

# Efficient Verification of Halting Properties for MPI Programs with Wildcard Receives

Stephen F. Siegel\*

Laboratory for Advanced Software Engineering Research  
Department of Computer Science, University of Massachusetts  
Amherst MA 01003, USA  
siegel@cs.umass.edu  
<http://laser.cs.umass.edu>

**Abstract.** We are concerned with the verification of certain properties, such as freedom from deadlock, for parallel programs that are written using the Message Passing Interface (MPI). It is known that for MPI programs containing no “wildcard receives” (and restricted to a certain subset of MPI) freedom from deadlock can be established by considering only synchronous executions. We generalize this by presenting a model checking algorithm that deals with wildcard receives by moving back and forth between a synchronous and a buffering mode as the search of the state space progresses. This approach is similar to that taken by partial order reduction (POR) methods, but can dramatically reduce the number of states explored even when the standard POR techniques do not apply.

## 1 Introduction

It is well-known that finite-state verification techniques, such as model checking, suffer from the *state explosion problem*: the fact that the number of states of a concurrent system may – and often does – grow exponentially with the size of the system. Many different approaches have been studied to counteract this difficulty. These include partial order reduction (POR) methods, data abstraction, program slicing, and state compression techniques, to name only a few.

For the most part, these approaches have been formulated in very general frameworks. Their generality is both a strength and a weakness: the methods can be broadly applied, but may miss opportunities for reduction in specific situations. This observation has led to interest in more *domain-specific* approaches. The idea is to leverage knowledge of the restrictions imposed by a particular programming domain, or of common idioms used in the domain, in order to gain greater reductions than the generic algorithms allow. An example of this approach for concurrent Java programs is given in [2], where analysis that identifies common locking patterns, among other things, is exploited to dramatically improve the generic POR algorithms.

---

\* Research supported by the U.S. Army Research Laboratory and the U.S. Army Research Office under agreement number DAAD190110564.

This paper is concerned with the domain of parallel programs that employ the Message Passing Interface (MPI). The MPI Standard [9,10] specifies the syntax and semantics for a large library of message passing functions with bindings in C, C++, and Fortran. For many reasons – portability, performance, the broad scope of the library, and the wide availability of quality implementations – MPI has become the de facto standard for high-performance parallel computing. In addition, we focus on a particular class of properties of MPI programs, which we call *halting properties*: claims on the state of a program whenever execution halts, whether due to deadlock, or to normal termination. Freedom from deadlock is an example of a halting property; another would be an assertion on the values of variables when a program terminates.

Some explanation of the most essential MPI functions is required for what follows. The basic MPI function for sending a message to another process is `MPI_SEND`. To use it, one must specify the destination process and a message tag, in addition to other information. The corresponding function for receiving a message is `MPI_RECV`. In contrast to `MPI_SEND`, an `MPI_RECV` statement may specify its source process, or it may use the *wildcard* value `MPI_ANY_SOURCE`, indicating that this statement will accept a message from any source. Similarly, it may specify the tag of the message it wishes to receive, or it may use the wildcard value `MPI_ANY_TAG`. A receive operation that uses either or both wildcards is called a *wildcard receive*. The use of wildcards and tags allows for great flexibility in how messages are selected for reception.

Previous work has established that if a program (restricted to a certain subset of MPI) contains no wildcard receives, then a suitable model  $\mathcal{M}$  of that program can be constructed with the following property:  $\mathcal{M}$  is deadlock-free if, and only if, no synchronous execution of  $\mathcal{M}$  can deadlock [12, Theorem 7.4]. This is exactly the kind of result we are after, as the need to represent all possible states of message channels is often a significant source of state explosion. Unfortunately, wildcard receives are common in actual MPI programs, and the theorem may fail if the hypothesis on wildcard receives is removed [12, Sec. 7.3].

The approach of this paper generalizes the earlier result in three ways. First, it shows that the hypothesis forbidding wildcard receives may be relaxed to allow the use of `MPI_ANY_TAG`, with no ill effects. Second, the range of properties is expanded to include all halting properties. But most importantly, it provides a model checking algorithm that deals with `MPI_ANY_SOURCE` by moving back and forth between a synchronous and a buffering mode as the search of the state space progresses. This approach is similar to that taken by POR methods, but can dramatically reduce the number of states explored even when the standard POR techniques do not apply.

The discussion proceeds as follows. Section 2 establishes the precise definition of a model of an MPI program, and of the execution semantics of such a model. The definition of a halting property and the statement of the main theorem are given in Sec. 3. Section 4 deals with consequences of the main theorem. These include a bounded model checking algorithm for halting properties; the consequences for programs that do not use `MPI_ANY_SOURCE` are also ex-

plored. Section 5 discusses the relationship with the standard POR techniques. Results of an empirical investigation are presented in Sec. 6, and conclusions are drawn in Sec. 7. Proofs of the theorems, a description of the program and model for each example, complete MPI/C source code for the examples, and all experimental results can be downloaded from <http://laser.cs.umass.edu/~siegel/projects>.

## 2 Models of MPI Programs

For the purposes of this paper, an MPI program consists of a fixed number of concurrent processes, each executing its own code, with no shared variables, that communicate only through the MPI functions. The precise notion of a *model* of such a program is defined below. While there are many issues that arise in creating models from code, these are beyond the scope of this paper, and the reader is referred to [12] for a discussion of this subject and some examples. It is argued there that this notion of model suffices to represent MPI\_SEND, MPI\_RECV, MPI\_SENDRECV (which concurrently executes one send and one receive operation), as well as the 16 collective functions, such as MPI\_BCAST, MPI\_ALLREDUCE, etc. The definition of receiving states here is slightly more general, in order to accommodate a new way to deal with tags, explained below.

### 2.1 Definition of a Model of an MPI Program

An *MPI context* is a 7-tuple

$$\mathcal{C} = (\text{Proc}, \text{Chan}, \text{sender}, \text{receiver}, \text{msg}, \text{loc}, \text{com}).$$

The first two components are finite sets, representing the set of *processes*, and the set of communication *channels*, respectively. The next two components are functions from Chan to Proc; they define the sending and receiving process for each channel. The function msg assigns, to each  $c \in \text{Chan}$ , a nonempty set msg( $c$ ); this is the set of messages that can be sent over channel  $c$ . The final two components are functions of Proc. For  $p \in \text{Proc}$ , loc( $p$ ) is a finite set representing the set of *local events* for  $p$ , while com( $p$ ) is defined to be the set of *communication events* for  $p$ , namely, the set of send and receive symbols

$$\{c!x, d?y \mid c, d \in \text{Chan}, x \in \text{msg}(c), y \in \text{msg}(d), \text{sender}(c) = p = \text{receiver}(d)\}.$$

Finally, for all  $p, q \in \text{Proc}$ , we assume  $\text{loc}(p) \cap \text{com}(q) = \emptyset$ , and  $p \neq q \Rightarrow \text{loc}(p) \cap \text{loc}(q) = \emptyset$ .

Let  $p \in \text{Proc}$ . An *MPI state machine for  $p$  under  $\mathcal{C}$*  is a 6-tuple

$$M_p = (\text{States}_p, \text{Trans}_p, \text{src}, \text{des}, \text{label}, \text{start}_p)$$

where  $\text{States}_p$  and  $\text{Trans}_p$  are sets, src and des are functions from  $\text{Trans}_p$  to  $\text{States}_p$ , label is a function from  $\text{Trans}_p$  to  $\text{loc}(p) \cup \text{com}(p)$ , and  $\text{start}_p \in \text{States}_p$ .

We do not use the subscript  $p$  for the functions  $\text{src}$ ,  $\text{des}$ , and  $\text{label}$ , because the process  $p$  will always be clear from the context. Finally, we require that every state  $u$  must fall into one of 5 categories, which are determined by the transitions departing from  $u$ . First, we define the following:

$$\begin{aligned}
 R(u) &= \{(d, y) \mid d \in \text{Chan}, y \in \text{msg}(d), \exists t \in \text{Trans}_p: \text{src}(t) = u \wedge \text{label}(t) = d?y\} \\
 Q(u) &= \{d \in \text{Chan} \mid \exists y \in \text{msg}(d): (d, y) \in R(u)\} \\
 R_d(u) &= \{y \in \text{msg}(d) \mid (d, y) \in R(u)\} \quad (d \in Q(u)).
 \end{aligned}$$


Now the 5 possibilities for  $u$  are as follows:

1.  $u$  is a *final state*: there are no transitions departing from  $u$ ,
2.  $u$  is a *local-event state*: there is at least one transition departing from  $u$ , and the transitions departing from  $u$  are labeled by local events for  $p$ ,
3.  $u$  is a *sending state*: there is precisely one transition departing from  $u$  and it is labeled by a send event for  $p$ ,
4.  $u$  is a *receiving state*: there is at least one transition departing from  $u$ , and the transitions departing from  $u$  are labeled by distinct receive events for  $p$ , or
5.  $u$  is a *send-receive state* (see Fig. 1):  $R(u) \neq \emptyset$ , and there is a  $c \in \text{Chan}$  with  $\text{sender}(c) = p$ , an  $x \in \text{msg}(c)$ , a state  $u'$ , and states  $v(d, y)$  and  $v'(d, y)$  for all  $(d, y) \in R(u)$ , such that the following all hold:
  - (a) the set of transitions departing from  $u$  consists of one transition to  $u'$  whose label is  $c!x$ , and, for each  $(d, y) \in R(u)$ , one transition labeled  $d?y$  to  $v(d, y)$ ,
  - (b) for each  $(d, y) \in R(u)$ , there is precisely one transition departing from  $v(d, y)$ , it is labeled  $c!x$ , and it terminates in  $v'(d, y)$ , and
  - (c) for each  $(d, y) \in R(u)$ , there is a transition from  $u'$  to  $v'(d, y)$ , it is labeled  $d?y$ , and these make up all the transitions departing from  $u'$ .

The point of the send-receive state is to model the `MPI_SENDRECV` function, which executes one send and one receive operation concurrently. This is modeled by allowing the send and receive to happen in either order.

Finally, a *model*  $\mathcal{M}$  of an MPI program is a pair  $(\mathcal{C}, M)$ , where  $\mathcal{C}$  is a context and  $M$  is a function that assigns, to each  $p \in \text{Proc}$ , an MPI state machine  $M_p$  for  $p$  under  $\mathcal{C}$ , such that  $\text{States}_p \cap \text{States}_q = \emptyset = \text{Trans}_p \cap \text{Trans}_q$  for  $p \neq q$ .

Given an MPI program, one may construct a model using one channel  $c_{p,q}$ , with  $\text{sender}(c_{p,q}) = p$  and  $\text{receiver}(c_{p,q}) = q$ , for each  $(p, q) \in \text{Proc} \times \text{Proc}$ . To translate a receive statement  $r$  it suffices to specify the sets  $Q(u)$  and  $R_d(u)$  for



**Fig. 1.** A send-receive state  $u$  with  $Q(u) = \{d, e\}$ ,  $R_d(u) = \{1, 2\}$ ,  $R_e(u) = \{\alpha\}$ .

the receiving state  $u$  corresponding to  $r$ . If  $r$  occurs in process  $q$  and specifies its source  $p$ , then we let  $Q(u) = \{c_{p,q}\}$ . If  $r$  instead uses `MPI_ANY_SOURCE` then we let  $Q(u) = \{c_{p,q} \mid p \in \text{Proc}\}$ . We may assume that the tags have been encoded in the messages, so that to each message  $x$  is associated an integer  $\text{tag}(x)$ . Now if  $r$  specifies a tag  $t$ , we let

$$R_d(u) = \{x \in \text{msg}(d) \mid \text{tag}(x) = t\} \quad (d \in Q(u)).$$

If instead  $r$  uses `MPI_ANY_TAG`, we take  $R_d(u) = \text{msg}(d)$ . We will see below that the execution semantics in effect allow a receive operation to choose non-deterministically among the receiving channels  $Q(u)$ , but, for a given  $d \in Q(u)$ , it must pick out the oldest message in  $d$  with a matching tag. This corresponds exactly to the requirements of the MPI Standard [9, Sec. 3.5].

## 2.2 Execution Semantics of a Model of an MPI Program

Let  $\mathbf{N} = \{0, 1, \dots\}$  and  $\mathbf{N}^\infty = \mathbf{N} \cup \{\infty\}$ . A sequence  $S = (x_1, x_2, \dots)$  of elements of a set  $X$  may be either infinite or finite. We write  $|S|$  for the length of  $S$ . If  $A$  is a subset of a set  $B$ , and  $S$  is a sequence of elements of  $B$ , then *the projection of  $S$  onto  $A$*  is the sequence that results by deleting from  $S$  all elements that are not in  $A$ . If  $S$  is any sequence and  $n \in \mathbf{N}$ , then  $S^n$  denotes the sequence obtained by truncating  $S$  after the  $n^{\text{th}}$  element.

Let  $\mathcal{M}$  be a model of an MPI program. A *global state*  $\sigma$  of  $\mathcal{M}$  is a pair of functions  $(u, \alpha)$ , where  $u$  assigns, to each  $p \in \text{Proc}$ , a state  $u_p \in \text{States}_p$ , and  $\alpha$  assigns to each  $c \in \text{Chan}$  a finite sequence  $\alpha_c$  of elements of  $\text{msg}(c)$ . The sequence represents the *pending* messages for  $c$ : messages that have been sent but not yet received. We define  $\text{Pending}_c(\sigma) = \alpha_c$  and  $\text{state}_p(\sigma) = u_p$ . The *initial state* of  $\mathcal{M}$  is the global state for which  $u_p = \text{start}_p$  for all  $p$ , and  $\alpha_c$  is empty for all  $c$ .

Suppose  $\sigma = (u, \alpha)$  and  $\sigma' = (u', \alpha')$  are global states of  $\mathcal{M}$ ,  $p \in \text{Proc}$ ,  $t \in \text{Trans}_p$ , and that  $\text{src}(t) = u_p$ ,  $\text{des}(t) = u'_p$ ,  $u_q = u'_q$  for  $q \neq p$ , and one of the following holds:

1.  $\text{label}(t) \in \text{loc}(p)$  and  $\alpha = \alpha'$ ,
2. there exist  $c \in \text{Chan}$  and  $x \in \text{msg}(c)$  such that  $\text{label}(t) = c!x$ ,  $\alpha'_c$  is obtained by appending  $x$  to the end of  $\alpha_c$ , and  $\alpha'_d = \alpha_d$  for  $d \neq c$ , or
3. there exist  $d \in \text{Chan}$  and  $y \in \text{msg}(d)$  such that  $\text{label}(t) = d?y$ ,  $y$  is the first element of the projection of  $\alpha_d$  onto  $R_d(u_p)$ ,  $\alpha'_d$  is obtained by deleting the first occurrence of  $y$  from  $\alpha_d$ , and  $\alpha'_c = \alpha_c$  for  $c \neq d$ .

Then we call the triple  $\tau = (\sigma, \sigma', t)$  a *simple global transition* of  $\mathcal{M}$ , and we define  $\text{label}(\tau) = \text{label}(t)$ .

Suppose now that  $\sigma$ ,  $\sigma'$ , and  $\sigma''$  are global states,  $t_1, t_2$  are transitions,  $c \in \text{Chan}$ ,  $x \in \text{msg}(c)$ ,  $p = \text{receiver}(c)$ , and that the following all hold:

1.  $\text{label}(t_1) = c!x$  and  $\text{label}(t_2) = c?x$ ,
2.  $\text{Pending}_c(\sigma)$  contains no element of  $R_c(\text{state}_p(\sigma))$ , and
3.  $(\sigma, \sigma', t_1)$  and  $(\sigma', \sigma'', t_2)$  are simple global transitions.

In this case we will refer to the 4-tuple  $\tilde{\tau} = (\sigma, \sigma'', t_1, t_2)$  as a *synchronous global transition*, as it corresponds to a synchronous MPI communication: a message that is transferred directly from the sender to the receiver in one atomic step. We do *not* want to think of  $\tilde{\tau}$  as “passing through” the intermediate state  $\sigma'$ , but rather as leading directly from  $\sigma$  to  $\sigma''$ . In particular, since  $\text{Pending}_c(\sigma) = \text{Pending}_c(\sigma'')$ ,  $\tilde{\tau}$  leaves all of the channels unchanged. We define  $\text{label}(\tilde{\tau})$  to be the symbol  $c!x$ .

The *state graph* of  $\mathcal{M}$  is the ordered pair  $\mathcal{G} = (\mathcal{S}, \mathcal{T})$ , where  $\mathcal{S}$  is the set of all global states, and  $\mathcal{T}$  is the set of all (simple and synchronous) global transitions. Let  $\text{src}, \text{des}: \mathcal{T} \rightarrow \mathcal{S}$  be the projections onto the first and second coordinates, respectively. These give  $\mathcal{G}$  the structure of a directed graph.

An *event*  $\alpha$  is any element of  $\{\text{label}(\tau) \mid \tau \in \mathcal{T}\}$ . We say that  $\alpha$  is *enabled* at the global state  $\sigma$  if there exists  $\tau \in \mathcal{T}$  with  $\text{src}(\tau) = \sigma$  and  $\text{label}(\tau) = \alpha$ .

Given a path  $T = (\tau_1, \tau_2, \dots)$  in  $\mathcal{G}$ , we define the *atomic length* of  $T$  to be  $\|T\| = \sum_i \epsilon(\tau_i)$ , where  $\epsilon(\tau) = 1$  if  $\tau$  is simple and  $\epsilon(\tau) = 2$  if  $\tau$  is synchronous. This is sometimes a more natural measure of length than  $|T|$ . A *trace* of  $\mathcal{M}$  is any path in  $\mathcal{G}$  originating in the initial state. Finally, If  $T$  originates in the global state  $\sigma$  and  $c \in \text{Chan}$ , we define

$$\text{maxlen}_c(T) = \max_i \{|\text{Pending}_c(\sigma)|, |\text{Pending}_c(\text{des}(\tau_i))|\}.$$

### 3 The Main Theorem

The main theorem concerns *halting properties* so we first explain what these are. In general, a concurrent program is considered to be in a halted state if every process has become permanently blocked. A receive statement in an MPI program blocks, as one would expect, as long as there is no pending message that matches the parameters of that statement. However, the circumstances under which a sending statement blocks are more subtle. Typically, one would assume that each channel  $c$  has some fixed size  $\nu(c) \in \mathbf{N}$ , and declare that a send on  $c$  blocks whenever the length of  $c$  equals  $\nu(c)$ . The MPI Standard, however, imposes no such bounds, but instead declares that a send *may* block at any time, unless the receiving process is at a state from which it can receive the message synchronously [9, Sec. 3.4]. We thus make the following definition for a model  $\mathcal{M}$ :

**Definition 1.** A global state  $\sigma$  of  $\mathcal{M}$  is *potentially halted* if no receive, local, or synchronous event is enabled at  $\sigma$ .


We use the word “potentially” because a program in such a state may or may not halt, depending on the particular choices made by the MPI implementation.

For any predicate  $f$  on the global states of  $\mathcal{M}$ , and any subgraph  $\mathcal{H}$  of  $\mathcal{G}$  that contains the initial state  $\sigma_0$ , let  $\Pi(\mathcal{H}, f)$  denote the statement *for all states  $\sigma$  reachable in  $\mathcal{H}$  from  $\sigma_0$ ,  $f(\sigma)$* . Let  $\text{phalt}$  be the predicate defined by  $\text{phalt}(\sigma) \Leftrightarrow \sigma$  is potentially halted.

**Definition 2.** A *halting predicate* is a state predicate  $f$  of the form  $\text{phalt} \Rightarrow q$ , where  $q$  is any state predicate. A *halting property* is a statement of the form  $\Pi(\mathcal{H}, f)$ , where  $f$  is a halting predicate.

An example of a halting property is given by taking  $q = \text{false}$ , the predicate that holds at no state. For this  $q$ ,  $\Pi = \Pi(\mathcal{G}, f)$  states that  $\mathcal{M}$  never halts. One could also take  $q = \text{term}$ , the predicate that is true when all processes are at final states. Then  $\Pi$  states that whenever  $\mathcal{M}$  halts, all processes have terminated, i.e.,  $\mathcal{M}$  is deadlock-free. More generally, one could take  $q$  to be the predicate  $\text{term}_\Sigma$  that holds when all processes in a certain subset  $\Sigma$  are at final states. One could also let  $q$  be the conjunction of  $\text{term}_\Sigma$  with another predicate – for example, a predicate that holds when variables in the processes in  $\Sigma$ , whose values are encoded in the local states, have particular values. In this case  $\Pi$  would say that whenever the program halts, all processes in  $\Sigma$  have terminated and the variables have the specified values.

To motivate what follows, consider a model [12, Fig. 5] of three processes with state machines as follows:



Suppose we try to verify freedom from deadlock for this model by considering only synchronous executions. Then we only explore the sequence  $(c!?1, e!?1, d!?1)$ , which terminates normally, and miss the deadlocking sequence  $(c!1, e!?1, d!?1)$ . We can try to explain why we missed the deadlock in the following way. At the initial state, process  $p = \text{receiver}(c)$  is at a wildcard receive  $u$  with  $Q(u) = \{c, d\}$ . At this state,  $c$  is ready to receive a message (synchronously) but  $d$  is not. By pursuing only synchronous communication, we never get to see the state in which  $p$  is at  $u$  and a receive on  $d$  is enabled.

The solution is to consider all enabled events (not just synchronous ones) whenever a process  $p$  is at a wildcard receive  $u$ , unless  $u$  has become “urgent.” By this we mean that for each  $c \in Q(u)$ , either a (synchronous or buffered) receive on  $c$  is enabled or we know that a receive on  $c$  can never become enabled. Note that once a receive on  $c$  becomes enabled, it will remain enabled until  $p$  executes a transition, since  $p$  is the only process which may remove a message from  $c$ . Since no receive event can be enabled at a potentially halted state  $\sigma$ , the only way we can arrive at  $\sigma$  is if  $p$  eventually executes. Now if  $u$  is urgent, no new events in  $p$  can become enabled, and so one of the currently enabled events in  $p$  must occur if the system is to arrive at  $\sigma$ . Since those events are *independent* of events in other processes, we might as well explore the paths that result from scheduling each of those enabled events immediately. (If two events are independent then neither can disable the other and the effect of applying one and then the other does not depend on the order in which they are applied.) Local event states are similar, but they are always urgent since the local events are always enabled. The following definitions attempt to make all of this precise:

**Definition 3.** Let  $\sigma$  be a global state of  $\mathcal{M}$ ,  $p \in \text{Proc}$ , and  $u = \text{state}_p(\sigma)$ . We say  $p$  is *at an urgent state in  $\sigma$*  if either  $u$  is a local event state, or all of the following hold:

1.  $u$  is a receiving or send-receive state,
2. for all  $d \in Q(u)$ , at least one of the following holds:
  - (a) there is an event of the form  $d?y$  or  $d!y$  enabled at  $\sigma$ , or
  - (b)  $\text{state}_{\text{sender}(d)}(\sigma)$  is a final state,
 and
3. there is at least one  $d \in Q(u)$  for which 2(a) holds.

We define  $\text{Urgent}(\sigma)$  to be the set of all  $p \in \text{Proc}$  such that  $p$  is at an urgent state in  $\sigma$ . Finally, we say that  $\sigma$  is *urgent* if  $\text{Urgent}(\sigma) \neq \emptyset$ .

**Definition 4.** A global transition  $\tau$  is *urgent for process  $p$*  if  $\tau$  has the form  $(\sigma, \sigma', t)$  or  $(\sigma, \sigma', t', t)$ , where  $p \in \text{Urgent}(\sigma)$ ,  $t \in \text{Trans}_p$ , and  $\text{label}(t)$  is either a local event or a receive.

Condition 2(b) of Definition 3 can be relaxed somewhat: all that is really required is that  $\text{sender}(d)$  be in a state from which it can never reach a send on  $d$ . However, the version that we have stated has the advantage that it is very easy to check. Also, note that the third condition guarantees there is at least one enabled event at an urgent state.

We now fix a total order on  $\text{Proc}$ . The reason for this will become clear: we do not have to consider all urgent transitions departing from an urgent state, but only those for a single process, and so we will just choose the least one.

**Definition 5.** Let  $\tilde{T}$  denote the set of all global transitions  $\tau$  for which either  $\text{src}(\tau)$  is not urgent, or  $\tau$  is urgent for the minimal element of  $\text{Urgent}(\text{src}(\tau))$ . Let  $\tilde{\mathcal{G}} = (\mathcal{S}, \tilde{T})$ .

Now we can state the main theorem:

**Theorem 1.** *Given any path  $S$  in  $\mathcal{G}$  from a global state  $\sigma_0$  to a potentially halted global state  $\sigma$ , there exists a path  $T$  from  $\sigma_0$  to  $\sigma$  in  $\tilde{\mathcal{G}}$  such that  $\|T\| = \|S\|$ ,  $|T| \leq |S|$ , and  $\text{maxlen}_c(T) \leq \text{maxlen}_c(S)$  for all  $c \in \text{Chan}$ . In particular  $\Pi(\mathcal{G}, f) \Leftrightarrow \Pi(\tilde{\mathcal{G}}, f)$  for any halting predicate  $f$ .*

In light of the discussion above, it should come as no surprise that the proof of Theorem 1 relies on many of the restrictions imposed by our domain and property. For example, the fact that each channel has an exclusive receiving process was used to show that once a receive event becomes enabled, it must remain enabled until that process executes. The knowledge that the property could be violated only if no receive were enabled was also used. The fact that a sending state has exactly one outgoing transition also comes into play: if the sending state had outgoing transitions on two different channels then a synchronous event that was enabled on one channel could become disabled if the sending process were to send on the other channel. These arguments withstand the introduction of send-receive states only because the specific structure of those states guarantees that the send event is independent of the receive events. Remove any of these domain-specific restrictions, and Theorem 1 may fail.



## 4 Consequences of the Main Theorem

### 4.1 The Urgent Algorithm

In general, the number of reachable states in  $\mathcal{G}$  or  $\tilde{\mathcal{G}}$  may be very large (or even infinite). So it is common practice to place upper bounds on the channel sizes, or the search depth, in order to reach a conclusive result on at least a bounded region of the state space. For these reasons we define the following concepts. Let  $\nu : \text{Chan} \rightarrow \mathbf{N}^\infty$  and  $m \in \mathbf{N}^\infty$ . Let  $\mathcal{T}_{\nu,m}$  be the set of all global transitions that occur in traces  $T$  that satisfy (i)  $\|T\| \leq m$ , and (ii) for all global states  $\sigma$  through which  $T$  passes, and all  $c \in \text{Chan}$ ,  $|\text{Pending}_c(\sigma)| \leq \nu(c)$ . We let  $\mathcal{G}_{\nu,m} = (\mathcal{S}, \mathcal{T}_{\nu,m})$ .

Let  $\mathcal{T}_{\nu,m}^b$  be the set of all  $\tau \in \mathcal{T}_{\nu,m}$  such that  $\tau \in \tilde{\mathcal{T}}$  and


$$\text{if label}(\tau) = c!?x \text{ for some } c, x \text{ then } \sigma \text{ is urgent or } |\text{Pending}_c(\sigma)| = \nu(c), \quad (1)$$

where  $\sigma = \text{src}(\tau)$ . Condition (1) is not strictly necessary, but it may provide some additional reduction. The idea is that when  $\sigma$  is not urgent, it would be redundant to consider synchronous transitions since we are already pursuing all buffered sends and receives. An exception is made if a channel is full since then a buffered send would not be enabled. Let  $\mathcal{G}_{\nu,m}^b = (\mathcal{S}, \mathcal{T}_{\nu,m}^b)$ . We have the following consequence of Theorem 1:

**Corollary 1.** *Given any path in  $\mathcal{G}_{\nu,m}$  from a global state  $\sigma_0$  to a potentially halted global state  $\sigma$ , there exists a path in  $\mathcal{G}_{\nu,m}^b$  from  $\sigma_0$  to  $\sigma$ . In particular,  $\Pi(\mathcal{G}_{\nu,m}, f) \Leftrightarrow \Pi(\mathcal{G}_{\nu,m}^b, f)$  for any halting predicate  $f$ .*

If  $\text{States}_p$ ,  $\text{Trans}_p$ , and  $\nu(c)$  are finite for all  $p \in \text{Proc}$  and  $c \in \text{Chan}$ , then  $\mathcal{T}_{\nu,m}$  and  $\mathcal{T}_{\nu,m}^b$  are finite as well. It follows from Corollary 1 that we can verify a halting property in this case by performing a depth-first search of  $\mathcal{G}_{\nu,m}^b$ . Specifically, algorithm Urgent of Fig. 2 will find all reachable states in  $\mathcal{G}_{\nu,m}$  for which  $f$  does not hold. We assume  $\text{Proc} = \{p_1, \dots, p_N\}$  and  $p_1 < \dots < p_N$ . The search is initiated by setting the global variable  $R$  to the empty set and calling  $\text{search}(\sigma_0, 0)$ , where  $\sigma_0$  is the initial state. Function  $\text{urgent\_transitions}(\sigma, p)$  returns the set of all  $\tau \in \mathcal{T}$  such that  $\text{src}(\tau) = \sigma$  and  $\tau$  is urgent for  $p$ . Function  $\text{standard\_transitions}(\sigma, \nu)$  returns the set of all  $\tau \in \mathcal{T}$  that satisfy (i)  $\text{src}(\tau) = \sigma$ , (ii)  $|\text{Pending}_c(\text{des}(\tau))| \leq \nu(c)$  for all  $c$ , and (iii)  $\text{label}(\tau) = c!?x \Rightarrow |\text{Pending}_c(\sigma)| = \nu(c)$ . There is no need to specify  $\nu$  for  $\text{urgent\_transitions}$  since an urgent transition can never increase the length of a channel.

*Example.* In a model of a *client-server* system with  $n$  clients ( $n \geq 1$ ),  $\text{Proc} = \{0, 1, \dots, n\}$  with the natural order,  $\text{Chan} = \{c_1, d_1, \dots, c_n, d_n\}$ ,  $\text{msg}(c) = \{1\}$  for all  $c \in \text{Chan}$ , and  $\text{sender}(c_i) = i = \text{receiver}(d_i)$ ,  $\text{receiver}(c_i) = 0 = \text{sender}(d_i)$  for  $1 \leq i \leq n$ . For  $n = 2$ , the state machines for processes 0 (the server), 1, and 2, are respectively:



```

1  function selected_transitions( $\sigma$ )  /* returns  $\{\tau \in \mathcal{T}_{\nu,m}^b \mid \text{src}(\tau) = \sigma\}$  */
2      for  $i = 1$  to  $N$  do
3          if  $p_i \in \text{Urgent}(\sigma)$  then return urgent_transitions( $\sigma, p_i$ ) end if
4      end for;
5      return standard_transitions( $\sigma, \nu$ )
6  end function;

7  procedure search( $\sigma, l$ )
8      if  $l > m$  then return end if;
9       $R := R \cup \{\sigma\}$ ;
10     if not  $f(\sigma)$  then report_violation() end if;
11     for all  $\tau \in \text{selected\_transitions}(\sigma)$  do
12         if  $\text{des}(\tau) \notin R$  then search( $\text{des}(\tau), l + \epsilon(\tau)$ ) end if
13     end for all
14 end procedure

```

**Fig. 2.** The Urgent Algorithm: depth-first search of  $\mathcal{G}_{\nu,m}^b$ .

Let us see how the Urgent algorithm applies to this system for any  $\nu$  and  $m = \infty$ . We start with the initial state: this state is urgent for process 0, so we explore the states resulting from the global transitions labeled  $c_i!$  for all  $i$ . For any such  $i$ , the resulting state has process 0 in local state  $i$ , process  $i$  in local state 1, and all other processes and channels unchanged. This state is urgent for  $i$ , and so we explore the single transition  $d_i!$ . This returns us to the initial state, which is already in  $R$ . Hence the algorithm explores a total of  $n + 1$  global states, and  $2n$  transitions. Notice also that, in this case, the search does not explore any buffered communication, even though process 0 contains a wildcard receive.

## 4.2 Source-Specific Models and Synchronous Traces

We say that  $\mathcal{M}$  is *source-specific* if for every receiving and send-receive state  $u$  in  $\mathcal{M}$ ,  $|Q(u)| = 1$ ; this corresponds to an MPI program which never uses `MPI_ANY_SOURCE` (though it may use `MPI_ANY_TAG`). We say that a path in  $\mathcal{G}$  is *synchronous* if it consists solely of local and synchronous transitions.

Let  $\mathcal{M}$  be any model and  $\sigma$  a global state of  $\mathcal{M}$ . If  $\sigma$  is urgent, then clearly  $\sigma$  cannot be potentially halted. Now if  $\mathcal{M}$  is source-specific, the converse is also true. For if there is some  $c \in \text{Chan}$  and  $x \in \text{msg}(c)$  for which  $c?x$  or  $c!x$  is enabled at  $\sigma$ , then  $p = \text{receiver}(c) \in \text{Urgent}(\sigma)$ , since  $Q(\text{state}_p(\sigma)) = \{c\}$ .

Now suppose  $\mathcal{M}$  is source-specific and  $T$  is a trace terminating in a potentially halted state  $\sigma$ . By Theorem 1, there exists a trace  $\tilde{T} = (\tau_1, \dots, \tau_n)$  in  $\tilde{\mathcal{G}}$  terminating in  $\sigma$ , with  $n \leq |T|$  and  $|\tilde{T}| = |T|$ . Let  $\sigma_k = \text{des}(\tau_k)$  for  $1 \leq k \leq n$  and let  $\sigma_0$  be the initial state. Let  $i$  be the least integer for which  $\sigma_i$  is potentially halted. For  $0 \leq j < i$ ,  $\sigma_j$  is not potentially halted, which as we have seen means that  $\sigma_j$  is urgent. This implies that  $\tau_{j+1}$  is a local event, synchronous, or receive transition. But  $\tau_{j+1}$  cannot be a receive: if it were, there would have to be a preceding send. In other words,  $\tilde{T}^i$  is synchronous. We have proved:

**Corollary 2.** *Let  $\mathcal{M}$  be a source-specific model of an MPI program and  $T$  a trace terminating in a potentially halted state  $\sigma$ . Then there exist  $i \in \mathbf{N}$  and a trace  $\tilde{T}$  in  $\tilde{\mathcal{G}}$  that terminates in  $\sigma$  such that  $|\tilde{T}| \leq |T|$ ,  $\|\tilde{T}\| = \|T\|$ , and  $\tilde{T}^i$  is synchronous and terminates in a potentially halted state.*


This leads to the following, which generalizes [12, Theorem 7.4]. Note that 0 is used to denote the function on  $\mathbf{Chan}$  which is identically 0. Note also that all of the examples of halting predicates given in Sec. 3 satisfy the condition on  $q$ .

**Corollary 3.** *Suppose  $\mathcal{M}$  is a source-specific model of an MPI program, and  $q$  is a state predicate satisfying  $q(\sigma) \Rightarrow q(\sigma')$  for any simple global transition  $(\sigma, \sigma', t)$ . Let  $f$  denote the predicate  $\mathbf{phalt} \Rightarrow q$ ,  $\nu: \mathbf{Chan} \rightarrow \mathbf{N}^\infty$ , and  $m \in \mathbf{N}^\infty$ . Then  $\Pi(\mathcal{G}_{\nu,m}, f) \Leftrightarrow \Pi(\mathcal{G}_{0,m}^b, f)$ .*

### 5 Related Work

The literature on partial order reduction techniques is too large to summarize here, but [1, 4, 11] and the references cited cover much of the ground. *Persistent set* techniques [4, 5] form a family of POR methods that deal specifically with freedom from deadlock. Those techniques associate, to each global state  $\sigma$  encountered in the search, a subset  $T_\sigma$  of the set of all transitions enabled at  $\sigma$ , in such a way that the following condition holds: on any path in the full state graph departing from  $\sigma$ , no transition dependent on a transition on  $T_\sigma$  can occur without a transition in  $T_\sigma$  occurring first. (The word *transition* in this context corresponds to a set of our global transitions.) The reduced search explores only the transitions in  $T_\sigma$ , and so benefits whenever  $T_\sigma$  is a proper subset. If it is also the case that  $T_\sigma$  is empty only when there are no enabled transitions at  $\sigma$ , then we may conclude that the reduced search will explore all reachable deadlocked states [4, Thm. 4.3].

It should be emphasized, however, that here “deadlocked state” is used in the usual sense, to mean a state with no outgoing transitions. In our MPI context we call such states *absolutely halted*. An absolutely halted state is certainly potentially halted, but the converse is not always the case, and, in fact, the persistent set POR algorithms may miss potentially halted states. Consider, for example, the standard state graph (i.e., without the added synchronous transitions) arising from a model of an MPI program. For simplicity, let us assume the model has no send-receive states. Now for any global state  $\sigma$ , we could declare  $\sigma$  to be urgent if some process  $p$  were at either (i) a receiving state in which every receiving channel had at least one pending matching message, (ii) a local-event state, or (iii) a sending state. We could then let  $T_\sigma$  consist of the enabled events for the least urgent process (or all enabled events if no process is urgent), and this would satisfy the conditions of the previous paragraph. Consider a model with  $\mathbf{Proc} = \{0, 1\}$ , in the natural order, a local event  $\lambda$  in process 1, and with state machines as follows:



The model contains a potentially halted state  $\sigma_1$ , obtained from the initial state  $\sigma_0$  by executing  $\lambda$ . However, process 0 would be urgent at  $\sigma_0$ , so we would have  $T_{\sigma_0} = \{c!1\}$ . Hence the reduced search would not explore  $\sigma_1$ , and in fact would complete without encountering any potentially halted states.

We could attempt to correct this problem by simply not allowing the sending states to be urgent, and it can in fact be shown that this would lead to an algorithm that explored all potentially halted states. However, the algorithm would miss many of the opportunities for reduction. Consider, for example, a client server system with  $n$  clients. The system would not be in an urgent state until every client had sent at least one request. In particular, the reduced search would explore all possible states of the  $n$  request channels for which at least one channel is empty.

Our approach solves this problem by adding the synchronous transitions to the state graph and defining the  $T_\sigma$  to take advantage of those transitions under the appropriate circumstances. This solution cannot, strictly speaking, be characterized as a persistent set approach, since our  $T_\sigma$  do not necessarily satisfy the persistent set condition. Consider, for example, a client-server system with one client. At the initial state  $\sigma$ , our  $T_\sigma$  consists of the single transition labeled  $c!1$ . But the path  $c!1, c?1$  is also possible from  $\sigma$ , and both  $c!1$  and  $c?1$  are dependent on  $c!1$ .

Other POR techniques preserve more complex temporal properties. The *ample set* framework [1, Chap. 10] is an example of these. Here, the  $T_\sigma$  must satisfy several conditions in addition to those described above for persistent sets. If all the conditions are met, then the reduced state graph is guaranteed to be stutter-equivalent to the full state graph, and hence can be used to verify any  $LTL_X$  property [1, Cor. 2 and Thm. 12].

Returning to the MPI context, any halting property can be expressed as an  $LTL_X$  property of the form  $\Box(\mathbf{phalt} \Rightarrow q)$ , and therefore, in theory, is amenable to the ample set approach. Now, however, another problem arises. In the ample set framework, a transition is *invisible* if it can never change the truth value of a predicate (such as  $\mathbf{phalt}$ ) used in the LTL formula. The *invisibility condition* requires that whenever  $T_\sigma$  is a proper subset of the set of all transitions enabled at  $\sigma$ , every transition in  $T_\sigma$  is invisible. Since any local event, receive, or synchronous transition might change  $\mathbf{phalt}$  from false to true, these transitions are not necessarily invisible, and therefore an ample set algorithm would not include them in a  $T_\sigma$  (unless the  $T_\sigma$  consisted of all enabled transitions). This eliminates most, if not all, of the opportunities for reduction.

As it turns out, the ample set invisibility condition is unnecessarily strict, and all that is really required is that a transition never change  $\mathbf{phalt}$  from true to false, which is certainly the case for local event, receive, and synchronous transitions, since they are not even enabled at potentially halted states. (Notice that send transitions can change  $\mathbf{phalt}$  from true to false, which explains why it really would be a mistake to treat them as invisible.) Another condition, concerning cycles in the reduced graph, can also be safely ignored in our context. After these modifications, however, the ample set approach essentially reduces to the persistent set approach, discussed above.

Another deadlock-preserving reduction method is the *sleep set* technique [3, 4]. Sleep sets can be used in conjunction with the persistent set approach to further reduce the numbers of states and transitions explored. The idea is to associate to each state a dynamically changing set of transitions that can be safely ignored even if they appear in the persistent set for that state. It is possible that this method could be adapted to work with our urgent algorithm, an idea we hope to explore in future work.

## 6 Experimental Results

Eight scalable C/MPI programs were used for our empirical investigation. They range from standard toy concurrency examples to more complex programs from a well-known book on MPI [6]. For each, we constructed by hand an abstract model appropriate for verifying freedom from deadlock. These models were encoded as certain Java classes that can be read by the MPI-Optimized Verifier (MOVER), a Java tool developed for this project. Given the model and an object describing a halting property, MOVER can either (A<sub>1</sub>) execute a generic depth-first search of the state space to verify the property or report any violations, (A<sub>2</sub>) execute the Urgent algorithm to do the same, or (A<sub>3</sub>) produce a Promela model that can be used by the model checker SPIN [7] to do the same.

The processes and channels in the Promela model correspond exactly to those in the MPI model. There are no variables in the Promela, other than the channels. The local states of a process are encoded by labeled positions in the code. States with multiple departing transitions are encoded using the Promela selection construct (`if...fi`). A never claim is inserted corresponding to the LTL formula  $\langle \rangle !(\text{univenabled} \parallel \text{terminated})$ , where `univenabled` is defined to hold whenever a synchronous, local, or receive event is enabled (the definition refers to the lengths of the channels and the positions of the local processes), and `terminated` is defined to hold when all terminating processes are at final states. SPIN uses a POR algorithm that is similar to the ample set technique [8]. It might seem appropriate to use SPIN's `xr` and `xs` declarations, which declare a process to have exclusive read or write access to a channel and provide information to help the POR algorithm. However, this is not allowed, as the never claim makes reference to all the channels, and in fact an attempt to use those declarations causes SPIN to flag the error. This is SPIN's way of recognizing that the communication events may not be invisible with respect to the property.

(A different way to use SPIN to verify freedom from deadlock for MPI programs is described in [13]. In that approach, every send is immediately followed by a non-deterministic choice between blocking until the channel becomes empty and proceeding without blocking. Freedom from deadlock can then be checked in the usual way with SPIN, i.e., without a never claim. While we have not carried out an extensive comparison, it appears that the state-explosion is much worse for that approach than for the approach presented here, due to all the new states introduced by the non-deterministic choices.)

We applied all three approaches to each of the examples, increasing system size  $n$  until  $n = 200$  or we ran out of memory. In each case we recorded the numbers of states and transitions explored, and the time and memory used. We used the Java2 SDK 1.4.2 with options `-Xmx1900M` and SPIN 4.2.0, with options `-DCOLLAPSE -DMEMLIM=2800 -DSAFETY`; the maximum search depth also had to be increased in some cases. The experiments were run on a Linux box with a 2.2 GHz Xeon processor and 4 GB of memory. In the one case where a deadlock was found, the searches were stopped after finding the first counterexample.

Figures 3 and 4 show the number of states explored. We first observe that the numbers for  $A_1$  and  $A_3$  are exactly equal in all cases where both searches completed. Since  $A_1$  explores all reachable states, this means that SPIN's POR algorithm (on, by default) made no difference in the number of states explored. This is not surprising, since there are no invisible events for the algorithm to exploit. For the one case where a violation exists, SPIN did find the violation much sooner than either MOVER algorithm (Fig. 3(d)). This appears to be just a fluke related to process ordering: we ran the same problem but reversed the order in which the processes were declared (for both tools), and the results were almost exactly reversed.


For the Client-Server, Producer-Consumer, and the two exchange examples, the performance of  $A_2$  was the most impressive, reducing the complexity class from one that is apparently exponential to one that is linear. For Monte Carlo and Master-Slave, both functions appear to be exponential, but the exponent for the  $A_2$  function is lower (significantly so for Master-Slave), allowing it to scale further. In one case (Fig. 3(c)), the use of  $A_2$  makes almost no difference, but there the number of reachable states was quadratic to begin with so there was not much room for improvement. The Master Producer-Consumer proved the most difficult: there seemed to be a small constant reduction but no approach could scale beyond  $n = 4$ .

For Producer-Consumer, we give on one graph (Fig. 4, left) the results for various values of  $\nu$ . This graph demonstrates the impact of channel size on state explosion for systems that can buffer many messages. For  $\nu = 0$ , however, the number of reachable states for the system of size  $n$  is just  $n + 1$ , and  $A_2$  searches that number of states for any value of  $\nu$ , since the system contains no wildcard receives. We also give the time for the Master-Slave example; typical of these examples, the pattern is similar to that for the number of states.


In summary, the Urgent algorithm often dramatically reduced the number of states explored. It can never increase that number, as long as the search is carried to completion, nor did it appear to have a significant impact on the time required to complete the search. In contrast, the POR algorithm implemented in SPIN had no effect on the number of states explored.

## 7 Conclusions and Future Work


We have presented a POR-like optimization to the standard model checking algorithm for verifying halting properties of MPI programs. The algorithm seeks to control state explosion by limiting the number of transitions explored that




(a) Client-Server,  $\nu = 2$




(b) Master Producer-Consumer,  $\nu = 1$




(c) Chansize Deadlocker,  $\nu = n - 1$




(d) Chansize Deadlocker,  $\nu = n$  (violation)




(e) Exchange-Alternate,  $\nu = 2$



(f) Exchange-Sendrecv,  $\nu = 2$




(g) Master-Slave,  $\nu = 2$



(h) Monte Carlo,  $\nu = 2$

**Fig. 3.** Graphs of  $y = \log_{10}(f(n))$ , where  $f(n)$  is the number of states explored for the system of size  $n$ , with channel size bound  $\nu$ .



**Fig. 4.** Producer-Consumer states for  $\nu \in \{0, 1, \dots, 4\}$  ( $\log_{10}$  of number of states, left), and Master-Slave time ( $\log_{10}$  of number of seconds, right).

involve buffering messages. The technique also interacts well with the imposition of bounds on both the search depth and the sizes of the communication channels.

Earlier work showed that it suffices to consider only synchronous communication when verifying freedom from deadlock for certain MPI programs with no wildcard receives. We have shown how that result follows easily from the theorem that justifies our optimization. Moreover, we have shown that the restriction that forbids wildcard receives may be relaxed to allow the use of `MPI_ANY_TAG`.

We have demonstrated the effectiveness of our algorithm on several scalable examples, including some with wildcard receives. However, a better validation of effectiveness would utilize more “realistic” examples. There is no guarantee that scaling our simple examples presents the same kind of challenge to the Urgent algorithm that an actual production-level MPI code would. Due to the difficulty of creating models by hand, this task would benefit from an automated MPI model extractor. We intend to develop such a tool, and use it to verify not only freedom from deadlock, but also other halting properties. For example, we would like to model the arithmetic computations performed by an MPI program symbolically, and check that at termination the program has arrived at the correct arithmetic result.

Finally, the study of domain-specific approaches may also shed light on the general framework. It would be interesting to see if the ample set framework, for example, could be extended to incorporate the idea of switching between a synchronous and a buffering mode, generalizing our MPI-specific approach.

## Acknowledgments

The author is grateful to George Avrunin and to Reviewer 3 for helpful comments on earlier drafts of this paper.

## References

1. Clarke, Jr., E.M., Grumberg, O., Peled, D.A.: Model Checking. MIT Press, Cambridge (1999)



2. Dwyer, M.B., Hatcliff, J., Robby, Ranganath, V.P.: Exploiting object escape and locking information in partial-order reductions for concurrent object-oriented programs. *Formal Methods in System Design* **25** (2004) 199–240
3. Godefroid, P.: Using partial orders to improve automatic verification methods. In Clarke, E.M., Kurshan, R.P., eds.: *Computer-Aided Verification, 2nd International Conference, CAV '90*, New Brunswick, NJ, USA, June 1990, Proceedings. Volume 531 of *Lecture Notes in Computer Science*, Springer-Verlag (1990) 176–185
4. Godefroid, P.: *Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem*. Springer, Berlin (1996)
5. Godefroid, P., Pirottin, D.: Refining dependencies improves partial-order verification methods. In Courcoubetis, C., ed.: *Computer-Aided Verification, 5th International Conference, CAV '93*, Elounda, Greece, June/July 1993, Proceedings. Volume 697 of *Lecture Notes in Computer Science*, Springer-Verlag (1993) 438–449
6. Gropp, W., Lusk, E., Skjellum, A.: *Using MPI: Portable Parallel Programming with the Message-Passing Interface*. MIT Press, Cambridge, MA (1999)
7. Holzmann, G.J.: *The SPIN Model Checker*. Addison-Wesley, Boston (2004)
8. Holzmann, G.J., Peled, D.: An improvement in formal verification. In Hogrefe, D., Leue, S., eds.: *Formal Description Techniques VII, Proceedings of the 7th IFIP WG6.1 International Conference on Formal Description Techniques*, Berne, Switzerland, 1994. Volume 6 of *IFIP Conference Proceedings*, Chapman & Hall (1995) 197–211
9. Message Passing Interface Forum: *MPI: A Message-Passing Interface standard, version 1.1*. <http://www.mpi-forum.org/docs/> (1995)
10. Message Passing Interface Forum: *MPI-2: Extensions to the Message-Passing Interface*. <http://www.mpi-forum.org/docs/> (1997)
11. Peled, D.: Ten years of partial order reduction. In Hu, A.J., Vardi, M.Y., eds.: *Computer Aided Verification, 10th International Conference, CAV '98*, Vancouver, BC, Canada, June 28 – July 2, 1998, Proceedings. Volume 1427 of *Lecture Notes in Computer Science*, Springer (1998) 17–28
12. Siegel, S.F., Avrunin, G.S.: *Modeling MPI programs for verification*. Technical Report UM-CS-2004-75, Department of Computer Science, University of Massachusetts (2004)
13. Siegel, S.F., Avrunin, G.S.: Verification of MPI-based software for scientific computation. In Graf, S., Mounier, L., eds.: *Model Checking Software: 11th International SPIN Workshop, Barcelona, Spain, April 1–3, 2004, Proceedings*. Volume 2989 of *Lecture Notes in Computer Science*, Springer-Verlag (2004) 286–303