Thread-Modular Model Checking with Iterative Refinement

Wenrui Meng	Tsinghua University
Fei He	Tsinghua University
Bow-Yaw Wang	Academia Sinica
Qiang Liu	Tsinghua University

Shared-Memory Concurrent System

- Concurrent model becomes more and more important, especially with the trend of multi-core architecture

- Many processes run concurrently and interact by shared variables
- Problem in concurrent system design
  - Access conflict
  - Deadlock
  - Starvation

- Model checking ensure the correctness of system

- Challenge in model checking: state explosion

Related work

- Thread-Modular Model Checking[Flanagan, C., Qadeer, 2003]
  - Disadvantage: Incomplete
  - Advantage: Local Invariant

- Local Proof with Split Invariant[Cohen, A., Namjoshi, K 2007]
  - Disadvantage: Require for all information
  - Advantage: Complete (With efficient refinement)

Model Definition

- **Process** $P=(X, L, I, T)$ is a quadruple where
  - X is a set of global variables,
  - L is a set of local variables,
  - I is the initial condition,
  - T is the transition relation.

- **Composition** of $P_1(X, L_1, I_1, T_1)$ and $P_2(X, L_1, I_1, T_1)$ is $P_1 \| P_2=(X, L, I, T)$ where
  - $L_1 \cap L_2=\emptyset$, $L=L_1 \cup L_2$,
  - $I=I_1 \wedge I_2$,
  - $T= (T_1 \wedge (L_2` = L_2)) \vee (T_2 \wedge (L_1` = L_1))$.

Model Definition

- **Guarantee** of process $P(X, L, I, T)$ is $G(P)=(X, \emptyset, I_x, T_x)$ where
  - I_x is the initial condition;
  - T_x is the transition relation.

- S_1 is the states on V_1, S_2 is the states on V_2 and $Y=V_1 \cap V_2$, **join**($\bowtie$) is a binary operator on S_1 and S_2, $S=S_1 \bowtie S_2$: $t \in S$ *if*
  - $t\downarrow_{V1} \in S_1$;
  - $t\downarrow_{V2} \in S_2$.

Thread-Modular Model Checking With Refinement

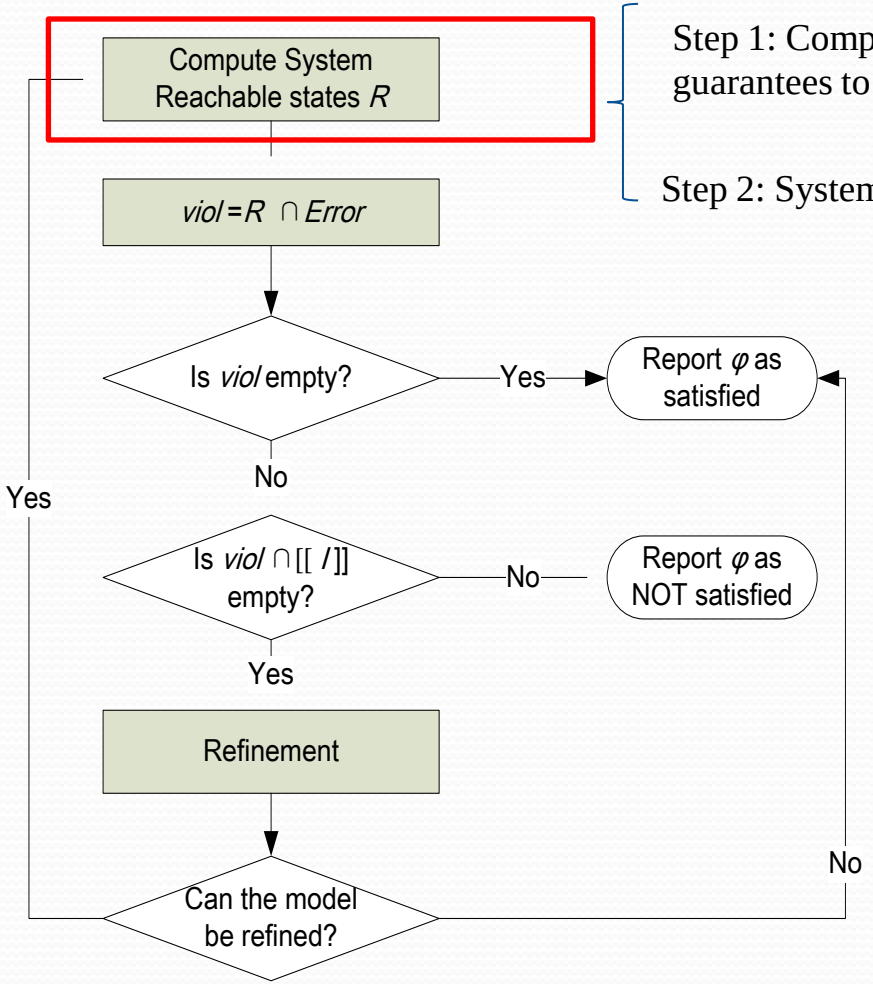

Step 1: Compose P_i with other processes' guarantees to calculate R_i for P_i;

Step 2: System reachable states $R = R_1 \bowtie \ldots \bowtie R_N$;

TMMCIR

Thread-Modular Model Checking With Refinement

Step 3

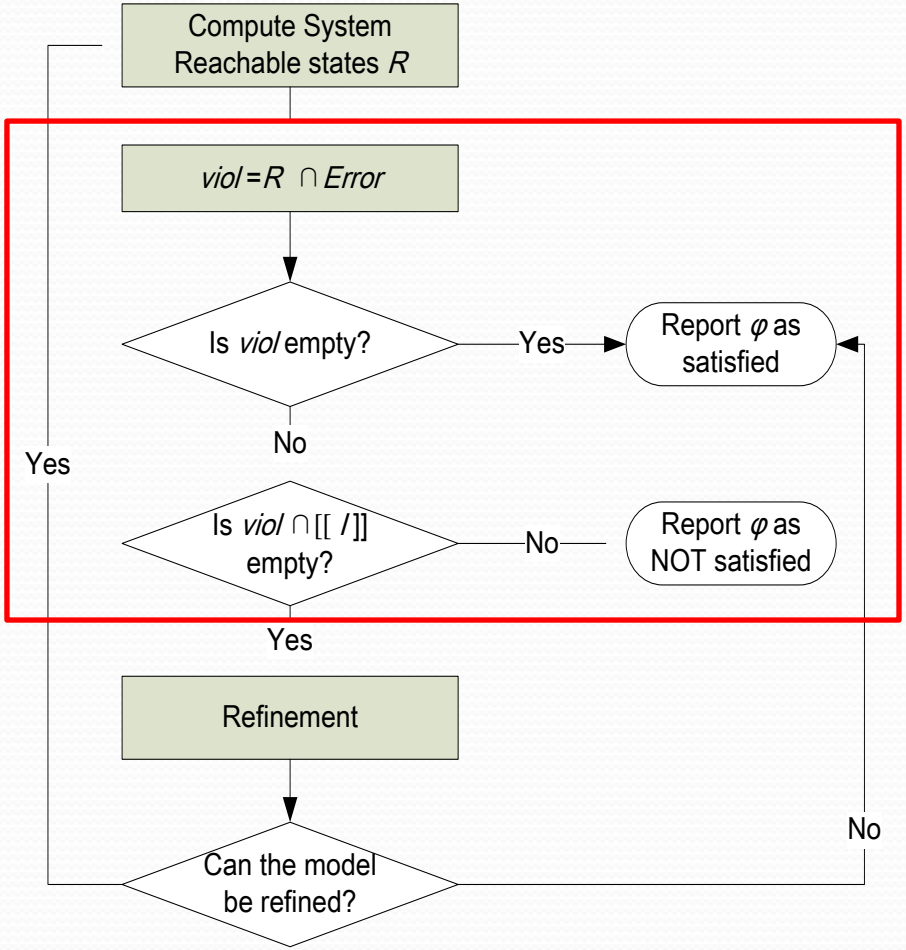

TMMCIR

Thread-Modular Model Checking With Refinement

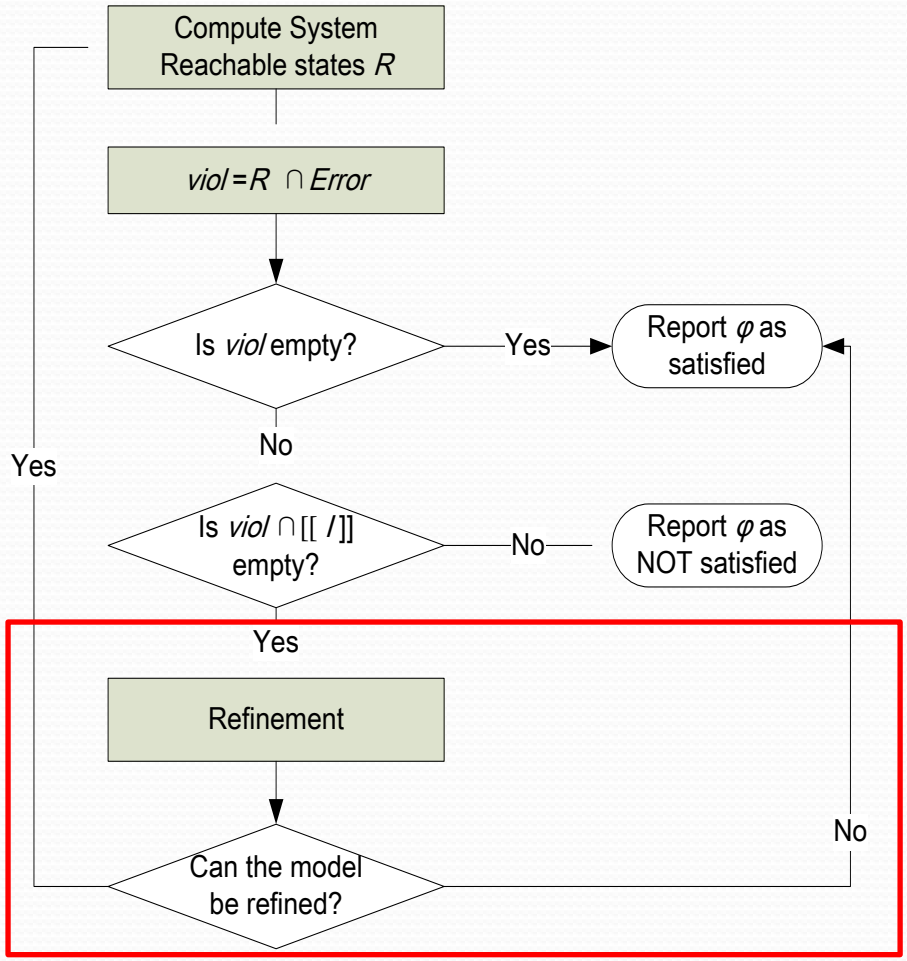

Step 4

TMMCIR

Case Analysis

$P_{i=1\sim N}$

in N : natural where N >1
x: boolean initially x =1
loop for ever
- $l = 0$: Non-Critical
- $l = 1$: request x
- $l = 2$: Critical
- $l = 3$: release x

Property φ:
For any $i \neq j$!((Pi.l=2,3)∧(Pj.l=2,3))

Mutual exclusion MUTEX-SEM

Case analysis- Step 1

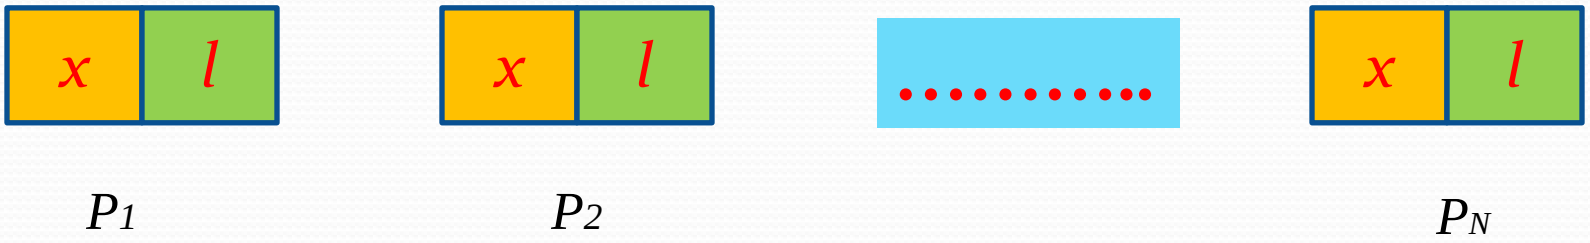

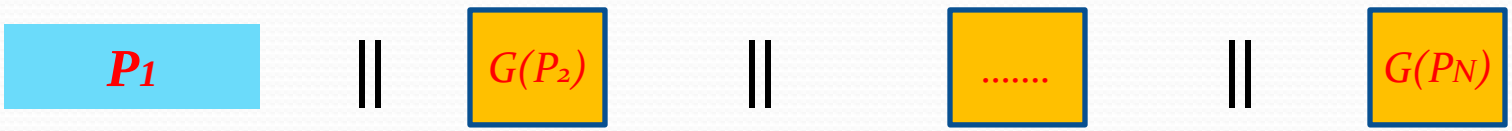

Reachable states
R1

Case analysis- Step 1

- P_i=({x}, {l}, (x=1, l=0), {(x=*1*, l=*0*, x`=*1*, $l$`=*1*), (x=*1*, l=*1*, x`=*0*, $l$`=*2*), (x=*0*, l=*2*, x`=*0*, $l$=*3*), ($x$=*0*, $l$=*3*, $x$`=*1*, l`=*0*)});

- $G(P_i)$=({x}, $\emptyset$, (x=*1*),{(x=*1*,x`=*1*), ($x$=*1*, $x$`=*0*), (x=*0*, x`=*1*)});

- R_1=[[(x=*0,1*)$\wedge$(l=*0,1,2,3*)]]

Case analysis- Step 2 & 3

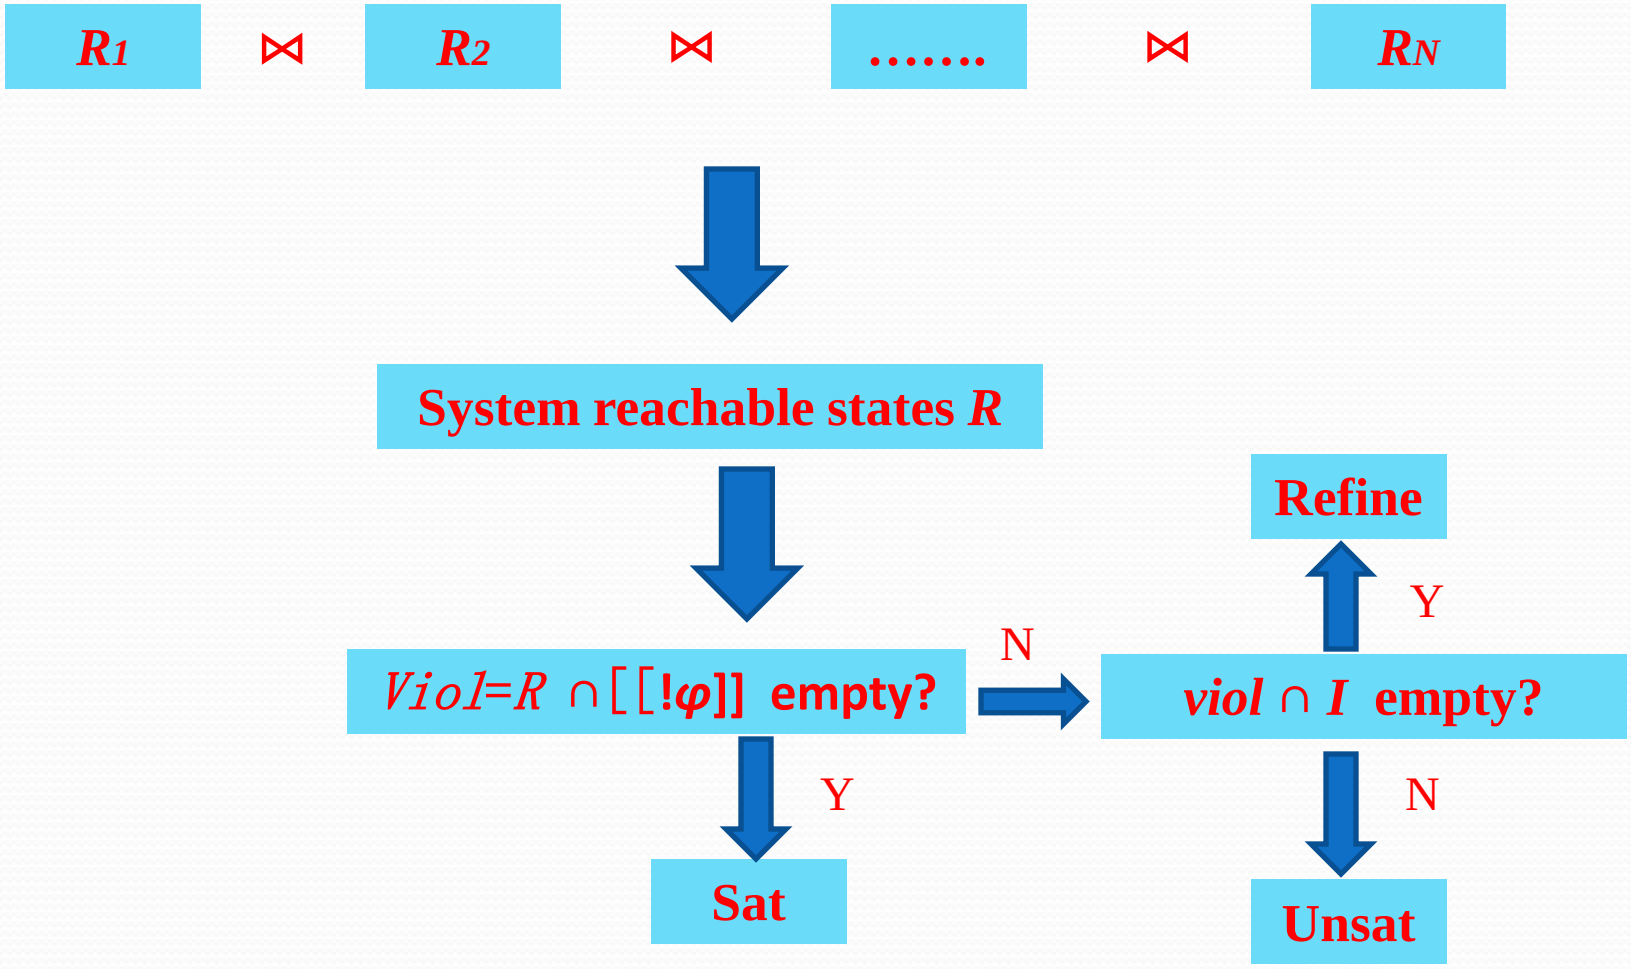

Case analysis- Step 2 & 3

- Suppose N=2, $R=R_1 \bowtie R_2$=[[(x=*0,1*) $\wedge$ ($P_1.l$=*0,1,2,3*) $\wedge$($P_2.l$=*0,1,2,3*)]];
- Error=[[($P_1.l$=*2,3*) $\wedge$($P_2.l$=*2,3*)]];
- $R \cap Error$!= $\emptyset$ and $R \cap Error \cap [[I]] = \emptyset$, so we call refinement.

Refinement

- Expose the relevant information for error states. For each state s in $R \cap error$, we will analyze the value of each local variable in *s* to find whether it's essential to the membership of error.

- If there is no essential value in *s*, we will add the predecessors of *s* into error.

Case analysis- Step 4

***s* ϵ *R* ∩ *Error* :**

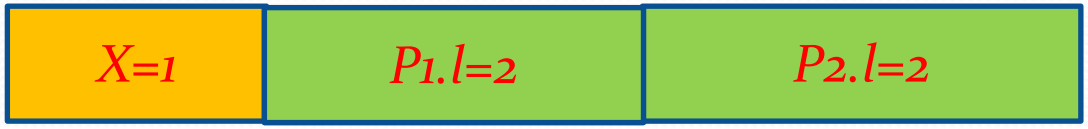

- $P_1.l=2$ is essential to the membership of error for s, we added it as an **essential predicate** ψ and added **auxiliary variable** b_1 for it. For each state s, $b_1 = \psi(s)$.

- Similarly we find other essential predicates $P_1.l=3$, $P_2.l=2$, $P_2.l=3$, their auxiliary variables b_2, b_3, b_4.

Case analysis

- After refinement, the global variable set $X=\{x, b_1, b_2, b_3, b_4\}$.
- Recalculate R_1 and R_2;
- System reachable states R is

$$\left\{\begin{array}{llll}(1,0,0,0,0|0,0), & (1,0,0,0,0|0,1), & (1,0,0,0,0|1,0), & (1,0,0,0,0|1,1),\\ (0,1,0,0,0|2,0), & (0,1,0,0,0|2,1), & (0,0,1,0,0|3,0), & (0,0,1,0,0|3,1),\\ (0,0,0,1,0|0,2), & (0,0,0,0,1|0,3), & (0,0,0,1,0|1,2), & (0,0,0,0,1|1,3)\end{array}\right\}$$

$R \bowtie Error = \emptyset$, so property φ is satisfied.

Sound & Complete

- Given processes $P_i(X, L_i, I_i, T_i)$ *for* $i=1,\ldots,N$ *and property* φ :
  - If TMMCIR return "*PASS*", then $P_1||\ldots||P_i \models \varphi$;
  - If TMMCIR return "*FAILURE*", then $P_1||\ldots||P_i \not\models \varphi$.

$S_0, S_1, \ldots, S_i, S_{i+1}, \ldots, S_k$

- Given processes $P_i(X, L_i, I_i, T_i)$ *for* $i=1,\ldots,N$ *and property* φ :
  - If $P_1||\ldots||P_i \models \varphi$, then TMMCIR return "*PASS*" ;
  - If $P_1||\ldots||P_i \not\models \varphi$, then TMMCIR return "*FAILURE*".

Heuristics

- After finding a essential predicate from state s, we will stop analysis of *s* and remove all the states which share the same predicate from $R \cap Error$.

- After refinement, we don't recalculate the reachable states for every process, but calculate the reachable states of the composition of all new guarantees, and then make join with previous reachable states of every process.
  - $R_G = Re(G(P_1) || \ldots || G(P_N))$;
  - $R = R_G \bowtie R_1 \bowtie \ldots \bowtie R_N$.

Experiment

Method	P	Time	BDD	R	A	Conclusive
TMMCIR	20	0.064	45990	1	40	Y
SPLIT	20	0.887	331128	1	38	Y
AFR	20	0.064	45990	/	/	Y
NuSMV	20	0.144	141036	/	/	Y
TMMC	20	0.032	20440	/	/	N
TMMCIR	50	0.580	401646	1	100	Y
SPLIT	50	12.187	4555054	1	98	Y
AFR	50	3.320	1242752	/	/	Y
NuSMV	50	3.412	2444624	/	/	Y
TMMC	50	0.228	203378	/	/	N
TMMCIR	100	5.536	1510344	1	200	Y
SPLIT	100	207.233	57265726	1	198	Y
AFR	100	208.561	3059868	/	/	Y
NuSMV	100	614.806	4762520	/	/	Y
TMMC	100	4.200	2057907	/	/	N
TMMCIR	200	27.966	2439514	1	400	Y
TMMCIR	300	145.093	5767146	1	600	Y

Mutual exclusion for MUTEX-SEM

Experiment

Method	P	Time	BDD	R	A	Conclusive
TMMCIR	10	0.020	14308	1	20	Y
SPLIT	10	0.321	100019	1	19	Y
AFR	10	5.996	617288	/	/	Y
NuSMV	10	2.288	1030176	/	/	Y
TMMC	10	0.016	10220	/	/	N
TMMCIR	20	0.104	122640	1	40	Y
SPLIT	20	2.520	930020	1	38	Y
AFR	20	1584.179	2634716	/	/	Y
NuSMV	20	3914.385	159986	/	/	Y
TMMC	20	0.044	47012	/	/	N

Mutual exclusion MUTEX-SEM-COUNT

Experiment

Method	P	Time	BDD	R	A	Conclusive
TMMCIR	4	0.100	96068	0	0	Y
SPLIT	4	0.267	218708	0	0	Y
AFR	4	0.084	91980	/	/	Y
NuSMV	4	0.140	106288	/	/	Y
TMMC	4	0.100	96068	/	/	Y
TMMCIR	8	26.246	2389436	0	0	Y
SPLIT	8	75.141	26776400	0	0	Y
AFR	8	240.555	4258674	/	/	Y
NuSMV	8	1282.984	25237268	/	/	Y
TMMC	8	25.780	2389436	/	/	Y

BAKERY

Experiment

Method	P	Time	BDD	R	A	Conclusive
TMMCIR	6	0.104	85848	3	12	Y
SPLIT	6	0.320	158410	3	6	Y
AFR	6	0.016	14308	/	/	Y
NuSMV	6	0.036	17374	/	/	Y
TMMC	6	0.008	7154	/	/	N
TMMCIR	8	1.028	367920	3	16	Y
SPLIT	8	16.442	1176322	5	10	Y
AFR	8	0.236	243236	/	/	Y
NuSMV	8	0.236	223818	/	/	Y
TMMC	8	0.020	24528	/	/	N
TMMCIR	10	15.981	1475768	3	20	Y
SPLIT	10	5274.488	4193266	6	11	Y
AFR	10	2.592	1815434	/	/	Y
NuSMV	10	3.556	1739444	/	/	Y
TMMC	10	0.052	48034	/	/	N

Dining Philosopher

Acknowledgement & Conclusion

- Thanks to Kedar Namjoshi for introducing the detail of the refinement algorithm

- Conclusion:

  In the experiment, our approach performs better than other complete verification algorithms. The main reason is that we compute the reachable states separately. The heuristics are also very helpful.

Q & A