BDD Variable Ordering for Interacting Finite State Machines

Adnan Aziz   Serdar Taşıran   Robert K. Brayton[*]
Email: {adnan,serdar,brayton}@cs.berkeley.edu
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA 94720, USA

Abstract

We address the problem of obtaining good variable orderings for the BDD representation of a system of interacting finite state machines (FSMs). Orderings are derived from the communication structure of the system. Communication complexity arguments are used to prove upper bounds on the size of the BDD for the transition relation of the product machine in terms of the communication graph, and optimal orderings are exhibited for a variety of regular systems. Based on the bounds we formulate algorithms for variable ordering. We perform reached state analysis on a number of standard verification benchmarks to test the effectiveness of our ordering strategy; experimental results demonstrate the efficacy of our approach. The algorithms described in this paper have been implemented in HSIS, a hierarchical synthesis and verification tool currently under development at Berkeley.

1 Introduction

Verification of a design is typically done by modelling it as a finite state machine. Properties to be verified can be specified in a temporal logic[1], or by a task automaton[2]. Verification algorithms proceed by performing some form of traversal of the state transition graph[3]. Large designs arising in practice are invariably the product of small interacting finite state machines. Industrial experience indicates that the largest component machines rarely have more than a hundred states[4]. However forming the product machine leads to the state explosion problem[5, 6]; given n Finite State Machines (FSMs) $\{M_1, M_2, \ldots, M_n\}$, the number of states in the product machine is the product of the number of states in each individual machine. As a result algorithms that explicitly operate on the state space of the product machine may have exponential time and space complexity.

A Binary Decision Diagram (BDD) [7] is a graph based structure used for representing logic functions. It can be used to represent the transition relation of a binary encoded sequential machine implicitly by forming the corresponding characteristic function [8, 9]. This representation can capture the regularity in the transition structure of the machine, and its canonicality makes it very useful in fixed point calculations. BDDs are now routinely used in formal verification[10, 11]. The success of such algorithms depends critically on the size of the resulting BDD's which is very sensitive to the variable ordering chosen. Given a logic function, the problem of finding the ordering which leads to a minimum sized BDD for the function is co-NP complete.

In this paper we address the variable ordering problem for building the transition relation of a set of interacting finite state machines. Using communication complexity[12], upper bounds on the size of the BDD for a specified ordering are derived. Similar results have been shown for combinational circuits in [13, 10] and our work was inspired by these. Indeed the transition relation for the product machine can be viewed as a logic circuit which takes the conjunction of the

transition relations of the component machines, and the techniques of [10] yield upper bounds on the BDD size. We derive a stronger bound and our derivation differs significantly from that of [10]. We stress that we obtain orderings that minimize the representation of the transition relation of the product machine, and as such may not be good for representations of the reached state sets, or equivalent state sets. Indeed there are examples of systems where orderings exist such that the BDD for the transition relation is linear sized, whereas the BDD for the reached state set is exponential sized under any ordering. However our experimental results indicate the orderings we obtain work well in reached state computations. Furthermore, dynamic variable reordering can be used if the BDDs for the reached state sets become unwieldy.

Previous work in this area deals with ordering strategies for combinational [14] and sequential logic circuits [8, 15]. Touati [8] suggests deriving an ordering on the next state variables first using a heuristic based on minimizing the cumulative variable support of the latches. The inputs and present state variables are interleaved with the next state variables; their ordering is derived by standard DFS ordering on the next state logic [14]. Jeong [15] gives efficient algorithms for finding BDD orderings based on the algebraic structure of the circuit. Another approach is based on dynamic ordering[16]. In this the BDD package automatically invokes a reordering routine which seeks to minimize the total number of BDD nodes by permuting small sets of adjacent variables. All these approaches are largely heuristic and do not yield a priori bounds.

In section 2 the basic notions of product machines and process communication graphs are defined. We prove upper bounds on the BDD size in terms of communication graph parameters. We characterize a large variety of interconnect structures for which asymptotically optimum orderings are derived. We also discuss interleaved orderings, and compare our approach with that of [8]. In section 3, we propose various algorithms based on these bounds for the variable ordering problem. In section 4 we present some results based on these algorithms. We conclude by discussing various extensions.

2 Theory

2.1 Definitions

In this section we define finite state machines and the semantics of their interaction. The definitions are motivated by the desire to model hardware designs. At the early stages of VLSI design, components may be incompletely specified, and the wires and states may not be encoded. Our definition allows this flexibility. Non-determinism is commonly used to abstract the environment or parts of the designs, and is reflected in the use of transition relations rather than functions to represent the state dynamics. Basically, we are dealing with non-deterministic Moore machines with transition predicates on the edges.

Definition 1 *A **finite state machine** with state space Q, inputs I, and output alphabet O is characterized by its transition relation $T \subset$*

$Q \times I \times Q$ *and output relation* $\Theta \subset Q \times O$

Systems are described as collections of hardware units, communicating through a set of wires, and driven in lockstep by a single clock; this is the basis for the definition of product machine. Consider a system of n interacting machines $M_1, M_2, \ldots, M_n$. We assume there are no external inputs. Any external inputs I_{ext} can be modeled by adding a one state FSM such that this FSM non-deterministically outputs any of the inputs from I_{ext}. Each component machine M_i has present state x_i, next state variable y_i, and takes as input some subset of outputs of the other machines.

Definition 2 *Given n machines $M_1, M_2, \ldots, M_n$, $M_i = (Q_i, T_i, \Theta_i)$, the* **product machine** $M = M_1 \otimes M_2 \otimes \ldots \otimes M_n$ *is the machine on state space $Q = Q_1 \times Q_2 \times \ldots \times Q_n$ and output space $O = O_1 \times O_2 \times \ldots \times O_n$, characterized by*

- *Output relation*

$$\Theta(x, o) = \prod_{k=1}^{n} \Theta_k(x_k, o_k)$$

- *Transition relation*

$$T(x, y) = \prod_{k=1}^{n} [T_k(x_k, i, y_k) \cdot \Theta_k(x_k, o_k)]$$

where the present state variable is $x = [x_1 x_2 \ldots x_n]$, the next state variable is $y = [y_1 y_2 \ldots y_n]$, and the output variable $o = [o_1 o_2 \ldots o_n]$.. We assume the sets O_i, O_j are disjoint for all i, j.

Binary Decision Diagrams may be used to represent the characteristic function of the transition relation, the output relation, and set of states encountered in verification.

The problem now addressed is: *Given machine M as previously defined, find a good variable ordering for T, using only the communication structure, i.e. without making use of the internal details of the component machines.* We start by restricting attention to the class of variable orderings where variables corresponding to different machines are not interleaved, i.e. only orderings which are permutations of $\{< \vec{x}_1, \vec{y}_1 >, \ldots, < \vec{x}_n, \vec{y}_n >\}$, allowing arbitrary permutations within the variables of a given machine. In section 2.3 we discuss the motivation for restricting our attention to these orderings and examine interleaved orderings.

As a first step towards solving this problem we use communication complexity to prove upper bounds on the size of the BDD corresponding to a permutation of the above form. The following definition is used to make the notion of bit communication complexity precise. To illustrate our approach we use a simplified notion of finite state machine, where the machine output is simply the state. Our results and algorithms are easily extended to machines with outputs that are functions of the state.

Definition 3 *Given machine M as above, a directed edge labeled graph G_M is defined as follows: There is a vertex v_j for each component machine M_j. If the transition relation for M_j depends on the state of M_i, we add edge $e_{i,j}$ labelled with $\vec{x}_i$. We refer to this graph as the* **process communication graph (PCG)**.


Figure 1: PCG for a system of interacting machines; edges are labelled with the present state of originating machine. The machines follow the permutation $\sigma = (413652)$. We are considering $k = 4$. $w_f^\sigma(4) = | x_1 | + | x_3 | + | x_2 |$, $w_r^\sigma(4) = | x_5 | + | x_4 |$. $| x_i |$ is the bit width of variable x_i.

2.2 Upper Bounds

We first develop some intuition behind the bounds derived in this section As an example let M, σ be as given in figure 2. $T(\vec{x}, \vec{y})$ is the product $\prod_{i=1}^{6} T_i$. The number of BDD nodes at level 4 is bounded as follows. Split T into the product of $(T_2 T_6 T_3 T_1)$ and $(T_5 T_4)$. The number of cofactors of T with respect to the variables in $V = \{x_2, y_2, x_6, y_6, x_3, y_3, x_1, y_1\}$ is no more than the number of cofactors of $(T_2 T_6 T_3 T_1)$ with respect to V times the number of cofactors of $(T_5 T_4)$ with respect to V.

$(T_5 T_4)$ has only a limited number of variables ($w_f^\sigma(4)$,defined as below) which are being cofactored in its support. There are at most $2^{w_f^\sigma(4)}$ possible for cofactors of $(T_5 T_4)$.

$(T_2 T_6 T_3 T_1)$ has only a limited number of variables ($w_r^\sigma(4)$, defined as below) remaining after cofactoring. There are at most 2^{2^n} functions on n boolean variables, so there are at most $2^{2^{w_r^\sigma(4)}}$ possible cofactors of $(T_2 T_6 T_3 T_1)$.

Hence the BDD for $T(\vec{x}, \vec{y})$ under the non-interleaved ordering derived from σ has no more than $2^{w_f^\sigma(4)} \cdot 2^{2^{w_r^\sigma(4)}}$ nodes at level 4. This approach can be extended to derive bounds at each level of the BDD and hence on the total size of the BDD.

Theorem 2.1 *Let M be a system of n interacting machines $M_1, M_2, \ldots, M_n$, and let σ be a permutation on $\{1, 2, \ldots, n\}$. Then the number of distinct cofactors of $T(\vec{x}, \vec{y})$ with respect to the variables in $\{\vec{x}_{\sigma(1)}, \vec{y}_{\sigma(1)}, \vec{x}_{\sigma(2)}, \vec{y}_{\sigma(2)}, \ldots, \vec{x}_{\sigma(k)}, \vec{y}_{\sigma(k)}\}$ is bounded by $2^{w_f^\sigma(k)} \cdot 2^{2^{w_r^\sigma(k)}}$, where*

$$
\begin{aligned}
w_f^\sigma(k) &= \quad \textit{number of distinct bits communicated from} \\
&\qquad \{M_{\sigma(1)}, \ldots, M_{\sigma(k)}\} \textit{ to } \{M_{\sigma(k+1)}, \ldots, M_{\sigma(n)}\} \\
w_r^\sigma(k) &= \quad \textit{number of distinct bits communicated from} \\
&\qquad \{M_{\sigma(k+1)}, \ldots, M_{\sigma(n)}\} \textit{ to } \{M_{\sigma(1)}, \ldots, M_{\sigma(k)}\}
\end{aligned}
$$

Proof: Refer to [17] for a detailed proof. ∎

As a corollary to the above theorem we get an upper bound on the number of nodes for the BDD for $T(\vec{x}, \vec{y})$:

Corollary 2.1 *The number of nodes in the BDD for $T(\vec{x}, \vec{y})$ for the ordering $\vec{x}_{\sigma(1)} \prec \vec{y}_{\sigma(1)} \prec \ldots \prec \vec{x}_{\sigma(n)} \prec \vec{y}_{\sigma(n)}$, is bounded above by :*

$$S^\sigma = \sum_{i=1}^{n} (2^{2(|x_{\sigma(i)}|)} \cdot 2^{w_f^\sigma(i+1)} \cdot 2^{2^{w_r^\sigma(i+1)}}) \tag{1}$$

A looser bound is

$$M^\sigma = n \cdot c \cdot 2^{w_f^\sigma} \cdot 2^{2^{w_r^\sigma}} \qquad (2)$$

where

- w_f^σ *is the maximum number of forward crossing bits across any partition induced by* σ

- w_r^σ *is the maximum number of reverse crossing bits across any partition induced by* σ

- c *depends only on the maximum number of state bits in some machine*

Proof:

From Theorem 2.1, it follows that the number of nodes in the BDD for $T(\vec{x}, \vec{y})$ at level $k + 1$, where the variables from the FSMs $M_{\sigma(1)}, \ldots, M_{\sigma(k)}$ have been cofactored out, is bounded by $2^{w_f^\sigma(k)} \cdot 2^{w_r^\sigma(k)}$. The total number of nodes at all the levels between level k and level $k + 1$ is no more than $2^{2(|x_{\sigma(i)}|)}$ times the number at level k. Summation over k yields the result. ∎

Remark: Given k, there are always choices of the individual machines such that there are at least $2^{w_f^\sigma(k)}$ distinct cofactors of $T(\vec{x}, \vec{y})$ with respect to $\{\vec{x}_{\sigma(1)}, \vec{y}_{\sigma(1)}, \vec{x}_{\sigma(2)}, \vec{y}_{\sigma(2)}, \ldots, \vec{x}_{\sigma(k)}, \vec{y}_{\sigma(k)}\}$, ie the bound of Theorem 2.1 is tight in the first term. Details are available in [17].

2.3 Interleaved Orderings

Typically, communication within a machine is dense, ie each bit of the next state depends on all the present state bits, since otherwise the machine would have a trivial factorization. Reasoning as in Corollary 2.1, it would seem that interleaving the variables leads to increased communication complexity and higher bounds. This suggests that variables corresponding to a single machine be ordered contiguously.

However, there are situations where interleaving state variables from different machines may be superior to a non-interleaved ordering. This can happen when the given partitioning of the system into interacting FSMs may not yield a sparse communication structure. Consider a product machine in which a component machine has a transition relation $T = [y \oplus (x_1^A x_1^B + x_2^A x_2^B + \ldots + x_n^A x_n^B)]$, where x_i^A, x_i^B are the present state bits from machine M^A and M^B. In this case, a non-interleaved ordering will result in an exponential sized BDD for T, where as the ordering $< x_1^A, x_1^B, \ldots, x_n^A x_n^B >$, is optimum and yields a linear sized BDD for T.

Interleaved orderings are also superior in the context of equivalent state computations. If states are equivalent only to themselves, the equivalence predicate is equality. For the equality relation a noninterleaved ordering is exponential, and an interleaved ordering is linear.

2.3.1 Comparison with Touati's heuristic

Touati et. al.[8] consider a BDD based approach to the problem of the equivalence of sequential hardware consisting of latches and logic gates. Their heuristic for variable ordering proceeds by first finding a permutation σ^* on the latches which minimizes the support, i.e. minimizes the following cost function:

$$\text{cost}(\sigma) = \sum_{1 \le j \le n} \left| \bigcup_{1 \le i \le j} \text{supp}(f_{\sigma_j}) \right|$$

where $|A|$ denotes the cardinality of A, and $\text{supp}(f_i)$ is the set of variables in the support of f_i. Malik's heuristic [14] is used to order the


Figure 2: An example on which Touati's heuristic fails

supports of the f_i, $\text{supp}(f_i)$ individually. Finally input and output variables are interleaved as follows: $\text{supp}(f_{\sigma(1)}), y_{\sigma(1)}, \ldots, \text{supp}(f_{\sigma(n)} - \bigcup_{1 \le i \le n-1} \text{supp}(f_{\sigma(i)})), y_{\sigma(n)}$.

Naturally, this ordering procedure can be used to derive orderings for systems of interacting FSMs. The component machines correspond to latches, and their fanins correspond to the support. $\text{cost}(\sigma)$ can be calculated directly from the process communication graph. However $\text{cost}(\sigma)$ is not directly correlated to the communication complexity of σ. Consider the system of FSMs interacting through a binary tree as shown in Figure 2. The ordering $<< M_1, M_2, \ldots, M_{2^n-1} >>$ is optimum for $\text{cost}(\sigma)$ but has high communication complexity as is seen in the partition $U = \{1, 2, \ldots, 2^{n-1} - 1\}$, $V = \{2^{n-1}, \ldots, 2^n - 1\}$, for which $w_f = 2^{n-1}$. In fact, there exists an ordering with low communication complexity, namely the ordering returned by lexicographically first DFS, for which $w_f = 2$. Our experimental results on actual BDD size in section 4.2 demonstrate the failure of Touati's heuristic on systems with a tree-like PCG.

2.4 Optimum Orderings

We can use the bounds derived above to get optimum variable orderings for a variety of sparse interconnect structures. The following lemma shows that for certain recursive structures whose communication complexity is independent of the number of nodes in the graph, the transition relation for the product machines grows linearly in n.

Lemma 2.1 *Let* $M(n)$ *be the composition of* n *component machines* $\{M_1, \ldots, M_n\}$. *Suppose there exist constants* w_f^*, w_r^*, *and* b *such that for for all n,*

1. *there exists an ordering* σ_n^* *of the vertices of the PCG such that for each k*

$$w_f^{\sigma_n^*}(k) \le w_f^*$$
$$w_r^{\sigma_n^*}(k) \le w_r^*$$

2. *the state of each* M_k^n *can be encoded in not more than* b *bits*

Then for all n *there is exists a variable ordering such that the BDD for* $T(n) = T_1 \times T_2 \ldots \times T_n$ *has at most* $c \cdot n$ *nodes, where* c *is independent of* n.

Proof: We use the bound of corollary 2.1, (equation 2)

$$M^\sigma = n \cdot 2^{2 \max_i (|x_i|)} \cdot \max_k (2^{w_f^\sigma(k)} \cdot 2^{2^{w_r^\sigma(k)}})$$

For each machine M_n using an ordering defined by σ_n^* leads to a BDD for T_n such that

$$|T_n| \le n \cdot 2^{2 \cdot \max_k |x_k|} \cdot \max_k (2^{w_f^{\sigma_n^*}(k)} \cdot 2^{2^{w_r^{\sigma_n^*}(k)}})$$

285

Ring   Linear Array   Binary Tree

Figure 3: A variety of sparse interconnect structures commonly encountered in verification

By the hypothesis of the lemma, $\max_k \mid x_k \mid \leq b$, $\max_k w_f^{\sigma_n^*}(k) \leq w_f^*$, and $\max_k w_r^{\sigma_n^*}(k) \leq w_r^*$.
Therefore

$$\mid T_n \mid \leq n \cdot 2^{2 \cdot b} \cdot 2^{w_f^*} \cdot 2^{2^{w_r^*}} \leq c \cdot n$$

where $c = 2^{2 \cdot b} \cdot 2^{w_f^*} \cdot 2^{2^{w_r^*}}$

∎

Armed with this lemma, we can find variable orderings for a variety of interconnect structures (parametrized by n) that yield linear sized BDDs, and so are optimum within constants. Figure 3 illustrates some communication graphs which satisfy the hypothesis of the lemma.

3 Algorithms

The upper bounds S^σ and M^σ of section 2 can be used to obtain a good variable ordering for the transition relation of a given product machine M composed of $\{M_1, \ldots M_n\}$. We first extract the process communication graph (PCG). The parameters in the bounds can be calculated directly from the PCG. We then find a permutation σ^* on the vertices of the graph that minimizes the bound. A non-interleaved variable ordering corresponding to σ is simply one in which all variables corresponding to the machine $M_{\sigma^*(1)}$ appear first (in any internal order amongst themselves), followed by all variables from $M_{\sigma^*(2)}$, etc.

Finding an optimum permutation on the vertices of the communication graph is a hard problem, akin to the Travelling Salesman Problem. We conjecture that it is NP-complete. Exhaustive search has factorial complexity, and dynamic programming yields an algorithm with complexity $O(n^3 2^n)$. We now discuss an exact branch and bound procedure and some heuristics for finding good permutations.

3.1 Branch and Bound

Consider the problem of obtaining a permutation σ^* which minimizes $S^\sigma = \sum_{i=1}^n g_i(\sigma)$. This problem lends itself to a branch and bound algorithm solution. This can be better understood by a closer look at the sum for S^σ. We have to minimize S^σ over all permutations. Let $S^{\hat{\sigma}}$ be the smallest sum obtained from the permutations considered so far. For the next permutation σ' being examined, if for some k $\sum_{i=1}^k g_i(\sigma') \geq S^{\hat{\sigma}}$, then the summation can be terminated. Moreover, any permutation that has $< \sigma'(1) \ldots \sigma'(k) >$ as a prefix can be eliminated. This provides a way to prune the search tree. In fact any permutation that has the FSMs above in that order can be pruned. But it is expensive to store and check for this information. The same formulation yields a branch and bound algorithm to find a permutation minimizing M^σ.

To increase the probability of pruning, we first calculate a good initial guess. This is explained in more detail in the next subsection. Then, we go over all permutations lexicographically, pruning wherever possible. The exhaustive search tree is never constructed, since this would require excessive memory and time. Instead, when a "bad" (in the sense described above) prefix is discovered, the algorithm branches to the next permutation in lexicographic order which does not have the bad prefix.

The input to the branch and bound algorithm is a process communication graph. Observe that all parameters in the bound can be calculated directly from the PCG. For each permutation examined, for each k, the number of bits crossing the cut in the forward and reverse directions are counted separately. Then the cost for the cuts so far for that permutation is calculated. Either pruning takes place before all cuts are examined, or the permutation considered is the best seen so far. The algorithm proceeds in this manner until all permutations are exhausted.

The memory used by this algorithm is linear in the size of the process communication graph. The algorithm can take factorial time in the worst case. It was necessary to experiment with the algorithm to see how effective the pruning is at reducing the search.

3.2 Heuristics

Several schemes, with varying degrees of sophistication, exist for picking the initial guess for the optimum permutation. The simplest is to choose vertices one at a time in the greedy sense, ie, choose the next vertex to minimize the partial cost. This greedy approach can be extended to an algorithm that has bounded look-ahead k. The algorithm proceeds recursively by computing all possible choices for the first k vertices in the permutation. It chooses the best among these and recursively completes the ordering. While improving the quality of the initial guess, look-aheads of k increase the time complexity of the algorithm by adding a factor of n^k. This is acceptable compared with the complexity of the branch and bound algorithm, since a good starting point can prune the search space drastically. When n is large, the initial guesses themselves can be used as approximate solutions. We implemented a simple greedy algorithm to minimize S^σ, and algorithms with look-aheads of 2 to minimize S^σ and M^σ.

4 Results

4.1 Computing an Optimum Permutation

The branch and bound algorithms that find permutations that minimize S^σ and M^σ were implemented using a greedily generated initial guess. Still for more than 10 to 15 vertices in the communication graph, the exact branch and bound algorithm was too slow. Our experiments show that a permutation which minimizes S^σ generated using a look-ahead of 2 yields a solution that is sufficiently close to being optimum and has negligible running time. When the communication graph has a regular structure (e.g. a mesh, tree, ring, etc), this permutation is often optimum. Thus, this is the method of choice.

4.2 Correlation between bound and actual size of BDD for Transition Relation

For assessing the usefulness of our bound as a measure of the actual BDD size, and checking the validity of assumptions, we needed a set of representative examples to test our algorithms. We artificially constructed several product machines according to various interconnection schemes. In each case the component machines had a small state space and a randomly chosen (possible nondeterministic) transition relation.

We extracted the process communication graphs, and found variable orderings as follows:

ran_vars All variables are ordered randomly

ran_comps Non-interleaved ordering where component machines are ordered randomly

touati_heur Interleaved ordering as described in section 2.3.1

min_comm Non-interleaved order where components are ordered to minimize communication complexity as described in section 2.1

286

EXAMPLE	ran_comps	$touati\_heur$	min_comm
Acyclic	$> 10^6$	5,336	1,188
Cyclic	$> 10^6$	21,174	3,185
Few_Ran_C	$> 10^6$	35,444	1,085
RT	$> 10^6$	2,455	421
Mesh20	$> 10^6$	21,573	9,203
Tree31	$> 10^6$	75,034	1,134
Ran_Mesh	43,727	17,435	6,511
AcyclicII	$> 10^6$	112,379	10,756
CyclicII	$> 10^6$	583,603	79,199
ran_vars: In all cases $> 10^6$ BDD nodes			

Table 1: Number of BDD nodes for the TR of the product machine; in all cases time to find the ordering was negligible
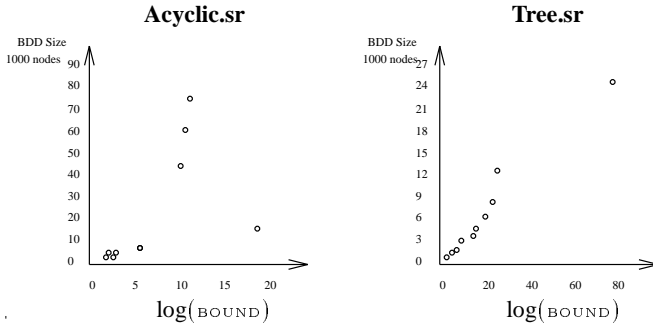

Figure 4: Correlation between bound S^σ and actual BDD size for several permutations σ.

4.3 Performance on Reachability Analysis

As mentioned in the introduction, an ordering which minimizes the size of the BDD representing the transition relation does not necessarily lead to an ordering which is good for general verification calculations. Reached state computation is a core routine in verification. A truer measure of the effectiveness of our ordering is given by our performance on reached state analysis for realistic verification benchmarks.

In Table 2 we describe results for reached state analysis using ordering derived from *min_comm* on common verification examples. We computed the reached state set in the most direct fashion, namely by directly computing the image of the reached state set under the global transition relation. There are more sophisticated ways of computing the image that involve using don't cares to minimize the BDDs, and partitioned transition relations to avoid building the full transition relation. However this simple experiment is enough to give us a good idea of the performance of our ordering.

Specifically we report the following three statistics—

TR size : Number of nodes in BDD for transition relation

Max-Reached : Number of nodes in largest BDD representing a set of states encountered during reached state computation

Time-Reached : Time taken to perform reached state computation

The ordering given by *touati_heur* was identical to that given by *min_comm* on all examples except for *DME15* where it was orders of magnitude worse. This is explained by the fact that *DinPhil32, MilSched16, Fis4* have ring like structure, and *touati_heur* will find the asymptotically best ordering for rings. However, as pointed out in section 2.3.1, *touati_heur* is especially bad for trees, and this is brought across in *DME15*. Also, the best ordering that could be obtained by hand for *Fis4* was much worse than that generated by *min_comm*.

EXAMPLE	min_comm		
	TR size	Max-reached	Time-reached (sec)
Dynachek3	332	37	2.0
DME16	4,982	45	1.2
DinPhil32	4,863	304	26.0
MilSched16	5,003	938	33.6
Fis5	22,227	27,901	1145.4
2MDLC	24,487	1,335	635

Table 2: Results on reached state analysis – Number of BDD nodes for TR, largest reached state set, and time to perform reached state analysis

As seen in Table 1, in all cases, our approach significantly reduces the BDD size as compared with random orderings. All algorithms took only a few seconds to compute the ordering. The running time was negligible when compared to the time taken to read in the example and build the BDD. Note that *Mesh20*, which has high communication complexity (as shown in [17]) had a large BDD compared to *Tree31* which has low communication complexity (as discussed in section 2.4). Also observe that Touati's heuristic does particularly badly on *Tree31* and *RT*, corroborating the discussion of section 2.3.1. We also experimented with some interleaving, similar to that given by Touati's heuristic. In many cases this allowed further reduction in BDD size. We concluded that ordering based on *minimizing the communication complexity, while allowing some heuristic interleaving works best in practice*.

The bound S^σ derived in Corollary 2.1 is an upper bound. To examine the relationship S^σ between the actual number of nodes in the BDD for the transition relation under the non-interleaved ordering derived from σ, we took several different permutations of the components and computed the bound. We then built the BDD and plotted the actual BDD size against the bound. The plots are seen in figure 4; they demonstrate a strong correlation between the bound and the actual BDD size.

5 Conclusion

We addressed the problem of deriving good variable orderings for the BDD representation of a system of interacting finite state machines for formal verification applications. Towards this end we introduced the notion of the process communication graph and proved results connecting BDD size to the communication graph. We justified the decision to derive orderings based only on knowledge of the communication graph. We use the bounds to formulate fast heuristic algorithms for variable ordering. The experimental results show good correlation between the BDD size predicted by the bound and the actual BDD size, and comparing our results on the transition relation size with orderings derived from Touati's heuristic validates our decision to use orderings which minimize the communication complexity through the communication graph. Our performance on reached state analysis for verification benchmarks further demonstrates the effectiveness of our approach.

6 Future Work

The BDD variable ordering problem for combinational logic circuits still bears scope for further research: consider the recent results of [16, 18]. As shown in the original inspiration for our work [13, 10], circuit structure can be used to bound the size of the BDDs for the outputs. As in the case of interacting FSMs, communication between gates can be analyzed to derive variable orderings that minimize the bound. We are investigating the correlation between the bound, and actual BDD size. Since the number of gates is large, even greedy algorithms such as those described in section 3 take prohibitively long; we plan to experiment with graph partitioning and hierarchical clustering procedures to derive orderings on the gates that minimize communication.

The *quantification ordering* problem arises in various contexts, including performing reached state analysis [3, 8] and composition of combinational relations [19], and is defined as follows: Given BDDs for a set of terms $T_1, T_2, \ldots, T_n$ and variables $x_1, x_2, \ldots, x_k$, efficiently compute the following expression:

$$T \;\; = \;\; \exists x_1, x_2, \ldots, x_k \, [\prod_{i=1}^{n} T_i \;]$$

Computing T by first forming $[\prod_{i=1}^{n} T_i \;]$ is frequently infeasible, as the BDDs for intermediate stages of the product can be very large. As a result, research has focussed on building T incrementally by forming products of subsets of the T_i's and smoothing out variables that do not appear in terms outside the subset.

Let α be a permutation on the variables $x_1, x_2, \ldots, x_k$. α leads to a natural schedule for forming T – take all terms depending on x_{α_1}, conjunct them and quantify out x_{α_1}, and recursively continue this procedure with the conjunctant and the remaining terms and variables. Given a variable ordering σ and a set of BDDs $T_1, T_2, \ldots, T_n$, the techniques given in this paper can be used to derive an upper bound on the size of the BDD for $[\prod_{i=1}^{n} T_i \;]$ under the variable ordering σ. Thus given any schedule we can a priori obtain a bound on the size of the largest BDD encountered in forming T according to this schedule. In the spirit of this paper, we can obtain a schedule that minimizes the bound. We plan to implement and experiment with schedules that minimize the bound.

References

[1] Z. MANNA AND A. PNEULI, "VERIFICATION OF CONCURRENT PROGRAMS: THE TEMPORAL FRAMEWORK," IN *The Correctness Problem in Computer Science* (R. S. BOYER AND J. S. MOORE, EDS.), INT. LECTURE SERIES IN COMPUTER SCIENCE, PP. 215–273, LONDON: ACADEMIC PRESS, 1981.

[2] Z. HAR'EL AND R. P. KURSHAN, "SOFTWARE FOR ANALYTICAL DEVELOPMENT OF COMMUNICATION PROTOCOLS," *AT&T Technical Journal*, PP. 45–59, JAN. 1990.

[3] J. R. BURCH, E. M. CLARKE, K. L. MCMILLAN, AND D. L. DILL, "SYMBOLIC MODEL CHECKING: 10^{20} STATES AND BEYOND," *Information and Computation*, VOL. 98, NO. 2, PP. 142–170, 1992.

[4] G. YORK. PERSONAL COMMUNICATION, FEB. 1993.

[5] O. GRÜMBERG AND D. E. LONG, "MODEL CHECKING AND MODULAR VERIFICATION," IN *Proc. of CONCUR '91: 2nd Inter. Conf. on Concurrency Theory* (J. C. M. BAETEN AND J. F. GROOTE, EDS.), VOL. 527 OF *Lecture Notes in Computer Science*, SPRINGER-VERLAG, AUG. 1991.

[6] A. AZIZ AND R. K. BRAYTON, "VERIFYING INTERACTING FINITE STATE MACHINES," TECH. REP. UCB/ERL M93/52, ELECTRONICS RESEARCH LAB, UNIV. OF CALIFORNIA, BERKELEY, CA 94720, JULY 1993.

[7] R. BRYANT, "GRAPH-BASED ALGORITHMS FOR BOOLEAN FUNCTION MANIPULATION," *IEEE Trans. Computers*, VOL. C-35, PP. 677–691, AUG. 1986.

[8] H. TOUATI, H. SAVOJ, B. LIN, R. K. BRAYTON, AND A. L. SANGIOVANNI-VINCENTELLI, "IMPLICIT STATE ENUMERATION OF FINITE STATE MACHINES USING BDD'S," IN *Proc. Intl. Conf. on Computer-Aided Design*, PP. 130–133, NOV. 1990.

[9] O. COUDERT AND J. C. MADRE, "A UNIFIED FRAMEWORK FOR THE FORMAL VERIFICATION OF SEQUENTIAL CIRCUITS," IN *Proc. Intl. Conf. on Computer-Aided Design*, PP. 126–129, NOV. 1990.

[10] K. L. MCMILLAN, *Symbolic Model Checking*. KLUWER ACADEMIC PUBLISHERS, 1993.

[11] A. AZIZ, F. BALARIN, R. K. BRAYTON, S.-T. CHENG, R. HOJATI, T. KAM, S. C. KRISHNAN, R. K. RANJAN, A. L. SANGIOVANNI-VINCENTELLI, T. R. SHIPLE, V. SINGHAL, S. TASIRAN, AND H.-Y. WANG, "HSIS: A BDD-BASED ENVIRONMENT FOR FORMAL VERIFICATION," IN *Proc. of the Design Automation Conf.*, JUNE 1994.

[12] J. E. HOPCROFT AND J. D. ULLMAN, *Introduction to Automata Theory, Languages and Computation*. ADDISON-WESLEY, 1979.

[13] C. BERMAN, "ORDERED BINARY DECISION DIAGRAMS AND CIRCUIT STRUCTURE," IN *Proc. Intl. Conf. on Computer Design*, OCT. 1989.

[14] S. MALIK, A. R. WANG, R. K. BRAYTON, AND A. SANGIOVANNI-VINCENTELLI, "LOGIC VERIFICATION USING BINARY DECISION DIAGRAMS IN A LOGIC SYNTHESIS ENVIRONMENT," IN *Proc. Intl. Conf. on Computer-Aided Design*, PP. 6–9, NOV. 1988.

[15] S.-W. JEONG, B. PLESSIER, G. D. HACHTEL, AND F. SOMENZI, "VARIABLE ORDERING FOR FSM TRAVERSAL," IN *Proc. Intl. Conf. on Computer-Aided Design*, 1991.

[16] R. RUDELL, "DYNAMIC VARIABLE ORDERING FOR BINARY DECISION DIAGRAMS," IN *Proc. Intl. Conf. on Computer-Aided Design*, PP. 42–47, NOV. 1993.

[17] A. AZIZ, S. TASIRAN, AND R. K. BRAYTON, "BDD VARIABLE ORDERING FOR INTERACTING FINITE STATE MACHINES," TECH. REP. UCB/ERL M93/71, ELECTRONICS RESEARCH LAB, UNIV. OF CALIFORNIA, BERKELEY, CA 94720, SEPT. 1993.

[18] H. FUJII, G. OOTOMO, AND C. HORI, "INTERLEAVING VARIABLE ORDERING METHODS FOR ORDERED BINARY DECISION DIAGRAMS," IN *Proc. Intl. Conf. on Computer-Aided Design*, PP. 38–41, NOV. 1993.

[19] R. K. BRAYTON, M. CHIODO, R. HOJATI, T. KAM, K. KODANDAPANI, R. P. KURSHAN, S. MALIK, A. L. SANGIOVANNI-VINCENTELLI, E. M. SENTOVICH, T. SHIPLE, K. J. SINGH, AND H.-Y. WANG, "BLIF-MV: AN INTERCHANGE FORMAT FOR DESIGN VERIFICATION AND SYNTHESIS," TECH. REP. UCB/ERL M91/97, ELECTRONICS RESEARCH LAB, UNIV. OF CALIFORNIA, BERKELEY, CA 94720, NOV. 1991.