbolic model checking to verify cache coherence in a distributed file system. Tech. Rep. CMU-CS-95-156, Carnegie Mellon Computer Science Department, 1995. CMU Electrical and Computer Engineering Bachelor's Thesis.

Appendix: State Transition Graphs for Coda

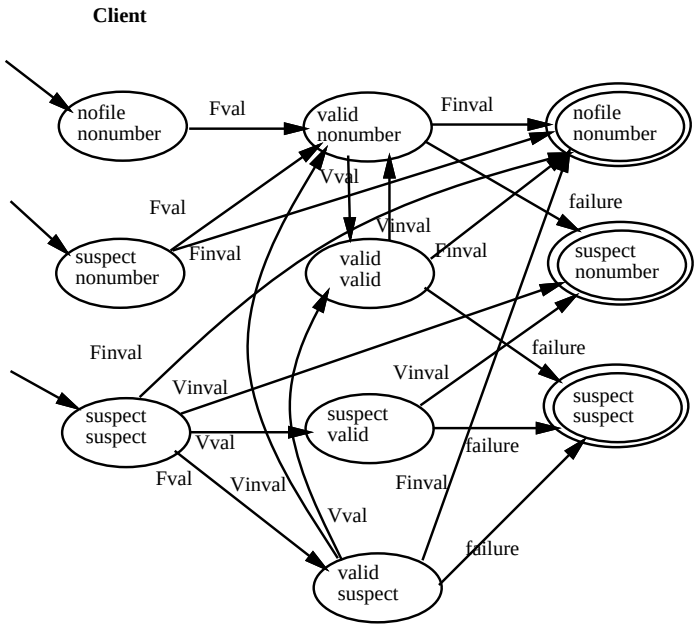
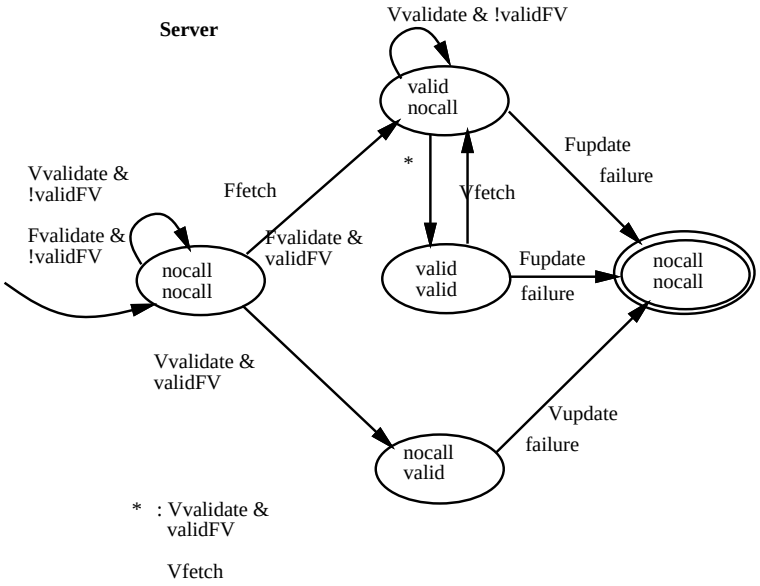
Each node is labeled by the value of the belief about the file and the value of the belief about the volume. Each arc is labeled by the name of the message received by the client or server.