

Statistical Model Checking for Timed Automata

Collaborators: Peter Bulychev,
Alexandre David
Axel Legay, Marius Mikucionis
Wang Zheng
Jonas van Vliet, Danny Poulsen



CAV 2011, PDMC 2011,
FORMATS 2011



UPPAAL

Safety ✓

$A[] \text{ forall } (i : \text{id_t}) \text{ forall } (j : \text{id_t})$
 $\text{Train}(i).\text{Cross} \ \&\& \ \text{Train}(j).\text{Cross} \ \text{imply } i == j$

Reachability ✓

$E \leftrightarrow \text{Train}(0).\text{Cross} \ \text{and} \ \text{Train}(1).\text{Stop}$

Liveness ✓

$\text{Train}(0).\text{Appr} \ \text{-->} \ \text{Train}(0).\text{Cross}$

$A \leftrightarrow .. E[] ..$ ✓

Limited quantitative analysis ✓

sup: .. inf: ..

Performance properties ✗

$\text{Pr}[\leftrightarrow \text{Time} \leq 500 \ \text{and} \ \text{Train}(0).\text{Cross}] \geq 0.7$
 $\text{Pr}[\text{Train}(0).\text{Appr} \ \text{-->}_{\text{Time} \leq 100} \ \text{Train}(0).\text{Cross}] \geq 0.4$

State-space explosion ✗



UPPAAL SMC

The screenshot displays the UPPAAL SMC interface with several components:

- Editor/Simulator/Verifier tabs:** Located at the top, with 'Simulator' selected.
- Simulation Trace:** A list of events such as 'Train(5) appr[0]: Train(0) --> Gate', 'Gate[0]: Gate --> Train(0)', 'Train(0).x - Train(5).x <= 10', 'Train(4).x - Train(0).x <= 10', 'Train(2).x - Train(3).x <= 10', 'Train(3).x - Train(5).x <= 10', and 'Train(4).x - Train(0).x <= 10'.
- State-space diagrams:** Multiple diagrams showing states like 'Safe', 'Cross', 'Start', and 'Stop' with transitions labeled with guard conditions and actions (e.g., 'x >= 3 leave[0]!', 'x <= 10 stop[1]?').
- Buttons:** 'Next', 'Reset', 'Open', 'Save', 'Random', and 'Slow' are visible.

Performance properties ✓

$\Pr[\leq 200] (\langle \rangle \text{Train}(5).\text{Cross})$

$\Pr[\leq 100] (\langle \rangle \text{Train}(0).\text{Cross}) \geq 0.8$

$\Pr[\leq 100] (\langle \rangle \text{Train}(5).\text{Cross}) \geq$

$\Pr[\leq 100] (\langle \rangle \text{Train}(1).\text{Cross})$

State-space explosion ✓

Generate runs

Performance properties

State-space explosion



Overview

- Statistical Model Checking in UPPAAL
 - Estimation
 - Testing
- Distributed SMC for Parameterized Models
 - Parameter Sweeps
 - Optimization
 - Nash Equilibria
- Distributing Statistical Model Checking
 - Estimation
 - Testing
- Parameter Analysis of DSMC
- Conclusion



Overview

- **Statistical Model Checking in UPPAAL**
 - Estimation
 - Testing
- Distributed SMC for Parameterized Models
 - Parameter Sweeps
 - Optimization
 - Nash Equilibria
- Distributing Statistical Model Checking
 - Estimation
 - Testing
- Parameter Analysis of DSMC
- Conclusion



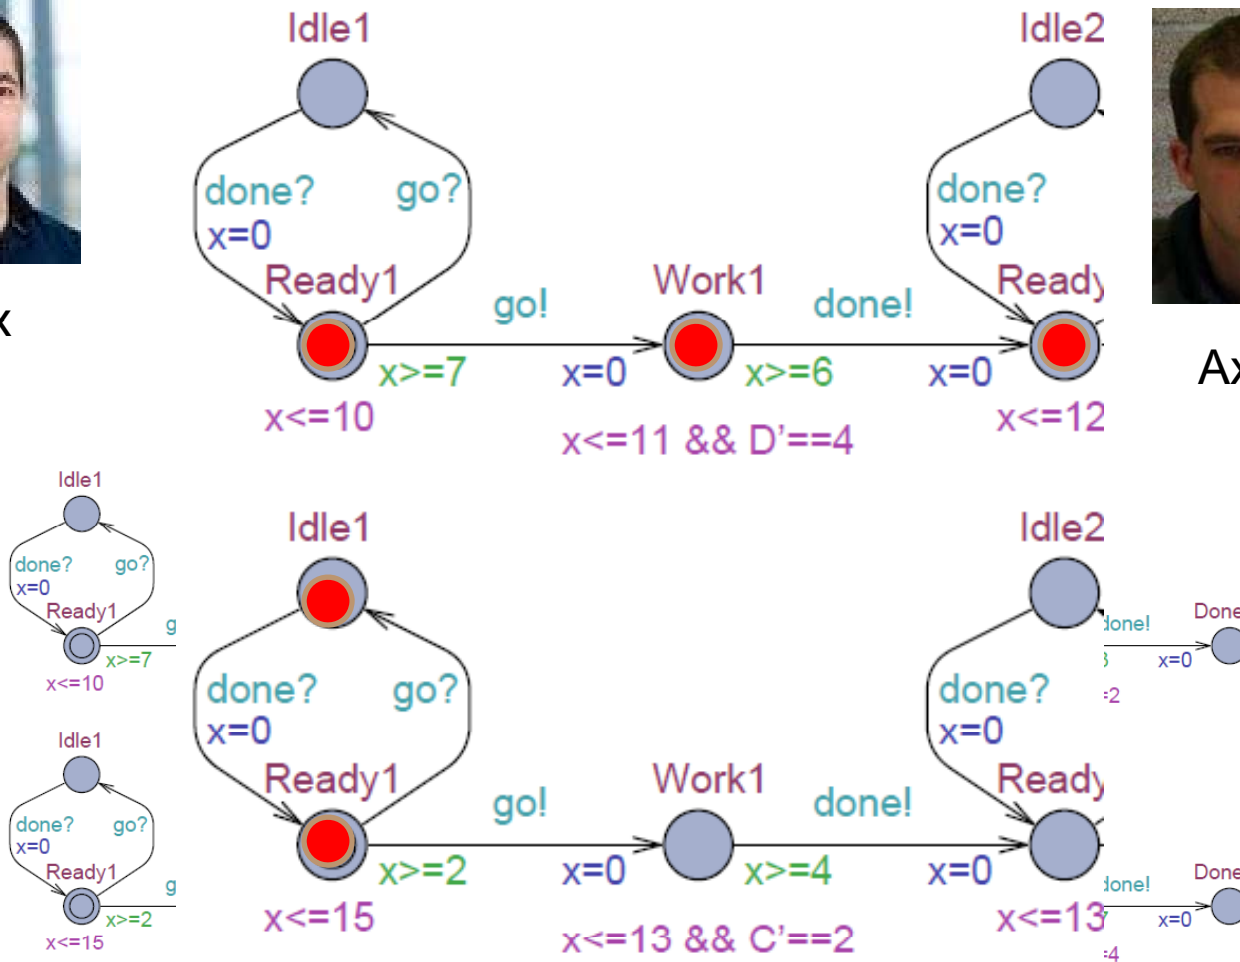
The Hammer Game



Alex



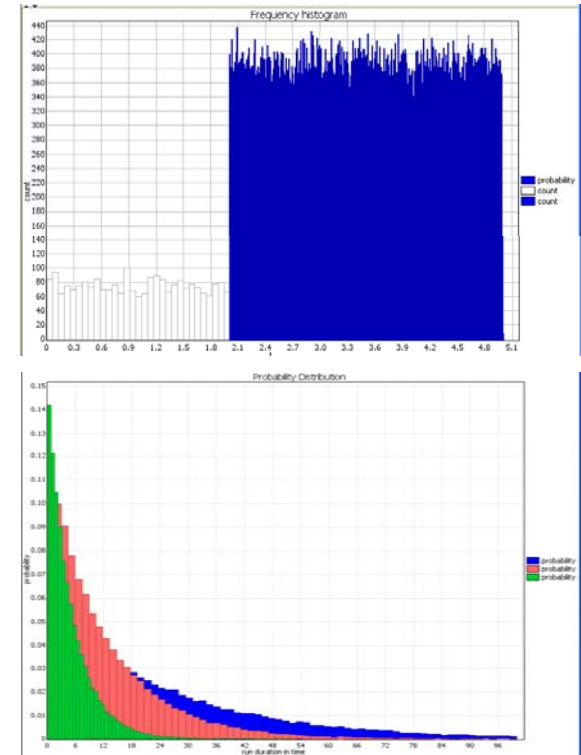
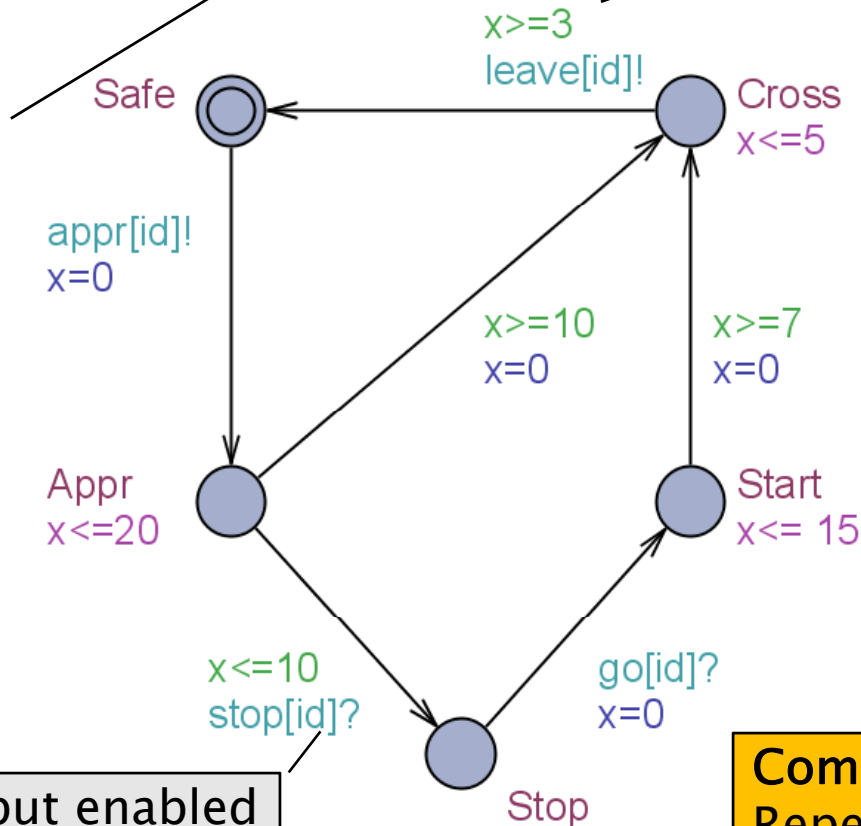
Axel



Stochastic Semantics of TA

Exponential Distribution

Uniform Distribution

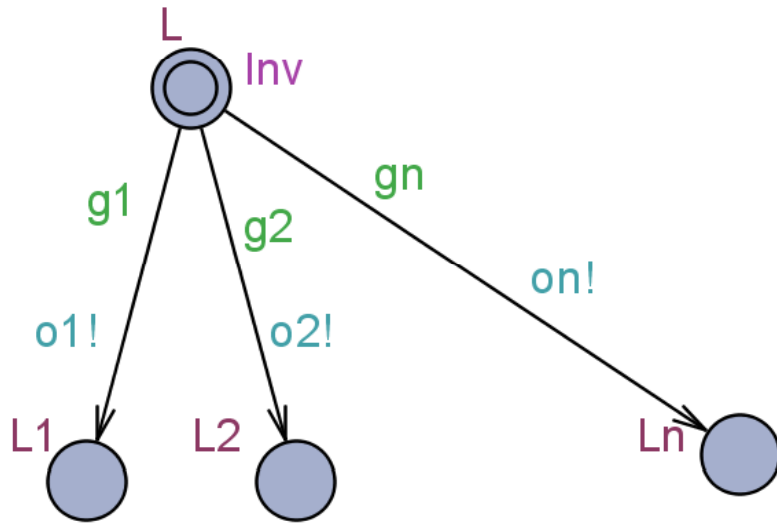


Input enabled

Composition =
Repeated races between components



Stochastic Semantics of Timed Automata

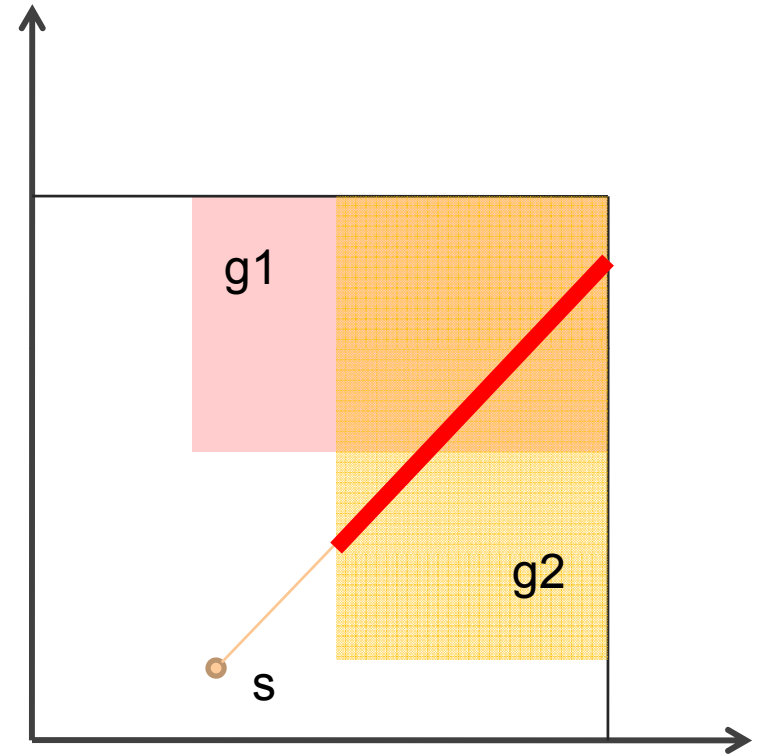


Delay Density Function

$$\mu_s: \mathbf{R} \rightarrow \mathbf{R}$$

Output Probability Function

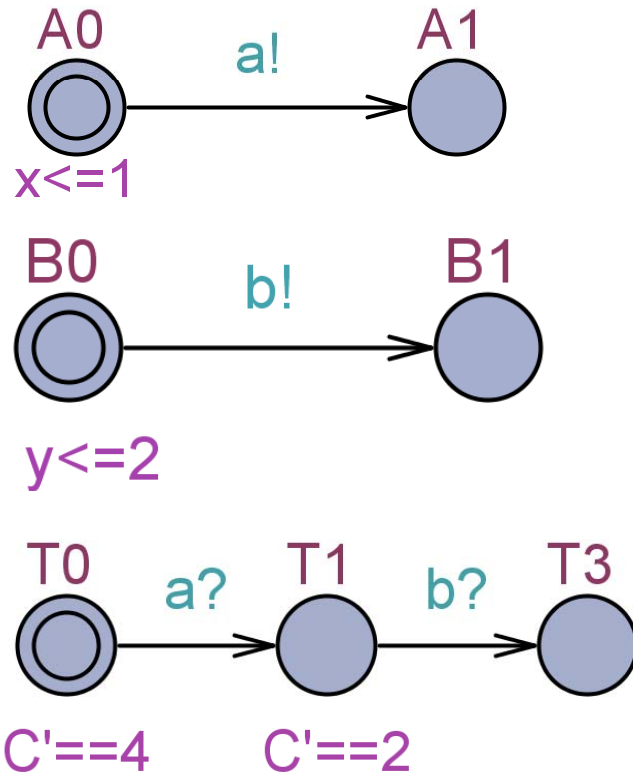
$$\gamma_s: \Sigma_o \rightarrow [0, 1]$$



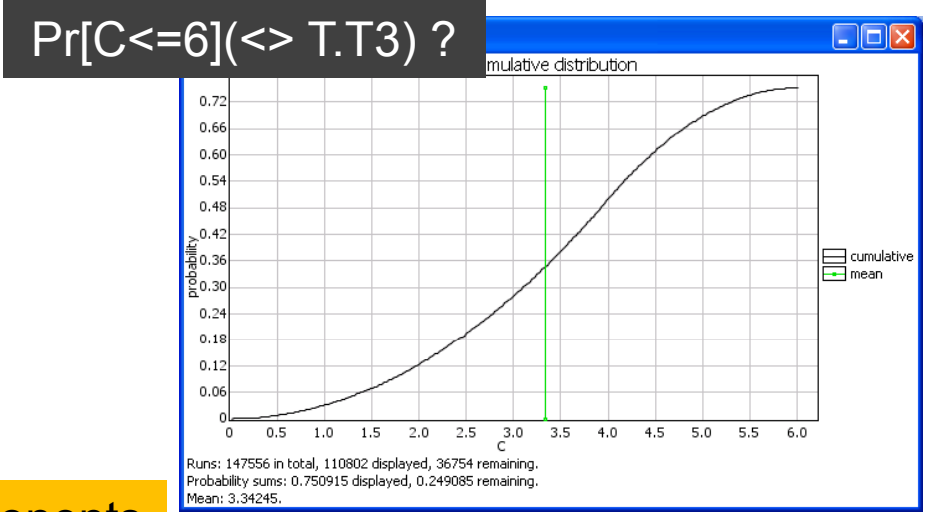
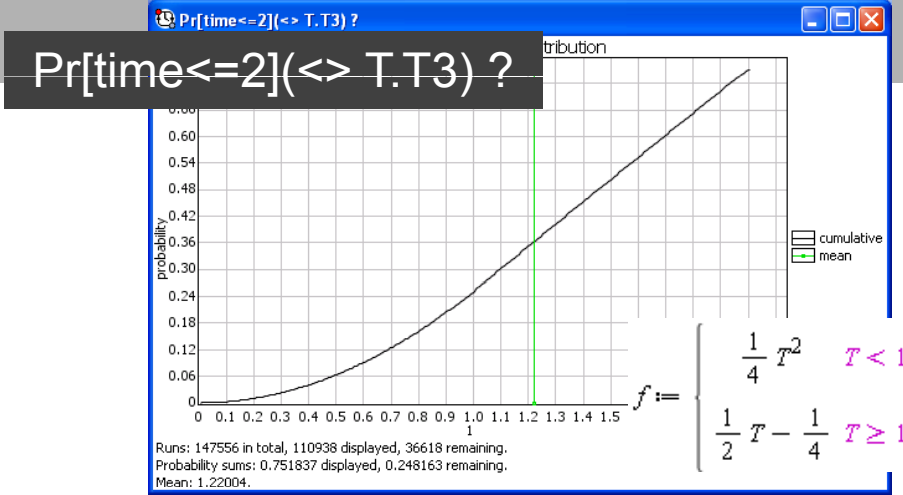
- μ_s uniform on $[d_{\min}, d_{\max}]$
- γ_s uniform over enabled outputs



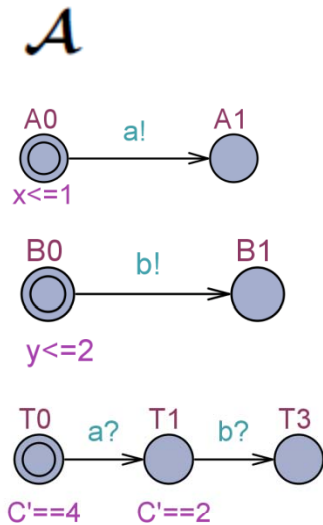
Stochastic Semantics of Timed Automata



Composition = Race between components for outputting



Stochastic Semantics of Timed Automata



Assumptions:

Component TAs are:

- Input enabled
- Deterministic
- Disjoint set of output actions

$\pi(\mathbf{s}, a_1 a_2 \dots a_n)$:
 the set of maximal runs from \mathbf{s} with a prefix
 $t_1 a_1 t_2 a_2 \dots t_n a_n$
 for some $t_1, \dots, t_n \in \mathbf{R}$.

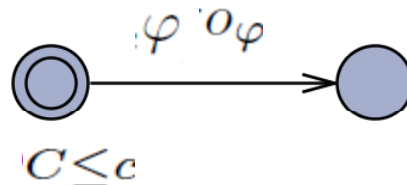
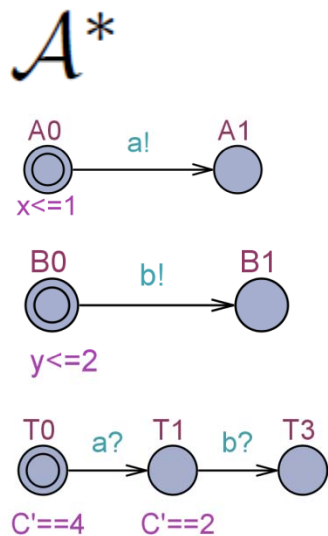
$$\mathbb{P}_{\mathcal{A}}(\pi(\mathbf{s}, a_1 a_2 \dots a_n)) = \int_{t \geq 0} \mu_{s_c}(t) \cdot \left(\prod_{j \neq c} \int_{\tau > t} \mu_{s_j}(\tau) d\tau \right) \cdot \gamma_{s_c}^t(a_1) \cdot \mathbb{P}_{\mathcal{A}}(\pi(\mathbf{s}^t)^{a_1}, a_2 \dots a_n) dt$$

where $c = c(a_1)$, and as base case we take $P_{\mathcal{A}}(\pi(\mathbf{s}), \varepsilon) = 1$.



Logical Properties

$$\psi ::= \mathbb{P}(\diamond_{C \leq c} \varphi) \sim p \mid \mathbb{P}(\square_{C \leq c} \varphi) \sim p$$



$$(A|B|T) \models \mathbb{P}(\diamond_{C \leq 6} T_3) = 0.75$$

$$(A|B|T) \models \mathbb{P}(\diamond_{t \leq 2} T_3) = 0.75$$

$$\mathcal{A} \models \mathbb{P}(\diamond_{C \leq c} \varphi) \sim p \text{ iff } \mathbb{P}_{A^*} \left(\bigcup_{\sigma \in \Sigma^*} \pi(s_0, \sigma o \varphi) \right) \sim p$$



SMC Algorithms in UPPAAL

Qualitative (Hypot)

$$p = \mathbb{P}_{\mathcal{A}}(\diamond_{C \leq c} \varphi)$$

α : prob of acc H_0 when H_1

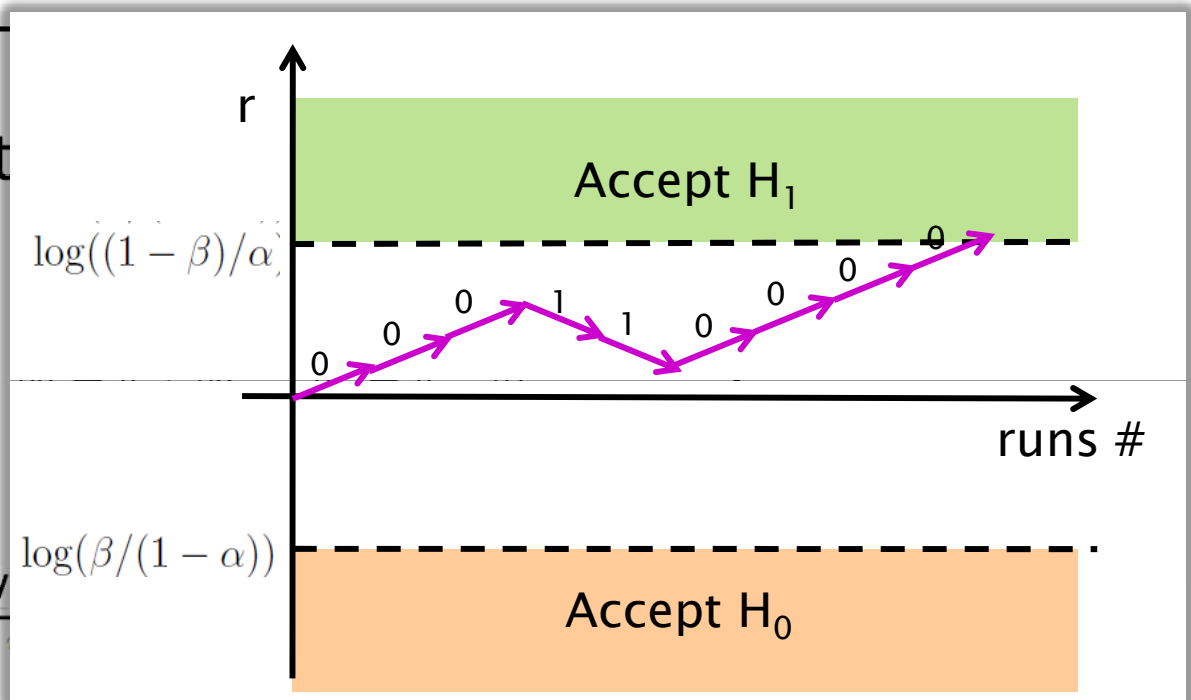
Algorithm II: Sequential Probability

```

function hypothesis( $S$ :model ,
1  $r := 0$ 
2 while true do
3   Observe the random variable  $x$  corresponding to  $\diamond_{C \leq c} \varphi$  for a run.
4    $r := r + x * \log(p_1/p_0) + (1 - x) * \log((1 - p_1)/(1 - p_0))$ 
5   if  $r \leq \log(\beta/(1 - \alpha))$  then accept  $H_0$ 
6   if  $r \geq \log((1 - \beta)/\alpha)$  then accept  $H_1$ 

```

11



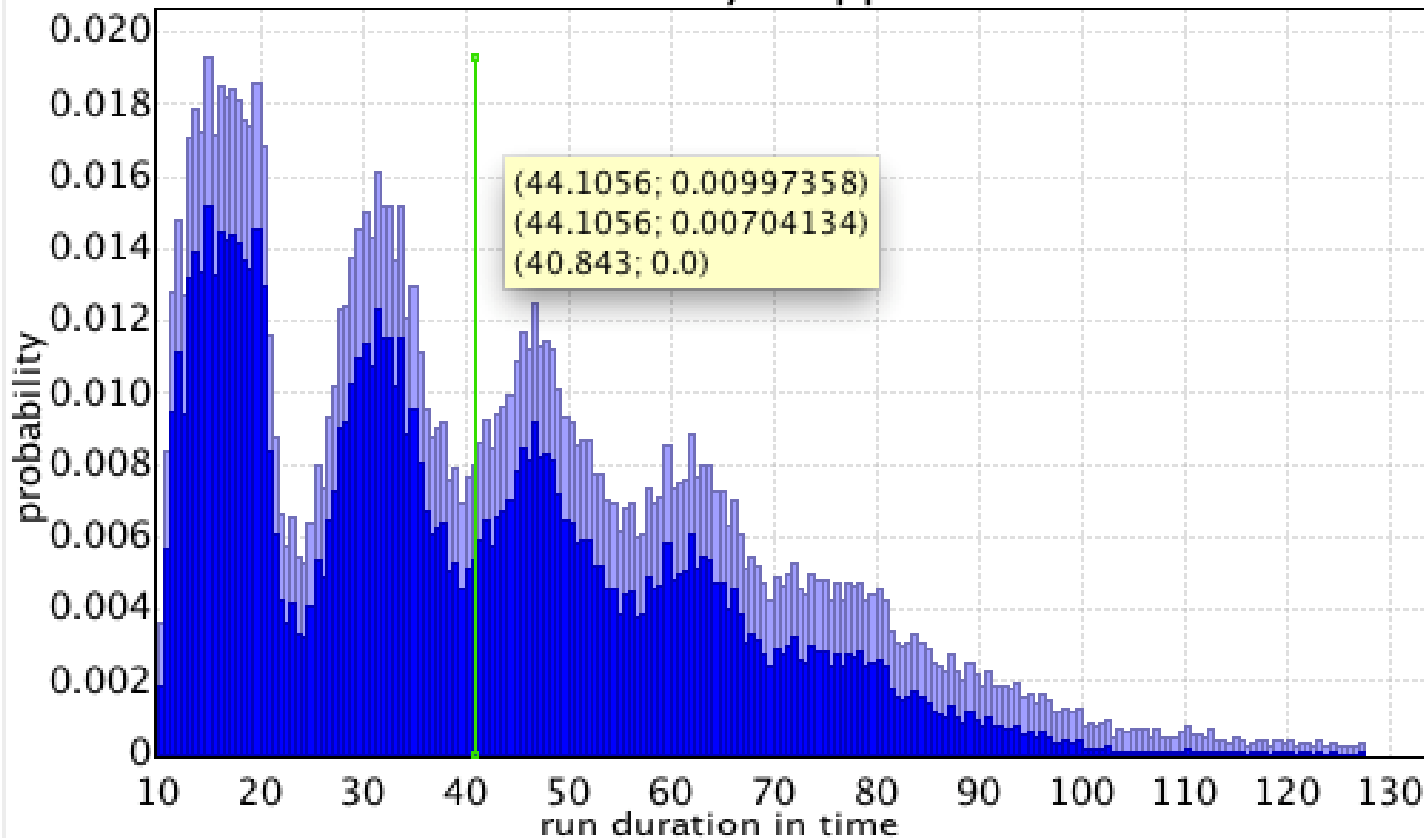
Queries in UPPAAL SMC

Pr[≤ 200]($\langle \rangle$ Train(5).Cross)

Message

Pr[≤ 200]($\langle \rangle$ Train(5).Cross)

Probability Clopper-Pearson CIs



Parameters: $\alpha=0.01$, $\epsilon=0.01$, bucket width=0.587972, bucket count=200.

Runs: 26492 in total, 26492 displayed, 0 remaining.

Probability sums: 1 displayed, 0 remaining.

Average: 40.843.

OK

x=0

Start
x \leq 15

?

Cross

Queries in UPPAAL SMC

$\text{Pr}[\leq 100] (\langle \rangle \text{Train}(0).\text{Cross}) \geq 0.8$

The screenshot displays the UPPAAL SMC interface. On the left, the 'Enabled Transitions' list includes 'appr[0]: Train(0) --> Gate'. The central area shows a list of variables: Gate.list[0] = 5, Gate.list[1] = 1, Gate.list[2] = 4, Gate.list[3] = 3, Gate.list[4] = 2, Gate.list[5] = 0, Gate.list[6] = 0, Gate.len = 5, Train(0).x >= 23, Train(1).x ∈ [13,60], and Train(2).x ∈ [0,15]. A 'Message' dialog box is open, displaying the results for 149 runs: 'Pr(⟨⟩ ...) <= 0.79 with confidence 0.99.' Below this, another 'Message' dialog box shows results for 651 runs: 'Pr(⟨⟩ ...) >= 0.51 with confidence 0.99.' The bottom of the interface features a 'Trace File' field, playback controls (Prev, Next, Replay, Open, Save, Random), and a speed slider from 'Slow' to 'Fast'. The bottom right corner shows a grid of train status indicators for Train(4) and Train(5), with variables like x >= 3 and leave[4]!

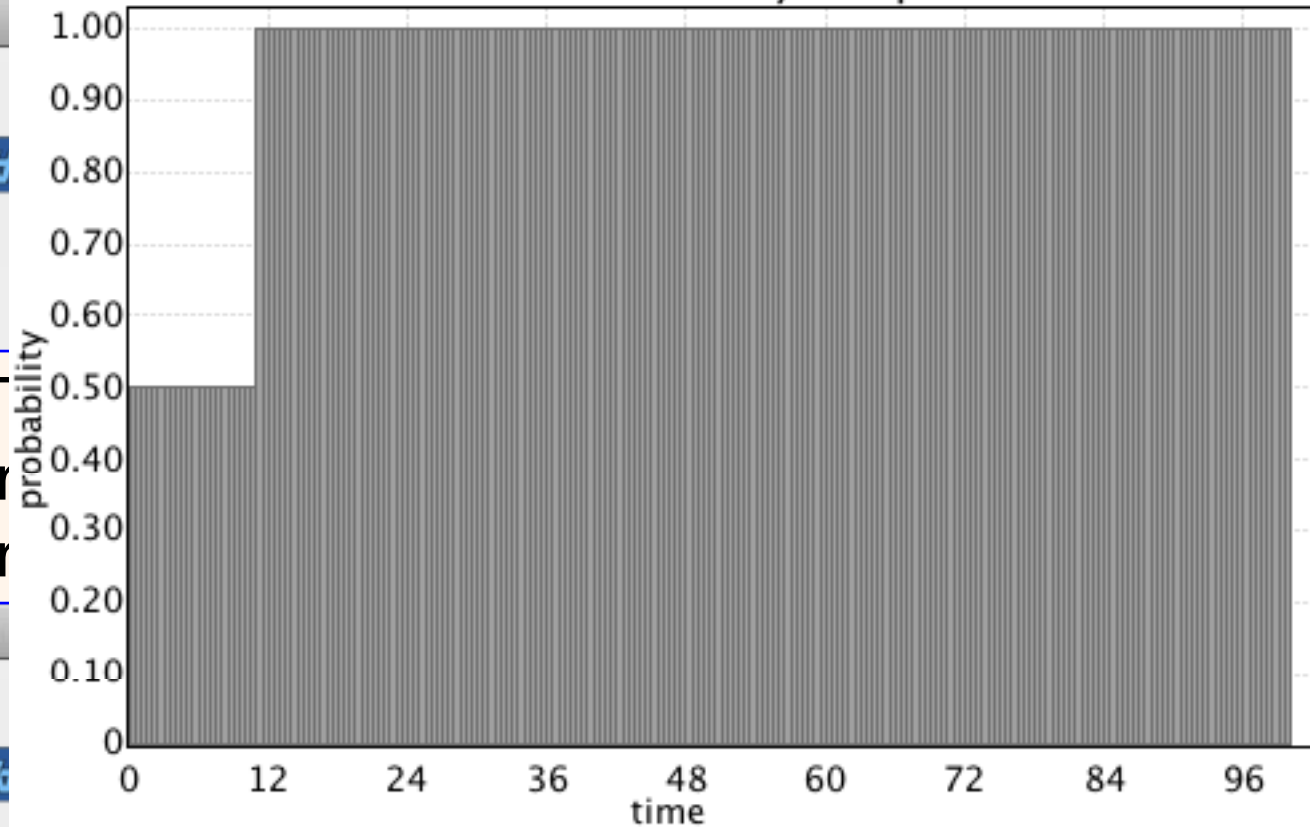
$\text{Pr}[\leq 100] (\langle \rangle \text{Train}(0).\text{Cross}) \geq 0.5$

Queries in UPPAAL SMC

$\Pr[\leq 100] (\langle \rangle \text{Train}(5).\text{Cross}) \geq$

$\Pr[\leq 100] (\langle \rangle \text{Train}(5).\text{Cross}) \geq \Pr[\leq 100] (\langle \rangle \text{Train}(1).\text{Cross})$

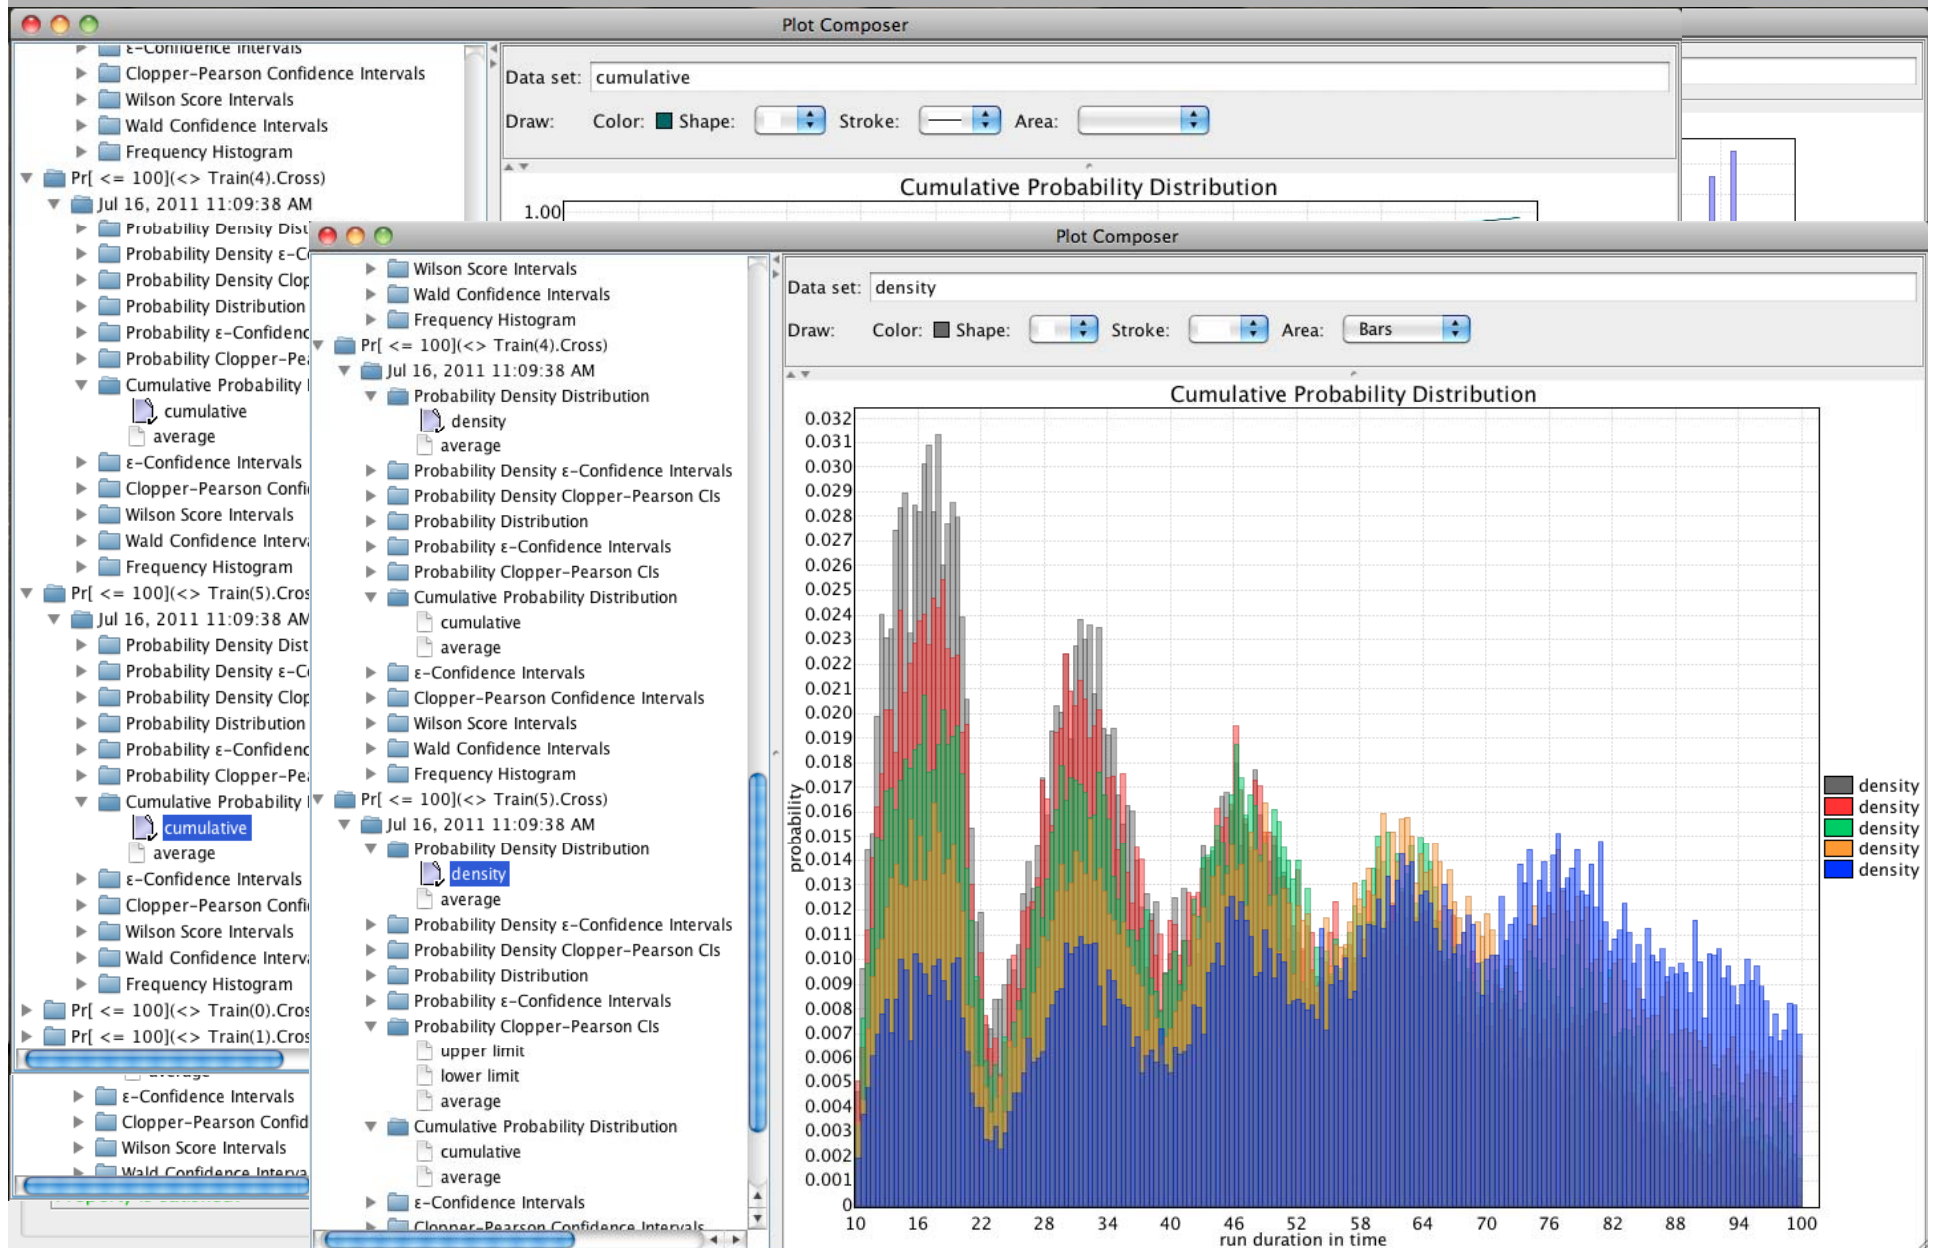
Probability comparison



value 0.0 means less-than is true.
value 0.5 means probabilities are indistinguishable.
value 1.0 means greater-than is true.



Analysis Tool: Plot Composer

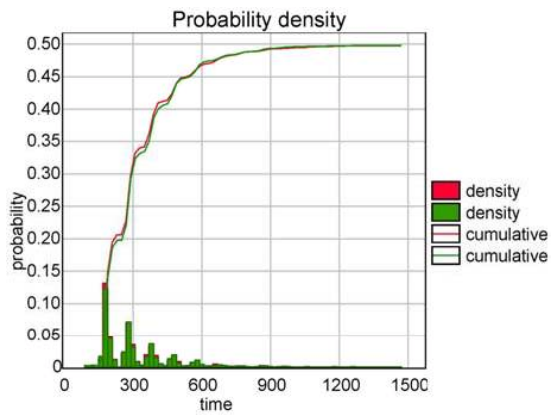


SMC in UPPAAL

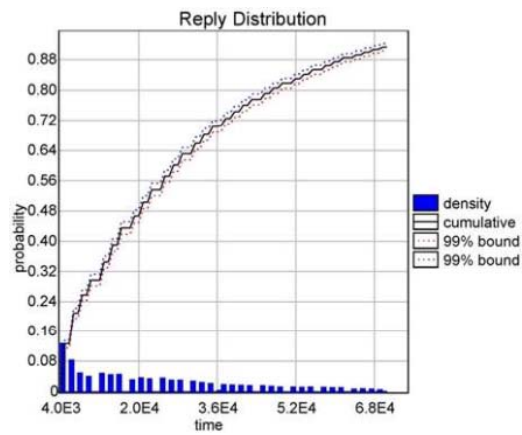
- Constant Slope Timed Automata
 - **Clocks** may have different (integer) **slope** in different locations.
 - **Branching edges** with discrete probabilities (weights).
 - **Beyond** Priced TA, Energy TA. Equal LHA in (non-stochastic) expressive power.
 - **Beyond** DTMC, beyond CTMC (with multiple rewards)
- All features of UPPAAL supported
 - User defined functions and types
 - Expressions in guards, invariants, clock-rates, delay-rates (rationals), and weights.
- New GUI for plot-composing and exporting.



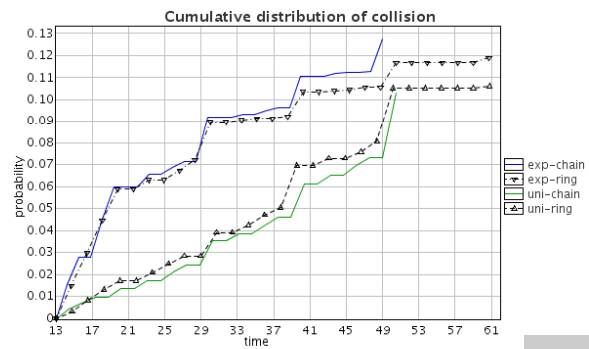
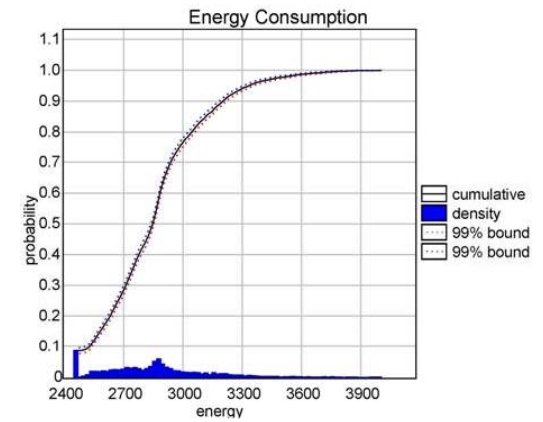
Case Studies



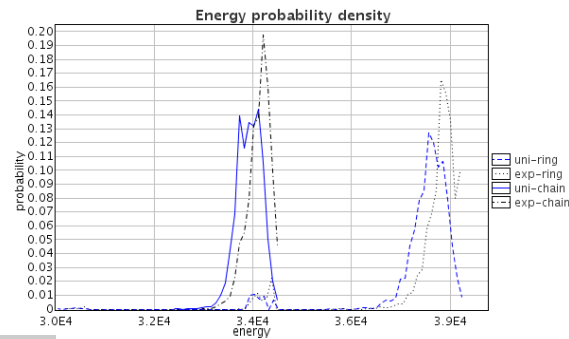
FIREWIRE



BLUETOOTH



LMAC



DPA

Benchmarking Duration Probabilistic Automata

		Param.			Estim.				Hyp. Testing			
		n	k	m	PRISM	Up_p	Up_d	Up_c	PRISM	Up_p	Up_d	Up_c
Fig res r_1	$\frac{n}{10}$	4	4	3	2.7	0.3	0.2	0.2	2.0	0.1	0.1	0.1
	$\frac{10}{10}$	6	6	3	7.7	0.6	0.5	0.4	3.9	0.2	0.2	0.3
	$\frac{10}{10}$	8	8	3	26.5	1.2	0.9	0.7	16.4	0.5	0.4	0.3
	$\frac{10}{10}$	20	40	20	>300				>300	35.5	26.2	20.7
	$\frac{10}{20}$	30	40	20	>300				>300	61.2	41.8	33.2
	$\frac{20}{20}$	40	40	20	>300				>300	92.2	56.9	59.5
	$\frac{20}{20}$	40	20	20	>300				>300	41.1	31.2	26.5
	$\frac{20}{20}$	40	30	20	>300				>300	68.8	46.7	46.1
	$\frac{20}{20}$	40	55	40	>300				>300	219.5		



Overview

- Statistical Model Checking in UPPAAL
 - Estimation
 - Testing
- **Distributed SMC for Parameterized Models**
 - Parameter Sweeps
 - Optimization
 - Nash Equilibria
- Distributing Statistical Model Checking
 - Estimation
 - Testing
- Parameter Analysis of DSMC
- Conclusion



UPPAAL & PDMC'05



PDMC'05
Gerd Behrman, Kim G Larsen



Architecture

GRID

Het. Cl

Hom. Cl

1-CPU

Fin
Time

Reach.

Liveness

Games

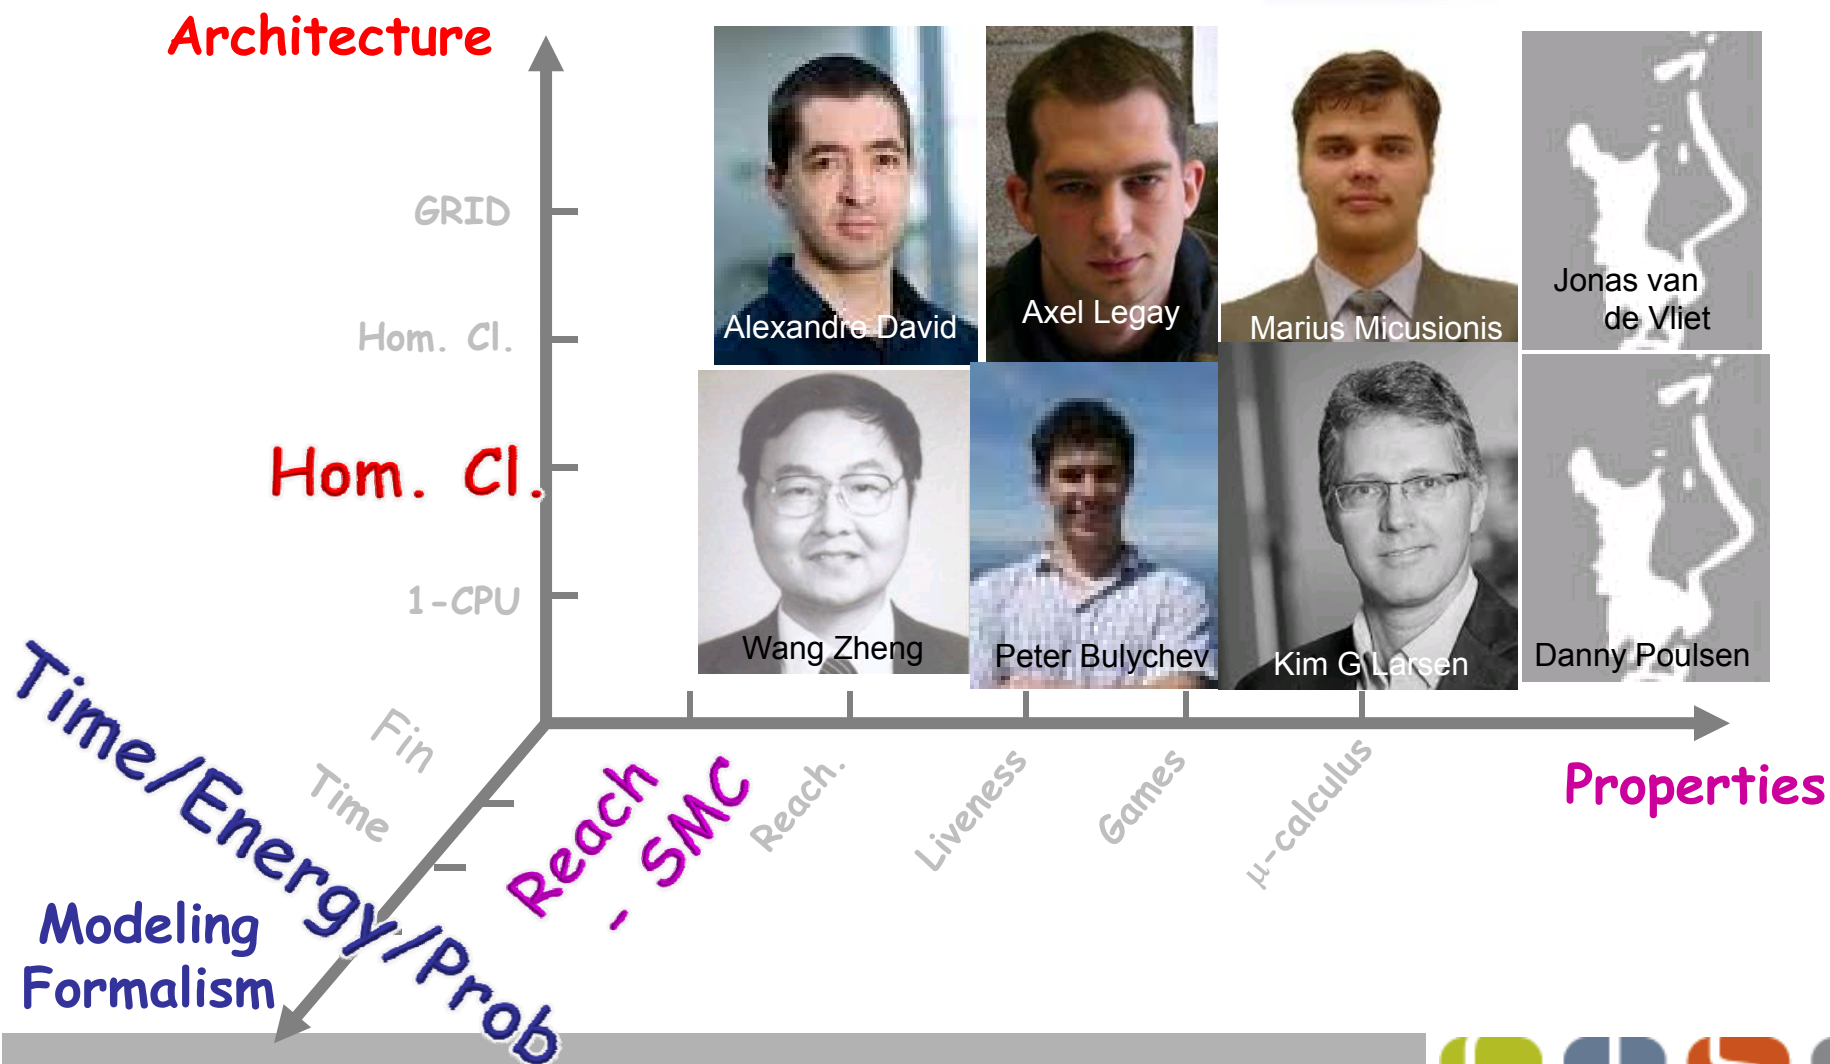
μ -calculus

Properties

Modeling
Formalism



UPPAAL & PDMC'11

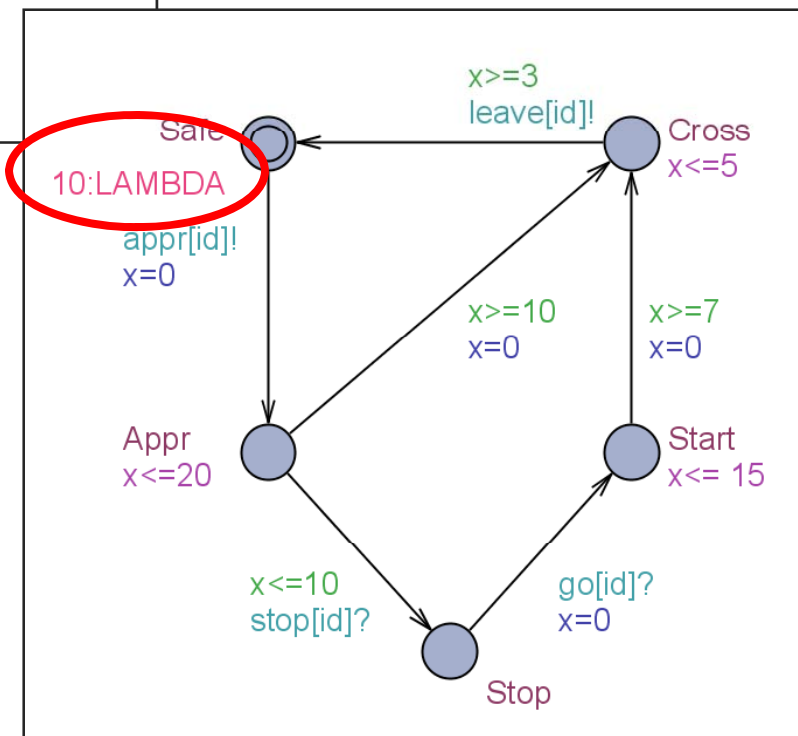


Parameterized Models in UPPAAL

```
const int N = #range(1, 10, ntrains);           // # trains
const int LAMBDA = 5*#range(1, 20, RATE);
typedef int [0,N-1] id_t;

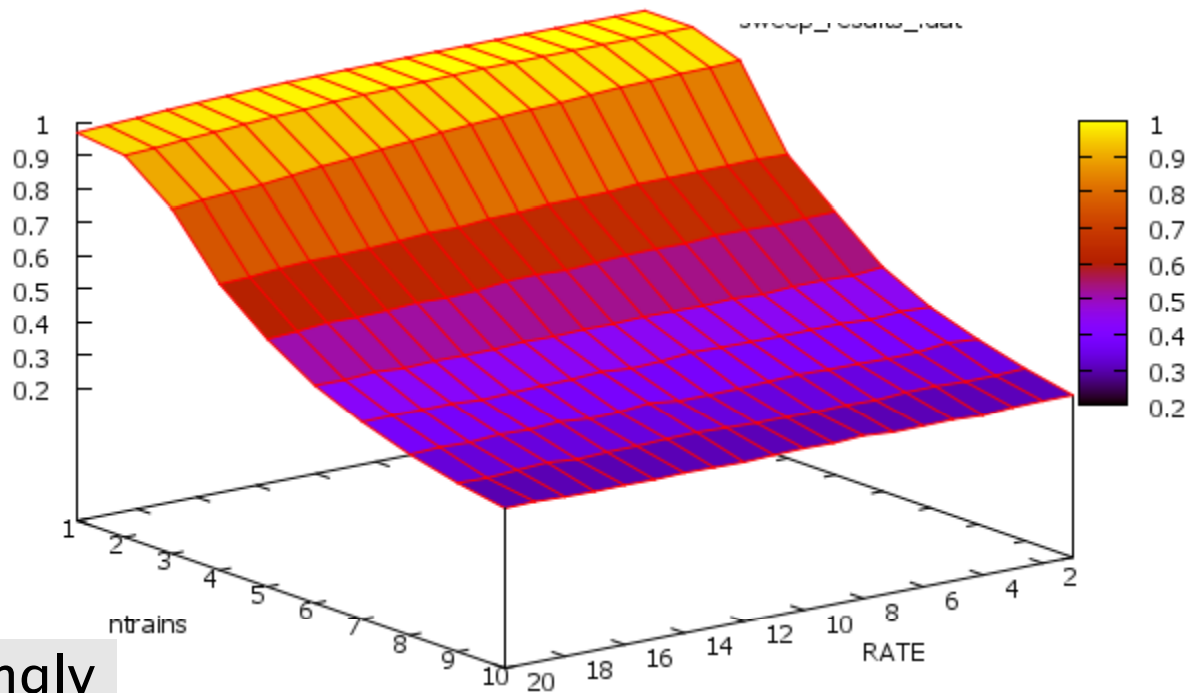
broadcast chan      appr[N], stop[N], leave[N];
urgent broadcast chan go[N];
clock time;
```

Extended Syntax
constants declared with a range
are treated as parameter



Parameterized Analysis of Trains

$\text{Pr}[\text{time} \leq 100](\langle \rangle \text{Train}(0).\text{Cross})$



“Embarrassingly
Parallelizable”



Lightweight Media Access Control

- Problem domain:
 - communication scheduling
- Targeted for:
 - self-configuring networks,
 - collision avoidance,
 - low power consumption
- Application domain:
 - wireless sensor networks
- **Initialization** (listen until a neighbor is heard)
- **Waiting** (delay a random amount of time frames)
- **Discovery** (wait for entire frame and note used slots)
- **Active**
 - choose free slot,
 - use it to transmit, including info about detected collisions
 - listen on other slots
 - fallback to Discovery if collision is detected
- Only neighbors can detect collision and tell the user-node that its slot is used by others



adopted from A.Fehnker, L.v.Hoesel, A.Mader

initialization

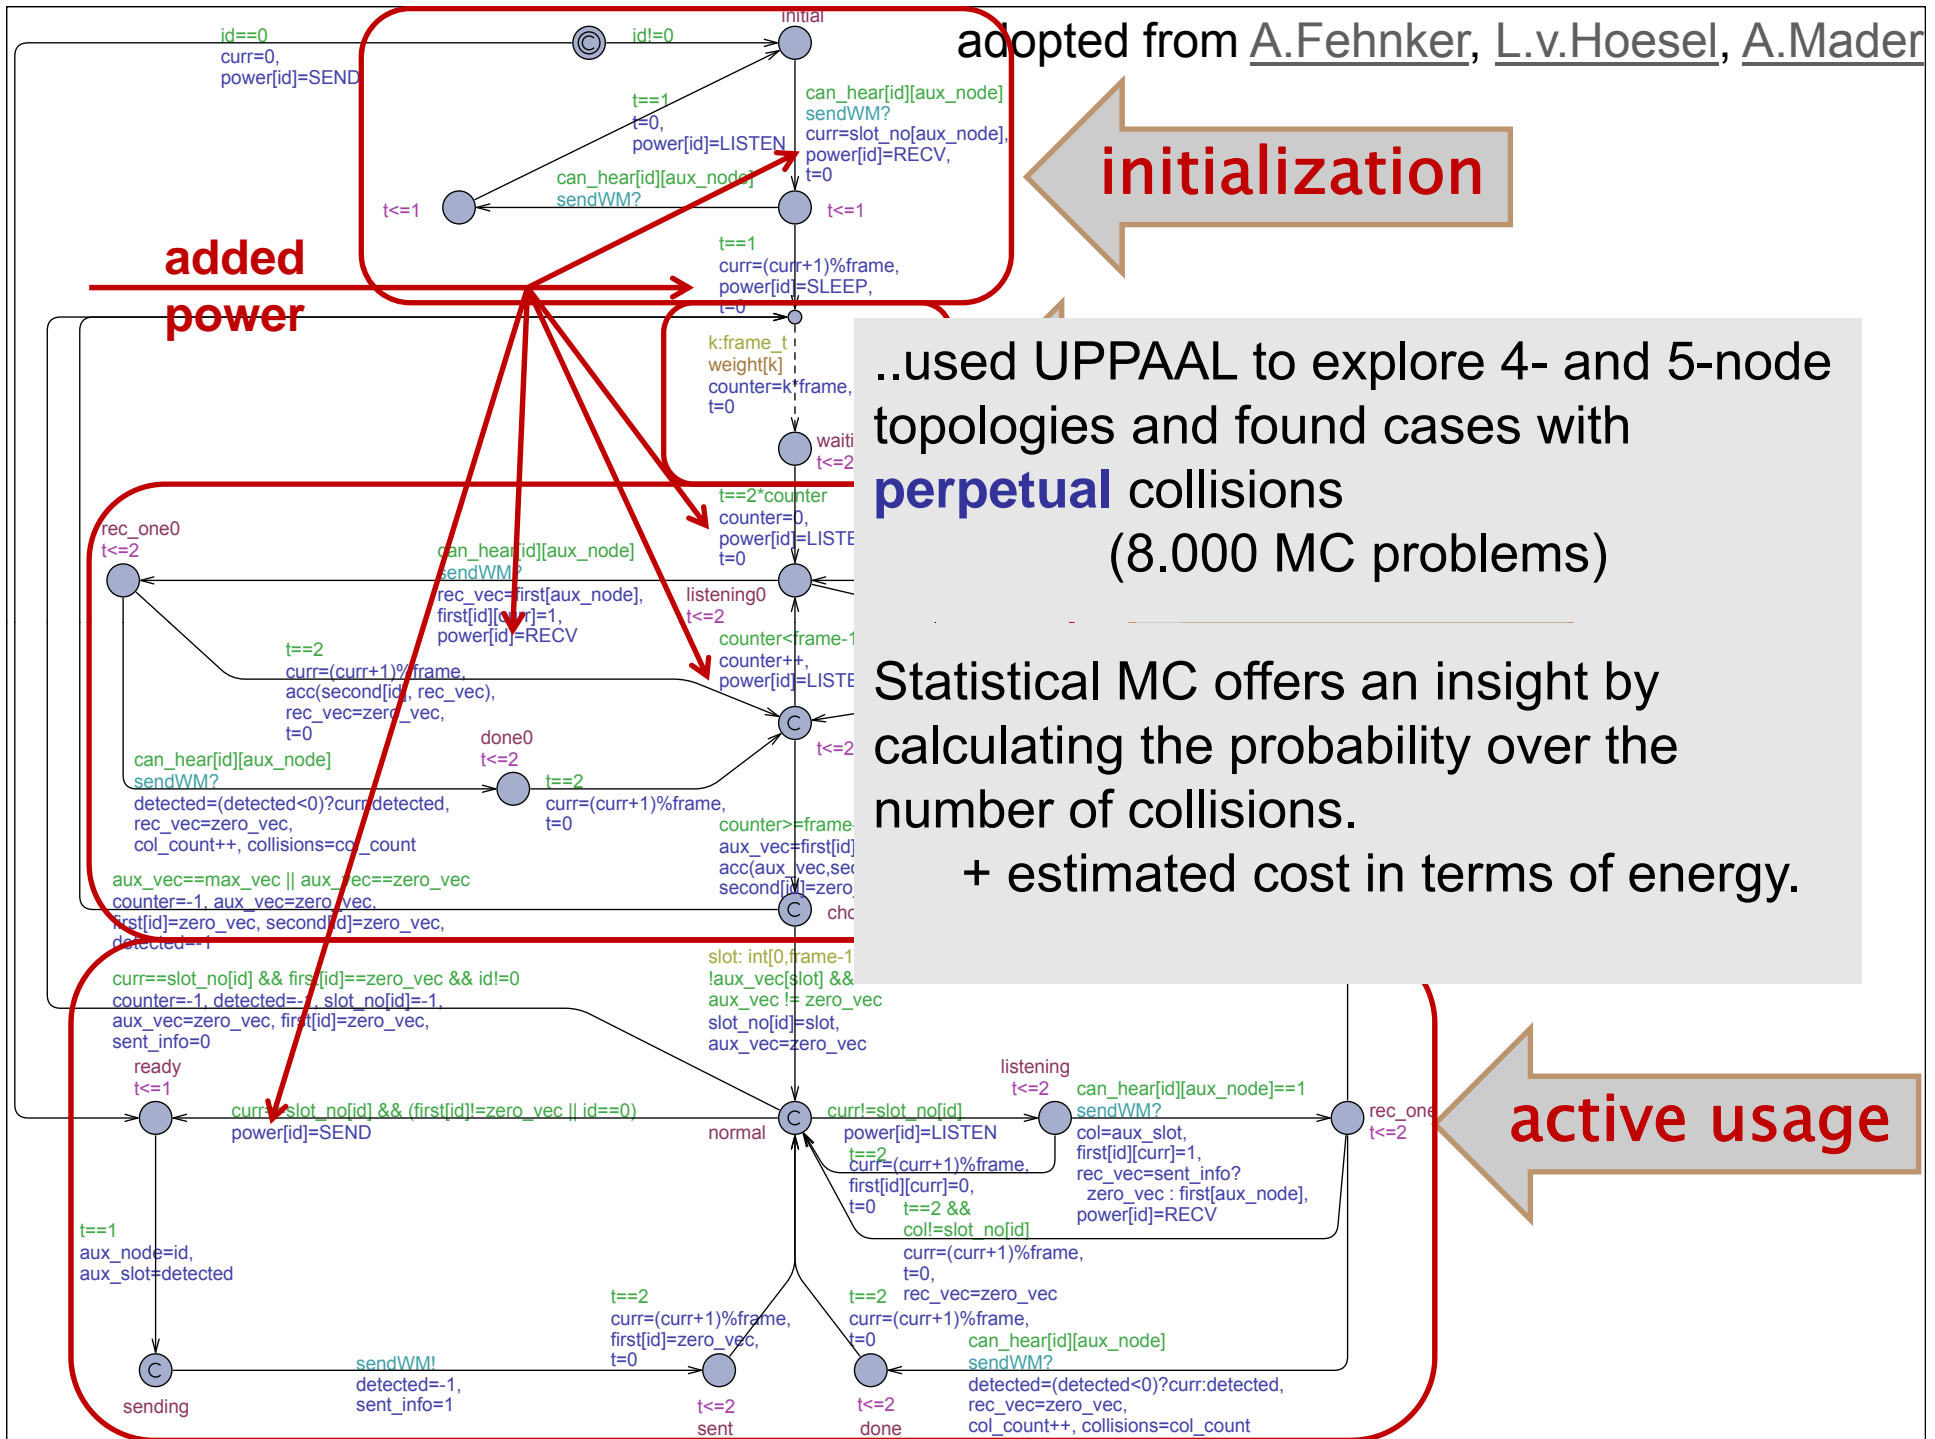
added power

..used UPPAAL to explore 4- and 5-node topologies and found cases with **perpetual collisions** (8.000 MC problems)

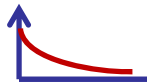


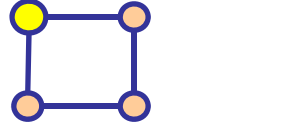
Statistical MC offers an insight by calculating the probability over the number of collisions.

+ estimated cost in terms of energy.

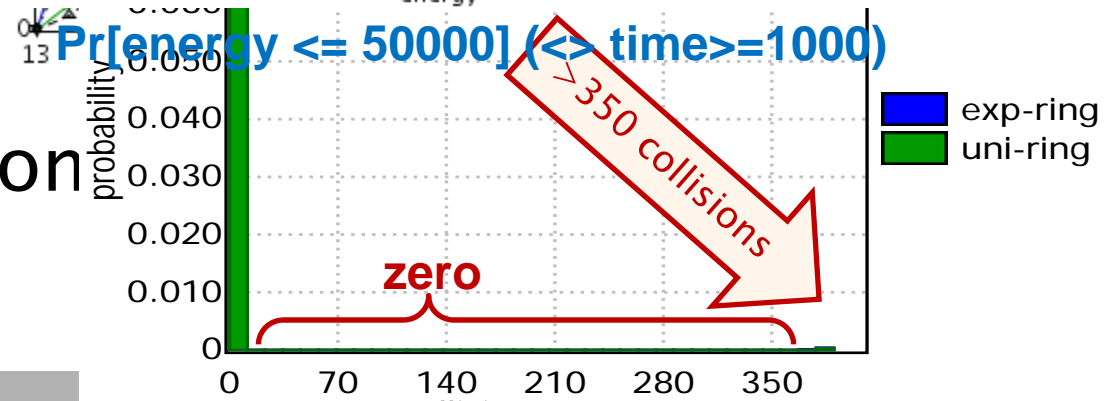
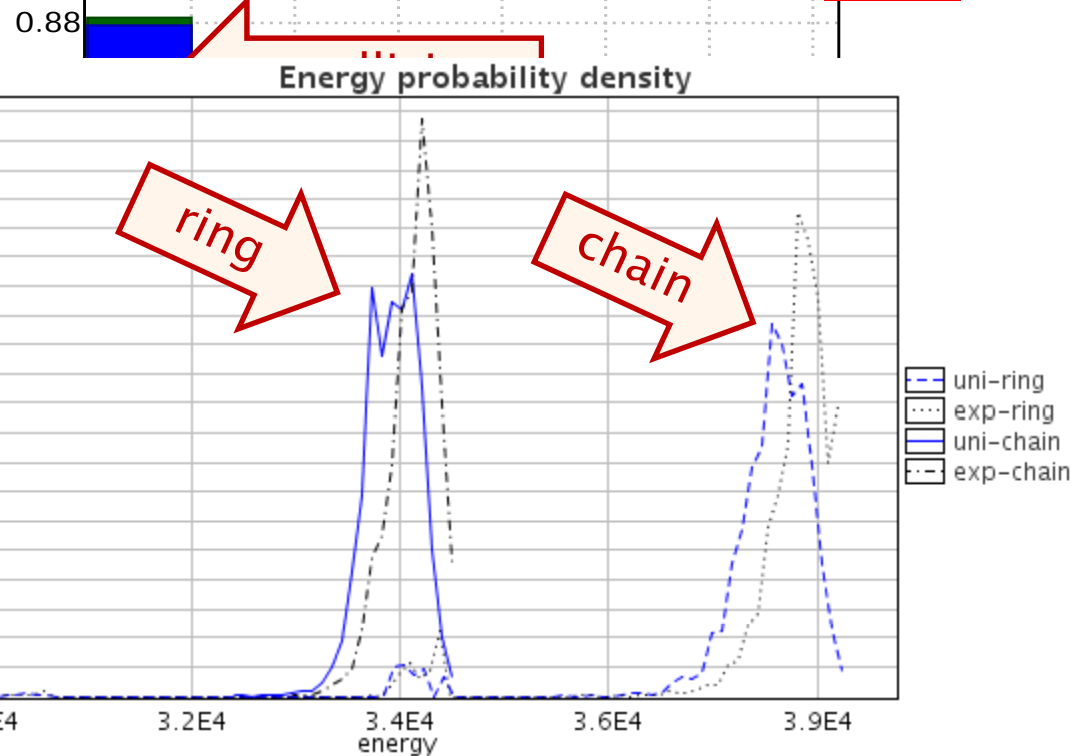
active usage



SMC of LMAC with 4 Nodes

- Wait distributio
 - geometric 
 - uniform 
- Network topolo
 - chain 
 - ring 
- Collision probal
- Collision count
- Power consumption

Probability density of Collision Count in a Chain



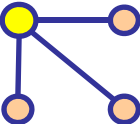
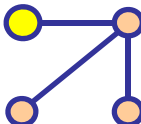
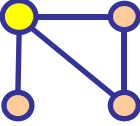
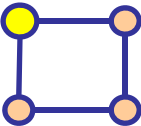
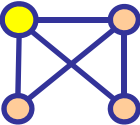
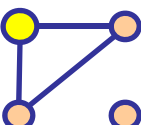
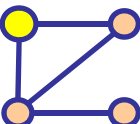
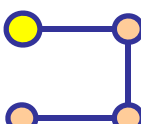
Kim Larsen [27]
 $\Pr[\text{energy} \leq 50000] (\langle \text{time} \rangle = 1000)$
 $\Pr[\text{collisions} \leq 50000] (\langle \text{time} \rangle = 1000)$

LMAC with Parameterized Topology

Distributed SMC

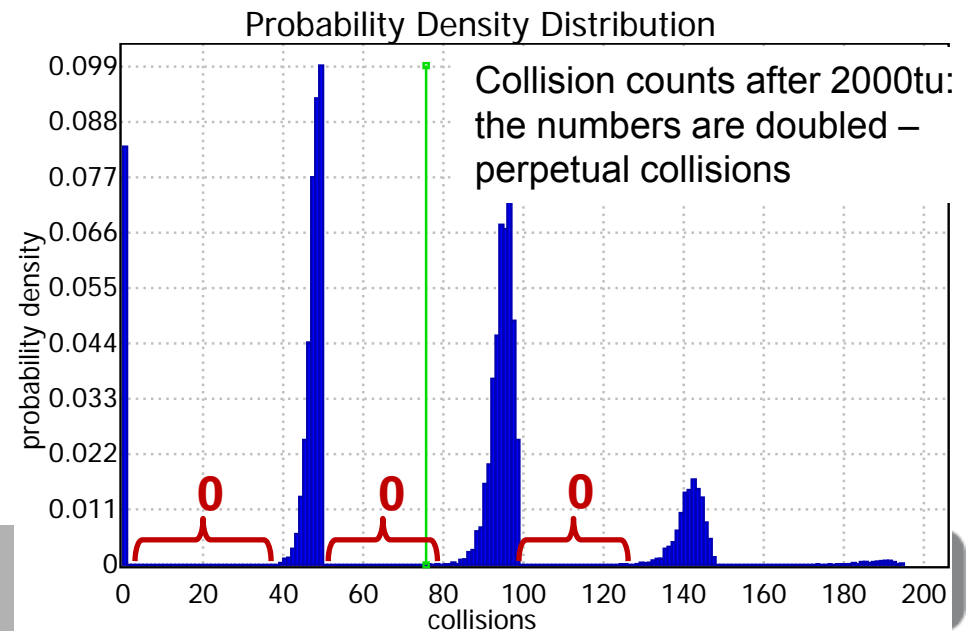
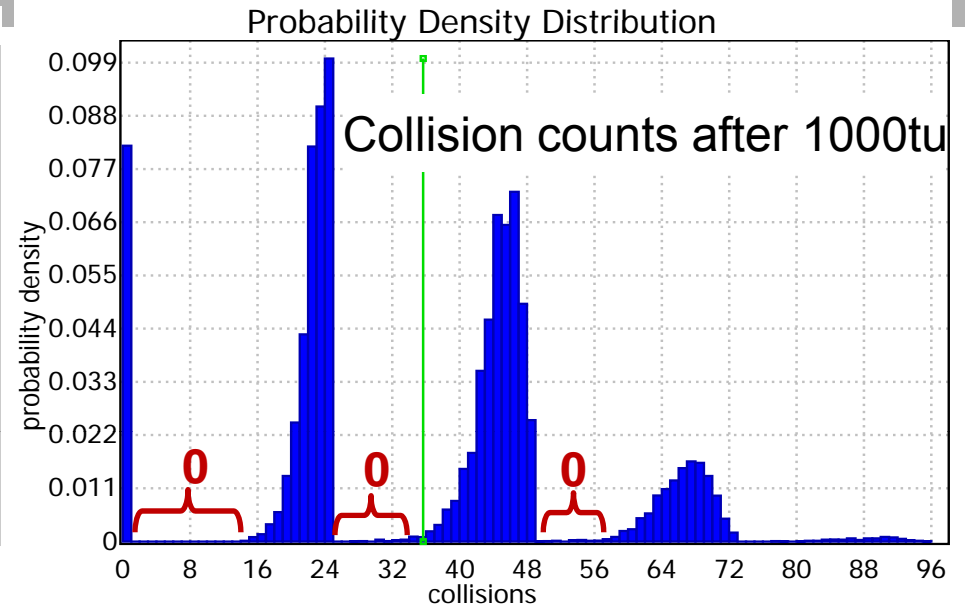
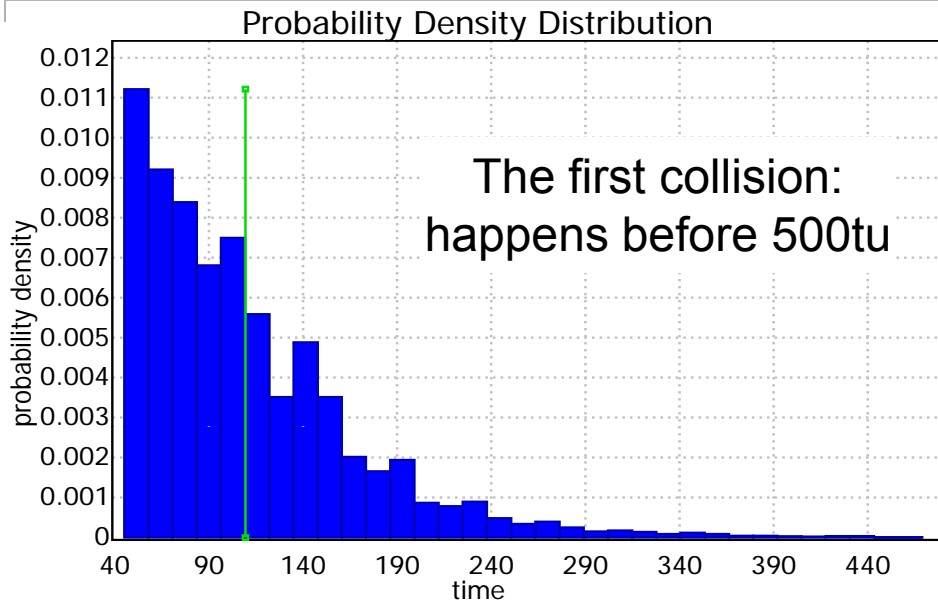
Collision probability in a 4 node network: sweep over all topologies.
32 core cluster: - 8xIntel Core2 2.66GHz CPU

$\Pr[\text{time} \leq 200] (\Leftrightarrow \text{col_count} > 0)$

topology	collision probability	topology	collision probability
(star)	 [0.36; 0.39]	 [0.08; 0.19]	
 [0.29; 0.36]	(ring)	 [0.11; 0.13]	
 [0.26; 0.30]	 [0.08; 0.15]		
 [0.19; 0.21]	(chain)	 [0.049; 0.050]	



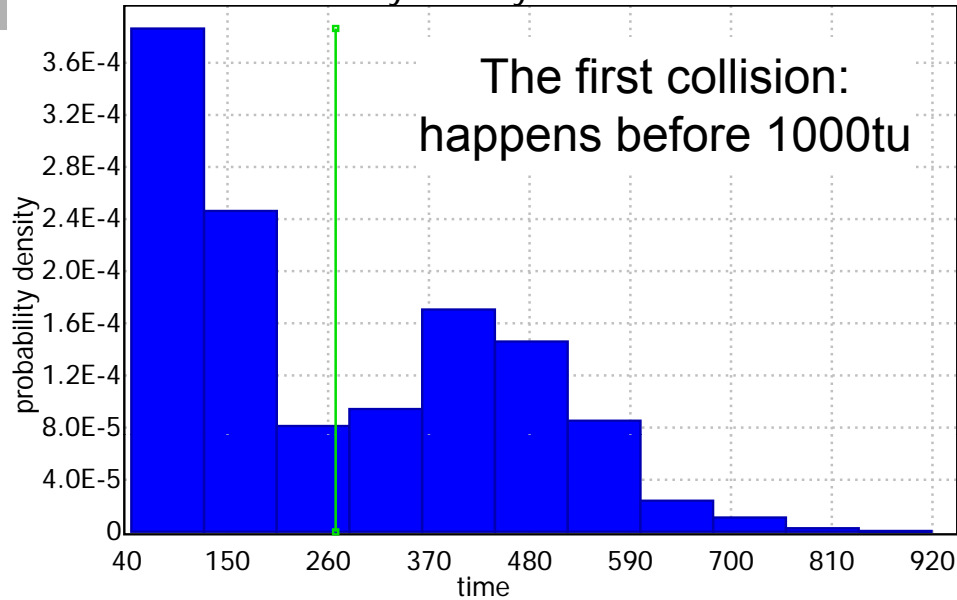
10-Node Star



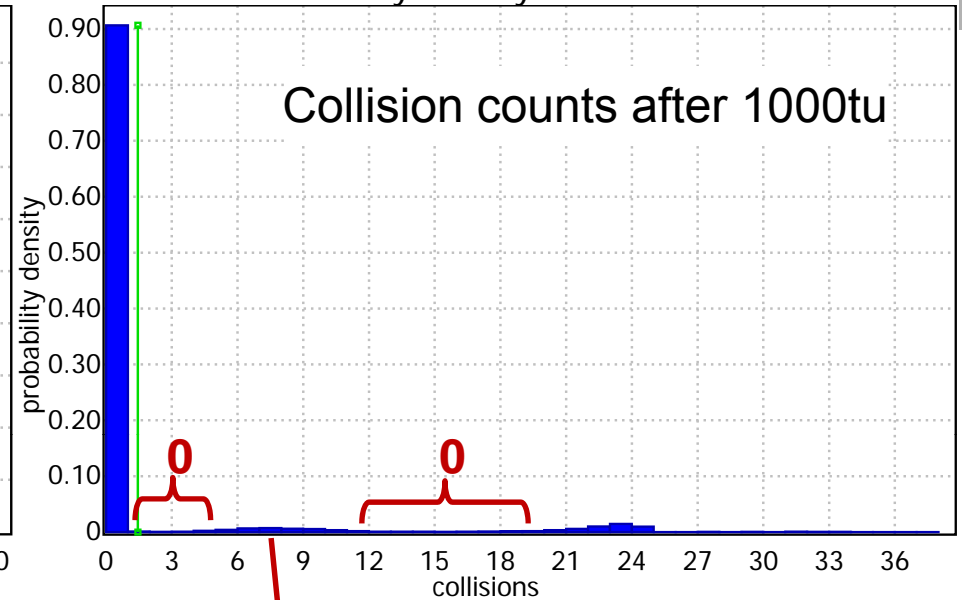
- The first collisions happen before **500tu**.
- It is unlikely (**8.2%**) that there will be **0** collisions.
- And if they happen, they are perpetual.

10-Node Ring

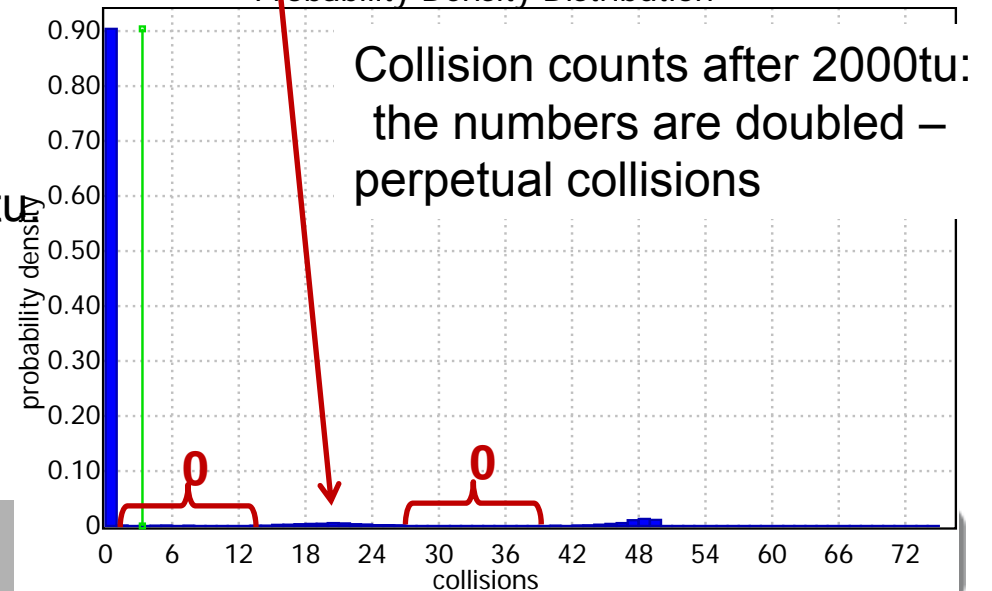
Probability Density Distribution



Probability Density Distribution

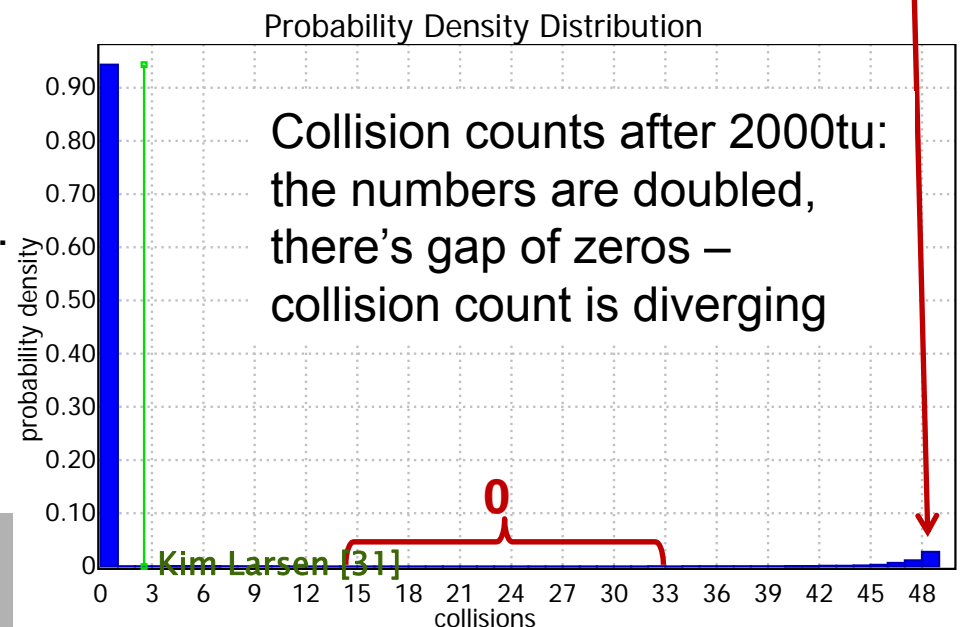
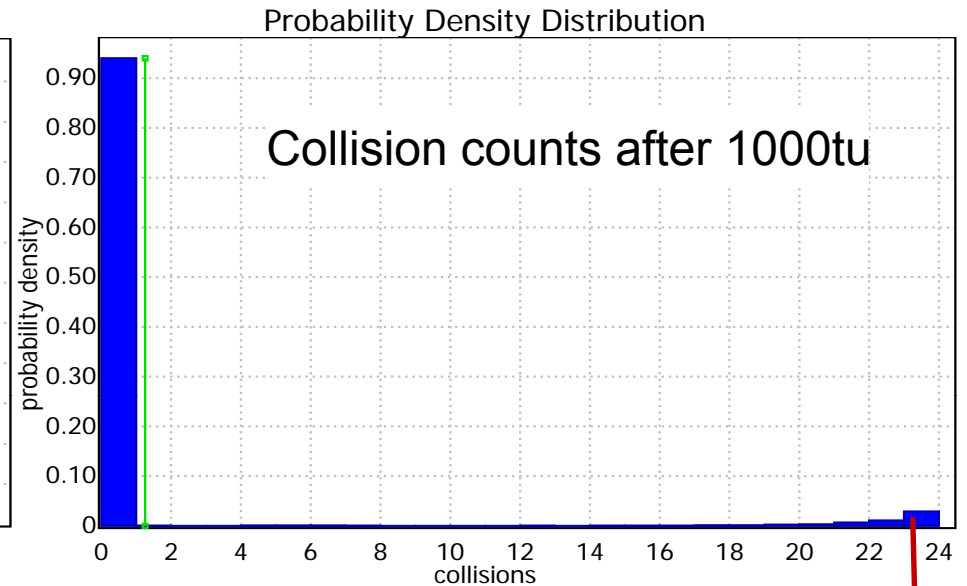
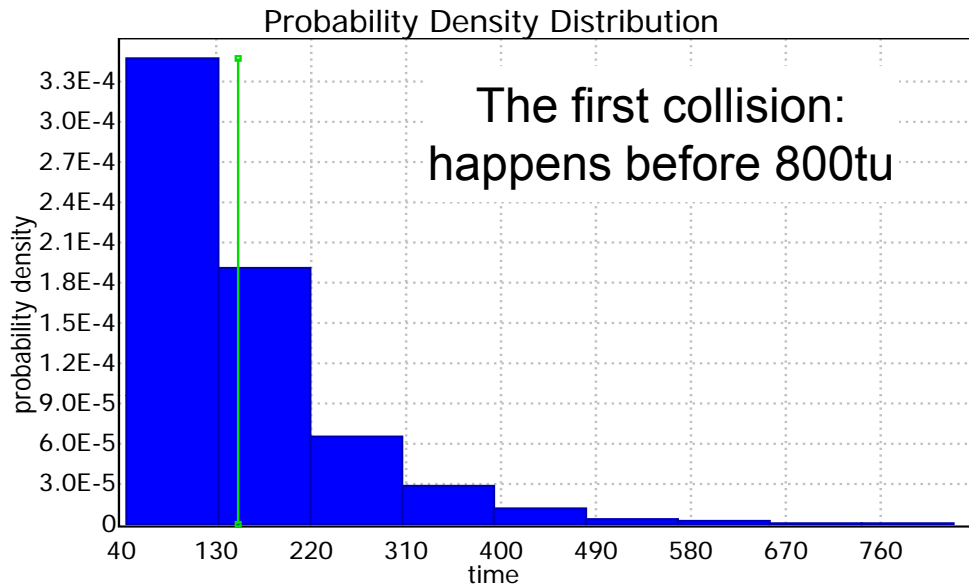


Probability Density Distribution



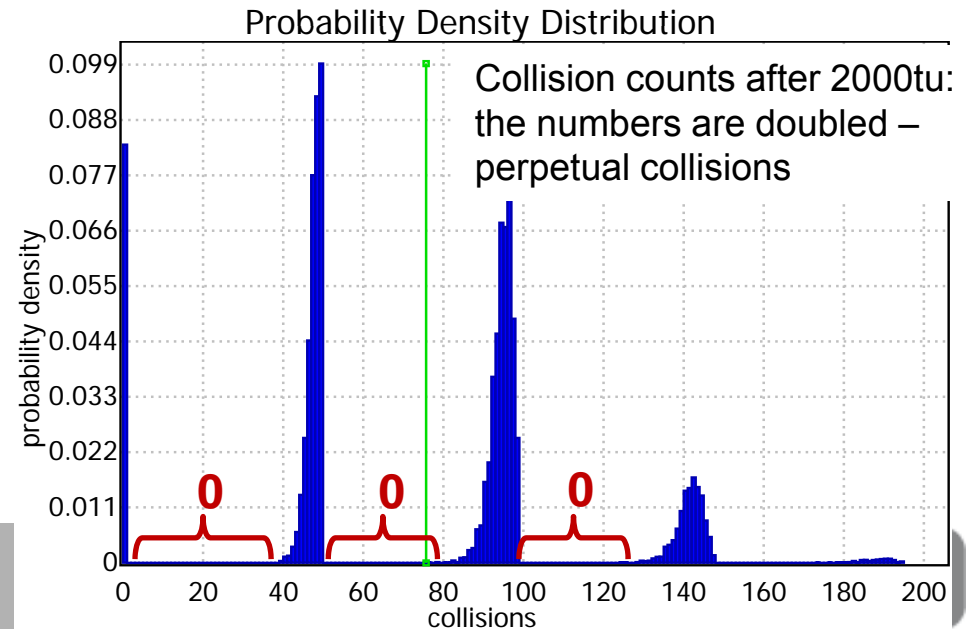
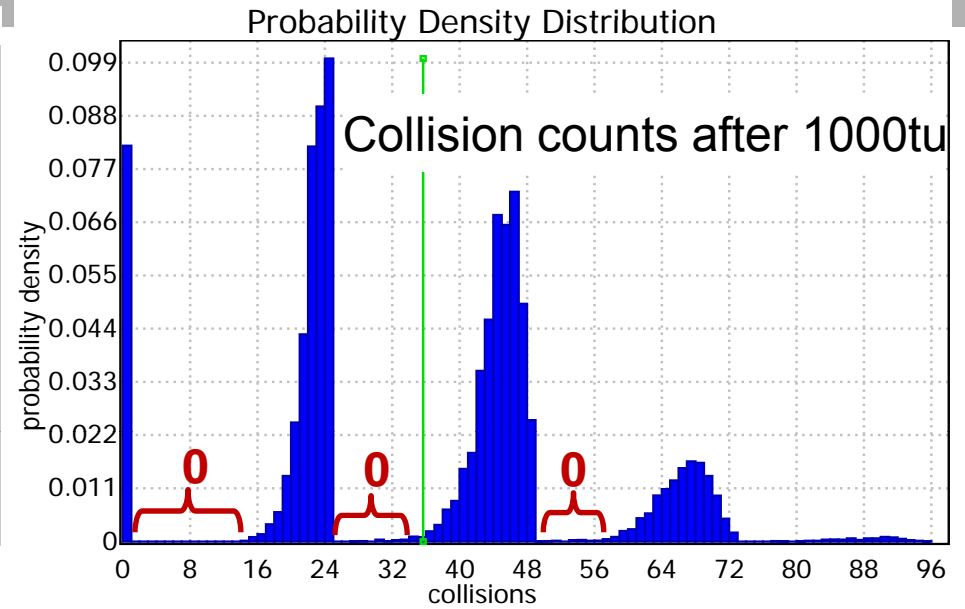
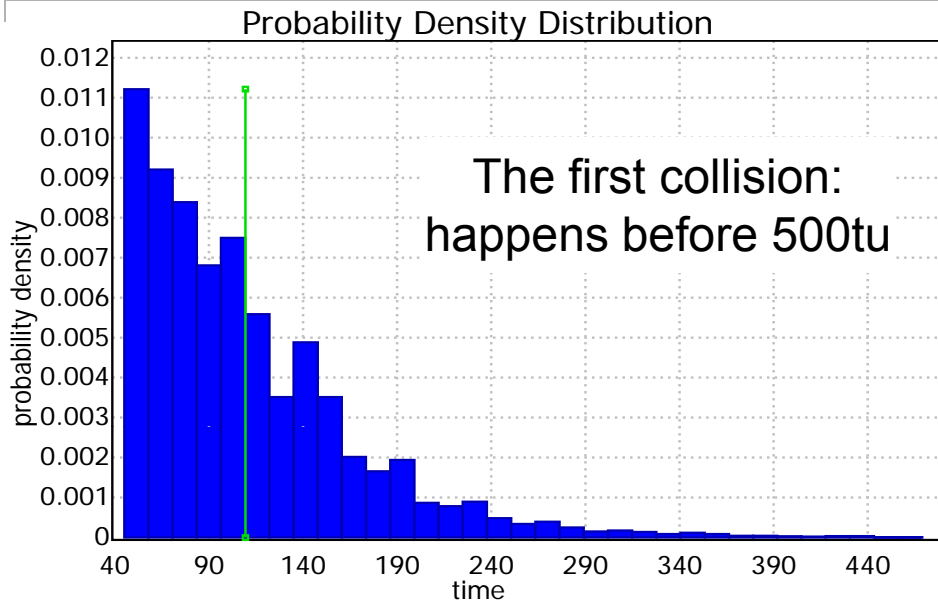
- The first collisions can be as late as 920tu
- It is very likely (>90%) that there will be 0 collisions.
- But if they happen, they are perpetual.

10-Node Chain



- The first collisions can be as late as **800tu**.
- It is very likely (**>94%**) that there will be **0** collisions.
- But if they happen, some are perpetual.

10-Node Star

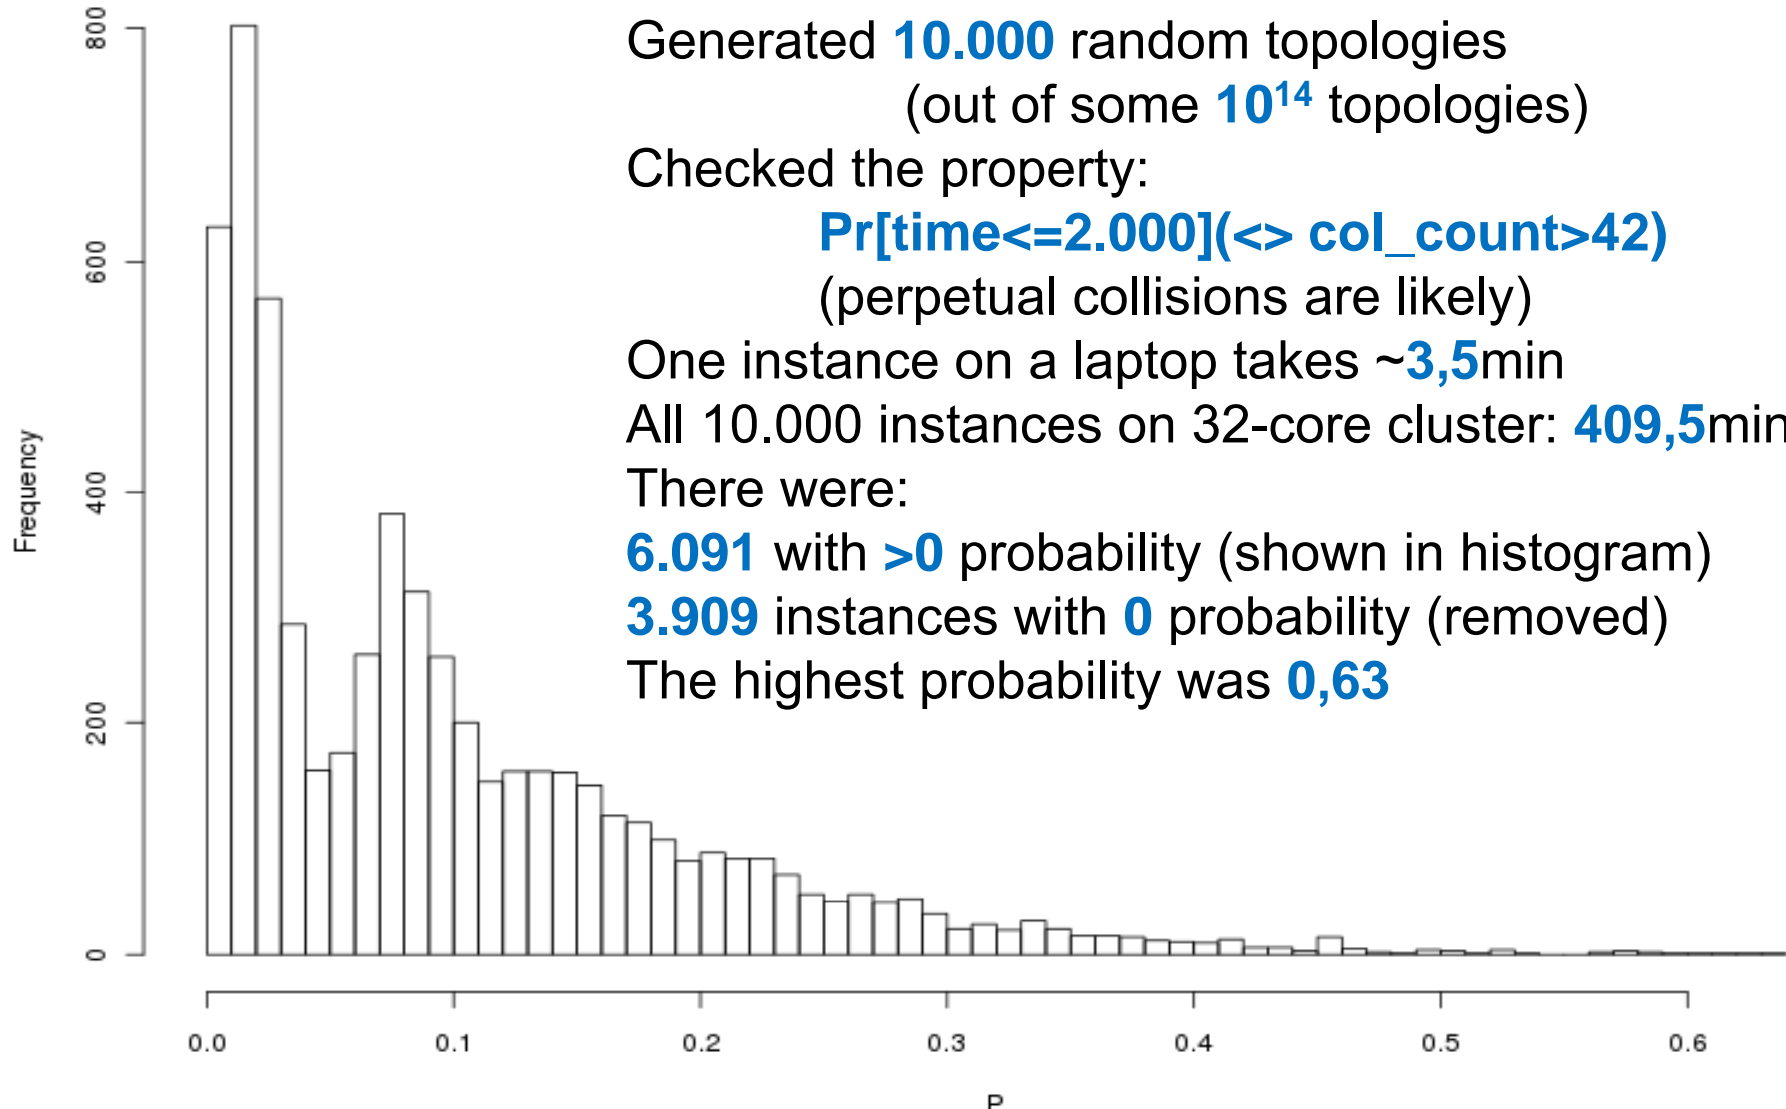


- The first collisions happen before **500tu**.
- It is unlikely (**8.2%**) that there will be **0** collisions.
- And if they happen, they are perpetual.

10-Node Random Topologies

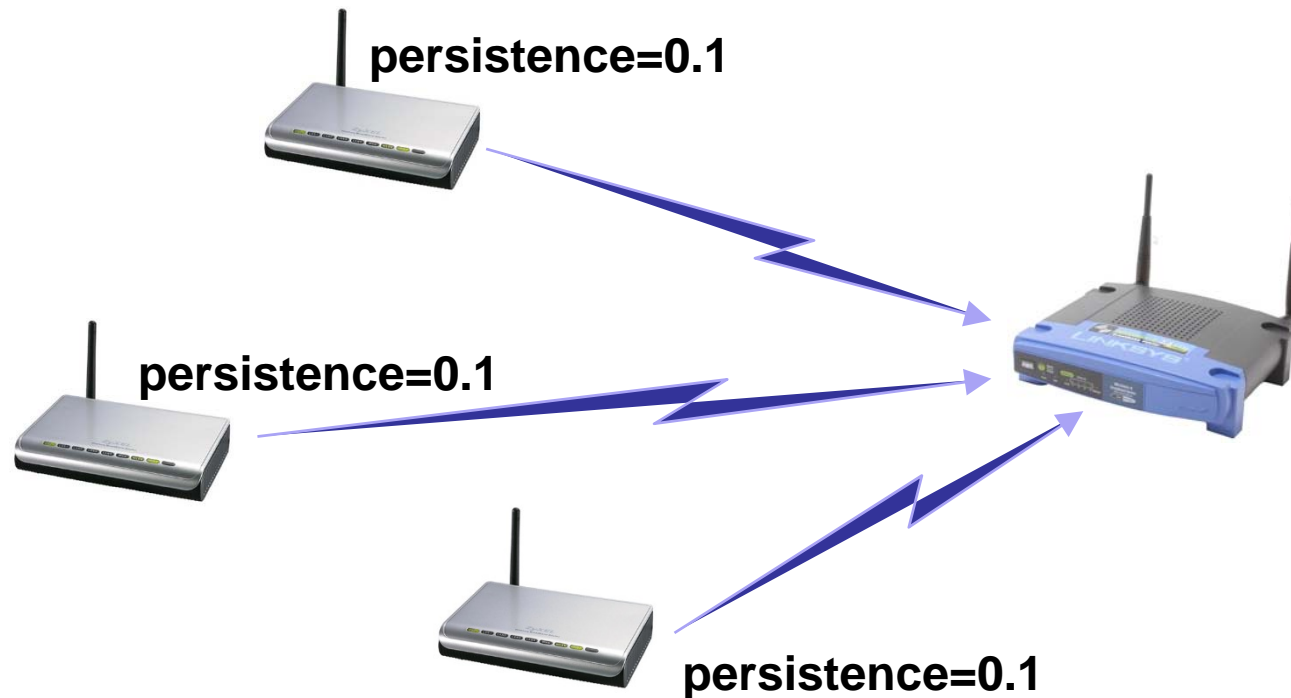
Distributed SMC

Histogram of P



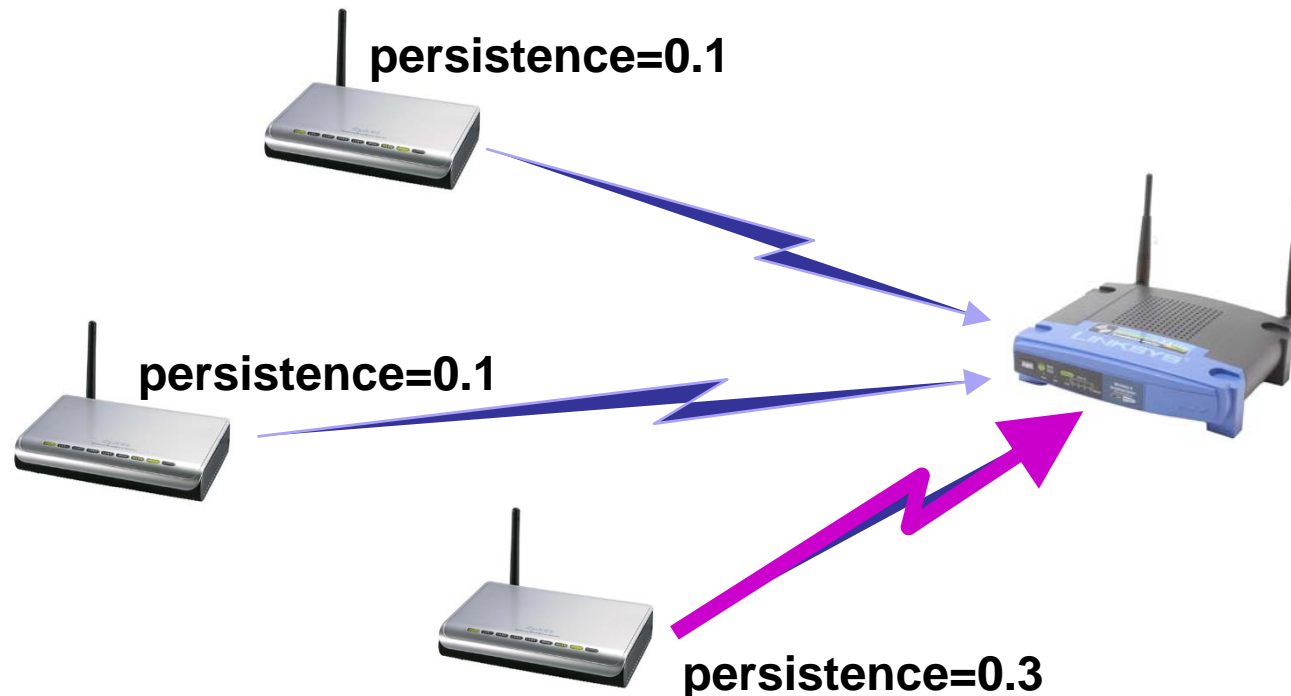
Nash Eq in Wireless Ad Hoc Networks

Consider a wireless network, where there are nodes that can independently adapt their parameters to achieve better performance



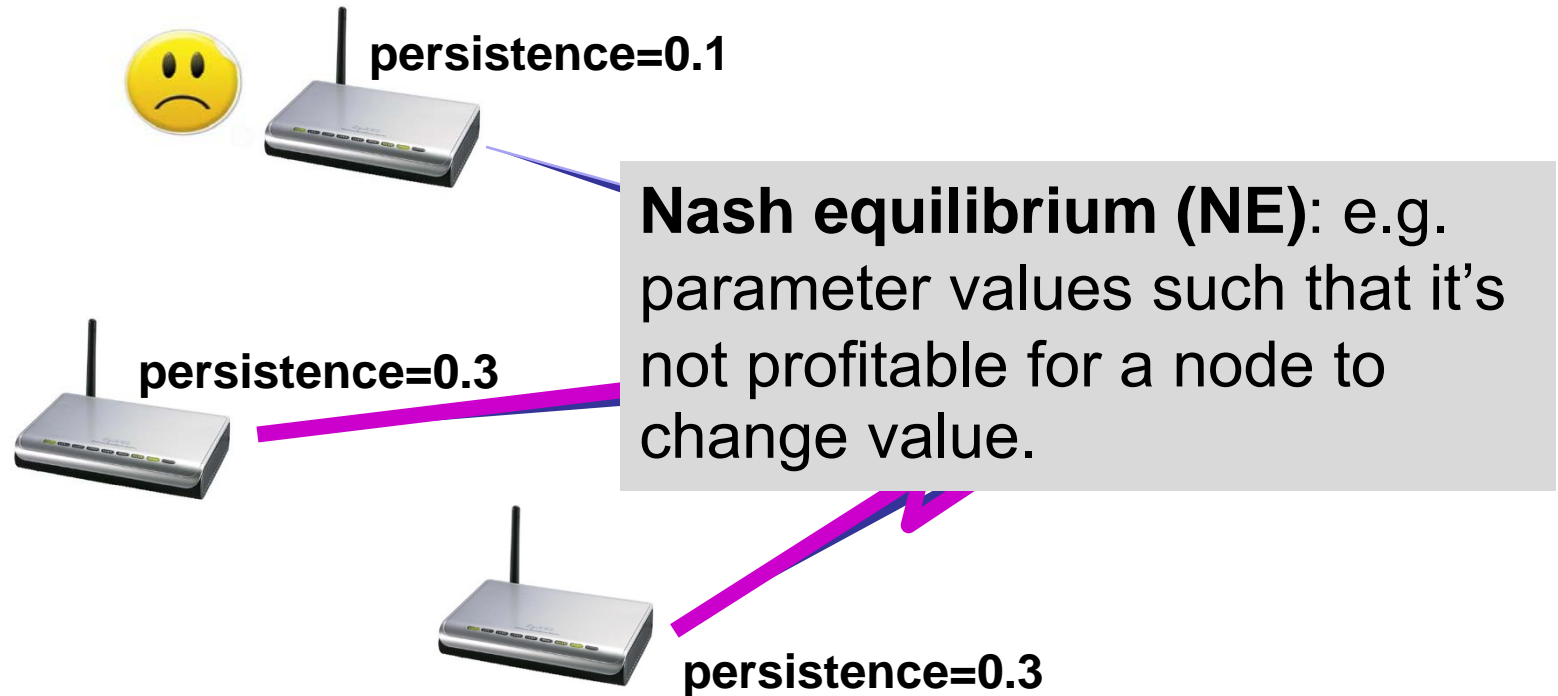
Nash Eq in Wireless Ad Hoc Networks

Consider a wireless network, where there are nodes that can independently adapt their parameters to achieve better performance

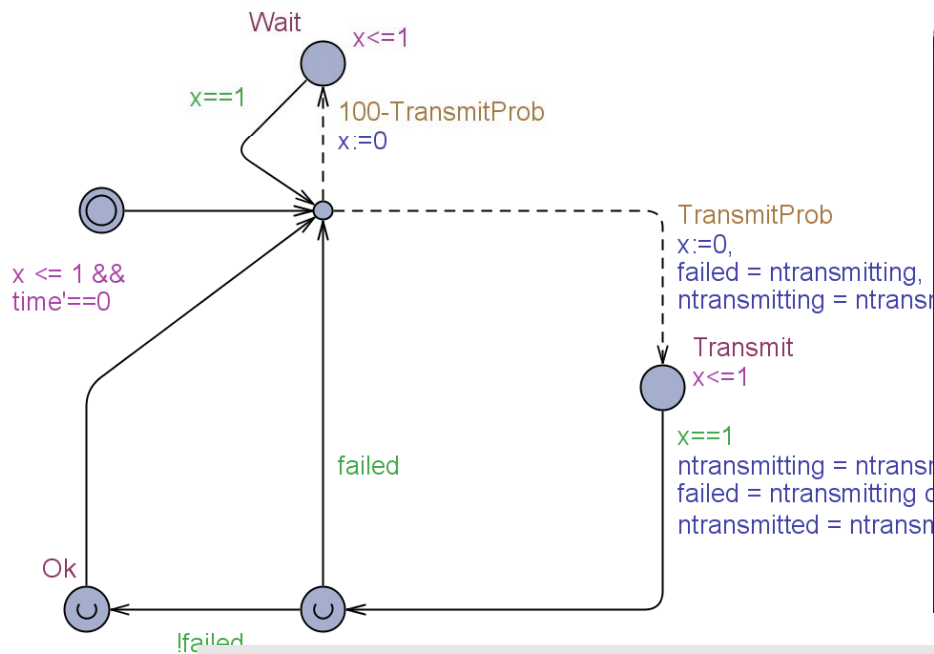


Nash Eq in Wireless Ad Hoc Networks

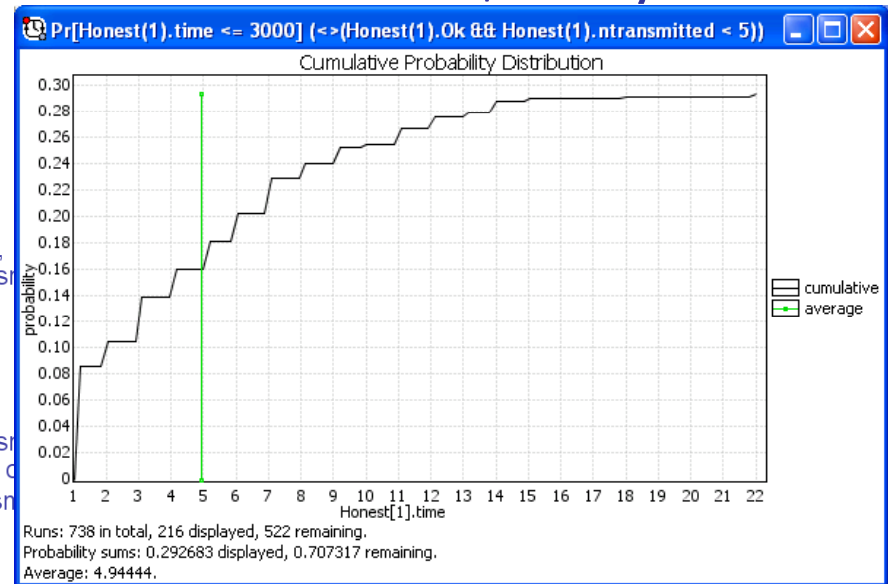
Consider a wireless network, where there are nodes that can independently adapt their parameters to achieve better performance



Aloha CSMA/CD protocol



TransmitProb=0.5 / Utility=0.29



Pr[Node.time <= 3000](<>(Node.Ok && Node.ntransmitted <= 5))

- Simple random access protocol (based on p-persistent ALOHA)
 - several nodes sharing the same wireless medium
 - each node has always data to send, and it sends data after a random delay
 - delay geometrically distributed with parameter $p = \text{TransmitProb}$



Distributed Algorithm for Computing Nash Equilibrium

Input: $S=\{s_i\}$ – finite set of strategies, $U(s_i, s_k)$ – utility function

Goal: find s_i s.t. $\forall s_k U(s_i, s_i) \geq U(s_i, s_k)$, where $s_i, s_k \in S$

Algorithm:

1. **for every** $s_i \in S$ **compute** $U(s_i, s_i)$
2. candidates := S
3. **while** $\text{len}(\text{candidates}) > 1$:
4. **pick** some unexplored pair $(s_i, s_k) \in \text{candidates} \times S$
5. **compute** $U(s_i, s_k)$
6. **if** $U(s_i, s_k) > U(s_i, s_i)$:
7. **remove** s_i from candidates
8. **if** $\forall s_k U(s_i, s_k)$ is already computed:
9. **return** s_i

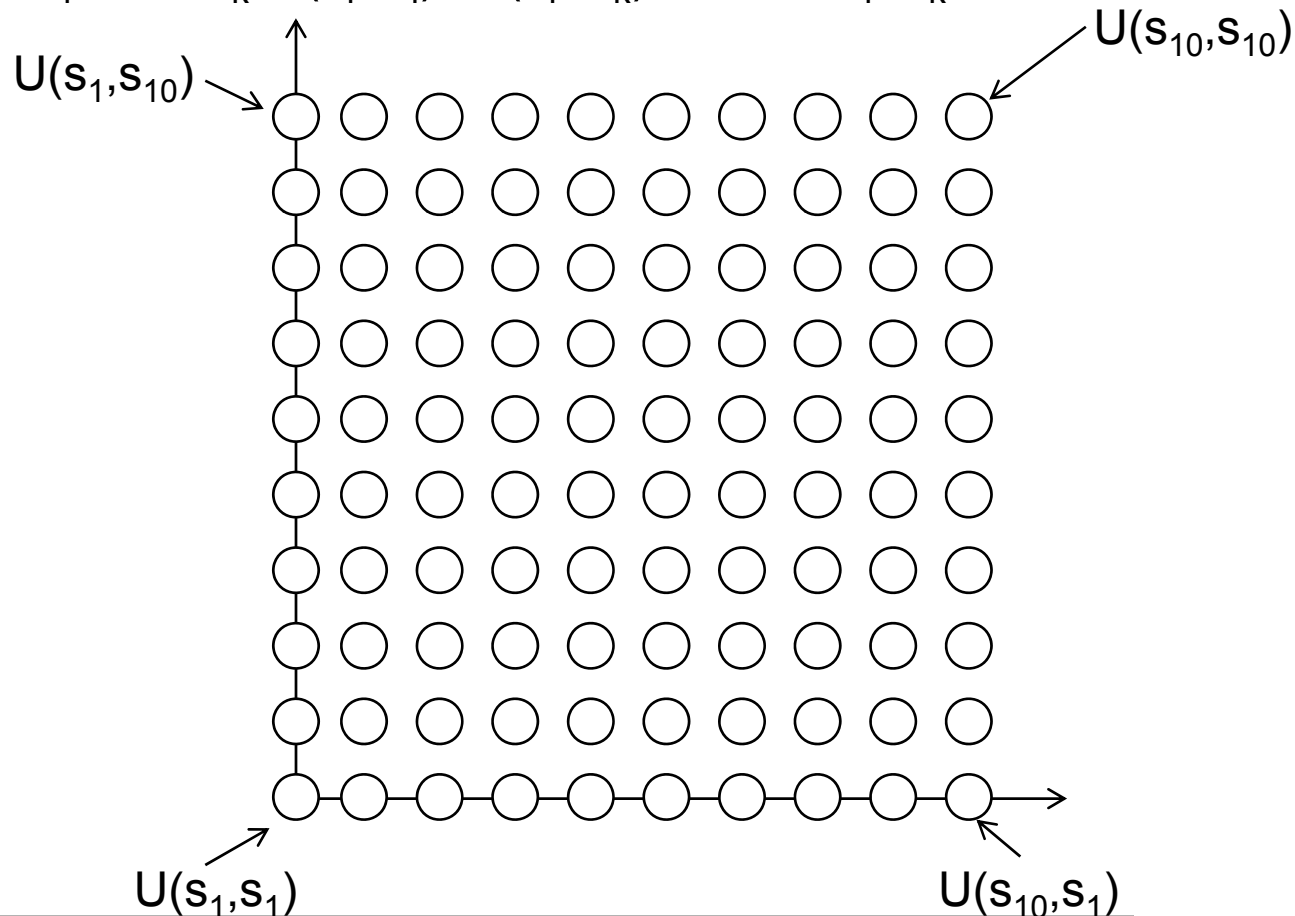
We can apply statistics to *prove* that (s_i, s_i) satisfies Nash equilibrium



Distributed algorithm for computing Nash equilibrium

Input: $S = \{s_1, s_2, \dots, s_{10}\}$ – finite set of strategies, $U(s_i, s_j)$ – utility function

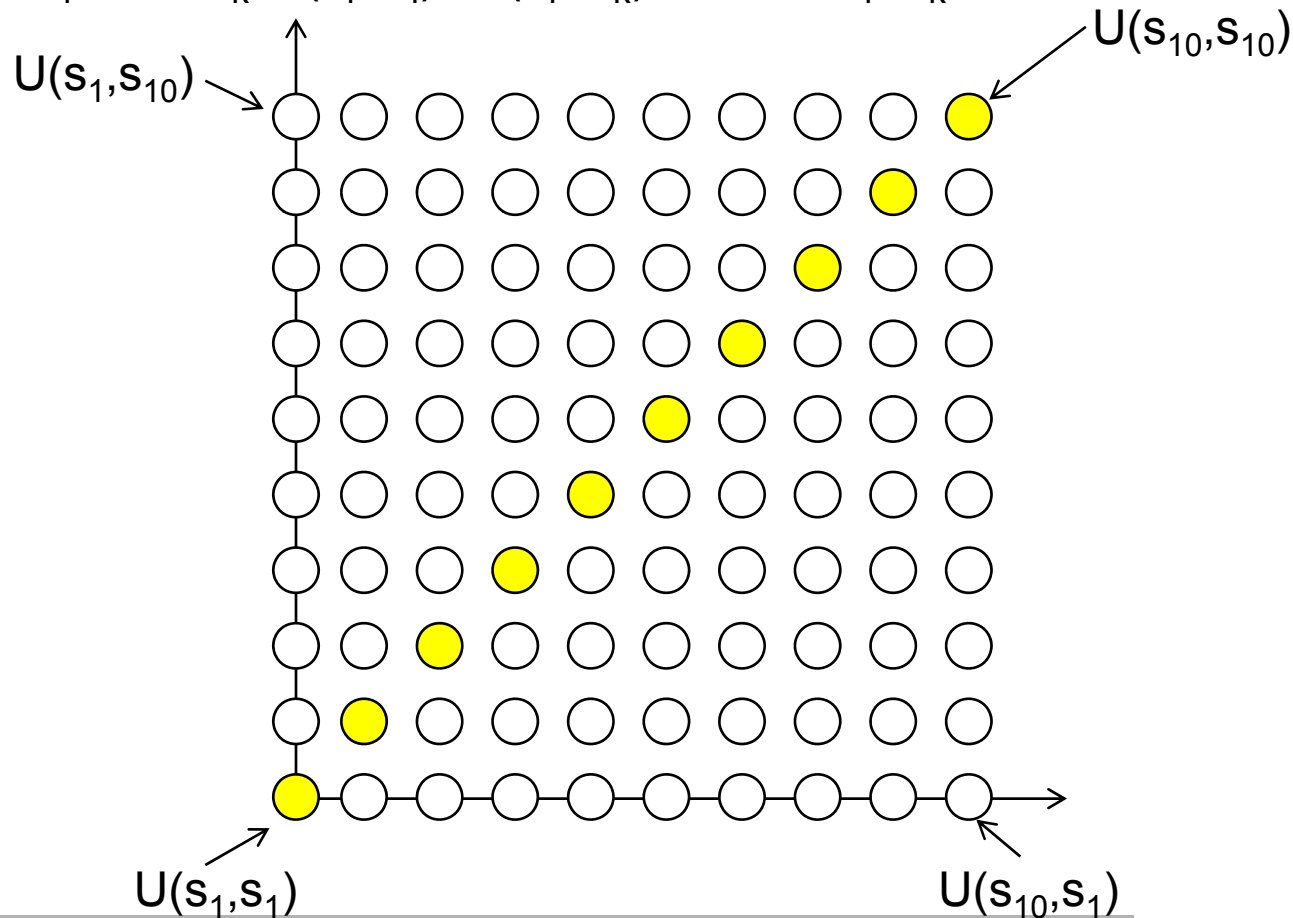
Goal: find s_i s.t. $\forall s_k U(s_i, s_i) \geq U(s_i, s_k)$, where $s_i, s_k \in S$



Distributed algorithm for computing Nash equilibrium

Input: $S = \{s_1, s_2, \dots, s_{10}\}$ – finite set of strategies, $U(s_i, s_j)$ – utility function

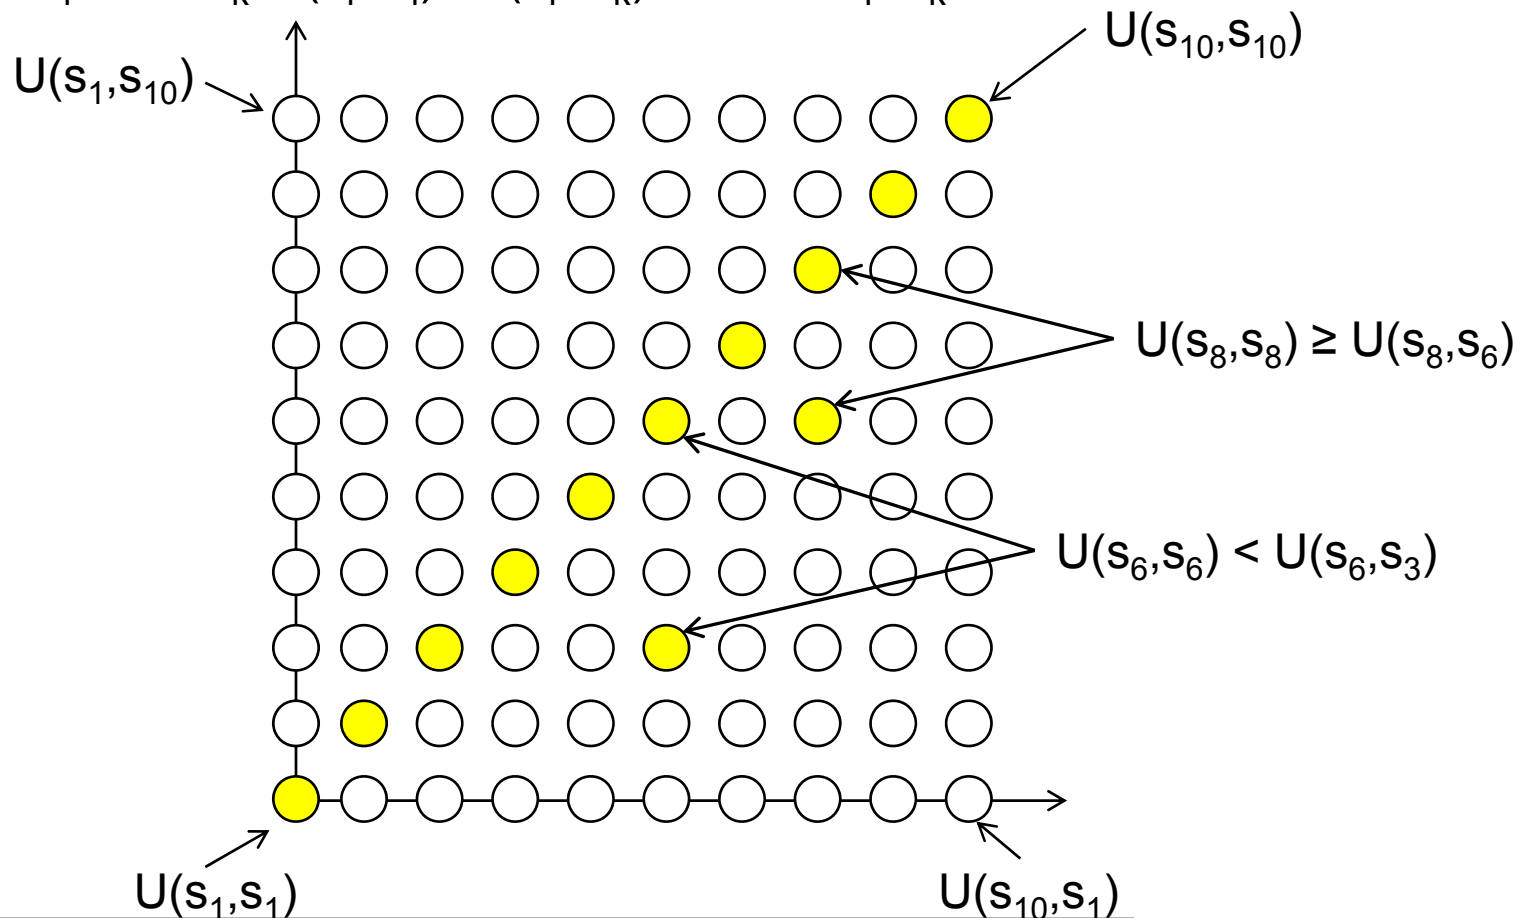
Goal: find s_i s.t. $\forall s_k U(s_i, s_i) \geq U(s_i, s_k)$, where $s_i, s_k \in S$



Distributed algorithm for computing Nash equilibrium

Input: $S = \{s_1, s_2, \dots, s_{10}\}$ – finite set of strategies, $U(s_i, s_j)$ – utility function

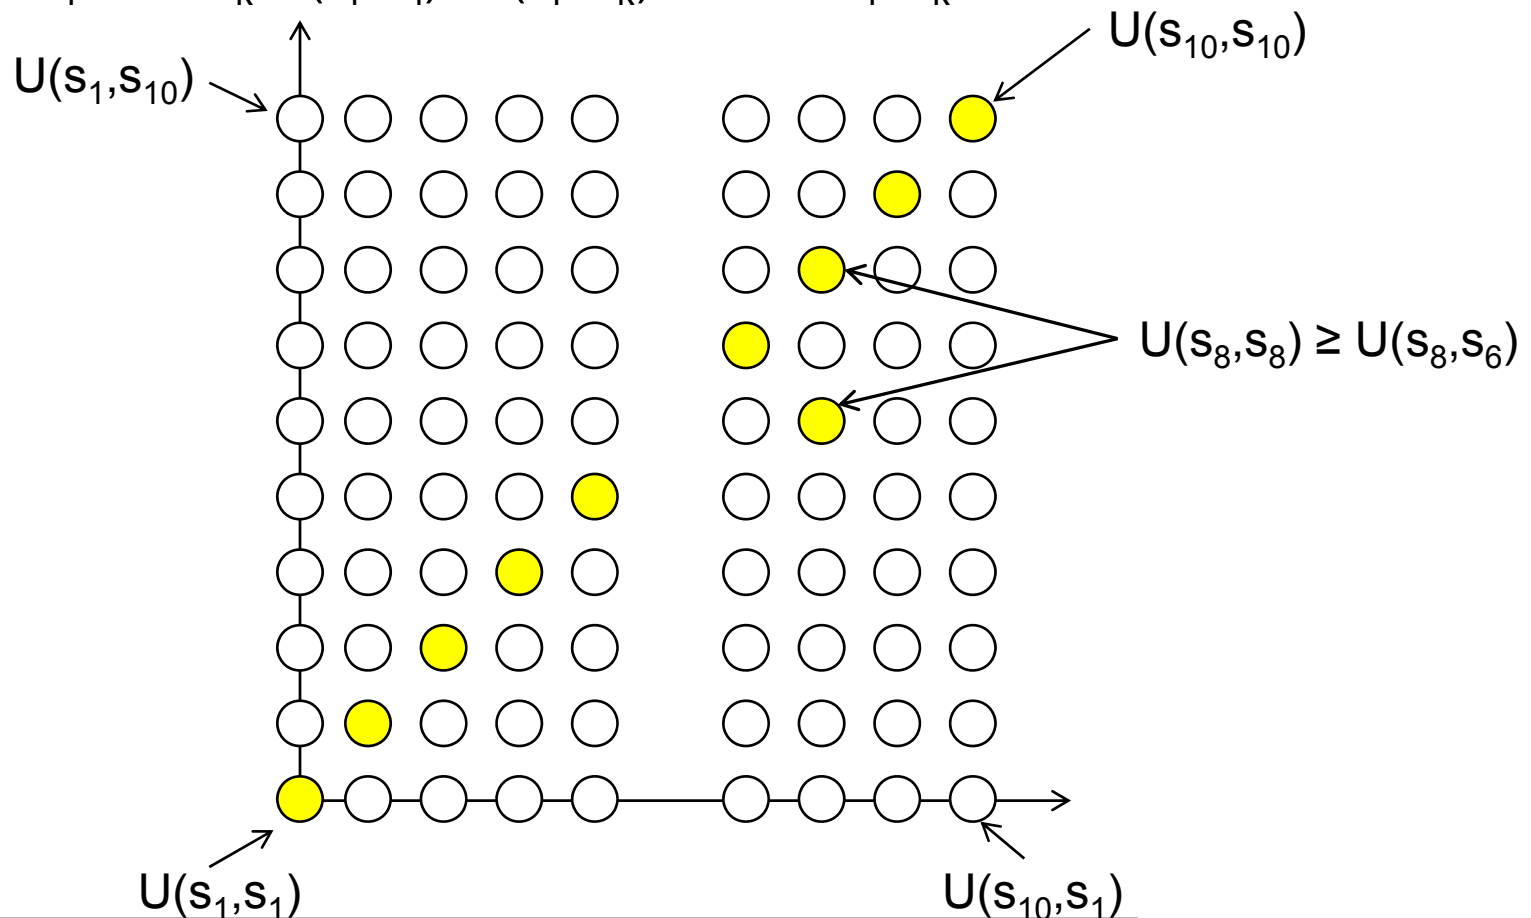
Goal: find s_i s.t. $\forall s_k U(s_i, s_i) \geq U(s_i, s_k)$, where $s_i, s_k \in S$



Distributed algorithm for computing Nash equilibrium

Input: $S = \{s_1, s_2, \dots, s_{10}\}$ – finite set of strategies, $U(s_i, s_j)$ – utility function

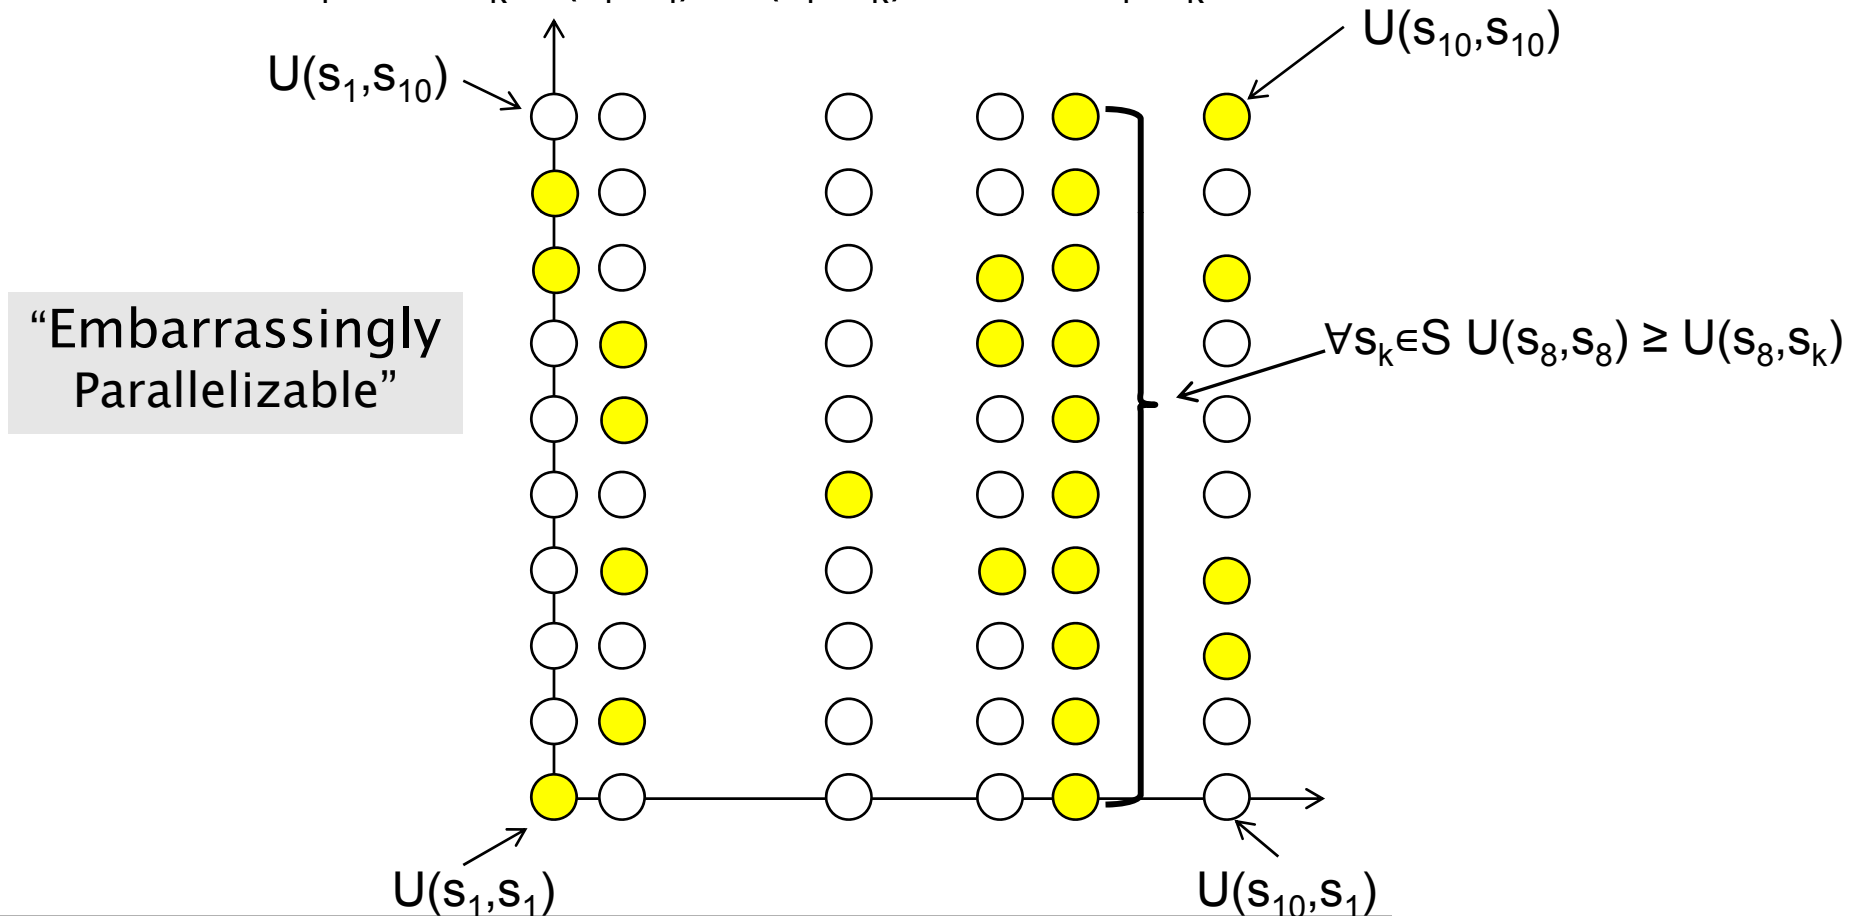
Goal: find s_i s.t. $\forall s_k U(s_i, s_i) \geq U(s_i, s_k)$, where $s_i, s_k \in S$



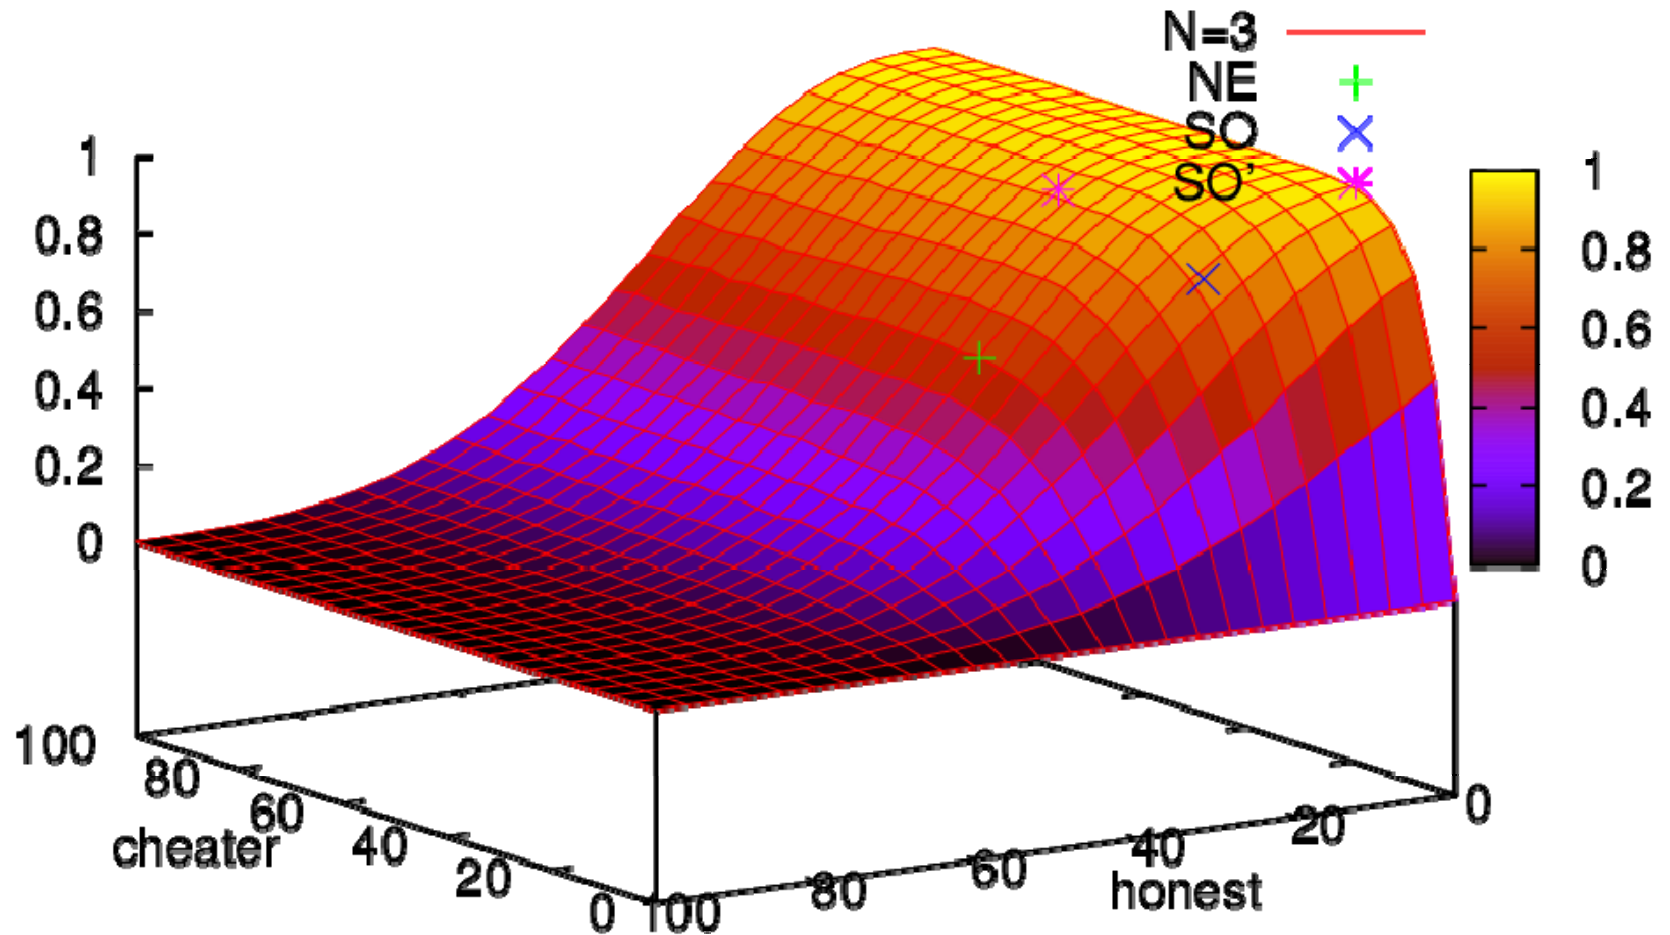
Distributed algorithm for computing Nash equilibrium

Input: $S = \{s_1, s_2, \dots, s_{10}\}$ – finite set of strategies, $U(s_i, s_j)$ – utility function

Goal: find s_i s.t. $\forall s_k \in S, U(s_i, s_i) \geq U(s_i, s_k)$, where $s_i, s_k \in S$



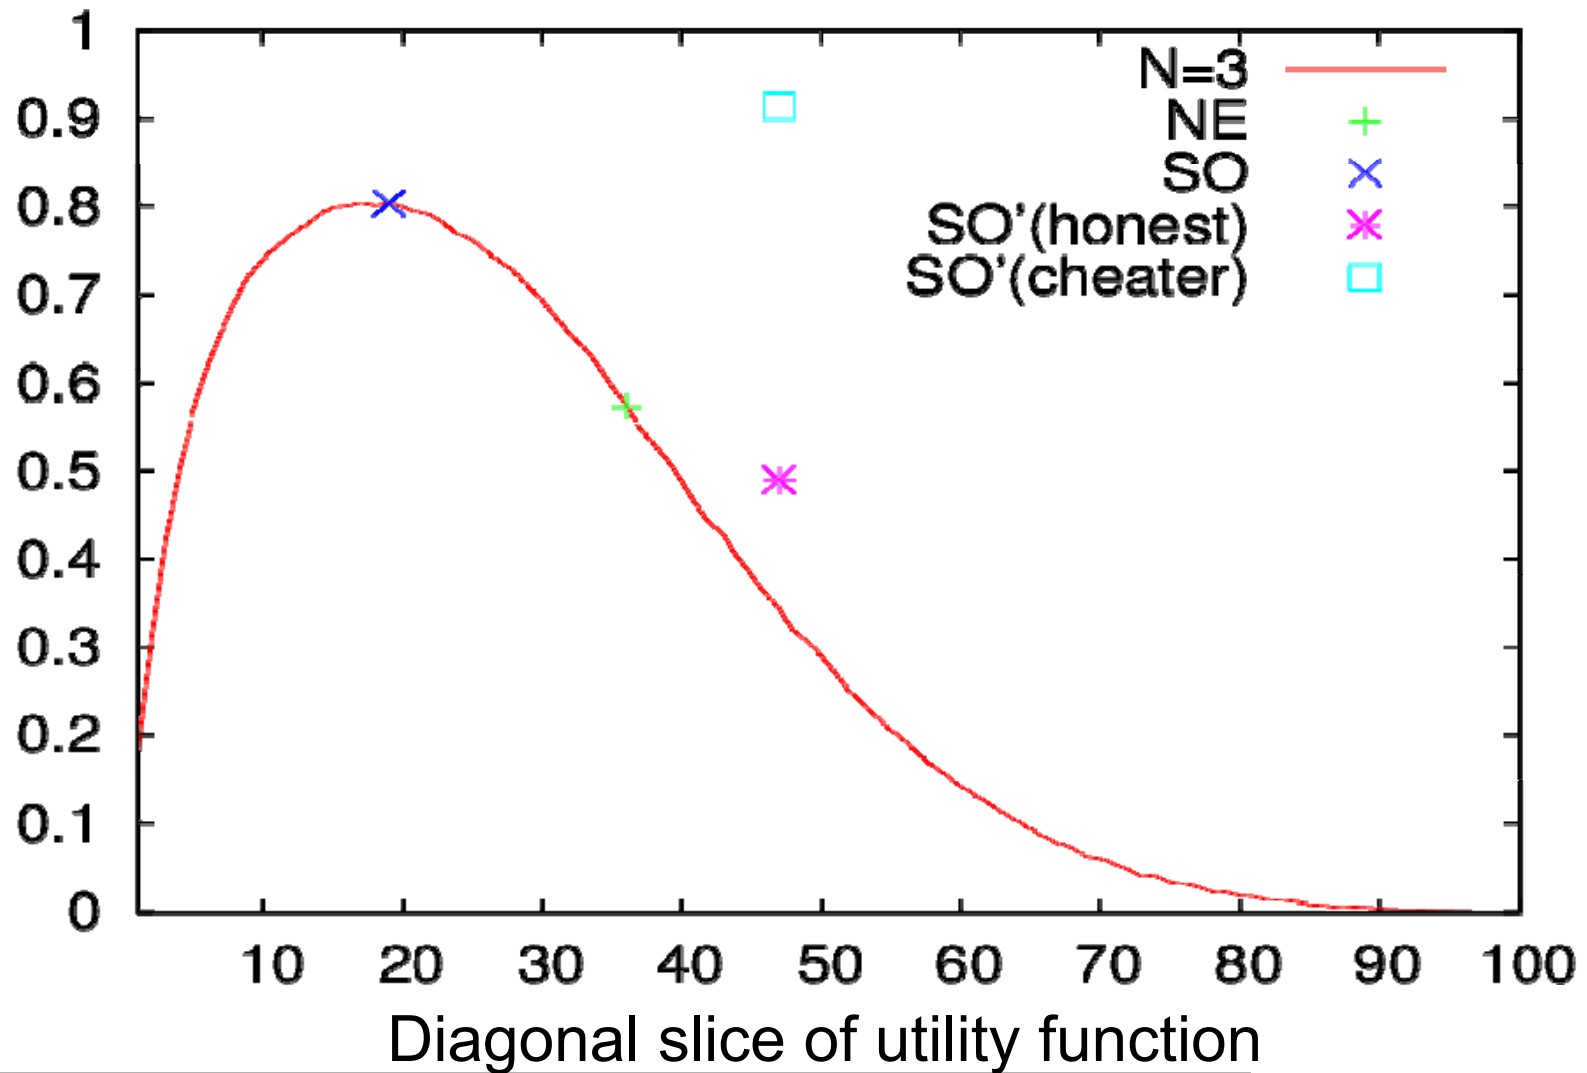
Results (3 nodes)



Value of utility function for the cheater node



Results (3 nodes)



Results

	N=2	N=3	N=4	N=5	N=6	N=7
Nash Eq (TrnPr)	0.32	0.36	0.36	0.35	0.32	0.32
$U(s_{NE}, s_{NE})$	0.91	0.57	0.29	0.15	0.10	0.05
Opt (TrnPr)	0.25	0.19	0.14	0.11	0.09	0.07
$U(s_{opt}, s_{opt})$	0.93	0.80	0.68	0.58	0.50	0.44

Symmetric Nash Equilibrium and Optimal strategies for different number of network nodes

#cores	4	8	12	16	20	24	28	32
Time	38m	19m	13m	9m46s	7m52s	7m04s	6m03s	5m

Time required to find Nash Equilibrium for N=3
100x100 parameter values
(8xIntel Core2 2.66GHz CPU)



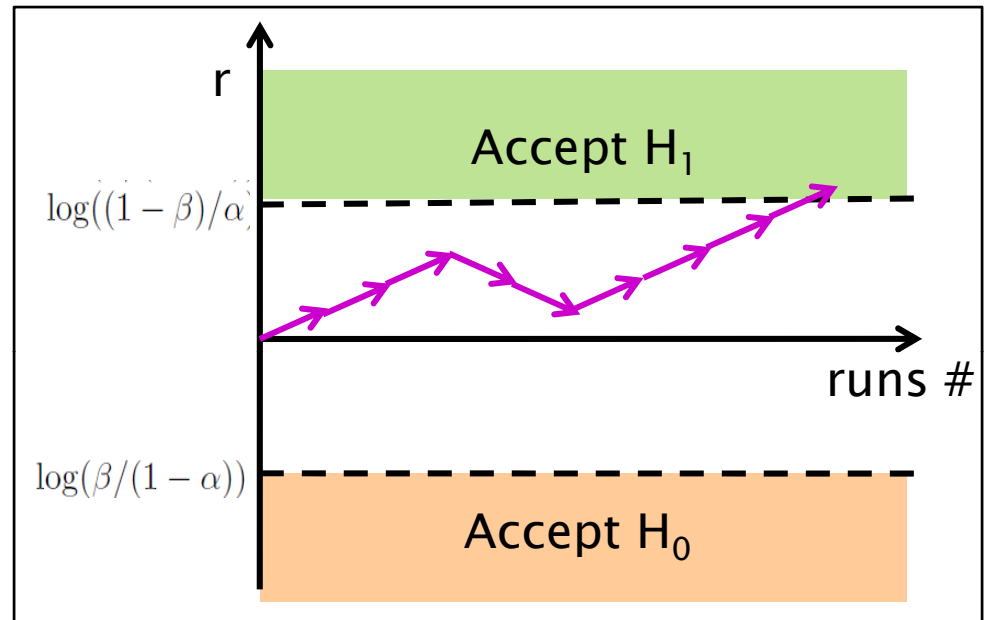
Overview

- Statistical Model Checking in UPPAAL
 - Estimation
 - Testing
- Distributed SMC for Parameterized Models
 - Parameter Sweeps
 - Optimization
 - Nash Equilibria
- **Distributing Statistical Model Checking**
 - Estimation
 - Testing
- Parameter Analysis of DSMC
- Conclusion



Bias Problem

- Suppose that generating accepting runs is fast and non-accepting runs is slow.
- 1-node exploration:
 - Generation is sequential, only the outcomes count.
- N-node exploration:
 - There may be an unusual peak of accepting runs generated more quickly by some nodes that will arrive long before the non-accepting runs have a chance to be counted!
 - The decision will be biased toward accepting runs.



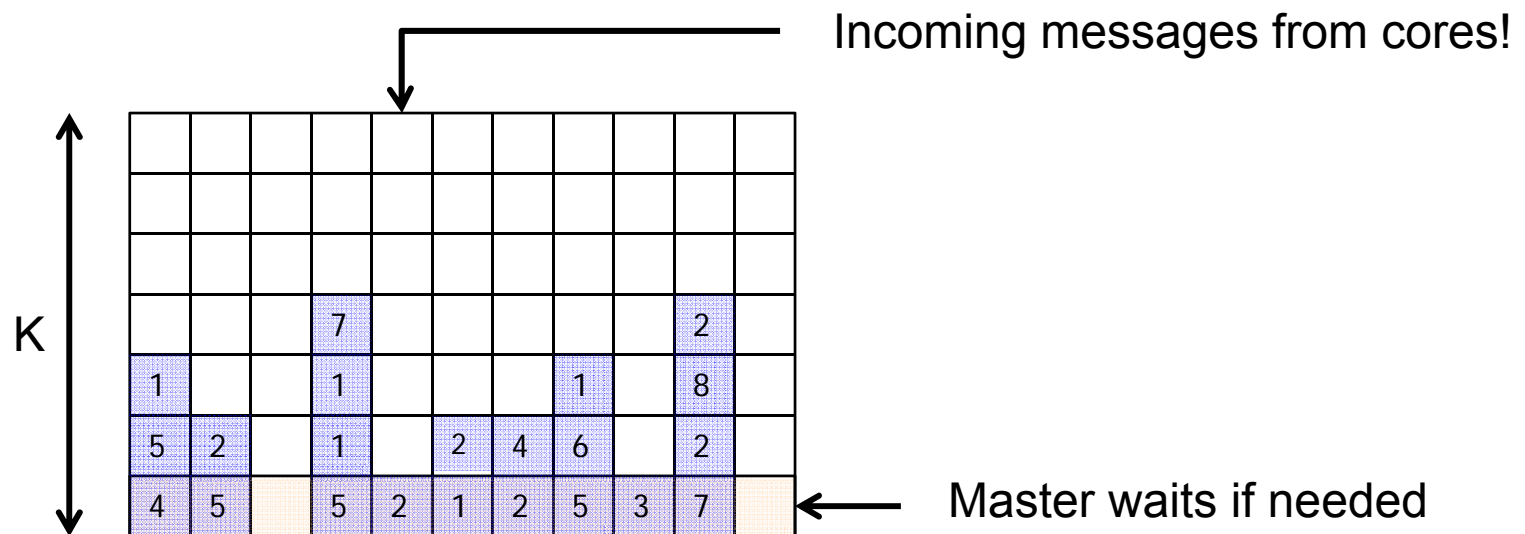
Solving Bias [Younes'05]

Queue the results at a master, use Round-Robin between nodes to accept the results.



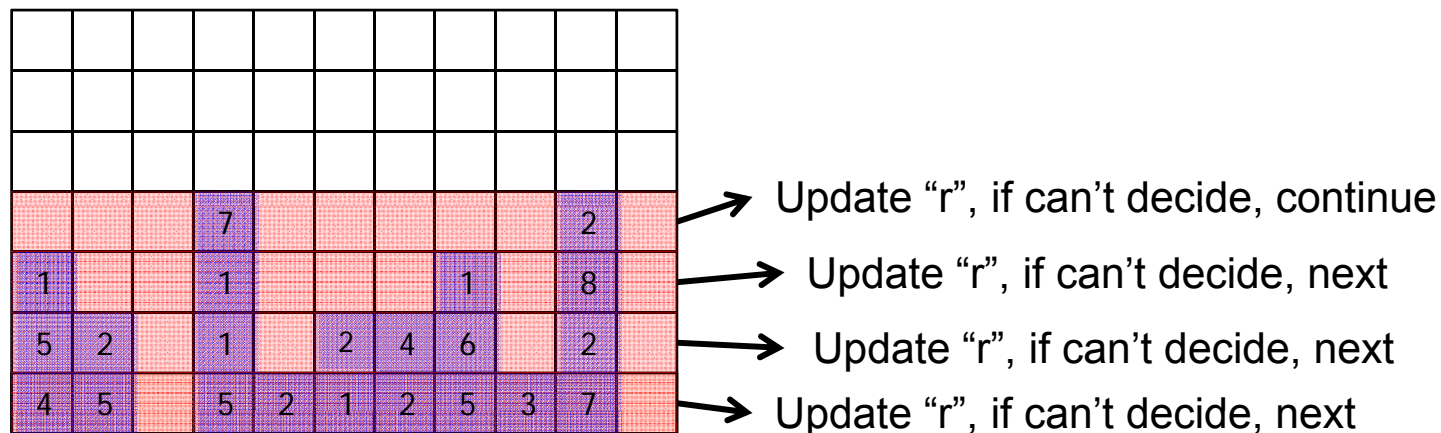
Our Implementation

- Use a batch of B (e.g 10) runs, transmit one count per batch.
- Use asynchronous communication (MPI)
- Queue results at the master and wait only when the buffer (size= K) is full.



Our Implementation

- Senders have a buffer of (K) asynchronously sent messages and blocks only when the buffer is full.
- The master periodically add results in the buffer.



Experiment on Multi-Core

- Machine: i7 4*cores HT, 4GHz.
 - Hyperthreading is an interesting twist:
 - threads share execution units,
 - have **unpredictable** running times
(may run on same physical core if < 8 threads).
- Model: Train Gate with 20 trains.
- Configuration – $B=40$, $K=64$
- Property:
“mutual exclusion on the bridge within time ≤ 1000 ”

H0: accept if $\text{Pr} \geq 0.9999$

H1: accept if $\text{Pr} \leq 0.9997$

$\alpha=0.001$, $\beta=0.001$.



Performance

- Compared to base non-MPI version.
- Min, average, max *.
- 4.99 max speedup on a **quad**-core.

	<u>Speedup</u>			<u>Efficiency</u>		
1	0.95	0.98	1.00	95%	98%	100%
2	1.86	1.94	1.98	93%	97%	99%
3	2.78	2.89	2.96	93%	96%	99%
4	3.33	3.76	3.90	83%	94%	98%
5	2.97	3.22	3.66	59%	64%	73%
6	3.61	3.74	3.87	60%	62%	65%
7	4.09	4.31	4.47	58%	62%	64%
8	3.65	4.73	4.99	46%	59%	62%

Base time
Min=44.35s
Avg=44.62s
Max=45.49s



Early Cluster Experiments

- Xeons 5335, 8 cores/node.

- Estimation
Firewire protocol
22 properties
node x cores
speed-up, efficiency

1x1	1x2	1x4	1x8		
1, 100%	1.8, 92%	3.5, 88%	6.7, 84%		
16min					
	2x1	2x2	2x4	2x8	4x8
	1.8, 92%	3.9, 98%	6.8, 85%	12.3, 77%	19.6, 61%

1x1	1x2	1x4			
1, 100%	1.7, 86%	3.3, 83%			
6m30s					
	2x1	2x4	4x2		
	1.7, 84%	5.9, 74%	7.1, 89%		

- Estimation
Lmac protocol
1 property
- Encouraging results despite simple distribution.

Thanks to Jaco van de Pol, Axel Belifante, Martin Rehr, and Stefan Blom for providing support on the cluster of the University of Twente



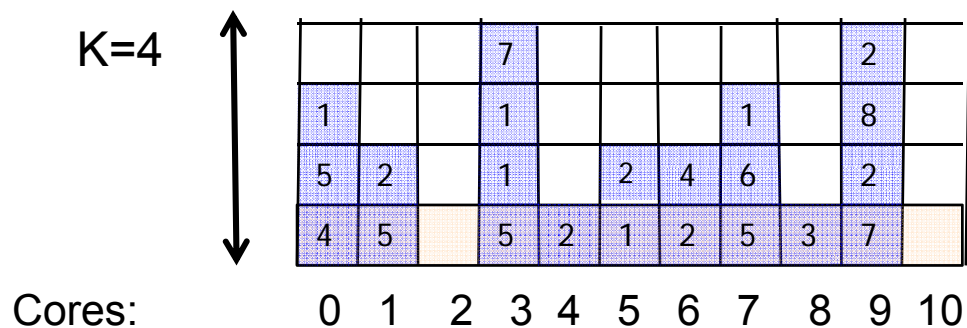
Overview

- Statistical Model Checking in UPPAAL
 - Estimation
 - Testing
- Distributed SMC for Parameterized Models
 - Parameter Sweeps
 - Optimization
 - Nash Equilibria
- Distributing Statistical Model Checking
 - Estimation
 - Testing
- **DSMC of DSMC**
- Conclusion



Distributed SMC

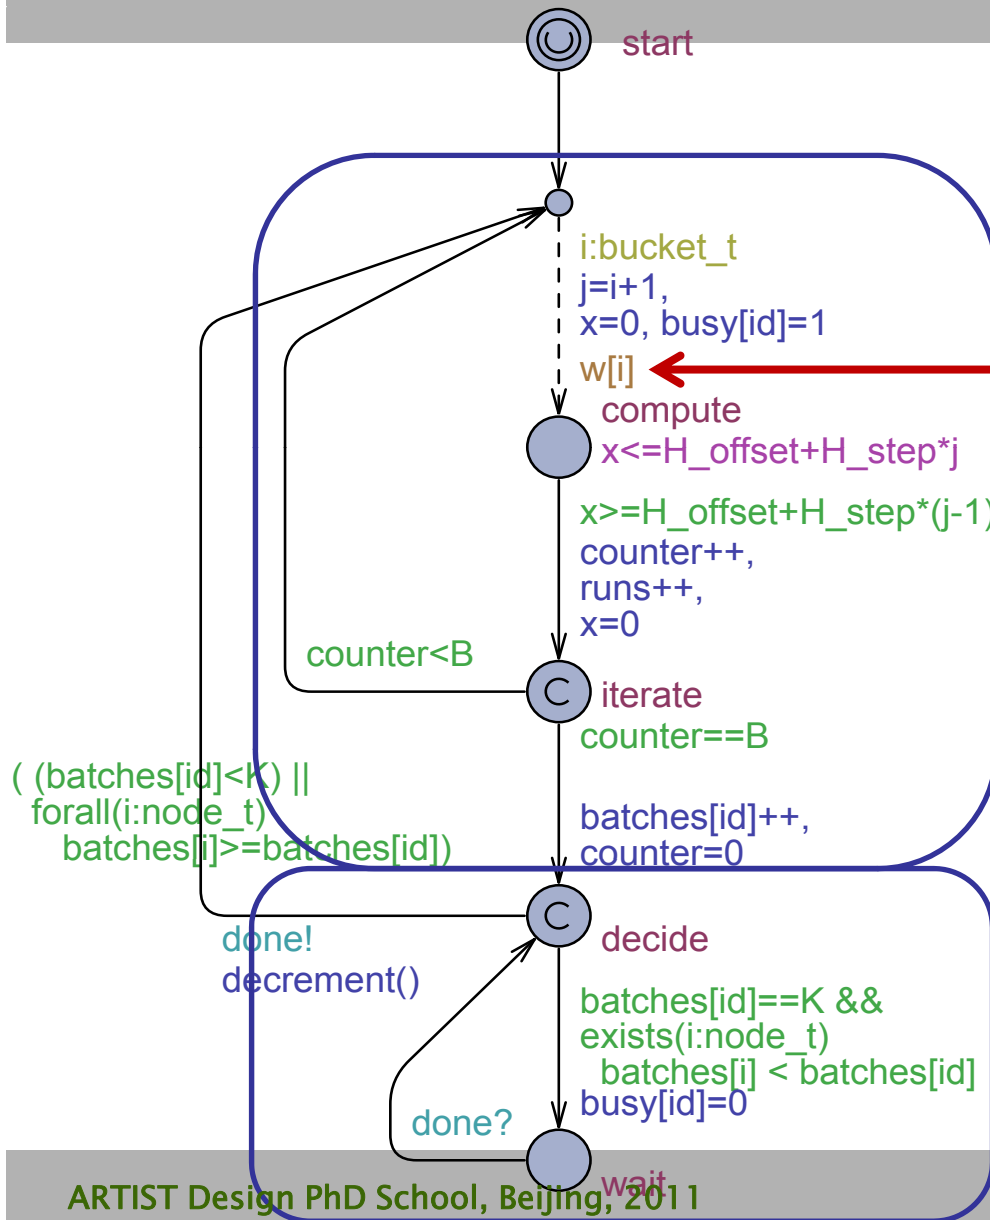
- SMC simulations can be distributed across a cluster of machines with **N** number of **cores**.
- The simulations are grouped into batches of **B** number of **simulations** in each to avoid bias.
- Each core is not allowed to be ahead by more than **K** batches than any other core.



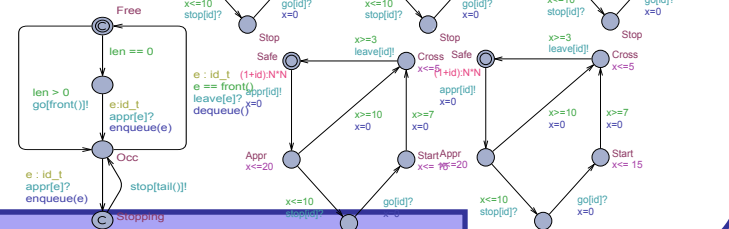
Core0 is computing **4th** batch
 Core2 is computing **1st** batch
 Core1 is computing **3rd** batch
 Core9 is blocked, **waiting** for Core2+10
 Core3 is blocked, **waiting** for Core2+10
 Only **complete** row of batches is used



Distributed SMC: Model of a Core



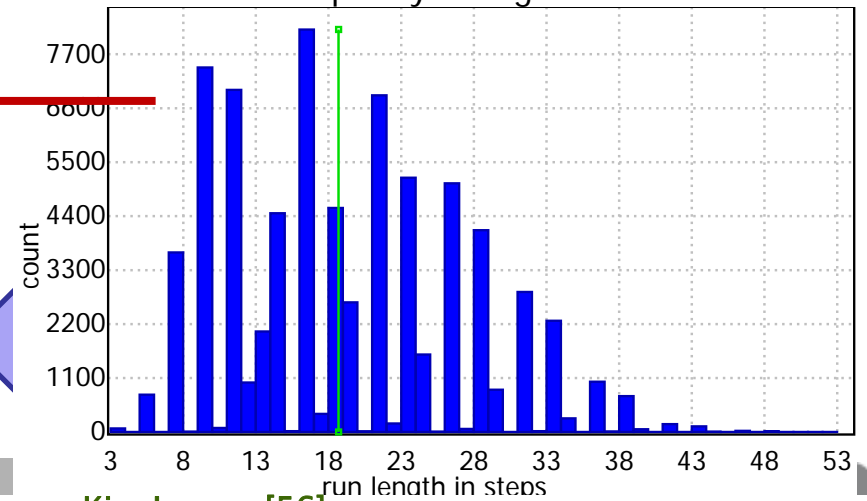
train gate
model



Computing one batch

$Pr[\# \leq 100](\langle \rangle Train(5).Cross)$

Frequency Histogram

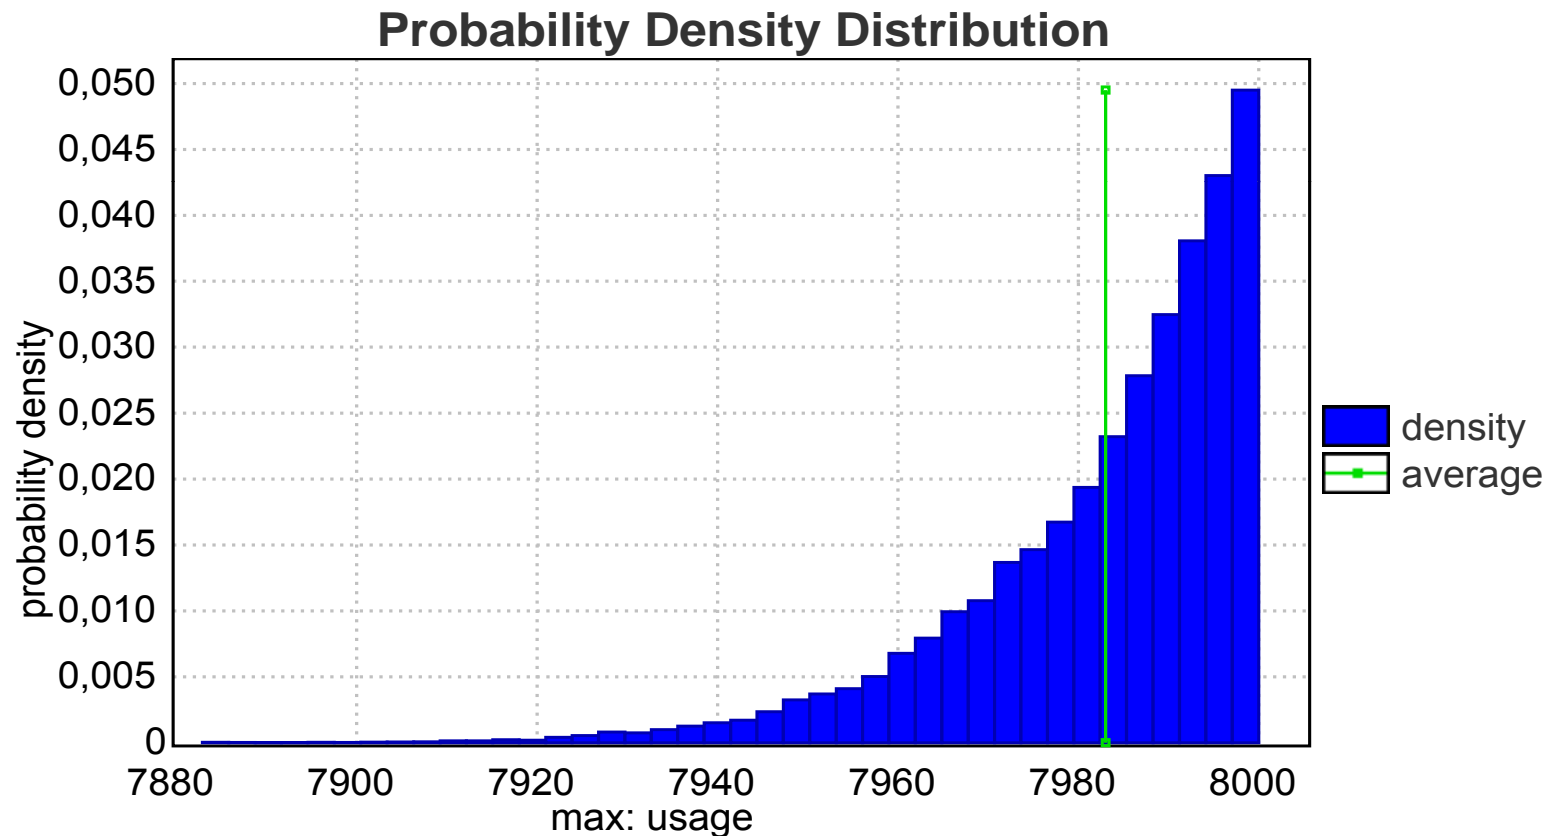


Kim Larsen [56]
 generation time ~ simulation steps

DSMC: CPU Usage Time

Parameter instantiation:
N=8, B=100, K=2

Property used:
E[time <= 1000; 25000] (max: usage)



Runs: 25000 in total, 25000 displayed, 0 remaining.

Probability sums: 1 displayed, 0 remaining.

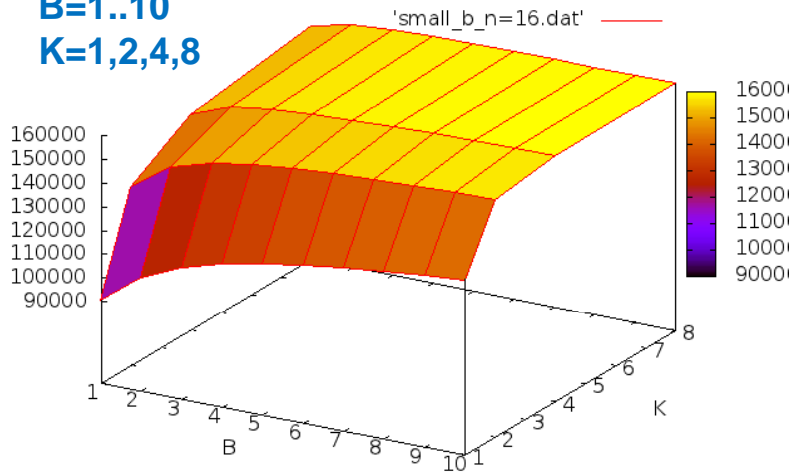
Average: 7983.02.

Kim Larsen [57]

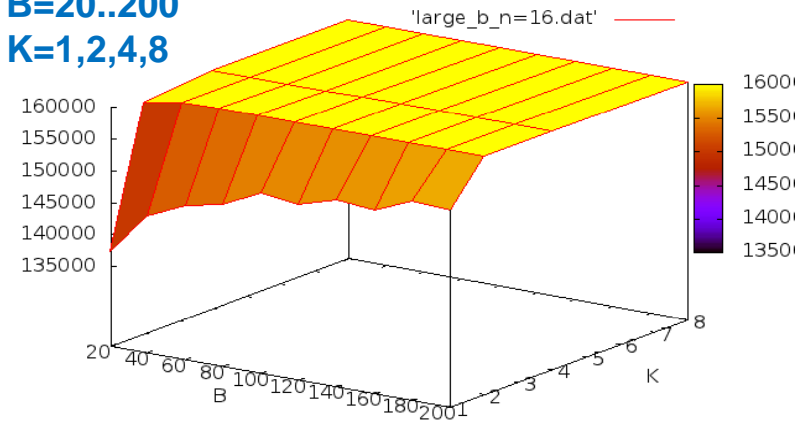


DSMC Performance Analysis

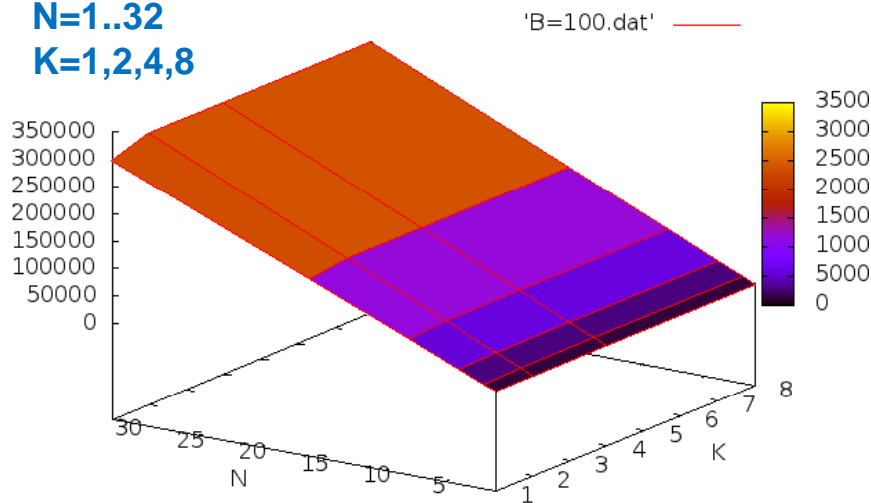
N=16
B=1..10
K=1,2,4,8



N=16
B=20..200
K=1,2,4,8



B=100,
N=1..32
K=1,2,4,8



Property used:

$E[\text{time} \leq 1000; 1000]$ (max: usage)

Conclusions:

K=1 has huge effect and should be avoided.

K=2 has effect if $B < 20$.

K > 2 are indistinguishable on homogeneous cluster.

K > 2 and $B > 20$: number of simulations scale linearly to the number of cores used.

Conclusion

- Preliminary experiments indicate that distributed SMC in UPPAAL scales very nicely.
- More work to identify impact of parameters for distributing individual SMC?
- How to assign statistical confidence to parametric analysis, e.g. optimum or NE?
- UPPAAL 4.1.4 available
(support for SMC, DSMC, 64-bit,..)

