

# Proving MCAPI Executions are Correct

## Applying SMT Technology to Message Passing

Yu Huang, Eric Mercer, and Jay McCarthy \*

Brigham Young University  
{yuHuang, egm, jay}@byu.edu

**Abstract.** Asynchronous message passing is an important paradigm in writing applications for embedded heterogeneous multicore systems. The Multicore Association (MCA), an industry consortium promoting multicore technology, is working to standardize message passing into a single API, MCAPI, for bare metal implementation and portability across platforms. Correctness in such an API is difficult to reason about manually, and testing against reference solutions is equally difficult as reference solutions implement an unknown set of allowed behaviors, and programmers have no way to directly control API internals to expose or reproduce errors. This paper provides a way to encode an MCAPI execution as a Satisfiability Modulo Theories (SMT) problem, which if satisfiable, yields a feasible execution schedule on the same trace, such that it resolves non-determinism in the MCAPI runtime in a way that it now fails user provided assertions. The paper proves the problem is NP-complete. The encoding is useful for test, debug, and verification of MCAPI program execution. The novelty in the encoding is the direct use of match pairs (potential send and receive couplings). Match-pair encoding for MCAPI executions, when compared to other encoding strategies, is simpler to reason about, results in significantly fewer terms in the SMT problem, and captures feasible behaviors that are ignored in previously published techniques. Further, to our knowledge, this is the first SMT encoding that is able to run in *infinite-buffer* semantics, meaning the runtime has unlimited internal buffering as opposed to no internal buffering. Results demonstrate that the SMT encoding, restricted to zero-buffer semantics, uses fewer clauses when compared to another zero-buffer technique, and it runs faster and uses less memory. As a result the encoding scales well for programs with high levels of non-determinism in how sends and receives may potentially match.

**Keywords:** Abstraction, refinement, SMT, message passing

## 1 Introduction

Embedded devices fill all sorts of crucial roles in our lives. They exist as medical devices, as network infrastructure, and they control our automobiles. Embedded devices continue to become more powerful as computing hardware becomes

---

\* Special thanks to Christopher Fischer, formally at BYU and now at Amazon

smaller and more modular. It is now becoming commonplace to find multiple processing units inside a single device. The Multicore Association (MCA) is an industry group that has formed to define specifications for low-level communication, resource management, and task management for embedded heterogeneous multicore devices [9].

One specification that the MCA has released is the Multicore Association Communications API (MCAPI) [10]. The specification defines types and functions for simple message passing operations between different computing entities within a device. Messages can be passed across persistent channels that force an ordering of the messages, or they can be passed to specific *endpoints* within the system. The specification places few ordering constraints on messages passed from one endpoint to another. This freedom introduces the possibility of a race between multiple messages to common endpoints thus giving rise to non-deterministic behavior in the runtime [14]. If an application has non-determinism, it is not possible to test and debug such an application without a way to directly (or indirectly) control the MCAPI runtime.

There are two ways to implement the MCAPI semantics: infinite-buffer semantics (the message is copied into a runtime buffer on the API call) and zero-buffer semantics (the message has no buffering) [18]. An infinite-buffer semantics provides more non-deterministic behaviors in matching send and receives because the runtime can arbitrarily delay a send to create interesting (and unexpected) send reorderings. The zero-buffer semantics follow intuitive message orderings as a send and receive essentially rendezvous.

Sharma *et al.* propose a method to indirectly control the MCAPI runtime to verify MCAPI programs under zero-buffer semantics [16]. As the work does not address infinite-buffer semantics, it is somewhat limited in its application. The work does provide a dynamic partial order reduction for the model checker, but such a reduction is not sufficient to control state space explosion in the presence of even moderate non-determinism between message sends and receives. A key insight from the approach is its direct use of match pairs–couplings for potential sends and receives.

Wang *et al.* propose an alternative method for resolving non-determinism for program verification using symbolic methods in the context of shared memory systems [19]. The work observes a program trace, builds a partial order from that trace called a concurrent trace program (CTP), and then creates an SMT problem from the CTP that if satisfied indicates a property violation.

Elwakil *et al.* extend the work of Wang *et al.* to message passing and claim the encoding supports both infinite and zero buffer semantics. A careful analysis of the encoding, however, shows it to not work under infinite-buffer semantics and to miss behaviors under zero-buffer semantics [6]. Interestingly, the encoding assumes the user provides a precise set of match pairs as input with the program trace, and it then uses those match pairs in a non-obvious way to constrain the happens-before relation in the encoding. The work does not discuss how to generate the match pairs, which is a non-trivial input to manually generate

Task 0	Task 1	Task 2
00 initialize(NODE_0,&v,&s);	10 initialize(NODE_1,&v,&s);	20 initialize(NODE_2,&v,&s);
01 e0=create_endpoint(PORT_0,&s);	11 e1=create_endpoint(PORT_1,&s);	21 e2=create_endpoint(PORT_2,&s);
02 msg_rcv_i(e0,A,sizeof(A),&h1,&s);	12 e0=get_endpoint(NODE_0,PORT_0,&s);	22 e0=get_endpoint(NODE_0,PORT_0,&s);
03 wait(&h1,&size,&s,MCAPI_INF);	13 msg_rcv_i(e1,C,sizeof(C),&h3,&s);	23 e1=get_endpoint(NODE_1,PORT_1,&s);
04 a=atoi(A);	14 wait(&h3,&size,&s,MCAPI_INF);	24 msg_send_i(e2,e0,"4",2,N,&h5,&s);
05 msg_rcv_i(e0,B,sizeof(B),&h2,&s);	15 msg_send_i(e1,e0,"1",2,N,&h4,&s);	25 wait(&h5,&size,&s,MCAPI_INF);
06 wait(&h2,&size,&s,MCAPI_INF);	16 wait(&h4,&size,&s,MCAPI_INF);	26 msg_send_i(e2,e1,"Go",3,N,&h6,&s);
07 b=atoi(B);	17 finalize(&s);	27 wait(&h6,&size,&s,MCAPI_INF);
08 if(b > 0);		28 finalize(&s);
09 assert(a == 4);		
0a finalize(&s);		

**Fig. 1.** An MCAPI concurrent program execution

for large or complex program traces. An early proof claims that the problem of finding a precise set of match pairs given a program trace is NP-complete [15].

This paper presents a proof that resolving non-determinism in message passing programs in a way that meets all assertions is NP-complete. The paper then presents an SMT encoding for MCAPI program executions that works for both zero and infinite buffer semantics. The encoding does require an input set of match pairs as in prior work, but unlike prior work, the match-set can be over-approximated and the encoding is still sound and complete. The encoding requires fewer terms to capture all possible program behavior when compared to other proposed methods making it more efficient in the SMT solver. To address the problem of generating match pairs, an algorithm to generate the over-approximated set is given. To summarize, the main contributions in this paper are

1. a proof that the problem of matching sends to receives in a way that meets assertions is NP-complete;
2. a correct and efficient SMT encoding of an MCAPI program execution that detects all program errors under zero or infinite buffer semantics given the input set of potential match pairs contains at least the precise set of match pairs; and
3. an  $O(N^2)$  algorithm to generate an over-approximation of possible match pairs, where  $N$  is the size of the execution trace in lines of code.

The rest of the paper is organized as follows: Section 2 presents an MCAPI program illustrating the non-determinism in the runtime specification. Section 3 is the proof that the problem is NP-complete. Section 4 defines the SMT encoding using potential match-pairs. The encoding is shown to be sound and complete even under an over-approximated set of match pairs. Section 5 provides a solution to the outstanding problem of generating feasible match pairs. Section 6 presents the experimental results that show the encoding to be efficient. Section 7 discusses related work. And Section 8 presents conclusions and future work.

## 2 Example

It is a challenge to explain intended behavior in simple scenarios consisting of a handful of calls when dealing with concurrency. Consider the MCAPI program execution in Figure 1 that includes three tasks that use send (`mcapi_msg_send_i`) and receive (`mcapi_msg_recv_i`) calls to communicate with each other. Line numbers appear in the first column for each task with the first digit being the task ID, and the declarations of the local variables are omitted for space.

Picking up the scenario just after the endpoints are defined, lines 02 and 05 receive two messages on the endpoint  $e0$  in variables  $A$  and  $B$  which are converted to integer values and stored in variables  $a$  and  $b$  on lines 04 and 07; task 1 receives one message on endpoint  $e1$  in variable  $C$  on line 13 and then sends the message “1” on line 15 to  $e0$ ; and finally, task 2 sends messages “4” and “Go” on lines 24 and 26 to endpoints  $e0$  and  $e1$  respectively. The additional code (lines 08 - 09) asserts properties of the values in  $a$  and  $b$ . The `mcapi_wait` calls block until the associated send or receive buffer is able to be used. Given the scenario, a developer might ask the question: “*What are the possible values of  $a$  and  $b$  after the scenario completes?*”

```

24 S2,4(0, &h5)
25 W(&h5)
-----
02 R0,2(2, &h1)
03 W(&h1)
-----
26 S2,6(1, &h6)
27 W(&h6)
-----
04 a = atoi(A);
13 R1,3(2, &h3)
14 W(&h3)
15 S1,5(0, &h4)
16 W(&h4)
-----
05 R0,5(1, &h2)
06 W(&h2)
07 b = atoi(B);
08 assume(b > 0);
09 assert(a == 4);

```

**Fig. 2.** A feasible execution traces of the MCAPI program execution in Figure 1

The intuitive trace is shown in Figure 2 using a shorthand notation for the MCAPI commands: send (denoted as **S**), receive (denoted as **R**), or wait (denoted as **W**). The shorthand notation further preserves the thread ID and line number as follows: for each command  $O_{i,j}(k, \&h)$ ,  $O \in \{\mathbf{S}, \mathbf{R}\}$  or  $\mathbf{W}(\&h)$ ,  $i$  represents the task ID,  $j$  represents the source line number,  $k$  represents the destination endpoint, and  $h$  represents the command handler. A specific destination task ID is indicated in the notation when a trace is fully resolved, otherwise “\*” notation indicates that a receive has yet to be matched to a specific send. The lines in the trace indicate the context switch where a new task executes.

From the trace, variable  $a$  should contain 4 and variable  $b$  should contain 1 since task 2 must first send message “4” to  $e0$  before it can send message

“Go” to  $e1$ ; consequently, task 1 is then able to send message “1” to  $e0$ . The `assume` notation asserts the control flow taken by the program execution. In this example, the program takes the true branch of the condition on line 08. At the end of execution the assertion on line 09 holds and no error is found.

```

24 S2,4(0, &h5)
25 W(&h5)
26 S2,6(1, &h6)
27 W(&h6)
-----
13 R1,3(2, &h3)
14 W(&h3)
15 S1,5(0, &h4)
16 W(&h4)
-----
02 R0,2(1, &h1)
03 W(&h1)
04 a = atoi(A);
05 R0,5(2, &h2)
06 W(&h2)
07 b = atoi(B);
08 assume(b > 0);
09 assert(a == 4);

```

**Fig. 3.** A second feasible execution traces of the MCAPI program in Figure 1

There is another feasible trace for the example shown in Figure 3 which is reachable under the infinite-buffer semantics. In this trace, the variable  $a$  contains 1 instead of 4, since the message “1” is sent to  $e0$  after sending the message “Go” to  $e1$  as it is possible for the send on line 24 to be buffered in transit. The MCAPI specification indicates that the wait on line 25 returns once the buffer is available. That only means the message is somewhere in the MCAPI runtime under the infinite-buffer semantics; it does not mean the message is delivered. As such, it is possible for the message to be buffered in transit allowing the send on line 15 to arrive at  $e0$  first and be received in variable “ $a$ ”. Such a scenario is a program execution that results in an assertion failure at line 09.

From the discussion above, it is important to consider non-determinism in the MCAPI runtime when testing or debugging an MCAPI program execution. The next section presents a proof that the problem of matching sends to receives in a way that meets all assertions is NP-complete. The proof justifies the encoding and SMT solver. Following the proof, the algorithm to generate the encoding is presented. It takes an MCAPI program execution with a set of possible send-receive match pairs and generates an SMT problem that if satisfied proves that non-determinism can be resolved in a way that violates a user provided assertion (the assertions are negated in the encoding) and if unsatisfiable proves the trace correct (meaning the user assertions hold on the given execution under all possible runtime behaviors). The encoding can be solved by an SMT solver such as Yices [4] or Z3 [12].

### 3 NP Completeness Proof

The complexity proof is inspired by the NP-completeness proof for memory coherence and consistency by Cantin *et al.* that uses a similar reduction from SAT only in the context of shared memory [2]. The complexity proof is on a new decision problem: *Verifying Assertions in Message Passing* (VAMP).

**Definition 1. *Verifying Assertions in Message Passing.***

*INSTANCE:* A set of constants  $D$ , a set of variables  $X$ , and a finite set  $H$  of task histories consisting of send, receive, and assert operations over  $X$  and  $D$ .

*QUESTION:* Is there a feasible schedule  $S$  for the operations of  $H$  that satisfy all the assertions?

The VAMP problem is NP-complete. The proof is a reduction from SAT. Given an instance  $Q$  of SAT consisting of a set of variables  $U$  and set of clauses  $C$  over  $U$ , an instance  $V$  of VAMP is constructed such that  $V$  has a feasible schedule  $S$  that meets all the assertions if and only if there is a satisfying truth assignment for  $Q$ . Feasible in this context means the schedule is allowed by the MCAPI semantics.

The reduction is illustrated in Figure 4. The figure elides the explicit calls to wait which directly follow each send and receive operation, and it elides the subscript notation as it is redundant in the figure. The figure also adds the value sent and the variable that receives the value to the notation as that information is pertinent to the reduction.

The reduction relies on non-determinism in the message passing to decide the value of each variable in  $U$ . The tasks  $h_{d_0}$  and  $h_{d_1}$  repeatedly send the constant value  $d_0$  (false valuation) or  $d_1$  (true valuation) to task  $h_C$ . The key intuition is that these tasks are synchronized with  $h_C$  so they essentially wait to send the value until asked.

The task  $h_C$  sequentially requests and receives  $d_0$  and  $d_1$  values for each variable in the SAT instance  $Q$ . It does not request values for a new variable until the current variable is resolved. As the values come from two separate tasks upon request, the messages race in the runtime and may arrive in either order at  $h_C$ . As a result, the value in each variable is non-deterministically  $d_0$  or  $d_1$ .

After the value of each variable  $u_i$  is resolved, the  $h_C$  task asserts the truth of each clause in the problem instance. As the clauses are conjunctive, the assertions are sequentially evaluated. If a satisfying assignment exists for  $Q$ , then a feasible schedule exists that resolves the values of each variable in such a way that every assert holds.

**Lemma 1.**  *$S$  is a feasible schedule for  $H$  that satisfies all assertions if and only if  $Q$  is satisfiable.*

*Proof. Feasible schedule for V implies Q is satisfiable:* proof by contradiction. Assume that  $Q$  is unsatisfiable even though there is a feasible schedule  $S$  for  $V$  that meets all the assertions. The reduction in Figure 4 considers all

<b>SAT:</b> $U \equiv \{u_0, u_1, \dots, u_m\}$		
$C \equiv \{c_0, c_1, \dots, c_n\}$		
$Q \equiv \{c_0 \wedge c_1 \wedge \dots \wedge c_n\}$		
<b>VAMPI:</b> $H \equiv \{h_{d_0}, h_{d_1}, h_C\}$		
$X \equiv \{u_0, \dots, u_m, g_0, g_1\}$		
$D \equiv \{d_0, d_1\}$		
$h_{d_0}$	$h_{d_1}$	$h_C$
$R(g_0, *)$	$R(g_1, *)$	$S(d_0, h_{d_0})$
$S(d_0, h_C)$	$S(d_1, h_C)$	$S(d_0, h_{d_1})$
		$R(u_0, *)$
		$R(u_0, *)$
$R(g_0, *)$	$R(g_1, *)$	$S(d_0, h_{d_0})$
$S(d_0, h_C)$	$S(d_1, h_C)$	$S(d_0, h_{d_1})$
		$R(u_1, *)$
		$R(u_1, *)$
...	...	...
		$assert(c_0)$
		$assert(c_1)$
		...

**Fig. 4.** General SAT to VAMPI reduction

truth values of the variables in  $Q$ , over every combination, by virtue of the non-determinism, and then asserts the truth of each of the clauses in  $Q$ . The complete set of possibilities is realized by sending in parallel from  $h_{d_0}$  and  $h_{d_1}$  the two truth valuations for a given variable to  $h_C$ . As these messages may be received in any order, each variable may assume either truth value. Further, each variable resolved is an independent choice so all combinations of variable valuations must be considered. This fact is a contradiction to the assumption of  $Q$  being unsatisfiable as the same truth values that meet the assertions would be a satisfying assignment in  $Q$ .

**Q is satisfiable implies feasible schedule for V:** the proof is symmetric to the previous case and proceeds in a like manner.

**Theorem 1 (NP-complete).** *VAMP is NP-complete.*

*Proof. Membership in NP:* a certificate is a schedule matching send and receives in each of the histories. The schedule is linearly scanned with the histories and checked that it does not violate MCAPI semantics. The next section constructs an operational model of MCAPI semantics that does just such a check given a schedule. The complexity is linear in the size of the schedule.

**NP-hard:** polynomial reduction from SAT. The correctness of the reduction is established by Lemma 1. The remainder of the proof is the complexity of the reduction. There are two tasks to send values  $d_0$  and  $d_1$  upon request. For each variable  $u_i \in U$ , each of these tasks,  $d_0$  and  $d_1$ , needs two operations: one to synchronize with  $h_C$  and another to send the value:  $O(2 * 2 * |U|)$ . The task  $h_C$

must request values from  $h_{d_0}$  and  $h_{d_1}$  and then receive both those values; it must do this for each variable:  $O(2 * 2 * |U|)$ . Once all the values are collected, it must then assert each clause:  $O(|C|)$ . As every term is linear, the reduction is linear.

<pre> <i>ctp</i> ::= (<i>t</i> ...) <i>t</i> ::= ([<math>\rho</math> <i>c</i>] ... <math>\perp</math>) <i>c</i> ::= (<b>assume</b> <i>e</i>)         (<b>assert</b> <i>e</i>)         (<b>x :=</b> <i>e</i>)         (<b>sndi</b> <i>a</i> <math>\alpha</math> <math>\beta</math> <i>e</i> <math>\rho</math>)         (<b>revi</b> <i>a</i> <math>\beta</math> <b>x</b> <math>\rho</math>)         (<b>wait</b> <i>a</i>) <i>m</i> ::= <math>\perp</math>   <math>\delta</math> <math>\delta</math> ::= ((<math>\beta</math> <math>\alpha</math>) ... <math>\perp</math>) <i>trace</i> ::= (<math>\sigma</math> ...) <math>\sigma</math> ::= (<math>\rho</math> <i>m</i>) <i>e</i> ::= (<b>op</b> <i>e</i> <i>e</i>)         <i>c</i>         <b>x</b>         <i>v</i> <i>v</i> ::= <b>number</b>         <i>bool</i> <i>bool</i> ::= <b>true</b>           <b>false</b> <math>\alpha</math> ::= <math>\gamma</math> <math>\beta</math> ::= <math>\gamma</math> </pre>	<pre> <i>mstate</i> ::= (<i>h</i> <math>\eta</math> <i>A</i> <i>Pnd.s</i> <i>Pnd.r</i> <i>Q.s</i> <i>ctp</i> <i>trace</i> <i>s</i>) <i>qstate</i> ::= (<i>Pnd.s</i> <i>Q.s</i> <i>m</i> <i>s</i>) <i>estate</i> ::= (<i>h</i> <math>\eta</math> <i>A</i> <i>Pnd.s</i> <i>Pnd.r</i> <i>Q.s</i> <i>e</i> <i>s</i> <i>k</i>)         <i>h</i> ::= <math>\emptyset</math>   (<i>h</i> [<i>l</i> <math>\rightarrow</math> <i>v</i>])         <math>\eta</math> ::= <math>\emptyset</math>   (<math>\eta</math> [<b>x</b> <math>\rightarrow</math> <i>l</i>])         <i>A</i> ::= ((<i>a</i> <math>\gamma</math>) ...) <i>Pnd.s</i> ::= <math>\emptyset</math>   (<i>Pnd.s</i> [<math>\beta</math> <math>\rightarrow</math> <i>frm</i>])         <i>frm</i> ::= <math>\emptyset</math>   (<i>frm</i> [<math>\alpha</math> <math>\rightarrow</math> <i>snd</i>])         <i>snd</i> ::= ([<i>a</i> <b>v</b>] ...) <i>Pnd.r</i> ::= <math>\emptyset</math>   (<i>Pnd.r</i> [<math>\beta</math> <math>\rightarrow</math> <i>rcv</i>])         <i>rcv</i> ::= ([<i>a</i> <b>x</b>] ...) <i>Q.s</i> ::= <math>\emptyset</math>   (<i>Q.s</i> [<math>\beta</math> <math>\rightarrow</math> <i>q</i>])         <i>q</i> ::= ([<i>a</i> <b>v</b>-<math>\perp</math>] ...) <i>v</i>-<math>\perp</math> ::= <math>\perp</math>   <b>v</b> <i>s</i> ::= <b>success</b>         <b>failure</b>         <b>infeasible</b>         <b>error</b> <i>k</i> ::= <b>ret</b>         (<b>assert</b> * <math>\rightarrow</math> <i>k</i>)         (<b>assume</b> * <math>\rightarrow</math> <i>k</i>)         (<b>x :=</b> * <math>\rightarrow</math> <i>k</i>)         (<b>op</b> * <i>e</i> <math>\rightarrow</math> <i>k</i>)         (<b>op v</b> * <math>\rightarrow</math> <i>k</i>) </pre>
(a)	(b)

**Fig. 5.** The trace language syntax with its evaluation syntax for the operational semantics—bold face indicates a terminal. (a) The input syntax with terminals **x**,  $\rho$  (which is unique), and *a* defined as strings and  $\gamma$  as a number. (b) The evaluation syntax with terminal *l* defined as a number.

$$\begin{array}{l}
\text{MACHINE STEP} \\
(Pnd.s \ Q.s \ m \ s) \rightarrow_q^* (Pnd.s_{p0} \ Q.s_{p0} \ m_p \ s_{p0}) \\
(h \ \eta \ A \ Pnd.s_{p0} \ Pnd.r \ Q.s_{p0} \ c_0 \ s_{p0} \ \mathbf{ret}) \rightarrow_e^* (h_p \ \eta_p \ A_p \ Pnd.s_{p1} \ Pnd.r_p \ Q.s_{p1} \ e \ s_{p1} \ \mathbf{ret}) \\
\hline
(h \ \eta \ A \ Pnd.s \ Pnd.r \ Q.s \ (t_0 \ \dots \ ([\rho_0 \ c_0] \ [\rho_1 \ c_1] \ [\rho_2 \ c_2] \ \dots) \ t_2 \ \dots) \ ([\rho_0 \ m] \ \sigma_1 \ \dots) \ s) \rightarrow_m \\
(h_p \ \eta_p \ A_p \ Pnd.s_{p1} \ Pnd.r_p \ Q.s_{p1} \ (t_0 \ \dots \ ([\rho_1 \ c_1] \ [\rho_2 \ c_2] \ \dots) \ t_2 \ \dots) \ (\sigma_1 \ \dots) \ s_{p1})
\end{array}$$

**Fig. 6.** Machine Reductions ( $\rightarrow_m$ )

## 4 Trace Language

The trace language is the theoretical framework for the match-pair encoding. The language syntax describes a CTP with a single execution trace on the same



$$\begin{array}{l}
\text{PROCESS QUEUE MOVEMENT} \\
([a_s v] [a_1 v_1] \dots) = \text{Pnd.s}(\beta)(\alpha) \\
\text{Pnd.s}_p = [\text{Pnd.s} \mid \beta \mapsto [\text{Pnd.s}(\beta) \mid \alpha \mapsto ([a_1 v_1] \dots)]] \quad ([a_1 v_1] \dots) = \text{Q.s}(\beta) \\
\text{Q.s}_p = [\text{Q.s} \mid \beta \mapsto ([a_s v] [a_1 v_1] \dots)] \quad s_p = \begin{cases} s & \text{if } |\text{Pnd.s}(\beta)(\alpha)| > 0 \\ \text{error} & \text{otherwise} \end{cases} \\
\hline
(\text{Pnd.s Q.s} ((\beta \alpha) (\beta_0 \alpha_0) \dots \perp) s) \rightarrow_q (\text{Pnd.s}_p \text{Q.s}_p ((\beta_0 \alpha_0) \dots \perp) s_p)
\end{array}$$

**Fig. 7.** Queue Reductions ( $\rightarrow_q$ )

$$\begin{array}{l}
\text{VARIABLE LOOKUP} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } x \ s \ k) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } h(\eta(x)) \ s \ k) \\
\\
\text{LEFT OPERAND} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } (op \ e_0 \ e) \ s \ k) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } e_0 \ s \ (op \ * \ e \rightarrow k)) \\
\\
\text{RIGHT OPERAND} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v \ s \ (op \ * \ e \rightarrow k)) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } e \ s \ (op \ v \ * \ \rightarrow k)) \\
\\
\text{BINARY OPERATION} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v_r \ s \ (op \ v_l \ * \ \rightarrow k)) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } op(v_l, v_r) \ s \ k) \\
\\
\text{ASSUME EXPRESSIONS EVALUATION} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } (assume \ e) \ s \ k) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } e \ s \ (assume \ * \ \rightarrow k)) \\
\\
\text{ASSUME COMMAND} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v \ s \ (assume \ * \ \rightarrow k)) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v \ assume(v, s) \ k) \\
\\
\text{ASSERT EXPRESSIONS EVALUATION} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } (assert \ e) \ s \ k) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } e \ s \ (assert \ * \ \rightarrow k)) \\
\\
\text{ASSERT COMMAND} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v \ s \ (assert \ * \ \rightarrow k)) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v \ assert(v, s) \ k) \\
\\
\text{ASSIGN EXPRESSIONS EVALUATION} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } (x \ := \ e) \ s \ k) \rightarrow_e \\
(h \eta A \text{Pnd.s Pnd.r Q.s } e \ s \ (x \ := \ * \ \rightarrow k)) \\
\\
\text{ASSIGN COMMAND} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } v \ s \ (x \ := \ * \ \rightarrow k)) \rightarrow_e \\
([h \mid \eta(x) \mapsto v] \eta A \text{Pnd.s Pnd.r Q.s } v \ s \ k) \\
\\
\text{SNDI COMMAND} \\
\frac{([a_1 v_1] \dots) = \text{Pnd.s}(\beta)(\alpha) \quad \text{Pnd.s}_p = [\text{Pnd.s} \mid \beta \mapsto [\text{Pnd.s}(\beta) \mid \alpha \mapsto ([a_0 h(\eta(x))] [a_1 v_1] \dots)]]}{(h \eta ([a \ \gamma] \dots) \text{Pnd.s Pnd.r Q.s } (\text{sndi } a_0 \ \alpha \ \beta \ x) \ s \ k) \rightarrow_e (h \eta ([a_0 \ \alpha] [a \ \gamma] \dots) \text{Pnd.s}_p \text{Pnd.r Q.s } \text{true } s \ k)} \\
\\
\text{RCVI COMMAND} \\
\frac{([a_1 x_1] \dots) = \text{Pnd.r}(\beta) \quad \text{Pnd.r}_p = [\text{Pnd.r} \mid \beta \mapsto ([a_0 x_0] [a_1 x_1] \dots)]}{(h \eta ([a \ \gamma] \dots) \text{Pnd.s Pnd.r Q.s } (\text{rcvi } a_0 \ \beta \ x_0) \ s \ k) \rightarrow_e (h \eta ([a_0 \ \beta] [a \ \gamma] \dots) \text{Pnd.s Pnd.r}_p \text{Q.s } \text{true } s \ k)} \\
\\
\text{WAIT (SNDI) COMMAND} \\
(h \eta A \text{Pnd.s Pnd.r Q.s } (\text{wait } a_s) \ s \ k) \rightarrow_e (h \eta A \text{Pnd.s Pnd.r Q.s } \text{true } s \ k) \\
\\
\text{WAIT (RCVI) COMMAND} \\
\frac{([a_0 \ \gamma_0] \dots [a_s \ \alpha] \dots [a_r \ \beta] \dots [a_1 \ \gamma_1]) = A \quad (h_p \ a_s \ \text{Pnd.r}_p \ \text{Q.s}_p \ s_p) = \text{getMarkRemove}(h, \eta, \text{Pnd.r}, \text{Q.s}, \beta, a_r, s) \quad A_p = ([a_0 \ \gamma_0] \dots [a_1 \ \gamma_1])}{(h \eta A \text{Pnd.s Pnd.r Q.s } (\text{wait } a_r) \ s \ k) \rightarrow_e (h_p \ \eta \ A_p \ \text{Pnd.s Pnd.r}_p \ \text{Q.s}_p \ \text{true } s_p \ k)}
\end{array}$$

**Fig. 8.** Expression Reductions ( $\rightarrow_e$ )

CTP. The evaluation syntax with its operational semantics define how to execute the CTP, following the specified trace, and define when that execution is a success (causes no assertion violation), a failure (causes an assertion violation), infeasible (causes an assume to not hold), or an error (uses a bogus match pair). Section 5 defines the encoding of a trace language program as an SMT problem and extends that encoding to capture a set of possible traces using match-pairs.

## 4.1 Syntax

Figure 5(a) is the syntax for a trace language program. This presentation uses ellipses (...) to represent zero or more repetitions, bold-face to indicate terminals, and omits commas in tuples for cleanliness. A trace language program is a CTP with a trace defining a sequential run of the CTP. The language defines a CTP ( $ctp$ ) as a list of threads. A thread ( $t$ ) is a list of pairs with each pair being a program location and a command. For simplicity, commands ( $c$ ) are restricted to assume, assert, assignment, non-blocking send ( $snd$ ) and receive ( $rcv$ ), and wait. We model a trace as an order of executed locations identified by ( $\rho$ ) and a series of queue movements that move a message from a source queue to a destination queue. The queue movement ( $m$ ) is either a special symbol indicating no movement or a list of move commands. The move list ( $\delta$ ) consists of several pairs of end-points where the first and second end-points are the destination and source end-points respectively. Each pair in the move list refers to a queue movement such that the first pending send of the source end-point should be moved to the send queue of the destination end-point. The non-terminal  $a$  in the grammar is a unique string identifier associated with a send or receive command referred to as an action ID in the text. The wait command takes a single action ID belonging to the associated send or receive action. The non-terminals  $\alpha$  and  $\beta$  are source and destination end-points respectively. The non-terminal  $\gamma$  is a number. The terminal  $x$  is any string not mentioned in the grammar definition and represents a program variable. Expressions ( $e$ ) are defined using prefix notation over binary operators. The bottom of Figure 1 is an example  $ctp$  in the trace language (omitting the first and second columns and trivially grouping each thread and its commands into an appropriate list using parenthesis).

A sequential trace of a CTP in the grammar is a list of trace entries ( $\sigma$ ). A trace entry is a pair consisting of a program location ( $\rho$ ) and either a queue movement list ( $m$ ) or a symbol ( $\perp$ ). An example of a trace can be seen in the bottom of Figure 1 in the second column by following the sequential order in the first column which starts on program location 20 of task 2. Notice that whenever the trace reaches a wait command on a receive action, the trace includes a queue movement list where the destination is listed followed by the source. Messages are delivered from the source queue to the destination queue. In other words, the trace resolves any non-determinism in scheduling or message buffering that is present in the CTP.

## 4.2 Operational Semantics

The operational semantics for the trace language are given by a term rewriting system using a *CESK* style machine<sup>1</sup> only the machine is augmented to include additional structure for modeling message passing. Figure 5(b) defines the machine state and other syntax relating to evaluation.

A machine step ( $\rightarrow_m$ ) in Figure 6 moves a thread forward by a single command. The rules operate on a machine state tuple (*mstate*) defined in Figure 5(b). The tuple can be partitioned into members relating to the *CESK* machine, members relating to the message passing model, and the trace status. The *CESK* machine members are *ctp* (the list of thread command sequences),  $\eta$  (an environment mapping a variable  $x$  to a location  $l$ ),  $h$  (a store mapping a location  $l$  to a value  $v$ ), and  $k$  (a continuation). Among the members of the message passing model,  $A$  is a dictionary mapping an action ID  $a$  to an end-point  $\gamma$ .  $Pnd\_s$  is a set of send queues where each queue is uniquely identified first by the source end-point and then by the destination end-point. The queue itself holds pairs consisting of an action ID and value  $(a,v)$ .  $Pnd\_r$  is a set of receive queues where each queue is uniquely identified by the destination end-point. The queue itself holds pairs consisting of an action ID and a variable  $(a,x)$ .  $Q\_s$  is also a set of queues where each queue is uniquely identified by the destination end-point. The queue itself holds pairs consisting of an action ID  $a$  and either a value  $v$  or a symbol  $\perp$ . Intuitively,  $Pnd\_s$ ,  $Q\_s$  and  $Pnd\_r$  are end-point queues tracking actions with associated values (sends) or variables (receives). Both  $Pnd\_s$  and  $Q\_s$  store the sends.  $Q\_s$  holds delivered messages that are ready to be received. A message moves from  $Pnd\_s$  to  $Q\_s$  through queue movement lists in the trace on the CTP.

The trace status in the *mstate* nine-tuple is given by  $s$  which ranges over a lattice:

$$\mathbf{success} \prec \mathbf{failure} \prec \mathbf{infeasible} \prec \mathbf{error}$$

The trace status only moves monotonically up the lattice starting from success. A success trace completes the entire trace, meets all the assume statements, and does not fail an assertion. A failure trace completes the entire trace, meets all the assume statements, but fails an assertion. An infeasible trace completes the entire trace but does not meet all the assume statements. An error is a trace that does not complete.

We have several *CESK* machines that handle different aspects of the CTP and trace. The machine step moves one step on the trace. In each machine step, the queue machine processes any queue movements in that step, and the expression machine handles any expressions in the command associated with the trace step.

The *Machine Step* inference in Figure 6 matches any *mstate* that has a thread whose first list entry matches the program location in the head of the trace. A match on the inference rewrites the *mstate* with new entries for each member of the nine-tuple by first applying the queue reduction relation until no more

<sup>1</sup> The *CESK* machine state is represented with a **C**ontrol string, **E**nvironment, **S**ore, and **K**ontinuation.

reductions apply (all queue movements are processed in the trace entry first) and then applying the expression reduction relation until no more reductions apply (as indicated by the asterisk). Note that queue reductions perform the queue movement such that all send actions are moved to the destination send queues. After completing all queue movements, messages can be delivered in the process of expression reduction.

The queue reduction for each command of the queue movement list in the trace entry is given in Figure 7. The definition of the *qstate* four-tuple is presented in the evaluation syntax in Figure 5(b). The symbol  $\perp$  in the queue movement  $m$  indicates that no more queue movement follows. A message moves from pending to delivered through queue movement lists. Each reduction step processes the first pair of the queue movement list and reduction steps follow until  $\perp$  is shown.

Expression reductions for each command in the trace language are given in Figure 8 and are defined over the *estate* nine-tuple in the evaluation syntax of Figure 5(b) which includes a continuation  $k$ . The **ret** continuation indicates that nothing follows, and an asterisk in a continuation is a place holder indication where evaluation is taking place. For example, the *Assume Expressions Evaluation* creates a continuation indicating that it is first evaluating the expression in the assume command. Once that expression reduces to a value, then the *Assume Command* inference matches to check the assumptions.

The expression reductions use several helper functions. The function  $op(v_l, v_r)$  applies the “op” to the left and right operands. The function `getMarkRemove` is explained in the later paragraph. The other helper functions are defined below:

$$\text{assert}(v, s) = \begin{cases} \mathbf{failure} & \text{if } s \prec \mathbf{failure} \wedge \\ & v = \mathbf{false} \\ s & \text{otherwise} \end{cases}$$

$$\text{assume}(v, s) = \begin{cases} \mathbf{infeasible} & \text{if } s \prec \mathbf{infeasible} \wedge \\ & v = \mathbf{false} \\ s & \text{otherwise} \end{cases}$$

Note that the status only moves monotonically up the lattice as mentioned previously. The notations  $h(\eta(x))$  (Sndi Command),  $Pnd\_r(\beta)$  (Rcvi Command), and  $Pnd\_s(\beta)(\alpha)$  (Sndi Command) in Figure 8 are used for lookup. For example,  $Pnd\_s(\beta)(\alpha)$  returns a list of pairs of action IDs and values as defined in Figure 5(b).

The *Sndi Command* and *Rcvi Command* in Figure 8 update  $Pnd\_r$  and  $Pnd\_s$  respectively with information to complete a message send or receive at a wait command. Consider a portion of the *Rcvi Command*:

$$Pnd\_r_p = [Pnd\_r \mid \beta \mapsto ([a_0 \ x_0] [a_1 \ x_1] \dots)]$$

$$([a_1 \ x_1] \dots) = Pnd\_r(\beta)$$

$Pnd\_r_p$  is a new set, just like the old set  $Pnd\_r$ , only the new set maps the destination end-point  $\beta$  to its contents in the old set plus the added entry to the front of the list of the action ID and variable for the receive command being

evaluated. Considering the entire rule in the figure, it also updates the action ID map,  $A = ([a_0 \ \beta] [a \ \gamma] \dots)$ , to include the coupling between the action ID and its destination end-point.

The function of the *Wait(SNDI) Command* rule in Figure 8 is to consume the wait command only. The function of the other rule for the wait command is more involved. The *Wait(RCVI) Command* rule is complicated because a CTP can wait on receive actions out of program order. Consider the following trace

$$(\dots (rcvA \perp) (rcvB \perp) (wait(rcvB) \perp) (wait(rcvA) \perp) \dots)$$

omitting the regular definition for the commands and the trace. It is perfectly valid to now call  $wait(rcvB)$  even though it appeared after the  $rcvA$  action. The *Wait(RCVI) Command* handles this very situation in function `getMarkRemove`. Consider the same trace above, we have the following structures of the *pending receive queue* and the *delivered queue* respectively,

$$Pnd\_r: (\dots (\beta \rightarrow ([rcvA \ x_0] [rcvB \ x_1] \dots)) \dots)$$

$$Q\_s: (\dots (\beta \rightarrow ([snd1 \ v_0] [snd2 \ v_1] \dots)) \dots)$$

Function `getMarkRemove` in the command  $wait(rcvB)$  works as the following steps: first finding the pair  $[rcvB \ x_1]$  in the pending receive queue; then getting the pair  $[snd2 \ v_1]$  in the delivered queue; then marking the pair  $[snd1 \ v_0]$  as first assigning  $v_0$  to variable  $x_0$  in the heap then updating the pair  $[snd1 \ v_0]$  to  $[snd1 \ \perp]$ ; finally removing the pair  $[rcvB \ x_1]$  and  $[snd2 \ v_1]$  in  $Pnd\_r$  and  $Q\_s$ , respectively.

The syntax with the operational semantics, as presented, are implemented directly in PLT Redex. PLT Redex is a language for testing and debugging semantics using term rewriting and is part of the Racket runtime.

The following definition is important to the proof of the match pair encoding in Section 5.

**Definition 2.** *A machine state  $(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ Q\_s \ ctp \ trace \ s \ k)$  is well formed if and only if it reduces to a final state where the CTP and trace run to completion, having matched every send and receive call, and the status is either success, failure, or infeasible:*

$$\begin{aligned} (h \ \eta \ A \ Pnd\_s \ Pnd\_r \ Q\_s \ ctp \ trace \ s \ k) &\rightarrow_m^* \\ (h_p \ \eta_p \ () \ Pnd\_s_p \ Pnd\_r_p \ Q\_s_p \ (( \dots) \ () \ s_p \ ret) \end{aligned}$$

such that

$$\begin{aligned} \forall \beta, \forall \alpha, Pnd\_s_p(\beta)(\alpha) = () \wedge \forall \beta, Pnd\_r_p(\beta) = () \wedge \\ \forall \beta, Q\_s_p(\beta) = () \quad \wedge s_p \prec \mathbf{error} \end{aligned}$$

For convenience, we define the function  $status(m)$  to return the final status after reduction of a well formed machine state  $m$ .

Intuitively, a well-formed machine state completes all the transitions of send and receive calls, meaning that there are no elements in  $A$ ,  $Pnd\_s$ ,  $Pnd\_r$  and  $Q\_s$ , the commands in the CTP and the execution trace are correctly executed, and the status of state never enters **error**.

$$\begin{array}{l}
\text{MACHINE STEP} \\
(Pnd\_s \ Q\_s \ m \ s) \rightarrow_q^* (Pnd\_s_{p0} \ Q\_s_{p0} \ m_p \ s_{p0}) \\
(h \ \eta \ A \ Pnd\_s_{p0} \ Pnd\_r \ Q\_s_{p0} \ c_0 \ s_{p0} \ \mathbf{ret} \ l \ smt) \rightarrow_e^* (h_p \ \eta_p \ A_p \ Pnd\_s_{p1} \ Pnd\_r_p \ Q\_s_{p1} \ e \ s_{p1} \ \mathbf{ret} \ l_p \ smt_{p0}) \\
smt_{p1} = \text{addHB}(smt_{p0} \ \rho_0 \ \rho_1) \\
\hline
(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ Q\_s \ (t_0 \ \dots \ ([\rho_0 \ c_0] \ [\rho_1 \ c_1] \ [\rho_2 \ c_2] \ \dots) \ t_2 \ \dots) \ ([\rho_0 \ m] \ \sigma_1 \ \dots) \ s \ l \ smt) \rightarrow_m \\
(h_p \ \eta_p \ A_p \ Pnd\_s_{p1} \ Pnd\_r_p \ Q\_s_{p1} \ (t_0 \ \dots \ ([\rho_1 \ c_1] \ [\rho_2 \ c_2] \ \dots) \ t_2 \ \dots) \ (\sigma_1 \ \dots) \ s_{p1} \ l_p \ smt_{p1})
\end{array}$$

**Fig. 9.** Machine Reductions to build the SMT model of a trace language program  
( $\rightarrow_{m\text{-smt}}$ )

$$\begin{array}{l}
\text{ASSUME EXPRESSIONS EVALUATION} \\
(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ (\mathit{assume} \ e) \ s \ k \ l \ (\mathit{defs} \ (\mathit{any} \ \dots))) \rightarrow_e (h \ \eta \ A \ Pnd\_s \ Pnd\_r \ e \ s \ (\mathit{assume} \ * \rightarrow k) \ l \ (\mathit{defs} \ (e \ \mathit{any} \ \dots))) \\
\text{ASSERT EXPRESSIONS EVALUATION} \\
(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ (\mathit{assert} \ e) \ s \ k \ l \ (\mathit{defs} \ (\mathit{any} \ \dots))) \rightarrow_e (h \ \eta \ A \ Pnd\_s \ Pnd\_r \ e \ s \ (\mathit{assert} \ * \rightarrow k) \ l \ (\mathit{defs} \ ((\mathit{not} \ e) \ \mathit{any} \ \dots))) \\
\text{ASSIGN EXPRESSIONS EVALUATION} \\
(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ (x \ := \ e) \ s \ k \ l \ (\mathit{defs} \ (\mathit{any} \ \dots))) \rightarrow_e (h \ \eta \ A \ Pnd\_s \ Pnd\_r \ e \ s \ (x \ := \ * \rightarrow k) \ l \ (\mathit{defs} \ ((= \ x \ e) \ \mathit{any} \ \dots))) \\
\text{SNDI COMMAND} \\
\begin{array}{l}
([a_1 \ v_1] \ \dots) = Pnd\_s(\beta)(\alpha) \\
Pnd\_s_p = [Pnd\_s \mid \beta \mapsto [Pnd\_s(\beta) \mid \alpha \mapsto ([a_0 \ h(\eta(x)) \ [a_1 \ v_1] \ \dots)]] \\
\mathit{any}_0 = (\mathit{define} \ a \ :: \ \mathit{send}) \\
\mathit{any}_1 = (\mathit{and} \ (= \ (\mathit{select} \ a \ \mathit{ep}) \ \beta) \ (= \ (\mathit{select} \ a \ \mathit{value}) \ x))
\end{array} \\
\hline
(h \ \eta \ ([a \ \gamma] \ \dots) \ Pnd\_s \ Pnd\_r \ Q\_s \ (\mathbf{sndi} \ a_0 \ \alpha \ \beta \ x) \ s \ k \ l \ ((\mathit{any}_d \ \dots) \ (\mathit{any}_a \ \dots))) \rightarrow_e \\
(h \ \eta \ ([a_0 \ \alpha] \ [a \ \gamma] \ \dots) \ Pnd\_s_p \ Pnd\_r \ Q\_s \ \mathbf{true} \ s \ k \ l \ ((\mathit{any}_0 \ \mathit{any}_d \ \dots) \ (\mathit{any}_1 \ \mathit{any}_a \ \dots))) \\
\text{RCVI COMMAND} \\
\begin{array}{l}
([a_1 \ x_1] \ \dots) = Pnd\_r(\beta) \\
Pnd\_r_p = [Pnd\_r \mid \beta \mapsto ([a_0 \ x_0] \ [a_1 \ x_1] \ \dots)] \quad \mathit{any}_0 = (\mathit{define} \ a \ :: \ \mathit{recv}) \\
\mathit{any}_1 = (\mathit{and} \ (= \ (\mathit{select} \ a \ \mathit{ep}) \ \beta) \ (= \ (\mathit{select} \ a \ \mathit{var}) \ x_0))
\end{array} \\
\hline
(h \ \eta \ ([a \ \gamma] \ \dots) \ Pnd\_s \ Pnd\_r \ Q\_s \ (\mathbf{rcvi} \ a_0 \ \beta \ x_0) \ s \ k \ l \ ((\mathit{any}_d \ \dots) \ (\mathit{any}_a \ \dots))) \rightarrow_e \\
(h \ \eta \ ([a_0 \ \beta] \ [a \ \gamma] \ \dots) \ Pnd\_s \ Pnd\_r_p \ Q\_s \ \mathbf{true} \ s \ k \ l \ ((\mathit{any}_0 \ \mathit{any}_d \ \dots) \ (\mathit{any}_1 \ \mathit{any}_a \ \dots))) \\
\text{WAIT (RCVI) COMMAND} \\
\begin{array}{l}
([a_0 \ \gamma_0] \ \dots \ [a_s \ \alpha] \ \dots \ [a_r \ \beta] \ \dots \ [a_1 \ \gamma_1]) = A \\
(h_p \ a_s \ Pnd\_r_p \ Q\_s_p \ s_p) = \mathit{getMarkRemove}(h, \eta, Pnd\_r, Q\_s, \beta, a_r, s) \\
A_p = ([a_0 \ \gamma_0] \ \dots \ [a_1 \ \gamma_1]) \quad (l_p \ \mathit{last\_as}) = \mathit{getlastsend/replace}(l, a_s, \alpha, \beta)
\end{array} \\
\hline
(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ Q\_s \ (\mathbf{wait} \ a_r) \ s \ k \ l \ (\mathit{defs} \ (\mathit{any}_a \ \dots))) \rightarrow_e \\
(h_p \ \eta \ A_p \ Pnd\_s \ Pnd\_r_p \ Q\_s_p \ \mathit{true} \ s_p \ k \ l_p \ (\mathit{defs} \ ((\mathbf{HB} \ (\mathit{select} \ \mathit{last\_as}) \ MP) \ (\mathit{select} \ a_s \ MP)) \ (\mathbf{MATCH} \ a_r \ a_s) \ \mathit{any}_a \ \dots))
\end{array}$$

**Fig. 10.** Expression Reductions to build the SMT model of a trace language program  
( $\rightarrow_{e\text{-smt}}$ )

## 5 SMT Encoding

The new SMT encoding is based on (1) a trace of events during an execution of an MCAPI program including control-flow assumptions and property assertions, such as Figure 2; and (2) a set of possible match pairs. A match pair is the coupling of a receive to a particular send. In the running example, the set admits, for example, that  $R_{0,2}$  can be matched with either  $S_{1,5}$  or  $S_{2,4}$ . This direct use of match pairs, rather than a state-based or indirect use of match pairs in an order-based encoding, [6] and [5], is novel.

The purpose of the SMT encoding is to force the SMT solver to resolve the match pairs for the system in such a way that the final values of program variables meet the assumptions on control flow but violate some assertion. In essence, the SMT solver completes a partial order on operations into a total order that determines the final match pair relationships.

### 5.1 Definitions

The encoding needs to express the partial order imposed by the MCAPI semantics as SMT constraints. The partial order is based on a *Happens-Before* relation over operations such as send, receive, wait, or assert:

**Definition 3 (Happens-Before).** *The Happens-Before (HB) relation, denoted as  $\prec_{\text{HB}}$ , is a partial order over operations.*

Given two operations,  $A$  and  $B$ , if  $A$  must complete before  $B$  in a valid program execution, then  $A \prec_{\text{HB}} B$  will be an SMT constraint.

The relation is derived from the program source and potential match pairs. In order to specify the constraints from the program source, each program operation is mapped to a set of variables that can be manipulated by the SMT solver.

**Definition 4 (Wait).** *The occurrence of a wait operation,  $W$ , is captured by a single variable,  $order_w$ , that constrains when the wait occurs.*

It is not enough to represent all events as simple numbers that will be ordered in this way. Such an encoding would not allow the solver to discover what values would flow across communication primitives. Instead, some events in the trace are modeled as a set of SMT variables that record the pertinent information about the event. For example,

**Definition 5 (Send).** *A send operation  $S$ , is a four-tuple of variables:*

1.  $M_s$ , the order of the matching receive event;
2.  $order_s$ , the order of the send;
3.  $e_s$ , the endpoint; and,
4.  $value_s$ , the transmitted value.

The most complex operation in MCAPI is a receive. Since receives are inherently asynchronous, it is not possible to represent them atomically. Instead, we

need to associate each receive with a wait that marks where in the program the receive operation is guaranteed to be complete. The MCAPAPI runtime semantics allow a single wait to witness the completion of many receives due to *the message non-overtaking property*. A wait that witnesses the completion of one or more receives is the *nearest-enclosing wait*.

**Definition 6 (Nearest-Enclosing Wait).** *A wait that witnesses the completion of a receive by indicating that the message is delivered and that all the previous receives in the same task issued earlier are complete as well.*

Task 0	Task 1
$R_{0,1}(*, \&h1)$	$S_{1,1}(0, \&h3)$
$R_{0,2}(*, \&h2)$	$W(\&h3)$
$W(\&h2)$	$S_{1,2}(0, \&h4)$
$W(\&h1)$	$W(\&h4)$

**Fig. 11.** Nearest-enclosing Wait example

Figure 11 shows that the wait  $W(\&h2)$  witnesses the completion of the receive  $R_{0,1}$  and  $R_{0,2}$  in task 0. Thus,  $W(\&h2)$  is their nearest-enclosing wait.

The encoding requires that every receive operation have a nearest-enclosing wait as it makes match pair decisions at the wait operation. The requirement is not a limitation of the encoding, as accessing a buffer from a receive that does not have a nearest-enclosing wait is an error. Rather, the wait is a convenience in the encoding to mark where a receive actually takes place. The same requirement can be made for sends for correctness but is not required for the encoding as send buffering is handled differently than receive buffering. The encoding effectively ignores wait operations for sends as will be seen.

**Definition 7 (Receive).** *A receive operation  $R$  is modeled by a five-tuple of variables:*

1.  $M_R$ , the order of the matching send event;
2.  $order_R$ , the order of the receive;
3.  $e_R$ , the endpoint;
4.  $value_R$ , the received value; and,
5.  $nw_R$ , the order of the nearest enclosing wait.

## 5.2 Assumptions, Assertions, and match pairs

The definitions so far merely establish the pertinent information about each event in the trace as SMT variables. It is necessary to now express constraints on those variables.

The most trivial kind of constraints are those for control-flow assumptions.



**Definition 8 (Assumption).** *Every assumption  $A$  is added as an SMT assertion.*

It may seem strange to turn *assumptions* into *assertions*, but from a constraint perspective, the assumption that we have already observed some property (during control-flow) is equivalent to instructing the SMT solver to treat it as inviolate truth, or an assertion.

The next level of constraint complexity comes from property assertions. These correspond to the invariants of the program. The goal is to discover if they can be violated, so we instruct the SMT solver to seek for a way to satisfy their *negation* given all the other constraints.

**Definition 9 (Property Assertion).** *For every property assertion  $P$ ,  $\neg P$  is added as an SMT assertion.*

Finally, we must express the relation in a given match pair as a set of SMT constraints. Informally, a match pair equates the shared components of a send and receive and constrains the send to occur before the nearest-enclosing wait of the receive. Formally:

**Definition 10 (Match Pair).** *A match pair,  $\langle R, S \rangle$ , for a receive  $R$  and a send  $S$  corresponds to the constraints:*

1.  $M_R = order_S$
2.  $M_S = order_R$
3.  $e_R = e_S$
4.  $value_R = value_S$  and
5.  $order_S \prec_{HB} nw_R$

The encoding is given a set of potential match pairs over all the sends and receives in the program trace. The constraints from these match pairs are not simply joined in a conjunctions. If we were to do that, then we would be constraining the system such that a single receive must be paired with all possible sends in a feasible execution rather than a single send. Therefore, we combine all the constraints for a given receive with all possible sends as specified by the input match pairs into a single disjunction:

**Definition 11 (Receive Matches).** *For each receive  $R$ , if  $\langle R, S_0 \rangle$  through  $\langle R, S_n \rangle$  are match pairs, then  $\bigvee_i \langle R, S_i \rangle$  is used as an SMT constraint.*

This encoding of the input ensures that the SMT solver can only use compatible send/receive pairs and ensures that sends happen before nearest-enclosing waits on receives.

### 5.3 Program Order Constraints

The encoding thus far is missing additional constraints on the *Happens-Before* relation stemming from program order. These constraints are added in four steps: we must ensure that sends to common endpoints occur in program order in a single task (step 1); similarly for receives (step 2); receives occur before their nearest-enclosing wait (step 3); and, that sends are received in the order they are sent (step 4).

*Step 1* For each task, if there are sequential send operations, say  $S$  and  $S'$ , from that task to a common endpoint,  $e_S = e_{S'}$ , then those sends must follow program order:  $order_S \prec_{HB} order_{S'}$ .

*Step 2* For each task, if there are sequential receive operations, say  $R$  and  $R'$ , in that task on a common endpoint,  $e_R = e_{R'}$ , then those receives must follow program order:  $order_R \prec_{HB} order_{R'}$ .

*Step 3* For every receive  $R$  and its nearest enclosing wait  $W$ ,  $order_R \prec_{HB} order_W$ .

*Step 4* For any pair of sends  $S$  and  $S'$  on common endpoints,  $e_S = e_{S'}$ , such that  $order_S \prec_{HB} order_{S'}$ , then those sends must be received in the same order:  $M_S \prec_{HB} M_{S'}$ .

For example, consider two tasks where task 0 sends two messages to task 1 as shown in Figure 12.

Task 0	Task 1
$S_{0,1}(1, \&h1)$	$R_{1,1}(*, \&h3)$
$S_{0,2}(1, \&h2)$	$R_{1,2}(*, \&h4)$
$W(\&h1)$	$W(\&h3)$
$W(\&h2)$	$W(\&h4)$

**Fig. 12.** Send Ordering Example

The  $M_S$  variables from the sends will be assigned to the orders for  $R_{1,1}$  and  $R_{1,2}$  by the match pairs selected by the SMT solver. The constraints added in this step force the send to be received in program order using the HB relation which for this example yields  $M_{S_{0,1}} \prec_{HB} M_{S_{0,2}}$ .

#### 5.4 Zero Buffer Semantics

The constraints presented so far correspond to an infinite-buffer semantics, because we do not constrain how many messages may be in transit at once. We can add additional, orthogonal, constraints to further restrict behavior and enforce a zero-buffer semantics. There are two kinds of such constraints.

First, for each task, if there are two sends  $S$  and  $S'$  such that  $order_S \prec_{HB} order_{S'}$ , and  $S$  and  $S'$  can both match a receive  $R$ , then we add the following constraint to the encoding:  $order_W \prec_{HB} order_{S'}$  where  $W$  is the nearest-enclosing wait that witnesses the completion of  $R$  in execution.

The second constraint relies on a dependence relation between two match pairs.

**Definition 12.** *To match pairs are dependent, denoted as  $\langle R, S \rangle \dashv \langle R', S' \rangle$ , if and only if*

```

...
01  $order_{R_{0,2}} \prec_{HB} order_{W(\&ch1)}$ 
02  $order_{R_{0,5}} \prec_{HB} order_{W(\&ch2)}$ 
03  $order_{R_{0,2}} \prec_{HB} order_{R_{0,5}}$ 
04  $order_{R_{1,3}} \prec_{HB} order_{W(\&ch3)}$ 
05 ( $assert (> b 0)$ )
06 ( $assert (not (= a 4))$ )
07  $\langle R_{0,2}, S_{2,4} \rangle \vee \langle R_{0,2}, S_{1,5} \rangle$ 
08  $\langle R_{0,5}, S_{2,4} \rangle \vee \langle R_{0,5}, S_{1,5} \rangle$ 
09  $\langle R_{1,3}, S_{2,7} \rangle$ 

```

**Fig. 13.** SMT Encoding

1. the nearest-enclosing wait  $W$  of  $R'$  issues before  $S$  on an identical endpoint;  
or
2.  $\exists \langle R''S'' \rangle$  such that  $\langle R, S \rangle \rightarrow \langle R'', S'' \rangle \wedge \langle R'', S'' \rangle \rightarrow \langle R', S' \rangle$ .

With the dependence relation, the second set of constraints for the zero-buffer semantics is give as: for each pair of sends  $S$  and  $S'$  that can both match an receive  $R$ , if there is a send  $S''$  issued after the issuing of  $S'$  by an identical endpoint, and a receive  $R'$  such that  $\langle R, S \rangle \rightarrow \langle R', S'' \rangle$ , then we add the following constraint to the encoding:  $order_W \prec_{HB} order_S$  where  $W$  is the nearest-enclosing wait that witnesses the completion of  $R$ .

## 5.5 Example

Figure 13 shows the encoding of Figure 1 as an SMT problem. We elide the basic definition of the variables discussed in Section 5.1. Lines 05 through 09 give the assumptions, assertions, and match pairs. The first four lines reflect the program order constraints: receives happen before corresponding wait operations and receives from a common endpoint follow program order. There are no constraints between sends because there are no sequential sends from a common endpoint to a common endpoint. To encode the zero-buffer semantics, the constraint  $order_{W(\&ch1)} \prec_{HB} order_{S_{1,5}}$  would need to be added to force the receive to complete before another send is issued.

## 5.6 Correctness

Before we can state our correctness theorem, we must define a few terms. We define our encoder as a function from programs and match pair sets to SMT problems:

**Definition 13 (Encoder).** For all programs,  $p$ , and match pair sets  $m$ , let  $SMT(p, m)$  be our encoding as an SMT problem.

We assume that an SMT solver can be represented as a function that takes a problem and returns a satisfying assignment of variables or an unsatisfiable flag:

**Definition 14 (SMT Solver).** *For all SMT problems,  $s$ , let  $SOL(s)$  be in  $\sigma + \text{UNSAT}$ , where  $\sigma$  is a satisfying assignment of variables to values.*

We assume that from a satisfying assignment to one of our SMT problems, we can derive an execution trace by observing the values given to each of the  $order_e$  variables. In other words, we can view the SMT solver as returning traces and not assignments.

We assume a semantics for traces that gives their behavior as either having an assertion violation or being correct. This formal semantics is presented in Section 4.

**Definition 15 (Semantics).** *For all programs,  $p$ , and traces  $t$ ,  $SEM(p, t)$  is either  $\text{BAD}$  or  $\text{OK}$ .*

Given this framework, our SMT encoding technique is sound if

**Theorem 2 (Soundness).** *For all programs,  $p$ , and match pair sets,  $m$ ,  $SOL(SMT(p, m)) = t \Rightarrow SEM(p, t) = \text{BAD}$ .*

Our soundness proof relies on the following lemma:

**Lemma 2.** *Any match pair  $\langle R, S \rangle$  used in a satisfying assignment of an SMT encoding is a valid match pair and reflects an actual possible MCAPI program execution.*

*Proof.* We prove this by contradiction. First, assume that  $\langle R, S \rangle$  is an invalid match pair (i.e. one that is not valid in an actual MCAPI execution). Second, assume that the SMT solver finds a satisfying assignment.

Since  $\langle R, S \rangle$  is not a valid match pair, match  $R$  and  $S$  requires program order, message non-overtaking, or no-multiple match to be violated. In other words, the *Happens-Before* constraints encoded in the SMT problem are not satisfied.

This is a contradiction: either the SMT solver would not return an assignment or the match pair was actually valid.

The correctness of our technique relies on completeness:

**Theorem 3 (Completeness).** *For all programs,  $p$ , and traces,  $t$ ,  $SEM(p, t) = \text{BAD} \Rightarrow \exists m. SOL(SMT(p, m)) = t$ .*

To completely prove completeness, we give the following modified framework that produces a correlating SMT problem for a given execution trace. The evaluation syntax for the machine reductions to build the SMT model of a trace

language program are largely those of the regular machine reductions with a few changes below

$$\begin{aligned}
mstate &::= (h \ \eta \ A \ Pnd\_s \ Pnd\_r \ Q\_s \ ctp \ trace \ s \ k \ l \ smt) \\
estate &::= (h \ \eta \ A \ Pnd\_s \ Pnd\_r \ Q\_s \ e \ s \ k \ l \ smt) \\
smt &::= (defs \ constraints \ ML) \\
defs &::= (any \ \dots) \\
constraints &::= (any \ \dots) \\
ML &::= (ML = \emptyset \mid (ML \ [\mu \rightarrow[(\mu \ \nu) \dots]])
\end{aligned}$$

The new syntax adds the *smt* member to the *mstate* and *estate* where the “any” term in *defs* and *constraints* matches any structure. The lists will be filled with definitions, HB entries, MATCH entries, etc. as defined by the SMT machine reductions. A match-list *ML* is a set of match pair lists uniquely identified by a receive action ID, and each list consists of a set of match pairs for this receive action ID and a send action ID. Note that we alpha renamed so that  $\mu$  and  $\nu$  are all unique labels. For convenience, we define function  $dom(ML)$  and  $range(ML)$  that return the set of receive action IDs in the set of match pairs of *ML* and the set itself, respectively. The symbol *l* is another member added to the *mstate* and *estate*. It records the last send operation that happens before the current one in the same thread. The *qstate* does not contain any SMT encodings in the reductions and it is not changed.

The changes for the reduction rules are presented in Figure 9 and Figure 10. The support function  $addHB((any_d\dots) (any_c\dots)) \ \rho_0 \ \rho_1$  adds program location to the definition list and adds a happens before relation

$$\equiv (((define \ \rho_0 :: int) \ any_d\dots)((HB \ \rho_0 \ \rho_1) \ any_c\dots))$$

The reduction rules in Figure 9 and Figure 10 add constraints to the SMT problem with respect to the semantic definition of the trace language. The Machine Step in Figure 9 adds the program order of actions in the same thread to the SMT problem. The evaluation of assume, assert and assign in Figure 10 adds assertions to the SMT problem. The *Sndi* and *Rcvi* commands define the send and receive actions in the SMT problem, and the *Wait(Rcvi)* reduction adds the “match” constraints to the SMT problem. Note that the function *getlastsend/replace* picks up the last send operation, *last\_as*, that happens before the current one in the same thread, and adds a HB relation between the *MP* of *last\_as* and that of the current send operation. The added relation ensures the messages from a common endpoint are non-overtaking (i.e., FIFO ordered).

The following definition supports the rest of the section.

**Definition 16.** *A machine state  $(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ ctp \ trace \ s \ k)$  is smt-enabled if it is well formed. The SMT problem is taken from the final SMT reduced state:*

$$\begin{aligned}
&(h \ \eta \ A \ Pnd\_s \ Pnd\_r \ ctp \ trace \ s \ k \ (()) \ ()) \rightarrow_{m-smt}^* \\
&(h_p \ \eta_p \ A_p \ Pnd\_s_p \ Pnd\_r_p \ ctp_p \ trace_p \ s_p \ k_p \ smt)
\end{aligned}$$

For convenience, we define the function  $\text{getSMT}(m) \mapsto \text{smt}$  to return the SMT problem in the final state, and  $\text{ANS}(\text{getSMT}(m)) \mapsto \{\mathbf{SAT}, \mathbf{UNSAT}\}$  to return the status of the SMT problem in the final state of  $m$ . Also, we define the relation of the range of function ANS such that  $\mathbf{UNSAT} > \mathbf{SAT}$ .

<pre> <i>defs</i> is not shown;  <i>constraints</i>; 00 <i>event</i><sub>R<sub>0,2</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch1) 01 <i>event</i><sub>R<sub>0,5</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch2) 02 <i>event</i><sub>R<sub>0,2</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>R<sub>0,5</sub></sub> 03 <i>event</i><sub>R<sub>1,3</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch3) 04 <i>event</i><sub>S<sub>1,5</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch4) 05 <i>event</i><sub>S<sub>2,4</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch5) 06 <i>event</i><sub>S<sub>2,7</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch6) 07 (<i>assert</i> (&gt; b 0)) 08 (<i>assert</i> (not (= a 4)))  <i>match</i>; 09 (R<sub>0,2</sub>, S<sub>2,4</sub>) 10 (R<sub>0,5</sub>, S<sub>1,5</sub>) 11 (R<sub>1,3</sub>, S<sub>2,7</sub>) </pre> <p style="text-align: center;">(a)</p>	<pre> <i>defs</i> is not shown;  <i>constraints</i>; 00 <i>event</i><sub>R<sub>0,2</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch1) 01 <i>event</i><sub>R<sub>0,5</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch2) 02 <i>event</i><sub>R<sub>0,2</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>R<sub>0,5</sub></sub> 03 <i>event</i><sub>R<sub>1,3</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch3) 04 <i>event</i><sub>S<sub>1,5</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch4) 05 <i>event</i><sub>S<sub>2,4</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch5) 06 <i>event</i><sub>S<sub>2,7</sub></sub> &lt;<sub>HB</sub> <i>event</i><sub>W</sub>(&amp;ch6) 07 (<i>assert</i> (&gt; b 0)) 08 (<i>assert</i> (not (= a 4)))  <i>match</i>; 09 (R<sub>0,2</sub>, S<sub>1,5</sub>) 10 (R<sub>0,5</sub>, S<sub>2,4</sub>) 11 (R<sub>1,3</sub>, S<sub>2,7</sub>) </pre> <p style="text-align: center;">(b)</p>	<pre> ... 01 <i>order</i><sub>R<sub>0,2</sub></sub> &lt;<sub>HB</sub> <i>order</i><sub>W</sub>(&amp;ch1) 02 <i>order</i><sub>R<sub>0,5</sub></sub> &lt;<sub>HB</sub> <i>order</i><sub>W</sub>(&amp;ch2) 03 <i>order</i><sub>R<sub>0,2</sub></sub> &lt;<sub>HB</sub> <i>order</i><sub>R<sub>0,5</sub></sub> 04 <i>order</i><sub>R<sub>1,3</sub></sub> &lt;<sub>HB</sub> <i>order</i><sub>W</sub>(&amp;ch3) 05 (<i>assert</i> (&gt; b 0)) 06 (<i>assert</i> (not (= a 4))) 07 (R<sub>0,2</sub>, S<sub>2,4</sub>) ∨ (R<sub>0,2</sub>, S<sub>1,5</sub>) 08 (R<sub>0,5</sub>, S<sub>2,4</sub>) ∨ (R<sub>0,5</sub>, S<sub>1,5</sub>) 09 (R<sub>1,3</sub>, S<sub>2,7</sub>) </pre> <p style="text-align: center;">(c)</p>
---	---	--

**Fig. 14.** SMT problems. (a) SMT problem based on the first trace. (b) SMT problem based on a second trace. (c) SMT problem built from the preceding two problems.

Figure 14 shows the SMT problem (minus the event definitions) generated from our new *CESK SMT* encoding machines for the program in Figure 1 on two different traces (parts (a) and (b)). Note that the encodings in Figure 14 (a) and (b) are identical except for the MATCH clauses, which are generated from the trace parts of the program. Furthermore, the encoding in Figure 14(a) violates the properties of interest encoded in the assert clauses, and the answer for the encoding in Figure 14 (a) is SAT. From this observation, we present the following lemma that indicates the equivalence between the status of our *CESK* machines and the satisfiability of the generated SMT encoding.

**Lemma 3.** *For a well-formed machine state  $m$ ,*

$$\text{status}(m) = \mathbf{failure} \iff \text{ANS}(\text{getSMT}(m)) = \mathbf{SAT}$$

*Proof.* The generated SMT problem in Figure 9 and Figure 10 captures the execution trace from the initial state up to the final state by looking at the following facts. First, the machine step in Figure 9 adds an HB relation for two consecutive program locations  $\rho_0$  and  $\rho_1$  in an identical thread, with respect to the program order constraint. Second, the assume, assert and assign expression evaluations in Figure 10 add assert constraints for variable  $x$  or expression  $e$

with respect to the trace constraint. Note that the negated system is used in the assert evaluation when generating the SMT statement. Third, the `sndi` and `rcvi` commands in Figure 10 add `any0` as the definition of a receive or send operation, followed by `any1` as the initialization of operation fields. Finally, the `wait(rcvi)` command in Figure 10 adds a match pair for the receive operation  $a_r$  and the send operation  $a_s$  where  $a_s$  is obtained by applying the process we described in Section 3.2. Furthermore, an HB relation is added with respect to the program order, such that the messages from a common endpoint are non-overtaking. The observations described above either adds program order constraints for the execution trace modeled in the machine states, or adds constraints for variables or expressions that are also modeled in the machine state. In a word, the generated SMT problem captures the execution trace in the machine.

Notice the fact above, we suppose the status of the final machine state is **failure**. From the reduction rules, we know that there exists at least one assert action that is evaluated false. In Figure 10, we know that expression “ $e$ ” is negated and added to the SMT encoding for the second reduction rule. Since  $e$  is evaluated false, ( $\text{not } e$ ) is evaluated true trivially in the generated SMT problem. Other statements are all evaluated true due to the trace capture fact that the program order is assigned correctly based on the execution trace modeled in the machine states, and the variables and expressions are also evaluated true in assume and assign commands. Thus, the SMT problem is evaluated **SAT**. On the other hand, suppose the SMT problem is evaluated **SAT**, indicating that all statements of the SMT problem are satisfiable. Since the reduction rules in Figure 10 add statements with respect to the transition flow of the trace, it implies that each transition of the trace is executed correctly except for some assert action. Because of the same reason above, we know that expression  $e$  for the assert action is evaluated false. Thus, the status of the final machine state is **failure**.  $\square$

As we discussed early, Figure 14 (a) and (b) are generated from the same program with two different traces. The encoding in Figure 14 (c) combines part (a) and (b) into one, the answer of which implies the non-deterministic behavior of the program. By solving the problem in Figure 14 (c), we implicitly solve two problems in Figure 14 (a) and (b) respectively. The following definition defines the “combination” behavior of two SMT problems.

**Definition 17.** A combination operator *COMB* for two SMT problems, represented as three-tuples  $smt_1$  and  $smt_2$ , where  $smt_1 = (\text{defs constraints } ML_1)$  and  $smt_2 = (\text{defs constraints } ML_2)$ , returns a new SMT tuple, such that,

$$COMB(smt_1, smt_2) = (\text{defs constraints } ML_{new})$$

where  $ML_{new}$  is a new relation such that  $\forall \mu \in \text{dom}(ML_1) \cup \text{dom}(ML_2)$ ,  $ML_{new}(\mu) = ML_1(\mu) \cdot ML_2(\mu)$ . Note that two SMT problems are assumed to hold the same defs and constraints excluding the match constraints.

To further finding the correlation between the combined SMT problem and the single SMT problems, we get the following lemma and proof.

**Lemma 4.** For a set of traces  $T_n$  for the same CTP, and a set of SMT problems  $SMT_n = \{smt_0, smt_1, \dots, smt_n\}$ , where each member in  $SMT_n$  encodes a trace in  $T_n$  according to our trace machine, there exists a new SMT problem  $smt_{total}$ , where

$$smt_{total} = COMB(smt_n, COMB(smt_{n-1}, COMB(smt_{n-2}, COMB(\dots))))$$

and

$$ANS(smt_{total}) = \begin{cases} \mathbf{SAT} & \text{iff } \exists smt_i \in SMT_n, \text{ s.t. } ANS(smt_i) = \mathbf{SAT} \\ \mathbf{UNSAT} & \text{otherwise} \end{cases}$$

*Proof.* We Prove it by induction.

We consider the base case as a SMT problem that encodes a single trace and the answer can be trivially proved by definition.

Induction. Assume we have combined  $n$  SMT problems and the combined SMT problem  $smt_{total}$  is evaluated **UNSAT**. We combine an additional SMT problem  $smt_{n+1}$ , which is different from any existent SMT problems. If  $ANS(smt_{n+1})$  is equal to **SAT**, the newly combined SMT problem,  $smt'_{total}$ , is evaluated **SAT** because the match pairs defined in  $smt_{n+1}$  are combined into  $smt'_{total}$ . If  $ANS(smt_{n+1})$  is equal to **UNSAT**, on the other hand,  $smt'_{total}$  is then evaluated **UNSAT** because all the traces implied in  $smt'_{total}$  are evaluated **UNSAT**. Additionally, it can be trivially proved if  $smt_{total}$  is evaluated **SAT** then the newly added SMT problem  $smt_{n+1}$  do not change the answer.  $\square$

From **Lemma 2**, the combined SMT problem is more powerful because it can find the violation of an assertion among several trace encodings. The following theorem states the relation between the ability of finding violation and the content of the match-list  $ML$ .

**Theorem 4.** For two SMT problems,  $smt_\phi = (defs \text{ constraints } ML_\phi)$  and  $smt = (defs \text{ constraints } ML)$ , if  $range(ML_\phi) \subseteq range(ML)$ ,

$$ANS(smt_\phi) \geq ANS(smt)$$

*Proof.* Since the range of  $ML_\phi$  is the subset of the range of  $ML$ , we can obtain  $smt$  by combining  $smt_\phi$  with other SMT problems. Suppose  $ANS(smt_\phi)$  is equal to **SAT**, by **Lemma 2**,  $ANS(smt)$  is equal to **SAT** as well. If  $ANS(smt_\phi)$  is equal to **UNSAT**,  $ANS(smt)$  is equal to either **UNSAT** or **SAT**, depending on the answers of other single SMT problems that combine  $smt$ . In either case,  $ANS(smt_\phi) \geq ANS(smt)$ .

These theorems above obscure an important problem: how do we know which match pair set to use? Soundness assumes we have one, while completeness merely asserts that one exists. Although Section 7 discusses our generation algorithm, we prove here an additional theorem that asserts that any conservative over-approximation of match pair sets is safe.



**Theorem 5 (Approximation).** *Give two match pair sets  $m$  and  $m'$ ,  $m \subseteq m' \Rightarrow SOL(SMT(p, m)) \sqsubseteq SOL(SMT(p, m'))$ , where  $UNSAT \sqsubseteq \sigma$ .*

Informally, this is true because larger match pair sets only allow *more* behavior, which means that the SMT solver has more freedom to find violations, but that all prior violations are still present. However, because of soundness, it is not possible that using a larger match pair set will discover false violations.

## 6 Experiments and Results

To assess the new encoding in this paper, three experiments with results are presented: a comparison to prior SMT encodings on a zero-buffer semantics, a scalability study on the effects of non-determinism in the execution time on infinite buffer semantics, and an evaluation on typical benchmark programs again with infinite buffer semantics. All of the experiments use the Z3 SMT solver ([12]) and are measured on a 2.40 GHz Intel Quad Core processor with 8 GB memory running Windows 7.

The initial program trace for the experiments is generated using the MCA provided reference solution with fixed input. In other words, the only non-determinism in the programs is that allowed by the MCAPI specification. As such, the experiments only consider one path of control flow through the program. Complete coverage of the program for verification purposes would need to generate input to exercise different control flow paths. Where appropriate, the time to generate the match pair sets from the input trace is reported separately.

### 6.1 Comparison to Prior SMT Encoding

To our best knowledge, the current most effective SMT encoding for verification of message passing program traces is the order-based encoding that describes the happens-before relation directly in the encoding and is only functional for zero-buffer semantics in its current form [6]. Although the tool to generate the encoding is not publicly available, the authors of the order-based encoding graciously encoded several contrived benchmarks used for correctness testing. These benchmarks are best understood as *toy* examples that plumb the MCAPI semantics to clarify intuition on expected behavior.

The zero-buffer encoding in this paper is compared directly to the order-based encoding on the contrived benchmarks. The order-based encoding yields incorrect answers for several programs. Where the order-based encoding returns correct answers, the new encoding, on average, requires 70% fewer clauses, uses half the memory as reported by the SMT solver, and runs eight times faster. The dramatic improvement of the new encoding over the order-based encoding is a direct result of the match pairs that simplify the happens-before constraints and avoids redundant constraints in the transitive closure of the happens-before relation.

## 6.2 Scalability Study

The intent of the scalability study is to understand how performance is affected by the number of messages in the program trace and the level of non-determinism in choosing match pairs where multiple sends are able to match to multiple receives. The programs for this study consist of a simple pattern of a single thread to receive messages and  $N$  threads to send messages. The single thread sequentially receives  $N$  messages containing integer values and then asserts that every message did not receive a specific value. In other words, a violation is one where each message has a specific value. The remaining  $N$  threads send a message, each containing a different unique integer value, to the single thread that receives. These programs represent the worst-case scenario for non-determinism in a message passing program as any send is able to match with any receive in the runtime, and the assertion is only violated when each send is paired with a specific receive. The SMT solver must search through the multitude of match pairs,  $N \times N$ , to find the single precise subset of match pairs that triggers the violation. In this program structure, there are  $N!$  feasible ways to match  $N$  sends to  $N$  receives.

**Table 1.** Scaling as a function of non-determinism

Test Programs		Performance	
$N$	Feasible Sets	Time (hh:mm:ss)	Memory(MB)
30	30!( $\sim 3E32$ )	00:00:36	20.11
40	40!( $\sim 8E47$ )	00:03:22	47.12
50	50!( $\sim 3E64$ )	00:16:11	102.65
60	60!( $\sim 8E81$ )	00:47:29	189.53
70	70!( $\sim 1E100$ )	02:00:30	364.25

The study takes an initial program of  $N = 30$ , so 31 threads, and varies  $N$  to see how the SMT solver scales. A small  $N$  is an easy program while a large  $N$  is a hard program. Table 1 shows how the new encoding scales with hardness. The first column is the number of messages, or  $N$ , and the second column is the number of feasible match pair subsets that correctly match every receive to a unique send. As expected, running time and memory consumption increase non-linearly with hardness.

The case where  $N = 70$  represents having 70 concurrent messages in flight from 70 different threads of execution. Such a scenario is not entirely uncommon in a high performance computing application, and it appears the new encoding is able to reasonably scale to such a level of concurrency. The result provides a bound on expected cost for analysis given the message passing behavior in a program. It is expected that the analysis of any program with fewer than 70! possible choices of feasible match pair resolutions will complete in a reasonable amount of time. Regardless, such a high-level concurrency seems unlikely in the embedded space to which MCAPI is targeted.

### 6.3 Typical Benchmark Programs

The results in the prior section suggest that the number of messages is not the deciding factor in hardness for the new encoding; rather, hardness is measured by the number of feasible match pair sets. This section further explores the observation to show that some programs are easy, even if there are many messages, while other programs are hard, even though there are only a few messages.

The goal of these experiments is to measure the new encoding on several benchmark programs. MCAPI is a new interface, and to date, the authors are not aware of publicly available programs written against the interface aside from the few toy programs that come with the library distribution. As such, the benchmarks in the experiments come from a variety of sources.

- *LE* is the leader election problem and is common to benchmarking verification algorithms.
- *Router* is an algorithm to update routing tables. Each router node is in a ring and communicates only with immediate neighbors to update the tables. The program ends when all the routing tables are updated.
- *MultiM* is an extension to an program in the MCAPI library distribution and is similar to the program in Figure 16. The extension adds extra iterations to the original program execution to generate longer execution trace.
- *Pktuse* is a benchmark from the MPI test suite [13]. The program creates 5 tasks—each of which randomly sends several messages to the other tasks.

The benchmark programs are intended to cover a spectrum of program properties. As before, the primary measure of hardness in the programs is not the number of messages but rather the size of the match pair set and the number of feasible subsets. The *LE* program is the easiest program in the suite. Although it sends 620 messages, there is only a single feasible match pair set. The programs *Router*, *MultiM*, and *Pktuse* respectively increase in hardness, which again is not related to the total number of messages but rather the total number of feasible match-sets that must be considered. For example, even though *Router* has 200 messages, it is an easier problem than *MultiM* that has 100 messages. The *Pktuse* program does have the most number of messages, 512, and in this case, the largest number of feasible match pair sets.

**Table 2.** Performance on selected benchmarks

Test Programs			Performance			
Name	# Mesg	Feasible Sets	EG(s)	MG(s)	Time (hh:mm:ss)	Memory(MB)
<i>LE</i>	620	1	1.49	0.051	<00:00:01	33.41
<i>Router</i>	200	~6E2	0.417	0.032	00:00:02	15.03
<i>MultiM</i>	100	~1E40	0.632	0.436	00:16:40	135.19
<i>Pktuse</i>	512	~1E81	10.190	9.088	02:06:09	1539.90

Table 2 shows the results for the benchmark suite. Other than the metrics used in Table 1, the time of generating the encoding and the match pairs is included in the third and fourth columns respectively. Note that the time shown

in the third column includes the time in the fourth column. As before, the running time tracks hardness and not the total number of messages. The table also shows the cost of match pair generation as it dominates the encoding time for the *Pktuse* program. Future work is to address the high-cost of match pair generation, which the authors believe to be NP-complete [15].

The benchmark suite demonstrates that a message passing program may have a large degree of non-determinism in the runtime that is prohibitive to verification approaches that directly enumerate non-determinism such as a model checker. The SMT encoding, however, pushes the problem to the SMT solver by generating the possible match pairs and then relying on advances in SMT technology to resolve the non-determinism in a way that violates the assertion. Of course, the SMT problem itself is NP-complete, so performance is only reasonable for small problem instances. The benchmark suite suggests that problem instances with astonishingly large numbers of feasible match pair sets are able to complete in a reasonable amount of time using the new encoding in this paper; though, the time to generate the match pairs may quickly become prohibitive.

## 7 Generating Match Pairs

The exact set of match pairs can be generated by simulating the program trace and using a depth-first search to enumerate non-determinism arising from concurrent sends and receives. Such an effort, however, solves the entire problem at once because if you simulate the program trace exploring all non-determinism, then you may as well verify all runtime choices for property violations at the same time.

In this section, we present an algorithm that does not require an exhaustive enumeration of runtime behavior in simulation. Our algorithm over-approximates the match pairs such that match pairs that can exist in the runtime are all included and some bogus match pairs that cannot exist in the runtime may or may not be included.

The algorithm generates the over-approximated match pair set by matching each pair of the send and receive commands at common endpoints and then pruning obvious matches that cannot exist in any runtime implementation of the specification.

Figure 15 presents the major steps of the algorithm. The algorithm proceeds by first linearly traversing each task of the program storing each receive and send command into two distinct structured lists. The receive list, `list_r`, is structured as in (1) and the send list, `list_s`, is structured as in (2).

$$\begin{aligned}
 (\mathbf{e}_0 &\rightarrow ((0, \mathbf{R}_{0,1}), (1, \mathbf{R}_{0,2}), \dots)) \\
 (\mathbf{e}_1 &\rightarrow ((0, \mathbf{R}_{1,1}), (1, \mathbf{R}_{1,2}), \dots)) \\
 \dots & \\
 (\mathbf{e}_n &\rightarrow ((0, \mathbf{R}_{n,1}), (1, \mathbf{R}_{n,2}), \dots))
 \end{aligned} \tag{1}$$

The list `list_r` groups receives by the issuing endpoint. The integer field merely records the order in which the receives are issued and increases by one on each

```

// initialization
input an MCAPI program
initialize list_r
initialize list_s

// check each receive and send with the same endpoint
for r in list_r
  for s in list_s
    let dest = destination endpoint(s)
    let src = source endpoint(s)
    // check matching criteria for r and s
    if
      1. endpoint(r) = dest
      2. index(r) >= index(s)
      3. index(r) =< (index(s)
                    + count(sends(dest=dest))
                    - count(sends(src=src, dest=dest)))
    then
      add pair (r, s) to match_set
    else
      continue
    end if
  end for
end for

output match_set;

```

**Fig. 15.** Pseudocode for generating over-approximated match pairs

receive. Similarly, the list `list_s` groups sends first by the destination endpoint and then by the source endpoint. Like `list_r`, an index increases by one to track the issue order. As the input is a program execution trace, any sends or receives in loops already have unique identifiers.

$$\begin{array}{l}
\begin{array}{ccc}
\text{“dest”} & \text{“src”} & \text{“src”} \\
(e_0 \rightarrow ((e_1 \rightarrow ((0, S_{1,1}), (1, S_{1,2}), \dots), (e_2 \rightarrow (\dots), \\
\dots)))) & & \\
(e_1 \rightarrow ((e_0 \rightarrow ((0, S_{0,1}), (1, S_{0,2}), \dots), (e_2 \rightarrow (\dots), \\
\dots)))) & & (2) \\
\dots & & \\
(e_n \rightarrow ((e_0 \rightarrow ((0, S_{0,3}), (1, S_{0,4}), \dots), (e_1 \rightarrow (\dots), \\
\dots)))) & &
\end{array}
\end{array}$$

Consider the program in Figure 16. The lists `list_r` and `list_s` for the program are

$$\begin{array}{l}
(0 \rightarrow ((0, R_{0,1}), (1, R_{0,2}), (2, R_{0,4}))) \\
(1 \rightarrow ((0, R_{1,2})))
\end{array} \tag{3}$$

$$\begin{array}{l}
(0 \rightarrow ((1 \rightarrow ((0, S_{1,1}), (1, S_{1,3}), (2 \rightarrow ((0, S_{2,1})))))) \\
(1 \rightarrow ((0 \rightarrow ((0, S_{0,3}))))))
\end{array} \tag{4}$$

The sends  $S_{1,1}$ ,  $S_{1,3}$ , and  $S_{2,1}$  have task 0 as an identical destination endpoint. The send  $S_{0,3}$  has task 1 as the destination endpoint. The list `list_s` in (4) reflects this partition. Receive  $R_{0,1}$  is the first receive operation in endpoint 0. This fact is again reflected in `list_r` in (3).

Task 0	Task 1	Task 2
$R_{0,1}(*, \&h1)$	$S_{1,1}(0, \&h5)$	$S_{2,1}(0, \&h8)$
$W(\&h1)$	$W(\&h5)$	$W(\&h8)$
$R_{0,2}(*, \&h2)$	$R_{1,2}(*, \&h6)$	
$W(\&h2)$	$W(\&h6)$	
$S_{0,3}(1, \&h3)$	$S_{1,3}(0, \&h7)$	
$W(\&h3)$	$W(\&h7)$	
$R_{0,4}(*, \&h4)$		
$W(\&h4)$		

**Fig. 16.** Another MCAPI concurrent program

The algorithm traverses the two lists in a nested loop to generate match pairs between send and receive commands. The function `index(r)` takes the endpoint of the receive and returns the issue order of that receive in the `list_r` structure. Similarly, the function `index(s)` takes the destination and source endpoints in the send and returns the issue order of that send in the `list_s` structure. These indexes essentially capture message non-overtaking.

The criteria to generate a match pair first requires the send and receive to be compatible (check 1), consistent with message non-overtaking (check 2), and that message non-overtaking does not preclude the match (check 3). A match is precluded by message non-overtaking when a receive cannot possibly match a send because by the time the program issues the receive, the send must have already been matched somewhere else. The function `count` gives the number of sends to a specific destination or the number of sends to a specific source and destination. As long as a receive is issued early enough to still match the send given the message non-overtaking rule, then the match is possible.

In our concrete example,  $R_{0,1}$  is matched with  $S_{1,1}$  or  $S_{2,1}$ , but it cannot be matched with  $S_{1,3}$  since the second rule is not satisfied such that the order of  $R_{0,1}$  is less than the order of  $S_{1,3}$  (i.e.,  $S_{1,3}$  would have to overtake  $S_{1,1}$  to satisfy the rule). The match between  $R_{0,4}$  and  $S_{1,1}$  is also precluded by check 3 as  $S_{1,1}$  must have already matched an earlier receive by message non-overtaking.

The generated set of match pairs for our example in Figure 16 is over-approximated by the algorithm because it includes pairs that cannot exist in any feasible execution. For example, the match pair  $(S_{2,1} R_{0,4})$  is not feasible because it is not possible to order  $S_{1,3}$  before  $R_{0,2}$  since  $R_{1,2}$  can only match with  $S_{0,3}$  that must occur after  $R_{0,2}$ . Fortunately, a satisfying solution is only possible using feasible match pairs. Non-feasible match pairs merely result in extra clauses in the encoding and potentially slow down the SMT solver.

The complexity of the algorithm is quadratic. Traversing the tasks to initialize the lists is  $O(N)$ , where  $N$  is the total lines of code of the program. Traversing the list of receives and the list of sends takes  $O(mn)$  to complete, where  $m$  is the total number of sends and  $n$  is the total number of receives. As  $m + n \leq N$ , the algorithm takes  $O(N + mn) \leq O(N + N^2) \approx O(N^2)$  to complete.

## 8 Related Work

Morse *et al.* provided a formal modeling paradigm that is callable from the C language for the MCAPI interface [11]. This model correctly captures the behavior of the interface and can be applied to model checking C programs that use the API. The work is a direct application of model checking and directly enumerates the non-determinism in the runtime to construct an exhaustive proof. The SMT encoding in this paper pushes that complexity to the SMT solver and leverages recent advances in SMT technology to find a satisfying assignment.

Sharma *et al.* present a dynamic model checker for MCAPI programs built on top of the MCA provided MCAPI runtime [16]. MCC systematically enumerates all non-determinism in the MCAPI runtime under zero-buffer semantics. It employs a novel dynamic partial order reduction to avoid enumerating redundant message orders. This work claims SMT technology is more efficient in practice in resolving non-determinism in a way to violate correctness properties.

Wang *et al.* present an SMT encoding for shared memory semantics for a given input trace from a multi-threaded program [19]. As mentioned previously, the program is partitioned into several concurrent trace programs, and the encoding for each program is verified using SMT technology. Elwakil *et al.* extend the encoding to message passing programs using the MCAPI semantics [5, 6]. The comparison to the encoding in this work is already discussed previously.

An important body of work is being pursued for MPI program verification [17, ?, ?, ?, 18, ?]. Highlights include an extension to the SPIN model checker for MPI programs, symbolic execution tools for MPI programs including new approaches to computing loop invariants, and various dynamic verification tools for MPI programs. Although MPI is more expressive than MCAPI, the correctness properties in MCAPI are similar to those in MPI. More importantly, the encoding in this work should be applicable to MPI programs that do not include collective operations. An important aspect of future work is to extend the encoding to collectives.

There is a rich body of literature for SMT/SAT based Bounded Model Checking. Burckhardt *et al.* exhaustively check all executions of a test program by translating the program implementation into SAT formulas [1]. The approach relies on counter-examples from the solvers to refine the encoding. The SMT encoding in this work is able to directly resolve the match-pair set over-approximation directly without needing to check a counter-example.

Dubrovin *et al.* give a method to translate an asynchronous system into a transition formula over three partial order semantics [3]. The encoding adds constraints to compress the search space and decrease the bound on the program unwinding. The encoding in this paper operates on a program execution and does not need to resolve a bound.

Kahlon *et al.* presented a partial order reduction, *MPOR*, that operates in the bounded model checking space [7]. It guarantees that exactly one execution is calculated per each Mazurkiewicz trace to reduce the search space. It would be interesting to see if *MPOR* is able to extend to message passing semantics.

Other work in bounded model checking explores heap-manipulating programs and challenges in sequential systems code [8, ?].

The application of static analysis is another interesting thread of research to test or debug message passing programs with some work in the MPI domain [20, ?,?]. The work is important as it lays the foundation for refining match-pair sets to only include those that cannot be statically pruned.

## 9 Conclusions and Future Work

The paper presents a proof that the problem of resolving non-determinism in message passing in a way that meets asserts is NP-complete. The paper then presents an SMT encoding of an MCAPI program execution that uses match pairs directly rather than the state-based or order-based encoding in the prior work. The encoding is generated from a given execution trace and a set of potential match pairs that can be over-approximated. The encoding takes extra care in the forming the SMT problem to preclude bogus match pairs in any over-approximation of the match pair input set. Critically, the encoding is the first to correctly capture the non-deterministic behaviors of an MCAPI program execution under infinite-buffer semantics.

The paper further defines an algorithm with  $O(N^2)$  time complexity to over-approximate the true set of match pairs, where  $N$  is the total number of code lines of the program. A comparison to prior work, [6], for a set of “toy” examples under zero-buffer semantics shows the new encoding capable and efficient in capturing correct behaviors of an MCAPI program execution. Experiments further show that the encoding scales to programs with significant levels of non-determinism in how sends are match to receives.

The results show that a large match-pair set does affect the runtime performance of the encoding in the SMT problem even if the encoding is sound under an over-approximation. Future work explores new methods for generating a much more precise set of match pairs. The encoding is dependent on an input execution trace of the program. Future work explores integrating the encoding into a model checker. The model checker generates a program trace that is encoded and verified. The result is then used to inform the model checker as to where it needs to backtrack to generate a new execution trace. The goal is to use the trace verification to construct a better partial order reduction in the model checker.

Finally, given the importance of high performance computing, future work looks to extend the encoding to account for MPI collective operations. This direction is motivated by the results where the encoding seems to scale to significant levels of concurrency. It should be possible to express MPI collectives as additional constraints in the encoding and apply the technique to MPI programs directly.



## References

1. Burckhardt, S., Alur, R., Martin, M.M.K.: Checkfence: Checking consistency of concurrent data types on relaxed memory models. In: ACM SIGPLAN PLDI. San Diego, California, USA (June 10–13, 2007)
2. Cantin, J.F., Lipasti, M.H., Smith, J.E.: The complexity of verifying memory coherence and consistency. *IEEE Trans. Parallel Distrib. Syst.* 16(7), 663–671 (Jul 2005), <http://dx.doi.org/10.1109/TPDS.2005.86>
3. Dubrovin, J., Junttila, T., Heljanko, K.: Exploiting step semantics for efficient bounded model checking of asynchronous systems. In: *Science of Computer Programming*. pp. 77(10–11):1095–1121 (2012)
4. Dutertre, B., de Moura Leonardo: A fast linear-arithmetic solver for DPLL(T). In: *CAV*. vol. 4144 of LNCS, pp. 81–94. Springer-Verlag (2006)
5. Elwakil, M., Yang, Z.: CRI: Symbolic debugger for MCAPI applications. In: *Automated Technology for Verification and Analysis* (2010)
6. Elwakil, M., Yang, Z.: Debugging support tool for mcapl applications. In: *PADTAD '10: Proceedings of the 8th Workshop on Parallel and Distributed Systems* (2010)
7. Kahlon, V., Wang, C., Gupta, A.: Monotonic partial order reduction: An optimal symbolic partial order reduction technique. In: *ACM CAV*. pp. 398–413. Springer, Berlin/Heidelberg, Grenoble, France (June 26–July 02, 2009)
8. Lahiri, S., Qadeer, S.: Back to the future revisiting precise program verification using SMT solvers. In: *POPL*. ACM, San Francisco, California, USA (2008)
9. MCA: The multicore association, <http://www.multicore-association.org>
10. MCA: The multicore association resource management API, <http://www.multicore-association.org/workgroup/mcapi.php>
11. Morse, E., Vrvilo, N., Mercer, E., McCarthy, J.: Modeling asynchronous message passing for C program. In: *Verification, Model Checking, and Abstract Interpretation*. vol. 7148 of LNCS, pp. 332–347. Springer-Verlag (2012)
12. de Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: *TACAS*. vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
13. MPPTest: MPPTest benchmark, <http://www.mcs.anl.gov/research/projects/mpi/mpptest/>
14. Netzer, R., Brennan, T., Damodaran-Kamal, S.: Debugging race conditions in message-passing programs. In: *ACM SIGMETRICS Symposium on Parallel and Distributed Tools*. pp. 31–40. Philadelphia, PA, USA (1996)
15. Sharma, S.: Private conversation on active research.
16. Sharma, S., Gopalakrishnan, G., Mercer, E., Holt, J.: MCC - a runtime verification tool for MCAPI user applications. In: *FMCAD* (2009)
17. Siegel, S.F., Zirkel, T.K.: Loop invariant symbolic execution for parallel program. In: Kuncak, V., Rybalchenko, A. (eds.) *Verification, Model Checking, and Abstract Interpretation: 13th International Conference, VMCAI 2012*. Lecture Notes in Computer Science, vol. 7148, pp. 412–427. Springer (2012)
18. Vakkalanka, S., Vo, A., Gopalakrishnan, G., Kirby, R.: Reduced execution semantics of MPI: From theory to practice. In: *FM*. pp. 724–740 (2009)
19. Wang, C., Chaudhuri, S., Gupta, A., Yang, Y.: Symbolic pruning of concurrent program executions. In: *ESEC/FSE*. pp. 23–32. ACM, New York, NY, USA (2009)
20. Zhang, Y., Evelyn, D.: Barrier matching for programs with textually unaligned barriers. In: *PPoPP*. pp. 194–204. ACM, San Jose, California, USA (2007)