

Using Partial Order Techniques to Improve Performance of Data Flow Analysis Based Verification*

Gleb Naumovich, Lori A. Clarke, and Jamieson M. Cobleigh

email: {naumovic|clarke|jcobleig}@cs.umass.edu
Laboratory for Advanced Software Engineering Research
Computer Science Department
University of Massachusetts
Amherst, Massachusetts 01003

1 Introduction

Finite state verification techniques automatically check that a software system conforms to a behavior specification or *property*. Such techniques are becoming extremely important with the proliferation of distributed systems. Distributed systems are more difficult to understand and reason about than sequential ones because of the potential non-deterministic interleaving of execution sequences from different threads of control, or *tasks*. While testing demonstrates the actual behavior of a system on selected test cases, distributed systems may not even produce the same results when re-executed with these same test cases. Finite state verification techniques, however, are capable of verifying a restricted, but interesting, class of properties for all possible executions of a program for all possible test cases. Unfortunately, in practice, finite state verification tools often require significant computing resources, and so there is a need for optimizations that improve the performance of such techniques.

One popular method for improving the performance of finite state verification techniques, without compromising the results of their analyses, is *partial order* optimization. Such an optimization is based on the observation that most representations of distributed systems, used in analysis, model the execution of the system as a total order between occurrences of events local to the tasks that execute in parallel. This means that multiple executions exist that differ from each other only by the relative order of appearance of events occurring in different tasks. In many cases, these differences are not important for checking the property of interest. In such cases, partial order reduction techniques choose and reason about a single representative ordering.

Necessarily, whether or not two interleavings can be considered equivalent depends on the property being checked. Thus, all partial order methods are defined for specific kinds of properties. The approaches of Godefroid and Wolper [6],

Valmari [12], and Katz and Peled [8] use partial orders for verifying safety properties. Peled [11] proposed a partial order method that allows checking stuttering-close Buchi automata properties.

Partial order methods have been shown to be successful in improving the performance of model checkers. One case study [1] showed that for many situations these techniques significantly improve both time and space requirements of the SPIN [7] model checker, thus enabling analysis of bigger problems. Godefroid, Peled, and Staskauskas [5] describe the design of a partial-order algorithm for a formal validation tool used for verification of several subsystems within Lucent Technologies 5ESS telephone switching system. This case study indicated that partial order approaches may significantly reduce analysis time for industrial software. Godefroid [3] uses partial order techniques in VeriSoft, a verification tool for distributed systems implemented in C or C++. VeriSoft was demonstrated to be effective in verifying an example of the Lucent Technologies' Heart-Beat Monitor of a telephone switch system [4].

To date, partial order reduction techniques have been applied in the context of model checking approaches. Such approaches enumerate all possible states of the system and reason about these states. In this paper we propose a partial order reduction for FLAVERS, a finite verification approach based on data flow analysis [2]. Although the program model that FLAVERS uses does not enumerate all possible states of the concurrent system under analysis, it uses a special type of edges to represent possible interleavings between events in different processes. The optimization that we propose in this paper uses partial orders to eliminate some of these edges, thereby improving efficiency of the FLAVERS analysis.

We evaluated the benefits of this optimization on a number of small, distributed programs. As expected, it was relatively easy to determine if the partial order optimization technique was applicable to a problem. For these 92 cases, the optimization was applicable to 35 of them. On average, for all applicable cases, the speedup of the FLAVERS analysis due to the use of this optimization was 21%. For one case, the optimization resulted in an analysis speedup of 91%.

In the next section we present a high-level overview of

*This research was partially supported by the Defense Advanced Research Projects Agency and the Air Force Research Laboratory/IFTD under agreement F30602-97-2-0032. The views, findings, and conclusions presented here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency, the Air Force Research Laboratory/IFTD, or the U.S. Government.

FLAVERS and give a detailed description of the system model that FLAVERS uses. Section 3 describes the partial order optimization approach that allows us to remove a significant number of edges from the FLAVERS program model. In Section 4 we present our experimental results. We conclude with observations and future research directions.

2 FLAVERS

Overview

FLAVERS (**FL**ow Analysis for **VER**ification of Systems) compactly represents a concurrent software system with a *Trace Flow Graph (TFG)* and uses an efficient fixed point data flow algorithm to determine if the behavior described by the TFG is consistent with a user-specified, event-based, safety property. The results of the FLAVERS analysis are conservative; in other words, the technique never claims that a property is verified when it is not. For efficiency reasons, similar to other finite state verification techniques, the TFG model used by FLAVERS over-approximates the potential executable sequences of events associated with the program. This leads to the possibility of *spurious results*, where FLAVERS reports a property violation when there is in fact no real executable behavior of the system that would violate that property. If FLAVERS detects a property violation, it provides the user with example paths that illustrate this property violation. By examining such paths users can often determine if the result of the analysis is spurious or not.

FLAVERS provides a flexible way for improving the precision of the analysis. Analysts do this by adding *feasibility constraints*, which specify additional semantic information about the system and which are used to limit the exploration of the TFG to only those paths that satisfy these feasibility constraints. If the constraints are well chosen by the user, spurious paths are eliminated and the subsequent analysis run will either verify the property or expose a counter example that corresponds to real executable behavior and, thus, exposes a bug in the system. FLAVERS provides automated support for creating several classes of feasibility constraints, for example constraints that model boolean, counter, or enumerated variables or model control flow through a specific thread of control.

Unfortunately, the use of constraints leads to the need to solve large and complex data flow problems, since the complexity of the FLAVERS analysis algorithm is exponential in the number of constraints, and so, if many constraints are used, the analysis algorithm has to deal with vast amounts of data. This led us to search for techniques for improving space and time characteristics of FLAVERS. In order to understand the technique proposed in this paper, in the remainder of this section we describe the TFG model in detail.

TFG Model of Ada Programs

The TFG for an Ada program is based on the control flow graphs (CFGs) for all tasks. Additional nodes and edges are added to the TFG to represent intertask communications.

Specifically, if the code region represented by node n in one task contains a synchronization statement that can correspond to one represented by node m in another task, a new node is added with incoming edges from n and m and outgoing edges to all successors of n and m . This is illustrated in Figure 1(c). A unique *initial* node that has no incoming edges and has outgoing edges to the start nodes of all CFGs and a unique *final* node that has no outgoing edges and has incoming edges from the end nodes of all CFGs are added to the TFG.

Formally, a TFG is a labeled directed graph $(N, E, n_{initial}, n_{final}, \Sigma, label)$, where N is the set of nodes, $E \subseteq N \times N$ is the set of edges, $n_{initial}, n_{final} \in N$ are unique initial and final nodes, Σ is the set of labels on the nodes of the graph, and *label* is a mapping from nodes to labels in Σ . The set of all nodes from the CFGs for all tasks forms the set of *local* TFG nodes. All other nodes, except the initial and the final nodes, represent intertask communications and thus are called *communication* nodes.

Let T be the set of tasks in the program. In the TFG for an Ada program, nodes may belong to one (local nodes), two (communication nodes), or no (initial and final nodes) tasks. We use function $task : N \rightarrow 2^T$ to associate with each TFG node the set of tasks it belongs to.

FLAVERS represents all properties and constraints as finite state automata (FSA). Transitions in these FSAs are labels from the TFG for the system; the FLAVERS analysis algorithm propagates the states of these automata throughout the TFG. Let \mathcal{A} be the set containing the property and all constraint automata used in the analysis. Let A be any automaton, either the property or a constraint. Let Σ_A be the alphabet of this automaton, that is, all events used in the transitions in this automaton. Define alphabet Σ_{aut} to contain all events in the property and constraint automata: $\Sigma_{aut} = \bigcup_{A \in \mathcal{A}} \Sigma_A$. Events from the TFG alphabet Σ that are not present in Σ_{aut} are not used by the property and constraints, and so they can be replaced with a single event τ . (In the following we assume that $\tau \in \Sigma$.)

Figure 1(a) shows a program that consists of two communicating Ada tasks, Figure 1(b) shows the CFGs for these two tasks with nodes labeled with the corresponding Ada program statements¹, and Figure 1(c) gives the TFG for this Ada program. The local nodes in this TFG have the same ID numbers associated with them as the corresponding nodes in the CFGs. The diamond-shaped nodes are the initial and the final nodes of this TFG and two communication nodes, labeled E1 and E2. The node labeled E1 represents the communication between the tasks at entry call T2. E1, and node labeled E2 represents the communication at entry

¹Only the two nodes labeled `x==true` and `x==false` do not directly correspond to any executable statements in the code. These two nodes represent the fact that the value of variable `x` is `true` on the first branch of the `if` statement and `false` on the other branch.

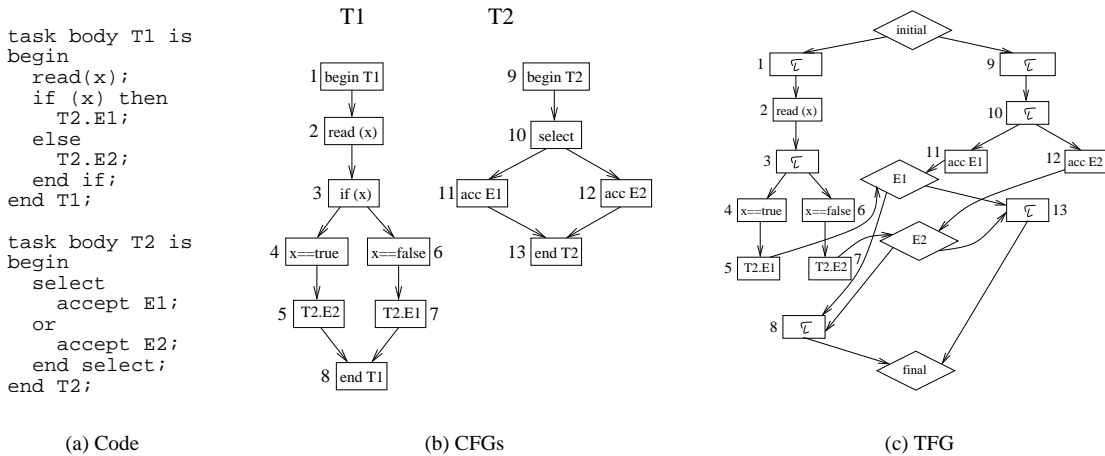


Figure 1: An example




Figure 2: The TFG after removal of some τ -labeled nodes

call $T2.E2$. Note that some nodes in the TFG are labeled τ . We assume that the labels on the corresponding nodes in the CFG are not in the alphabet Σ_{aut} for this example.

Theoretically, all τ -labeled nodes can be removed from the TFG in a safe manner, with the results of the FLAVERS analysis on the reduced TFG being the same as the results on the original TFG. When a τ -labeled node is removed, an edge is constructed from each predecessor of this node to each successor of this node. For nodes with multiple successors and predecessors this may lead to a quadratic blow-up in the number of control edges in the graph. We use a simple heuristic that removes τ -labeled nodes only where this does not lead to an increase in the number of edges in the TFG. In addition, we do not remove τ -labeled nodes in the situation where this would result in two communication nodes being direct successors of one another. Figure 2 shows the TFG from Figure 1(c) after the τ -labeled nodes that are not immediate predecessors of communication nodes are removed. The dashed edges between the nodes in Figure 2 are explained below.

In addition to edges that represent control flow within a single task, TFGs include edges that represent the fact that

during execution of a concurrent program, execution of a statement may immediately precede execution of a statement from another task. Thus, we divide the set of TFG edges E into two sets, one including edges that represent control flow within a task, and another including edges that represent the may immediately precede relation. The latter kind of edges is referred to as *MIP* edges. For Ada programs, we create a MIP edge from node n to node m and a MIP edge from node m to node n if there is a possibility that during some execution of the system regions of code that correspond to these two nodes may execute in parallel and if neither of nodes m and n is a τ -labeled node². We use an efficient data flow algorithm [9] that computes a conservative estimate of such node pairs (n, m) . Figure 2 shows the TFG where the MIP edges are shown as dashed lines. Note that each such line represents two edges, going in opposite directions.

The number of MIP edges in the TFG for a program depends on the synchronization pattern between the tasks in this program and the number of nodes with labels other than τ . In general, the number of MIP edges is quite large, far exceeding the number of control edges. This means that the analysis algorithm has to propagate information through a large number of edges in the graph, which may lead to poor performance. Thus, a reduction in the number of MIP edges would reduce the run time of the analysis algorithm. In addition, such a reduction could improve the precision of the analysis (the number of false negative results), because some paths through the graph that do not correspond to any real executions of the program would be eliminated. Note that such a reduction has to be *safe* with respect to a property, which means that the reduction does not eliminate sequences of events that correspond to real program executions that violate the property. In the next section we introduce such a safe MIP edge reduction based on partial orders.

3 Reduction of MIP Edges

²The traversal of such τ -labeled nodes from different tasks does not affect the property and constraints.

System	Nodes			Unoptimized			Optimized		
	τ	Local	Global	MIP Edges	MIP Time, s	State Prop, s	MIP Edges	MIP Time, s	State Prop, s
Chiron	45	68	51	236	0.02	1.51	166	0.19	1.37
Chiron	45	68	51	236	0.02	0.71	166	0.18	0.45
Chiron	45	70	51	236	0.02	0.69	166	0.18	0.46
Chiron	45	68	51	236	0.02	1.60	166	0.13	1.04
Chiron	45	68	51	236	0.01	0.70	166	0.19	0.46
Cyclic 2	17	16	9	190	0.01	0.25	50	0.09	0.11
Cyclic 2	17	12	9	116	0.01	0.11	59	0.10	0.09
Cyclic 4	35	16	21	608	0.02	0.91	273	0.10	0.21
Cyclic 4	33	32	21	1552	0.04	70.91	448	0.14	6.59
DPFM 2	8	18	8	136	0.01	0.08	64	0.09	0.08
DPFM 2	8	9	8	80	0.01	0.08	64	0.08	0.07
DPFM 3	11	27	10	216	0.01	0.10	104	0.09	0.09
DPFM 3	12	9	10	96	0.01	0.08	96	0.09	0.09
DPFM 7	23	63	18	536	0.02	2.02	264	0.13	0.92
DPFM 7	24	9	18	160	0.01	0.21	160	0.10	0.21
DPFM 10	32	90	24	776	0.04	66.52	384	0.18	25.22
DPFM 10	33	9	24	208	0.01	0.20	208	0.11	0.20
DPH 2	22	4	16	122	0.01	0.08	110	0.09	0.08
DPH 3	32	4	22	282	0.01	1.80	266	0.09	1.77
MMGT	47	16	78	1312	0.04	3.27	1280	0.15	3.19
TWH-P	5	11	84	4480	0.19	1.91	2933	0.25	0.76
TWH-P	11	7	16	232	0.01	0.09	202	0.09	0.08
TWH-I	5	11	102	6534	0.21	352.81	4644	0.37	221.94
TWH-I	13	7	18	276	0.01	0.09	243	0.09	0.09
RW 2	14	9	12	112	0.01	0.09	96	0.09	0.08
RW 2	15	5	12	88	0.01	0.09	88	0.09	0.08
RW 4	26	9	20	176	0.01	0.16	160	0.10	0.16
RW 4	27	5	20	152	0.01	0.16	152	0.10	0.16
RW 6	38	9	28	240	0.01	0.99	224	0.10	0.99
RW 6	39	5	28	216	0.01	0.92	216	0.10	0.92
RW 8	50	9	36	304	0.01	13.80	288	0.11	13.37
RW 8	51	5	36	280	0.01	9.56	280	0.12	9.57
Ring 2	13	26	20	556	0.02	0.31	388	0.10	0.24
Ring 3	19	39	28	1992	0.05	5.17	1270	0.15	3.25
Ring 4	25	52	36	4356	0.14	193.62	2692	0.29	117.82

Figure 4: Results of the experiment

examples had no MIP edges to start with³ or had no local events. Note that for such cases the partial order optimization does not have to be carried out in full, since one pass over the TFG, property, and constraints is sufficient to figure out if any TFG nodes are locally-labeled. Here we present the results only for the 35 examples where the TFG contained some MIP edges and local events.

Figure 4 shows the results of running the two versions of FLAVERS on these 35 examples. The first column of the table in gives the name of the program used in the example. For some examples we checked multiple properties, which explains the presence of the same program name in multiple rows. Most of these programs are well-known examples

³Although all examples are concurrent, in some of them all threads of control but one contained only τ -labeled nodes, and so no MIP edges were created, according to the simple optimization described in the beginning of Section 3.

from the concurrency literature, such as the dining philosophers and readers-writers examples.

The next three columns list the number of τ -, locally-, and globally-labeled nodes in the TFG. The columns 4-7 list the number of MIP edges, time for adding the necessary MIP edges to the TFG, and the time for the FLAVERS analysis algorithm in the unoptimized version of the algorithm. The last three columns are the same data for the optimized version of the algorithm, where the time for adding the necessary MIP edges to the TFG includes the time for performing the partial order optimization⁴.

It can be seen from the table in Figure 4 that in some cases

⁴Instead of implementing this approach in the way described in Section 3, where some of the existing MIP edges are removed as an optimization, in our implementation we create only those MIP edges that would not be removed by the optimization

this optimization improves the run time of the FLAVERS analysis algorithm by an order of magnitude. In the best case, we removed 74% of the edges and on another problem we saw a speedup of the analysis of 91%. More often, the improvements are not as striking but still quite significant. Only in 7 examples the optimization failed to remove any MIP edges. Out of the examples that benefited from the optimization, on average we removed 25% of the edges, with standard deviation .219. This led to an average speedup of the analysis algorithm of 21%, with standard deviation .2578. The extra overhead in performing the optimization is small, less than .2 seconds for all cases, which is small compared to the run time of the analysis algorithm.

We mentioned in Section 2 that this partial order optimization has the potential of improving the precision of the analysis. One benefit of this improved precision is that the user may not need to use some of the constraints that have to be used with the unoptimized version. (The constraints improve the analysis precision, and some of these improvements may not be necessary if the TFG itself is precise enough.) For each problem where the unoptimized version gave a conclusive result, we also ran the optimized version with a smaller number of constraints. It turned out that the precision improvement resulting from the removal of MIP edges was never sufficient to obtain conclusive results with some of the constraints missing.

5 Conclusion

We have shown how a simple optimization of the program model used by FLAVERS can significantly reduce time requirements of the data flow analysis of user-specified properties of concurrent software. This optimization is dependent on the property of interest and the feasibility constraints that the analysis uses. In particular, this optimization tends to work well with analyses that use feasibility constraints modeling variables local to tasks, because labels that represent operations on such variables tend to be local.

As presented in this paper, this optimization is specific to the Ada concurrency model. The only Ada-specific fact that this approach uses, however, is the presence of communication nodes in the TFG. The proposed approach for applying FLAVERS to concurrent Java programs [10] models the Java concurrency without using the Ada-style communication nodes. We believe that with some simple modifications, this partial order optimization can be easily extended to this Java-specific program model.

Potentially, the partial order optimization of FLAVERS described in this paper can be further improved. For example, all τ -labeled nodes can be considered as locally-labeled nodes. This could lead to a further reduction in the number of MIP edges, although it could have the undesired side-effect of increasing the size of some feasibility constraints necessary for the analysis. We plan to experiment with this trade-off. In addition, we plan to explore a number of directions for

further improvements of FLAVERS. For example, variables can be represented symbolically during the analysis, instead of being represented by finite state automata. This would remove the need to create and store variable automata and thus is likely to improve the analysis performance. We expect this and other optimizations to further improve both space and time requirements of the FLAVERS analysis, increasing its applicability to a wider range of concurrent programs.

REFERENCES

- [1] J. C. Corbett. Evaluating deadlock detection methods for concurrent software. *IEEE Transactions on Software Engineering*, 22(3):161–180, March 1996.
- [2] M. Dwyer and L. Clarke. Data flow analysis for verifying properties of concurrent programs. In *Proceedings of the Second ACM SIGSOFT Symposium on Foundations of Software Engineering*, pages 62–75, December 1994.
- [3] P. Godefroid. Model checking for programming languages using VeriSoft. In *Proceedings of the 24th ACM Symposium on Principles of Programming Languages*, pages 174–186, Jan. 1997.
- [4] P. Godefroid, R. S. Hanmer, and L. J. Jagadeesan. Model checking without a model: an analysis of the Heart-Beat Monitor of a telephone switch using VeriSoft. In *ACM SIGSOFT Proceedings of the 1998 International Symposium on Software Testing and Analysis*, pages 124–133, Mar. 1998.
- [5] P. Godefroid, D. Peled, and M. Staskauskas. Using partial-order methods in the formal validation of industrial concurrent programs. *IEEE Transactions on Software Engineering*, 22(7):496–507, July 1996.
- [6] P. Godefroid and P. Wolper. Using partial orders for the efficient verification of deadlock freedom and safety properties. pages 332–342.
- [7] G. J. Holzmann. *Design and Validation of Computer Protocols*. Prentice Hall Software Series, 1991.
- [8] S. Katz and D. Peled. Verification of distributed programs using representative interleaving sequences. *Distributed Computing*, (6):107–120, 1992.
- [9] G. Naumovich and G. S. Avrunin. A conservative data flow algorithm for detecting all pairs of statements that may happen in parallel. In *Proceedings of the Sixth ACM SIGSOFT Symposium on the Foundations of Software Engineering*, pages 24–34, Nov. 1998.
- [10] G. Naumovich, G. S. Avrunin, and L. A. Clarke. Data flow analysis for checking properties of concurrent Java programs. Technical Report 98-22, University of Massachusetts, Amherst, Apr. 1998. To appear in proceedings of the 1999 IEEE/ACM SIGSOFT International Conference on Software Engineering, May 1999. <http://laser.cs.umass.edu/abstracts/98-022.html>.
- [11] D. Peled. Combining partial order reductions with on-the-fly model checking. In *Proceedings of the 6th International Conference on Computer-aided Verification*, pages 377–390, June 1994.
- [12] A. Valmari. A stubborn attack on state explosion. In *Computer-Aided Verification 90*.