

# Exploiting partial orders and symmetries in efficient analysis of message-passing concurrency

**Subodh Sharma, Sept 6 2022, Presented in IARCS Verification Seminar Series**

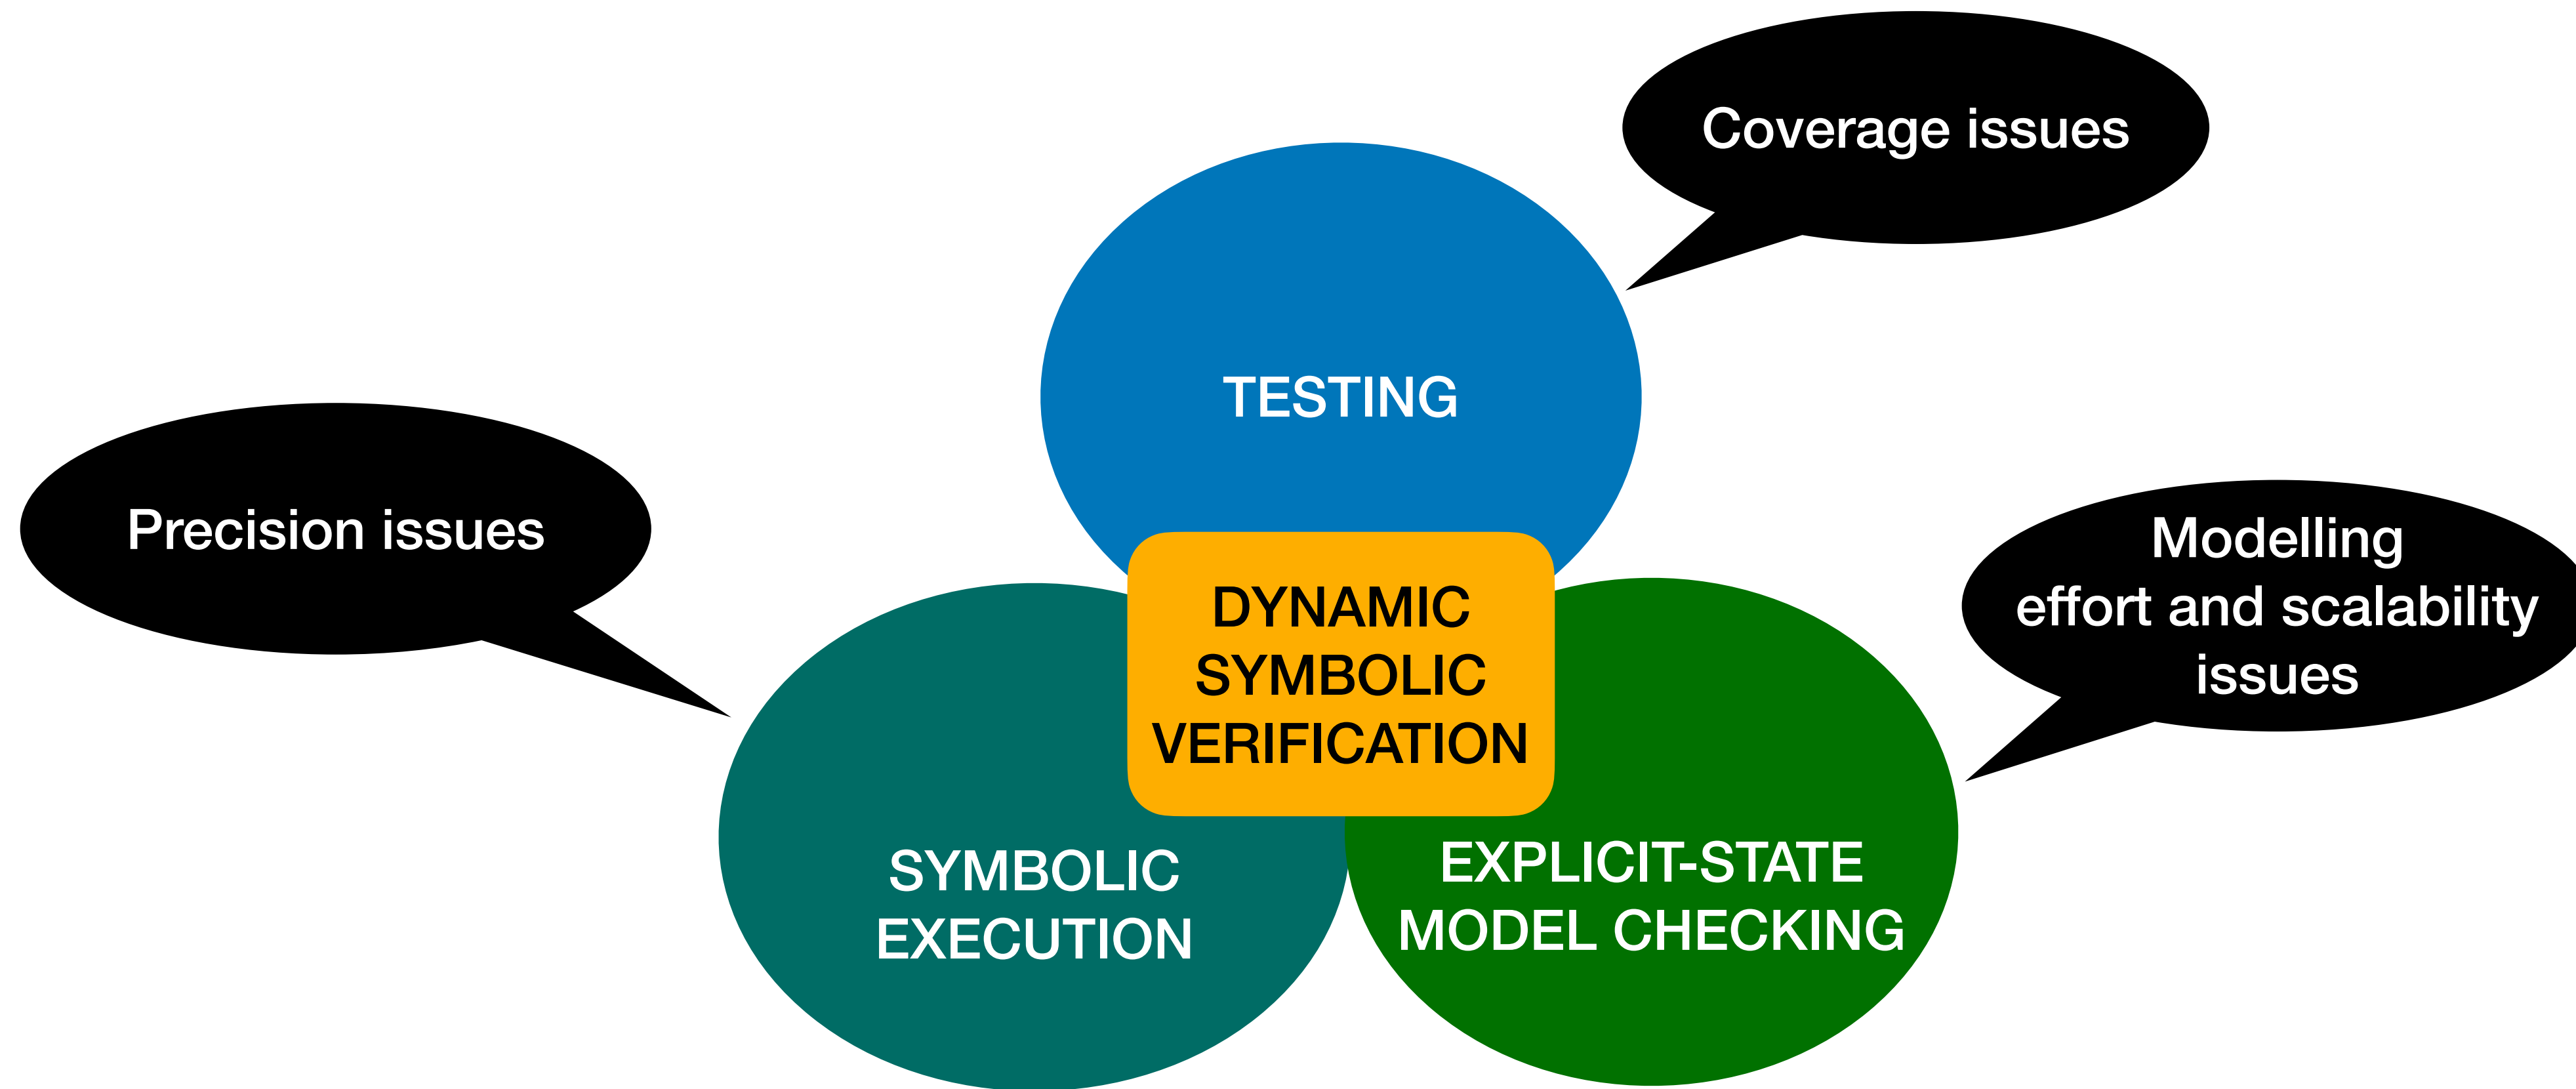
# Message Passing Systems

- Heavy use in **scientific computing**, client-server, peer-to-peer applications, consensus protocols etc.
  - Many languages and APIs — Go, Erlang, **MPI**, etc.
- Large and complex computation code interleaved with communication
- Many issues: communication races, deadlocks, assertion violations etc.
  - Reachability is **undecidable** when each process is a **FSM** with **FIFO buffers** [Brand and Zafiropulo, JACM'83]
- Lack of use of partial orders in practical verification techniques

# Soundness & Completeness

- Given a deductive system  $D$ :
  - $D$  is **sound** for a logic  $\mathcal{L}$ , if for every formula  $\phi \in \mathcal{L}$ ,  $\vdash_D \phi \Rightarrow \models \phi$ 
    - All formulae proven by the deductive system are indeed valid!
  - $D$  is **complete** for a logic  $\mathcal{L}$ , if for every formula  $\phi \in \mathcal{L}$ ,  $\models \phi \Rightarrow \vdash_D \phi$ 
    - All valid formulae are also proven by the deductive system!

# Dynamic Symbolic Verification



- **DSE: sound and complete w.r.t. a fixed input, relatively scalable**

# Syntax and Matches-before Relation

$C ::=$

- $x := e$
- $send(pid, x, h)$
- $recv(pid, x, h)$
- $wait(h)$
- $Barrier$

$pid \in PID$   
 $pid^* \in PID \cup \{\star\}$   
 $x, h \in \mathcal{V}$

- Assume the standard expression grammar for  $e$ .
- $h$  is the request handle.
- $wait$  is a **blocking** call (behaves differently under buffering)
- **Matches-before** ordering relation:
  - $send(\dots, h) \rightarrow wait(h); recv(\dots, h) \rightarrow wait(h)$
  - $send(pid, \dots) \rightarrow send(pid, \dots)$  (send to different destinations are **unordered**)
  - $wait(h) \rightarrow C$

# Formal MPI Semantics

- **State**  $s = \langle I, M \rangle$ 
  - $I$  is the set of issued calls
  - $M$  is the set of matched calls  $\subseteq I$
- **Issuable(x)**: Set of all  $x$ :  $x \notin I$  and for all  $y$  with  $y \prec_{po} x$ , we have  $y \in I$
- **Ready(m, q)**: If for all  $a \in m$  and every  $s \prec_{mo} a$ , we have  $s \in M$
- Semantics of  $P$  is a FSM  $S(P) = \langle Q, q_0, A, \delta \rangle$

# Modeling Message Passing Programs

- Each process is a finite state machine (FSM)
  - Deterministic computation
  - **FIFO** system buffers (0 — synchronous,  $k/\infty$  — asynchronous)
  - **Asynchronous non-blocking** (send/recv) calls.
  - Programs have finite control and finite data domain

# Nondeterminism: Reason for verification

```
void P1 () {  
  R1: Recv (*, x);  
  if (x==10)  
  R2: Recv (*, y);  
  else if (x==20)  
  R3: Recv (*, y);  
  else { }  
  R5: Recv (*, z);  
}
```

```
void P2 () {  
  Send (P1, 10);  
}
```

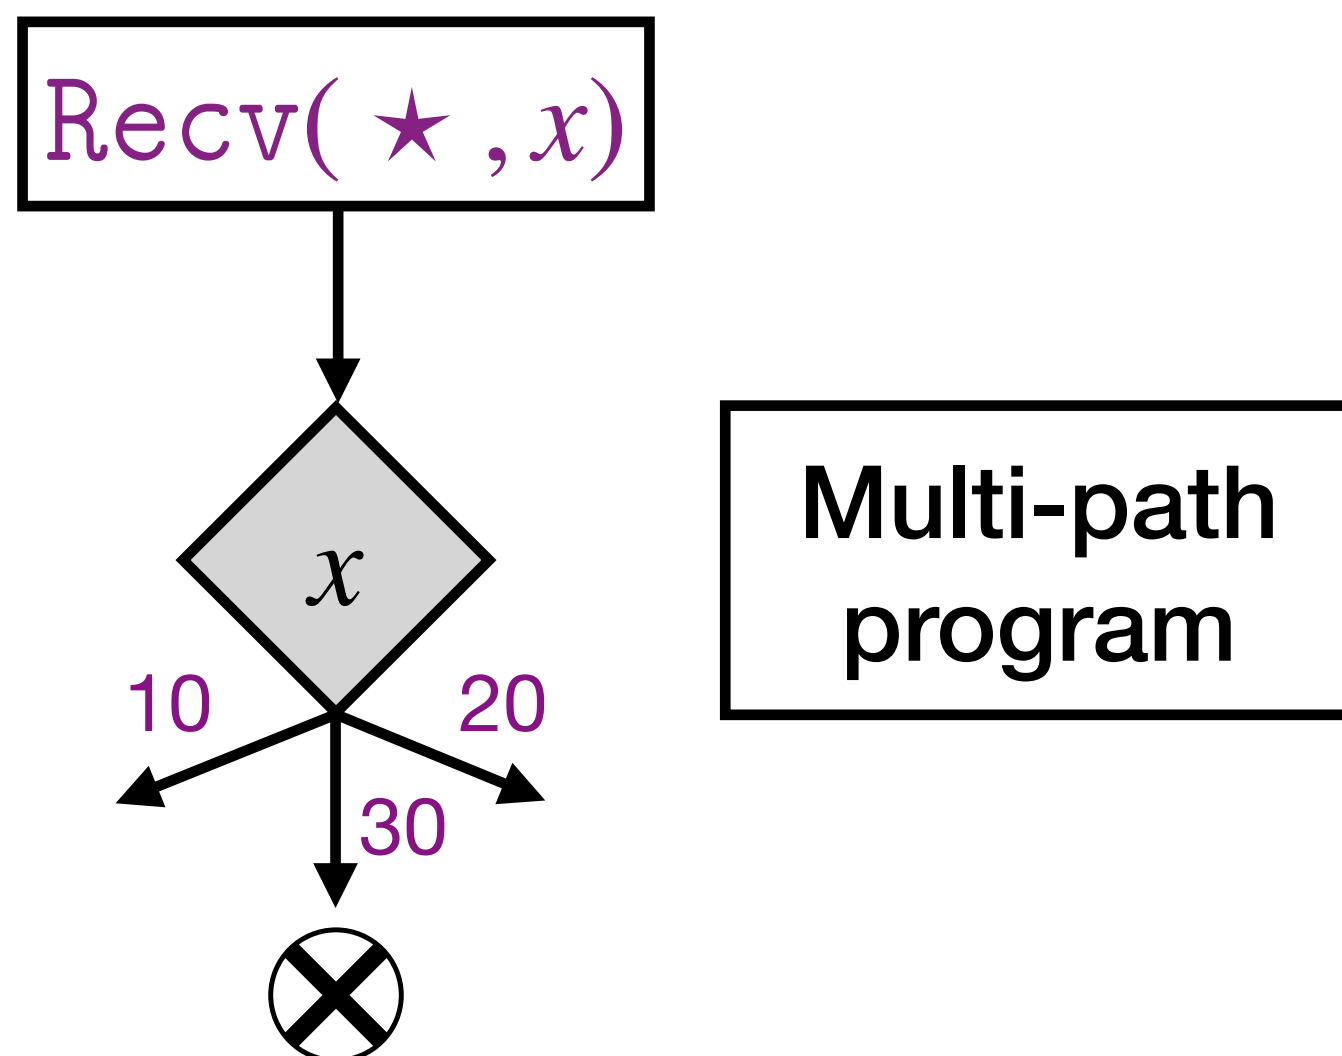
```
void P3 () {  
  Send (P1, 10);  
}
```

```
void P4 () {  
  Send (P1, 30);  
}
```

- Nondeterminism due to scheduling and control.
- Theoretical Results [\[TOPLAS'17, FM'18\]](#)

## Complexity

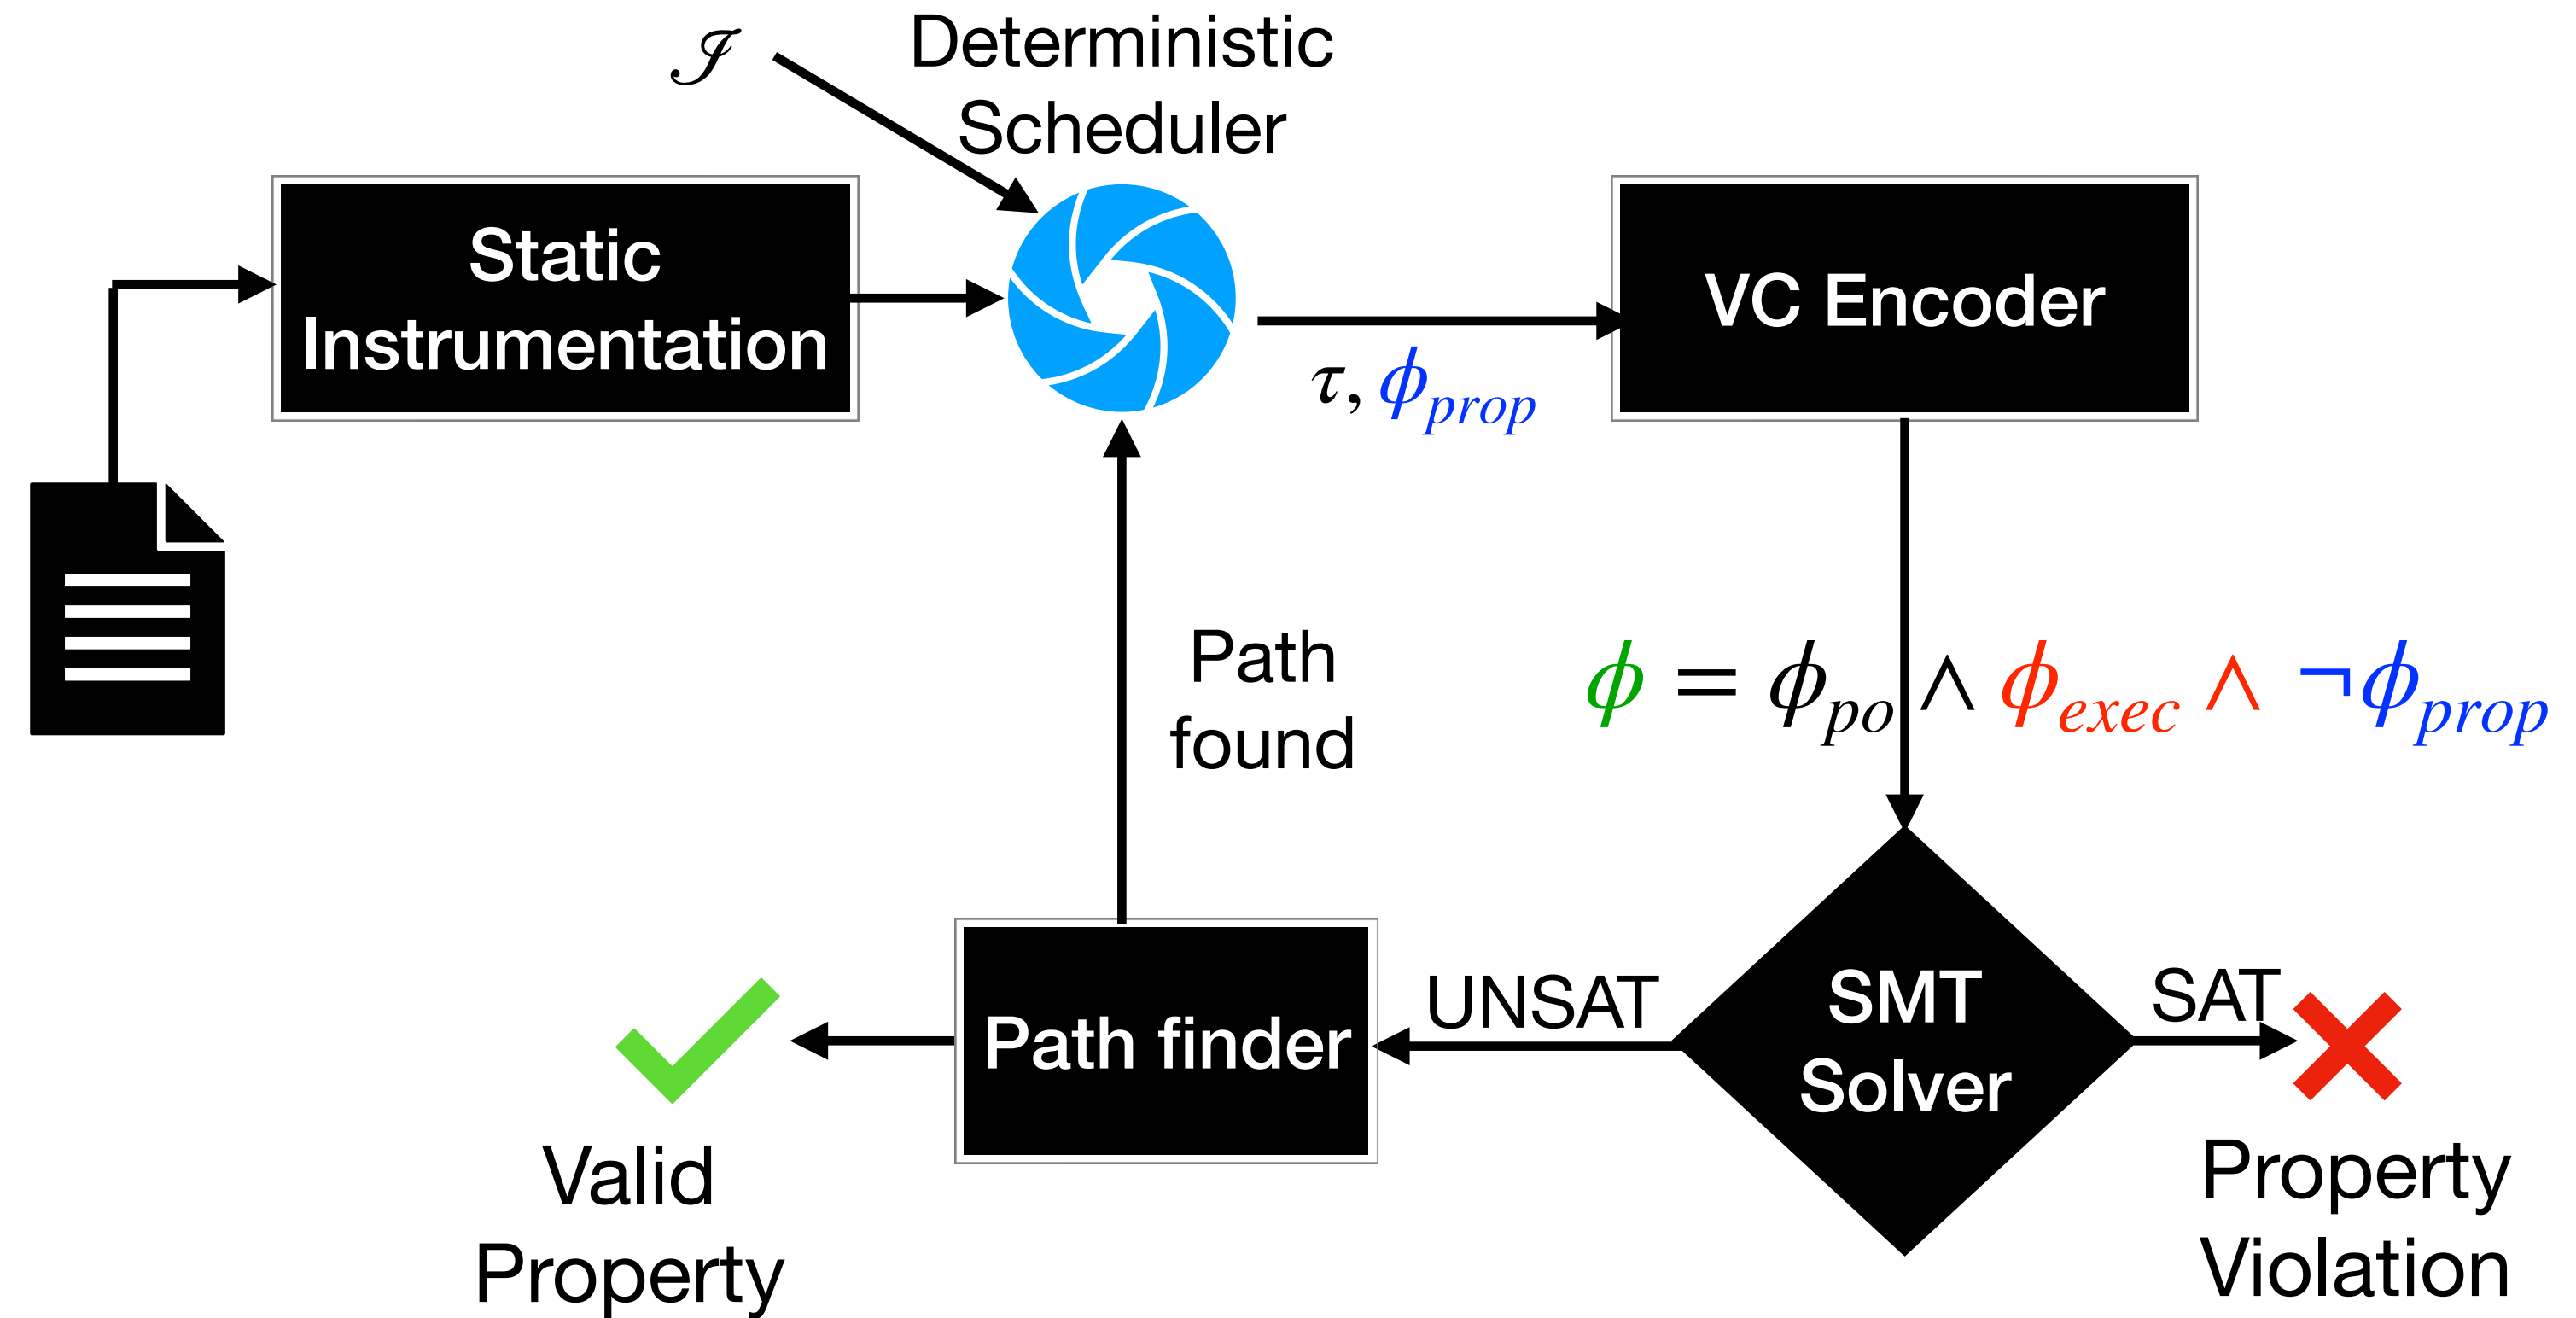
The deadlock discovery problem is NP-complete for both finite and infinite buffering models.





# Our technique using equivalence partitioning

- Propositional Encoding
  - Sound and complete
- Path finder:
  - nonchronological search for a new **equivalence class**
- Termination
  - Due to finite # of control decisions
- Validate results by concrete scheduling



## Soundness & Completeness

There is a deadlocking trace *iff* there is a satisfying assignment to the variables in the encoding.

# Propositional Encoding

## PO Constraints

$$\bigwedge_{b \in C} \bigwedge_{a < b} t_{ab} \quad (1)$$

$$\bigwedge_{a, b \in C} (t_{ab} \rightarrow (clk_a < clk_b)) \quad (2)$$

## Uniqueness & Clock Equality

$$\bigwedge_{\{a, b\} \in \mathbb{M}^+} (s_{ab} \rightarrow \bigwedge_{c \in \mathbb{M}^+(a), c \neq b} \neg s_{ac} \wedge \bigwedge_{c \in \mathbb{M}^+(b), c \neq a} \neg s_{cb}) \quad (3)$$

$$\bigwedge_{(a, b) \in \mathbb{M}^+} (s_{ab} \rightarrow (clk_a = clk_b)) \quad (4)$$

## Match correct & Matched only

$$\bigwedge_{a \in C} (m_a \rightarrow \bigvee_{b \in \mathbb{M}^+(a)} s_{ba} \wedge \bigvee_{b \in \mathbb{M}^+(a)} s_{ab}) \quad (5)$$

$$\bigwedge_{\alpha \in \mathbb{M}^+} (s_\alpha \rightarrow \bigwedge_{a \in \alpha} m_a) \quad (6)$$

## Match only issued

$$\bigwedge_{a \in C} (m_a \rightarrow r_a) \quad (7)$$

$$\bigwedge_{b \in C} (r_b \leftrightarrow \bigwedge_{a < b} m_a) \quad (8)$$

## Not all matched

$$\bigwedge_{\alpha \in \mathbb{M}^+} (\bigvee_{a \in \alpha} (m_a \vee \neg r_a)) \quad (9)$$

$$\bigvee_{a \in C} \neg m_a \quad (10)$$

# Savings through an example

```

void P1 () {
  R1: Recv (*, x);
  if (x==10)
  R2: Recv (*, y);
  else if (x==20)
  R3: Recv (*, y);
  else { }
  R5: Recv (*, z);
}
  
```

```

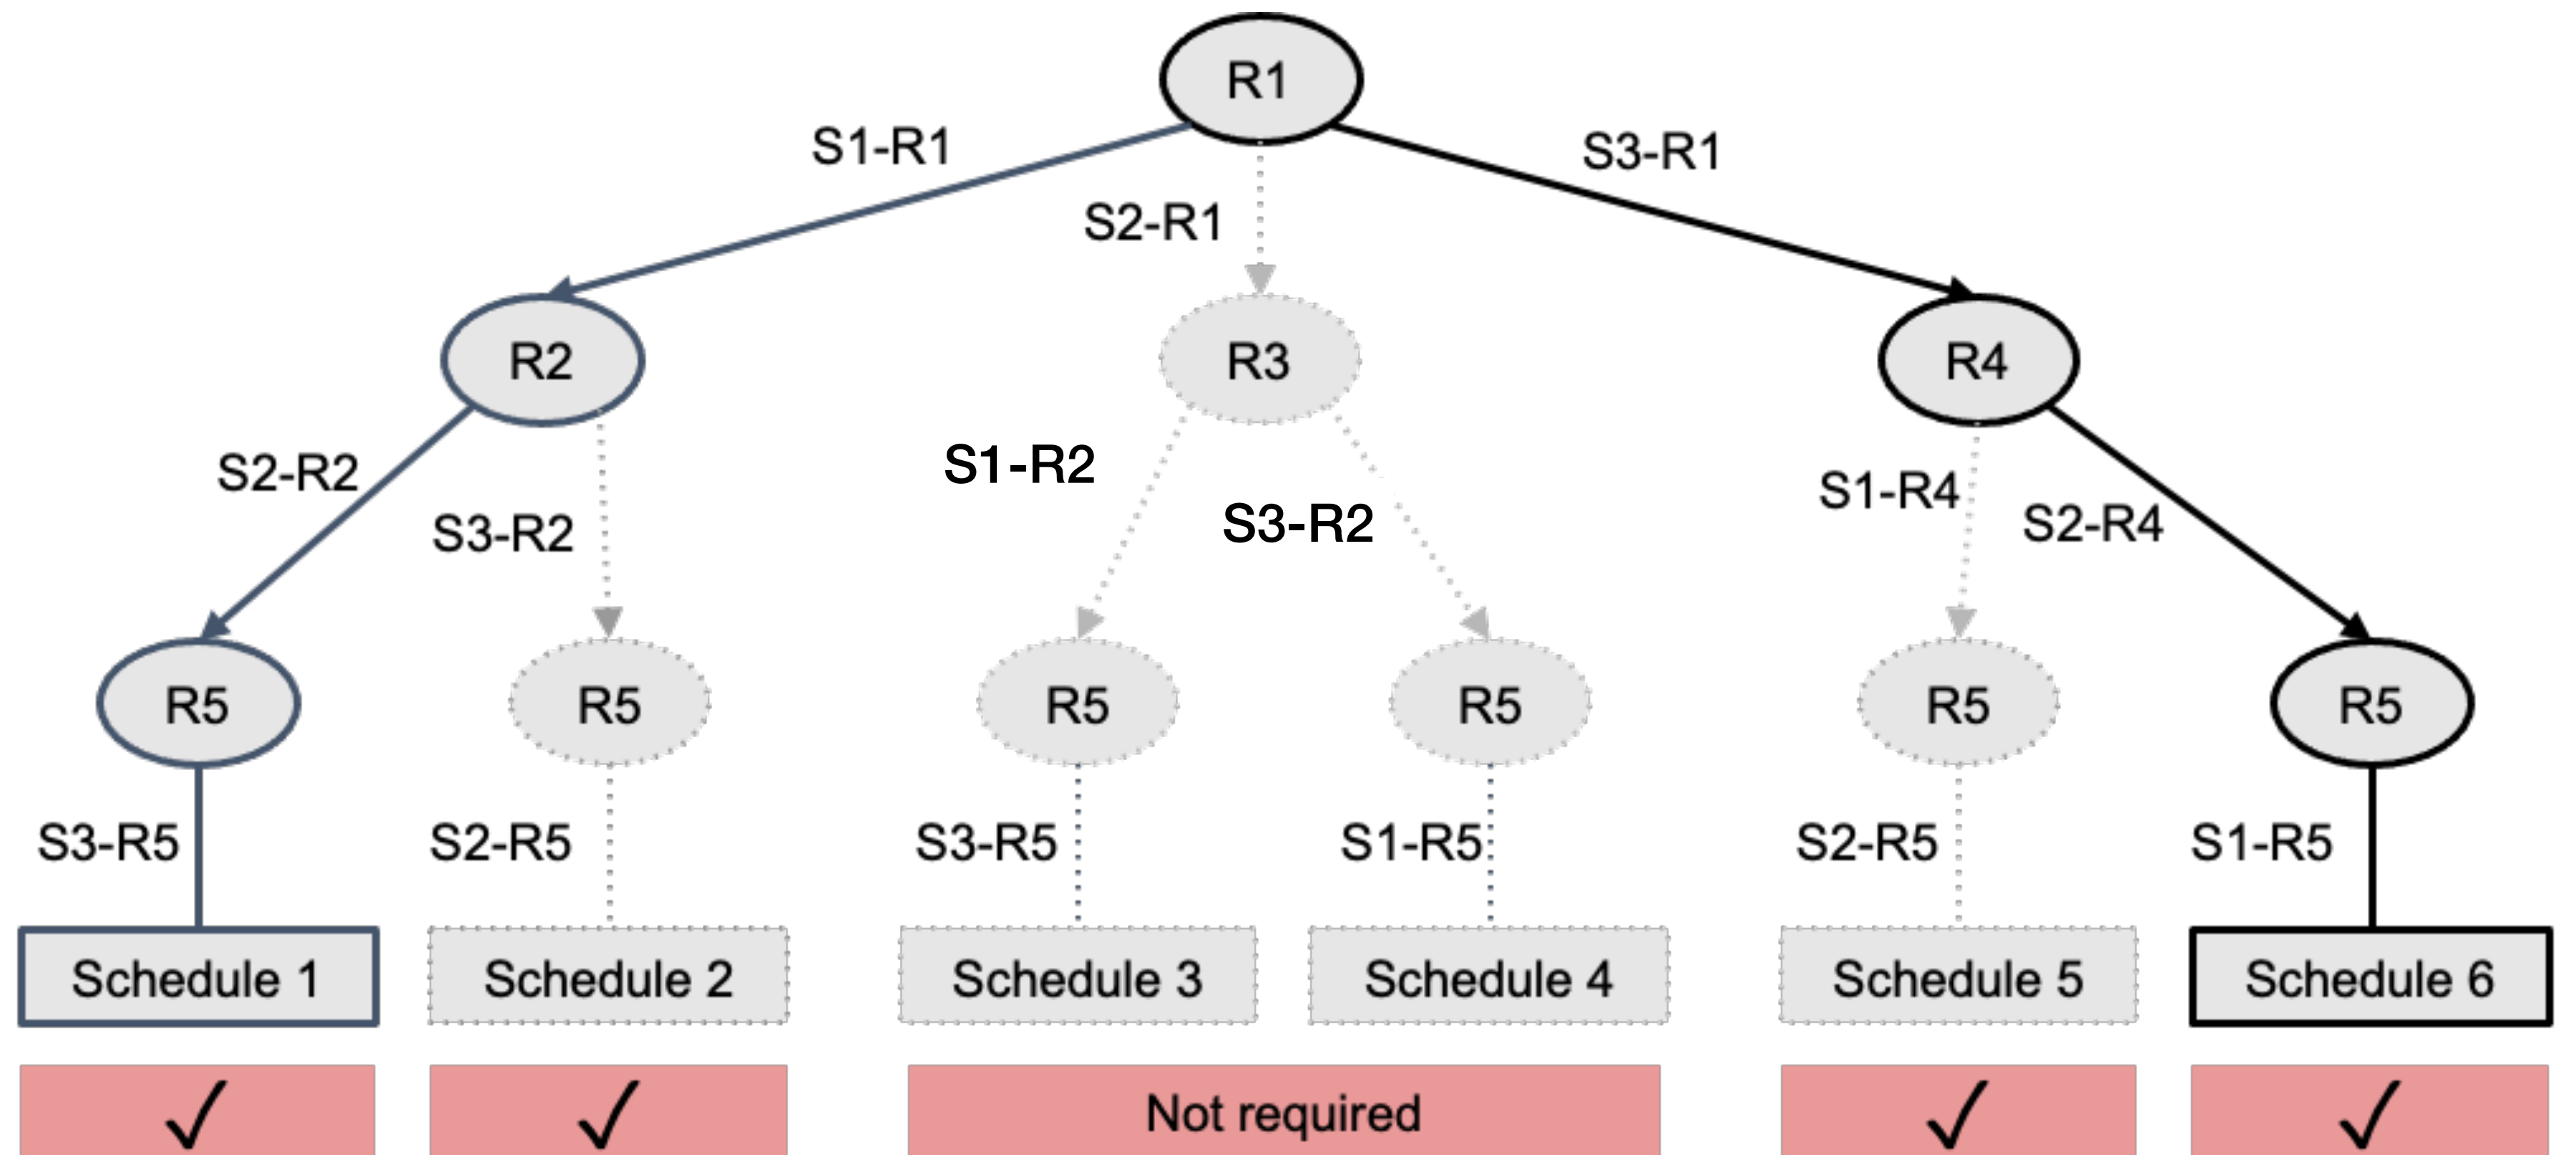
void P2 () {
  Send (P1, 10);
}
  
```

```

void P3 () {
  Send (P1, 10);
}
  
```

```

void P4 () {
  Send (P1, 30);
}
  
```



In general: exponential savings over runtime verification techniques

# Results

B'mark	# Proc	ISP		Aislinn Time (s)	CIVL Time (s)	HERMES	
		Runs	Time (s)			Runs	Time (s)
Matrix Multiply	6	720	837.64	1.33	48.07	1	12.93
	8	~1.5K	TO	2.21	258.86	1	19.67
Monte Carlo	6	120	151.60	5.03	*	5	4.82
	8	>1.2K	TO	10.17	*	7	7.43
Integrate	6	>3k	TO	27.84	157.63	1	0.49
	8	>3k	TO	TO	173.27	1	0.85

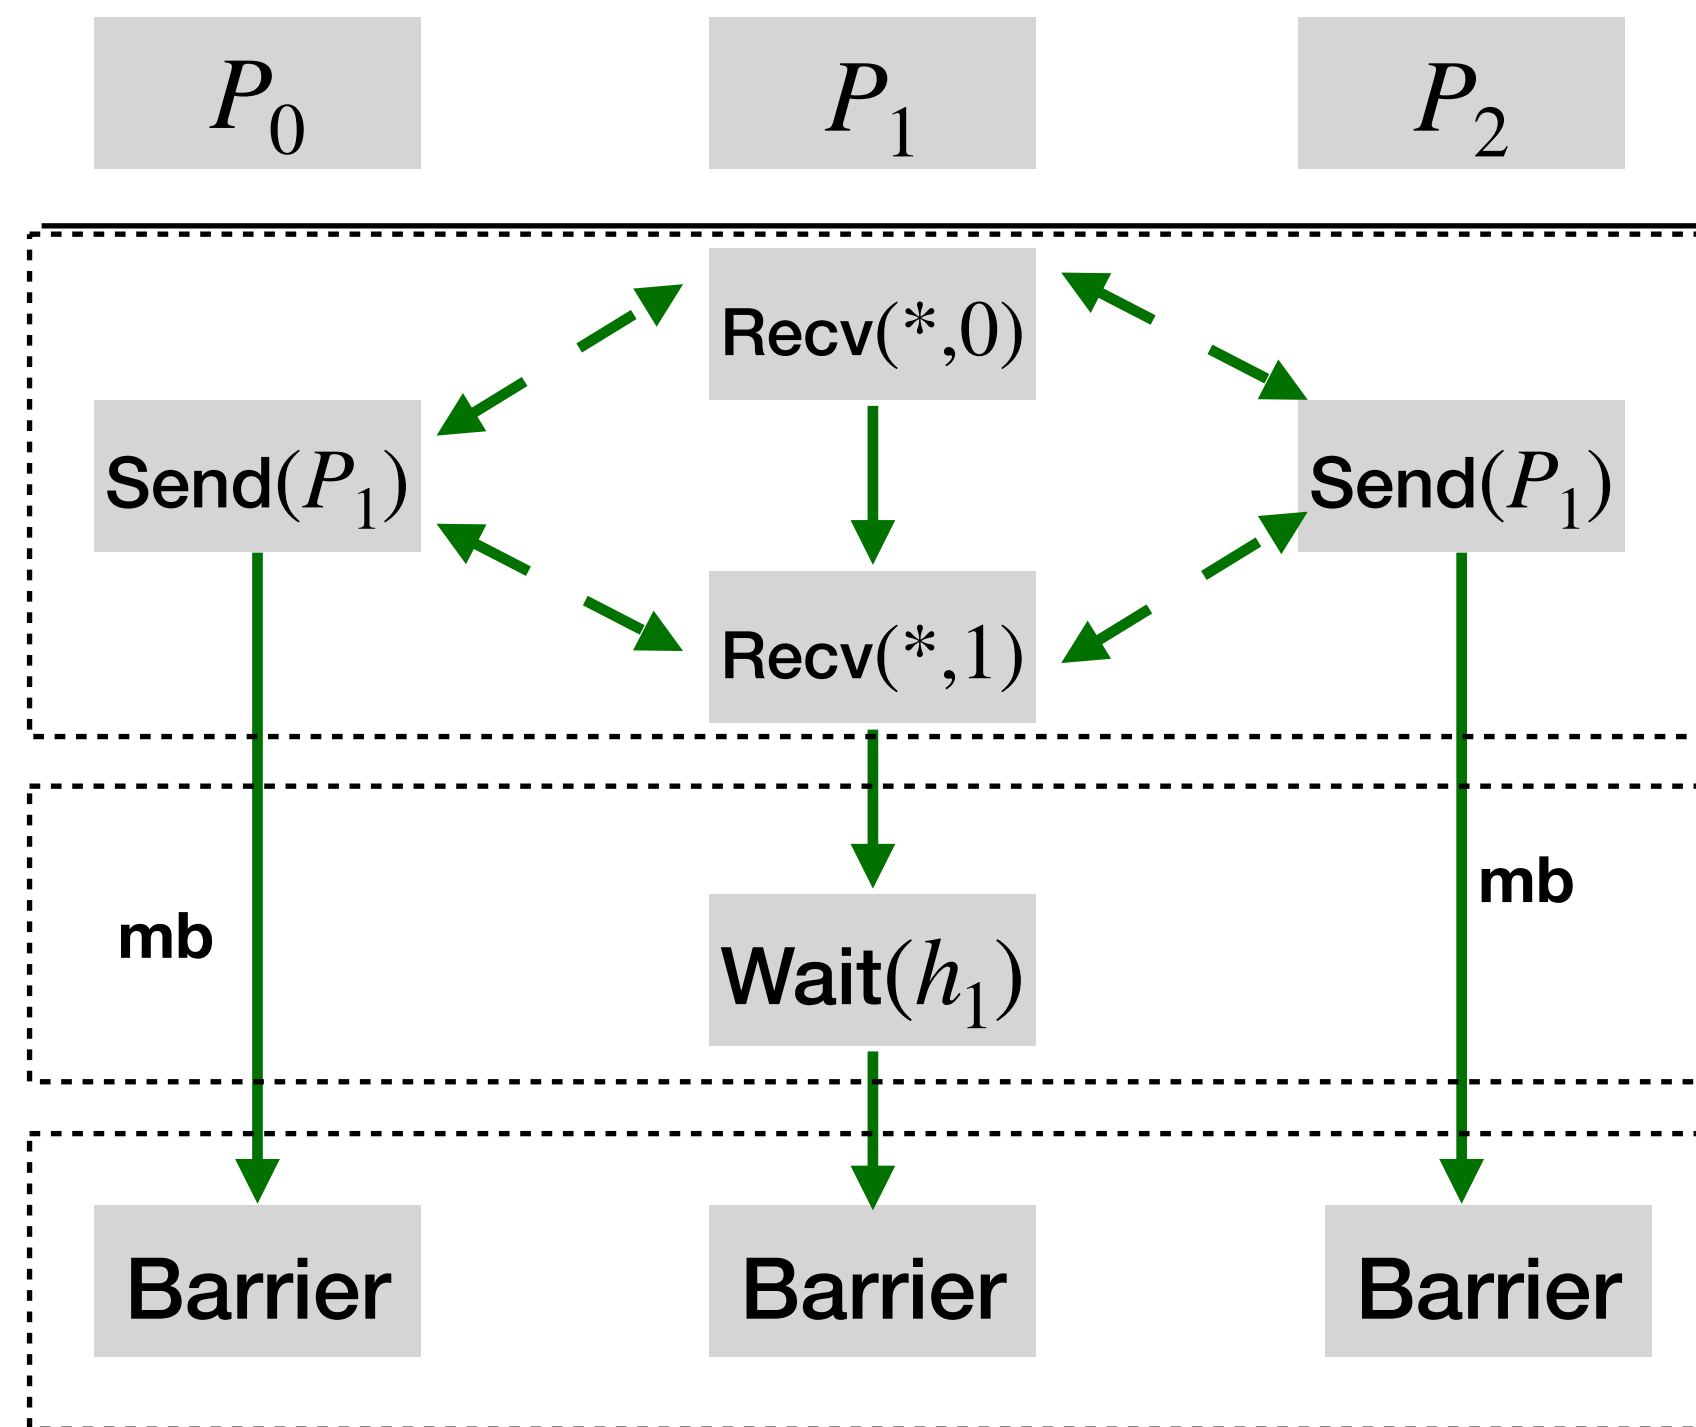
\* : Benchmark not supported  
 TO : Time out after 1800 secs

**Summary:** Orders of magnitude faster than explicit-state or full symbolic verifiers

**Generality:** Java 8 Multi-threaded Apache libraries for wait-notify deadlocks [ICSME20, ICST21]

# Exploiting Symmetries: Epochs

- Problem: Solvers are stuck on exploring symmetric runs of a program
- Problem: Symmetries may not be global

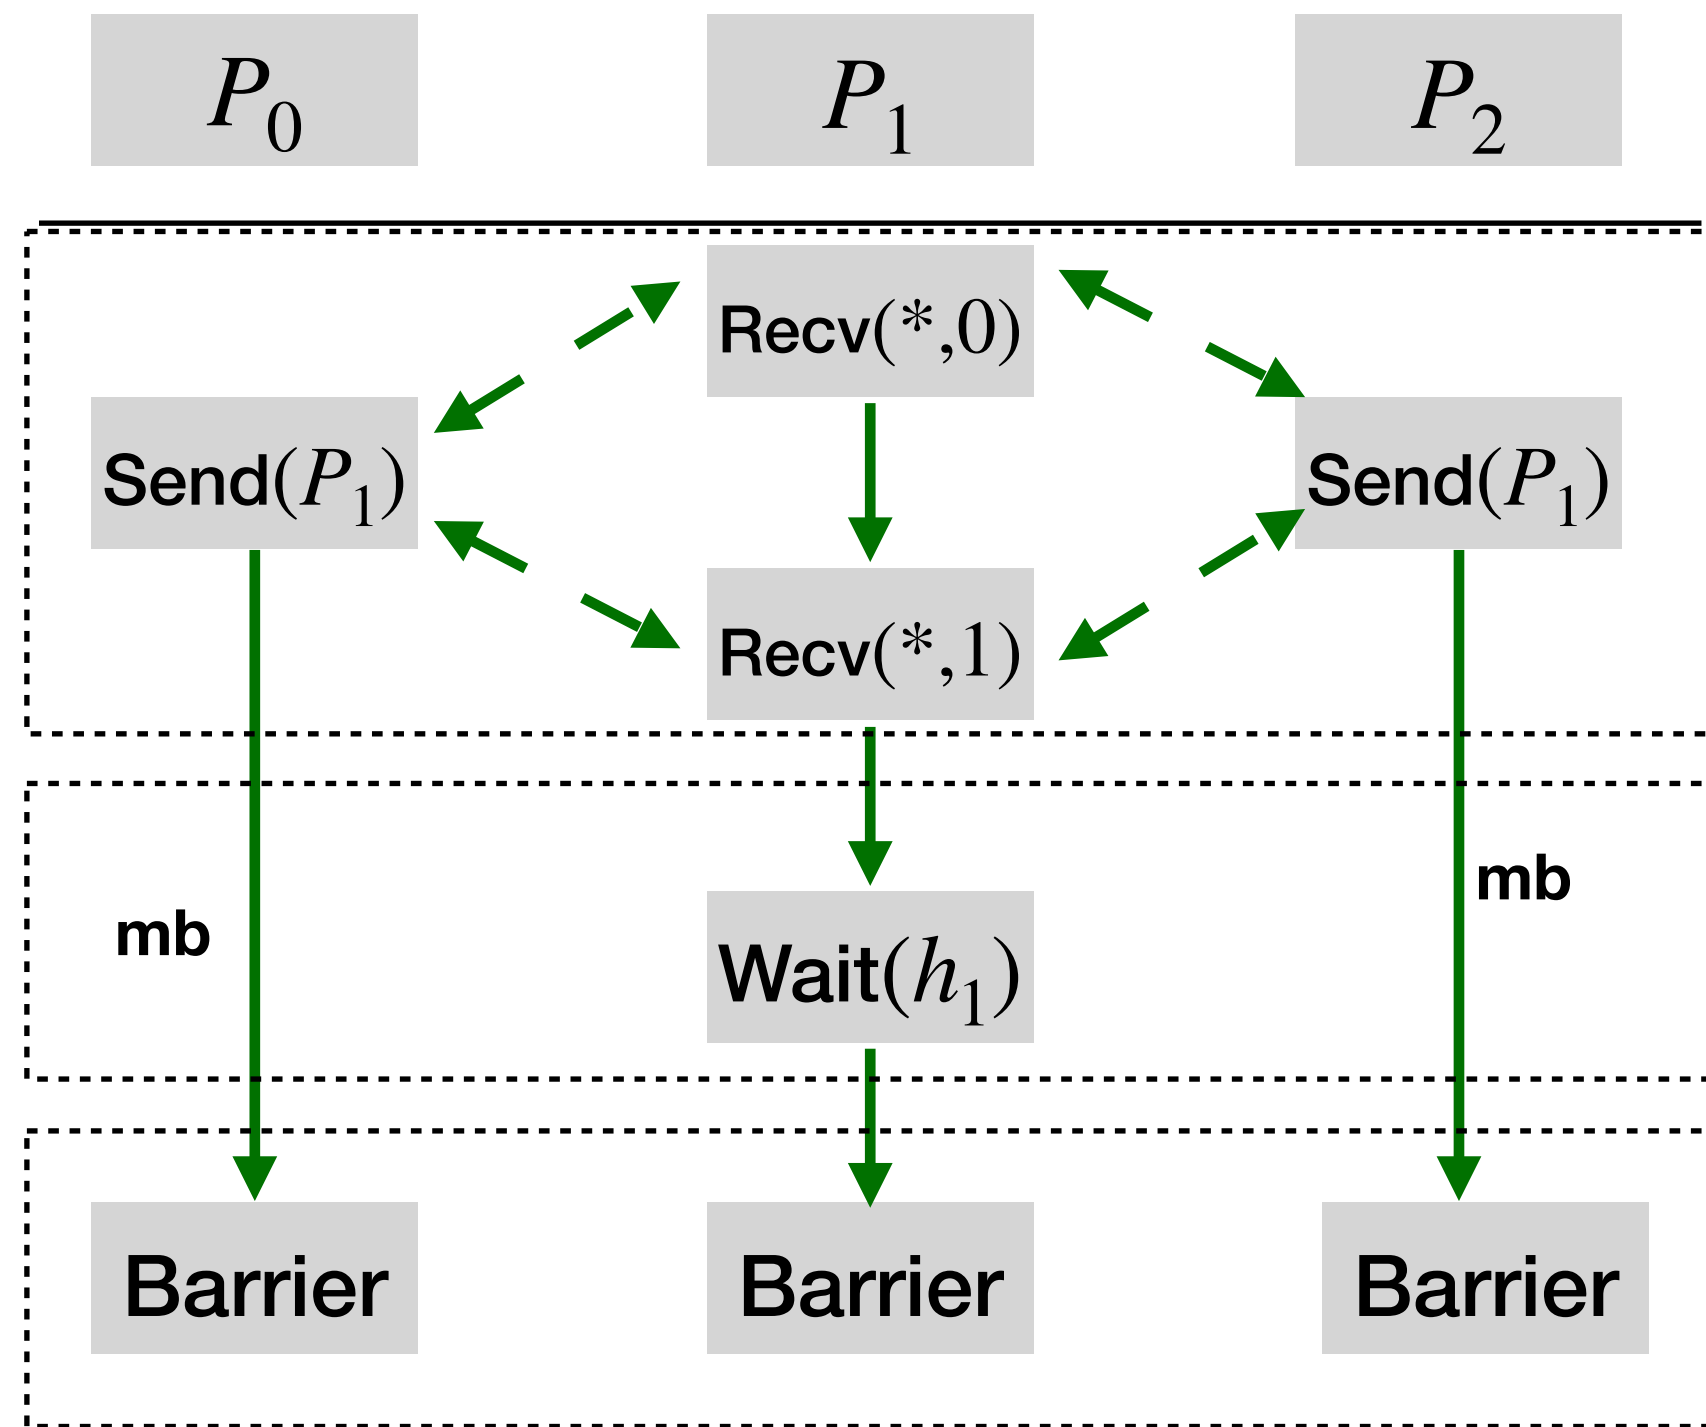


**Program Graph:**  $G = (C, E)$  s.t.  
 $\forall c_1, c_2 \in C, (c_1, c_2) \in E \iff c_1 < c_2$  and  
 $(c_1, c_2), (c_2, c_1) \in E \iff \{c_1, c_2\} \in M$

**Communication Epochs:** SCCs of program graph

# Detecting & Specifying Symmetries

- Automorphism  $\pi : C \rightarrow C$  if  $(c_1, c_2) \in E \iff (\pi(c_1), \pi(c_2)) \in E$
- **Lemma:**  $\pi$  preserves matches and deadlocks



$$\phi = \phi_{po} \wedge \phi_{exec} \wedge \neg \phi_{prop} \wedge \phi_{sbp}$$

- SBP (symmetry breaking predicate): **Lex-Leader** constraints

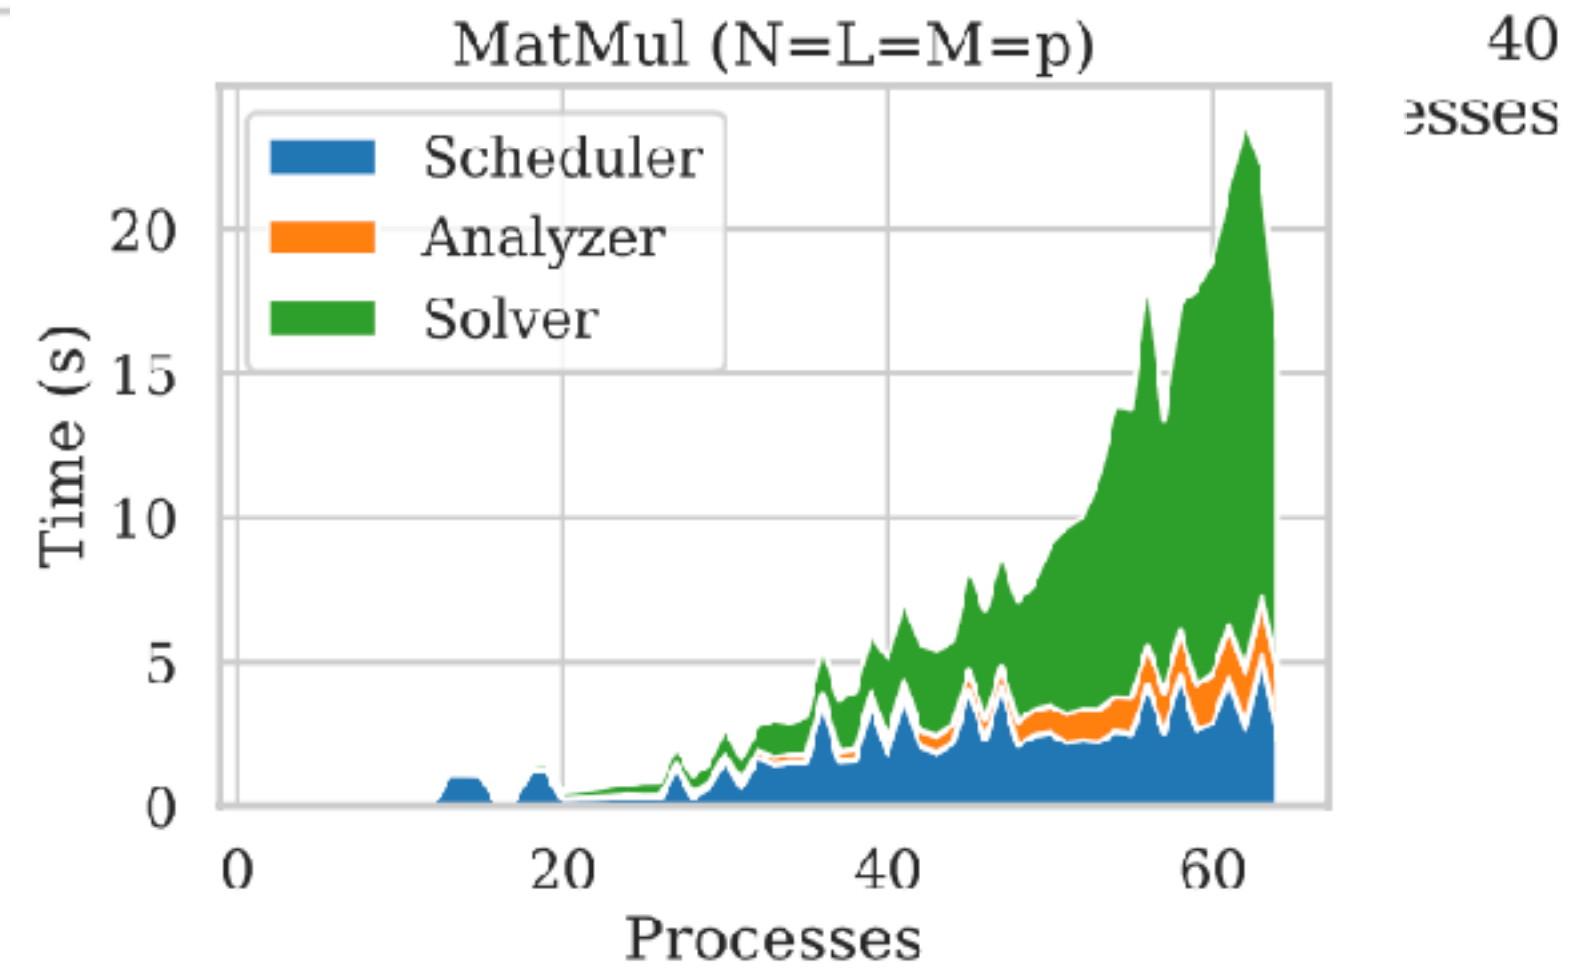
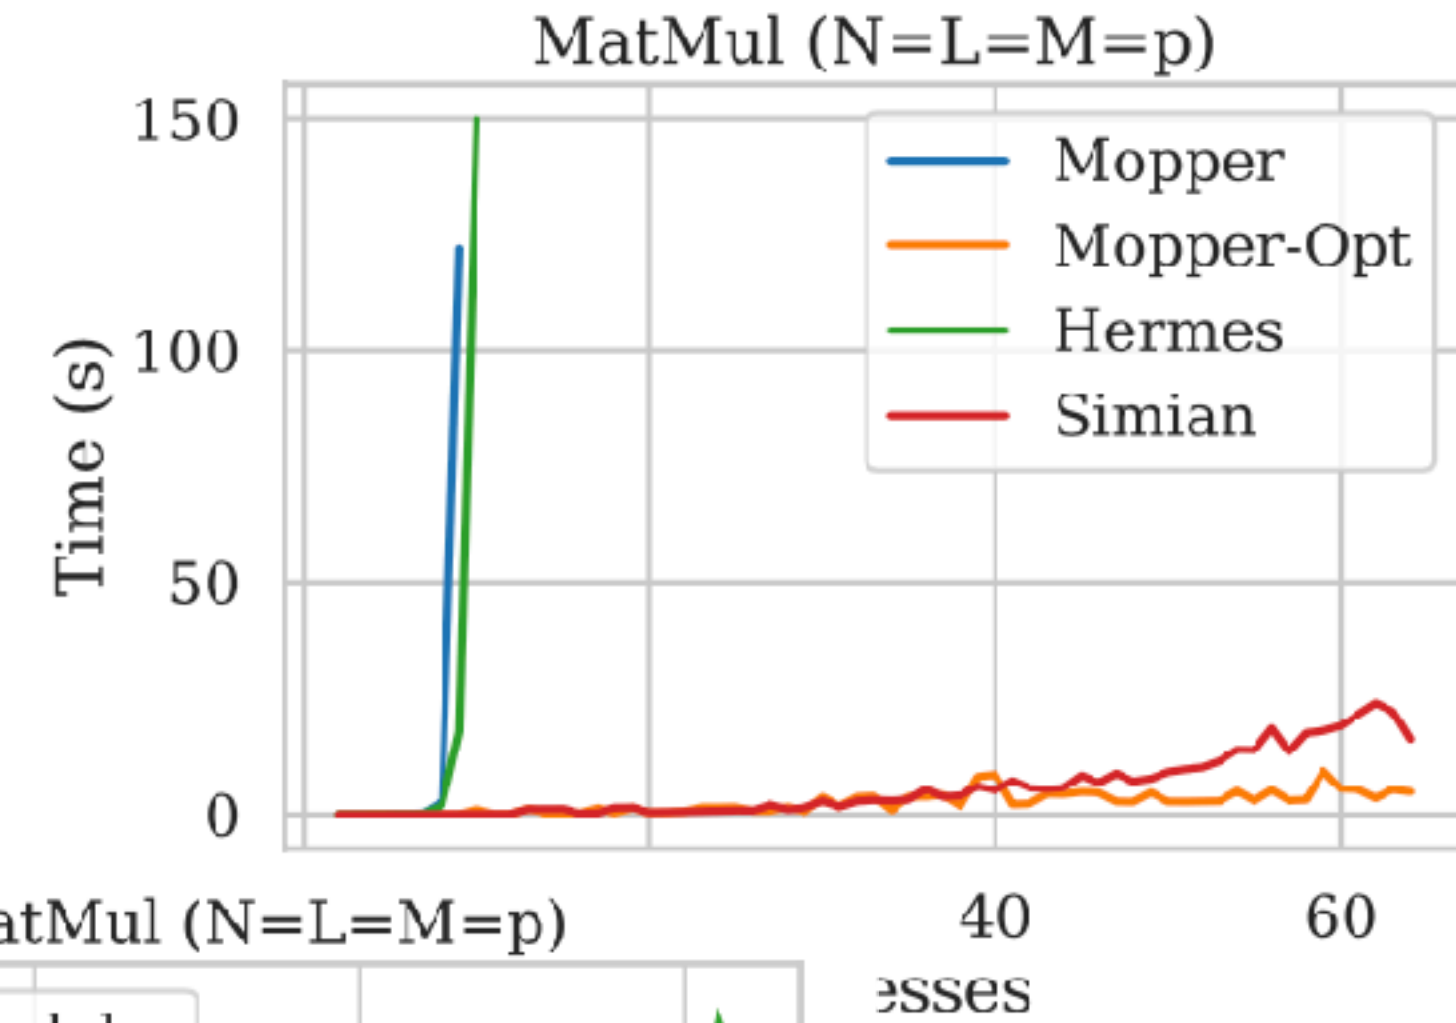
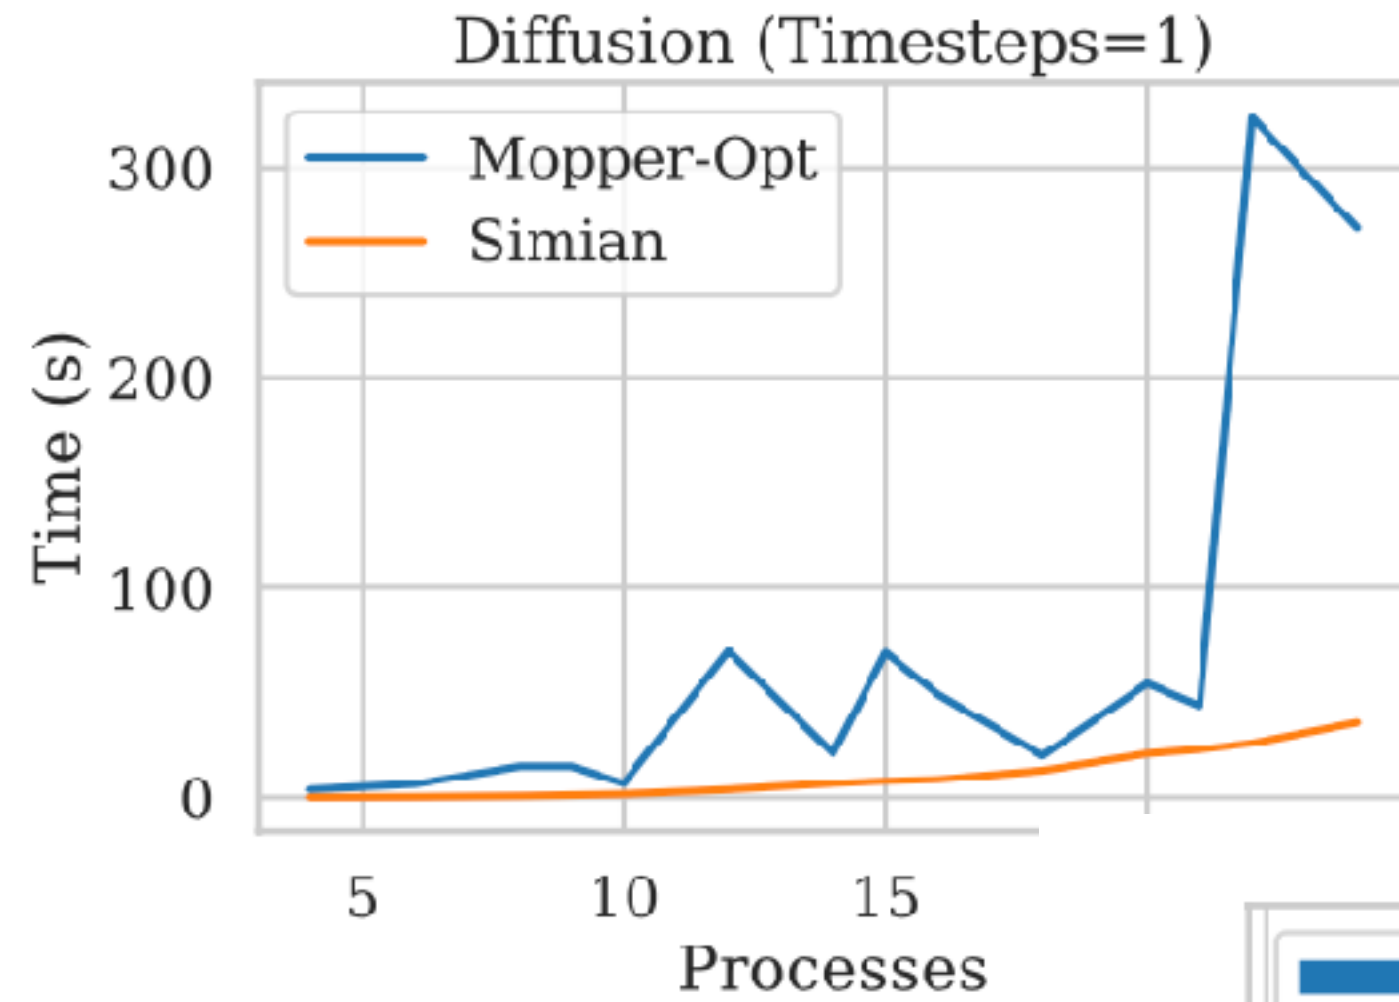
$$0 \leq i \leq n, P_i(V, W) \stackrel{\text{def}}{=} (V_{i-1} = W_{i-1} \rightarrow s_{\alpha_i} \leq s_{\beta_i})$$

$$[\tau_1] \leq [\tau_2] \stackrel{\text{def}}{=} \bigwedge P_i([\tau_1], [\tau_2])$$

$$\phi_{sbp} = \bigwedge_{\pi \in B} [\tau] \leq [\pi(\tau)]$$

# Exploiting Symmetries: Results

Name vs X	X	Deadlock	Hermes	Simian
Adder	8	No	0.212	0.06
	vs 16	No	TO	1.157
	Processes 32	No	TO	1.912
	64	No	TO	7.257
GaussElim	8	No	0.187	<b>0.186</b>
	vs 16	No	0.258	<b>0.283</b>
	Processes 32	No	1.993	<b>2.334</b>
	64	No	3.912	<b>3.226</b>
Heat	8	No	0.406	<b>0.325</b>
	vs 16	No	0.636	1.506
	Processes 32	No	2.005	4.255
	64	No	4.464	3.232
HeatErrors	8	Yes	<b>0.309</b>	0.353
	vs 16	Yes	0.662	<b>0.565</b>
	Processes 32	Yes	2.71	<b>2.19</b>
	64	Yes	6.435	5.051
Integrate	8	No	0.209	0.06
	vs 16	No	TO	<b>0.131</b>
	Processes 32	No	TO	<b>1.854</b>
	64	No	TO	7.583
Diffusion (Timesteps=1)	4	No	TO	<b>0.122</b>
	vs 8	No	TO	<b>0.716</b>
	Processes 16	No	TO	<b>8.566</b>
	24	No	TO	<b>36.308</b>
Diffusion (Grid=2x2)	2	No	TO	<b>0.225</b>
	vs 4	No	TO	<b>1.023</b>
	Processes 8	No	TO	7.197
	Timesteps 16	No	TO	<b>67.854</b>
MatMul (N=L=M=8)	8	No	1.815	<b>0.076</b>
	vs 16	No	3.032	<b>0.114</b>
	Processes 32	No	4.014	1.477
	64	No	5.48	<b>2.766</b>
MatMul (N=L=M=p)	8	No	1.918	<b>0.076</b>
	vs 16	No	TO	<b>0.224</b>
	Processes 32	No	TO	<b>2.822</b>
	64	No	TO	16.175
MatMul (p=8)	4	No	0.071	0.065
	vs 6	No	0.112	0.07
	Processes 8	No	1.79	<b>0.076</b>
	Size (N=L=M) 12	No	TO	<b>4.938</b>



- Summary:


- Exponential savings over the multi-path contribution
- Negligible analysis overheads (<10%)

# Thank You!

[svs@cse.iitd.ac.in](mailto:svs@cse.iitd.ac.in)

<https://subodhvsharma.github.io>

## Collaborators



Saurabh J

Ishita Agarwal  
Rishabh Ranjan  
Ganesh Narayanaswamy

