How Fast and Fat Is Your Probabilistic Model Checker?

an experimental performance comparison

David N. Jansen[3,1], Joost-Pieter Katoen[1,2], Marcel Oldenkamp[2], Mariëlle Stoelinga[2], Ivan Zapreev[1,2]

[1] MOVES Group, RWTH Aachen University
[2] FMT Group, University of Twente, Enschede
[3] ICIS, Radboud University, Nijmegen

Probabilistic Model Checking

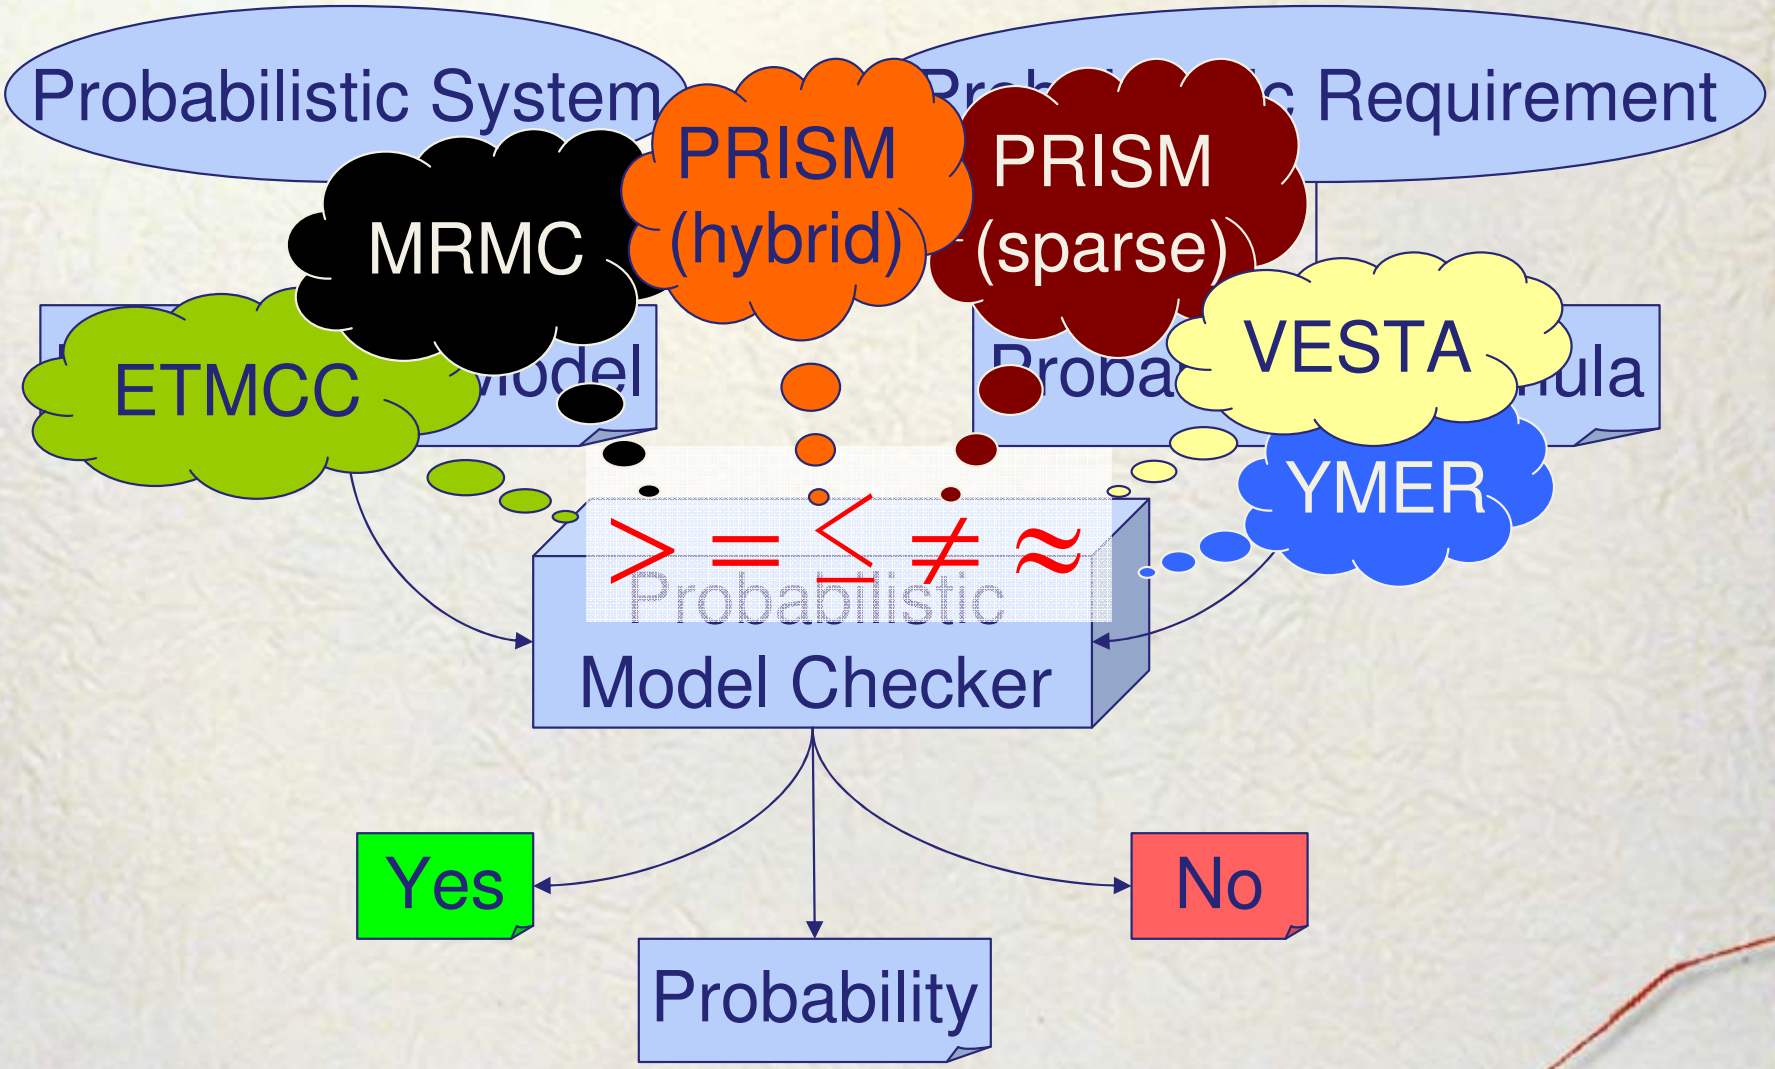

Why Are Probabilities Useful?

- system performance
- uncertainty in the environment
- randomized (networking) algorithms
- abstract from large populations

Probabilistic Model Checking...

- What is inside?
  - temporal logics + model checking
  - numerical and optimisation techniques from performance and operations research
- Where is it used?
  - powerful tools
  - applications: distributed systems, security, biology, quantum computing...
- Problem: Which tool to choose?

Probabilistic Models

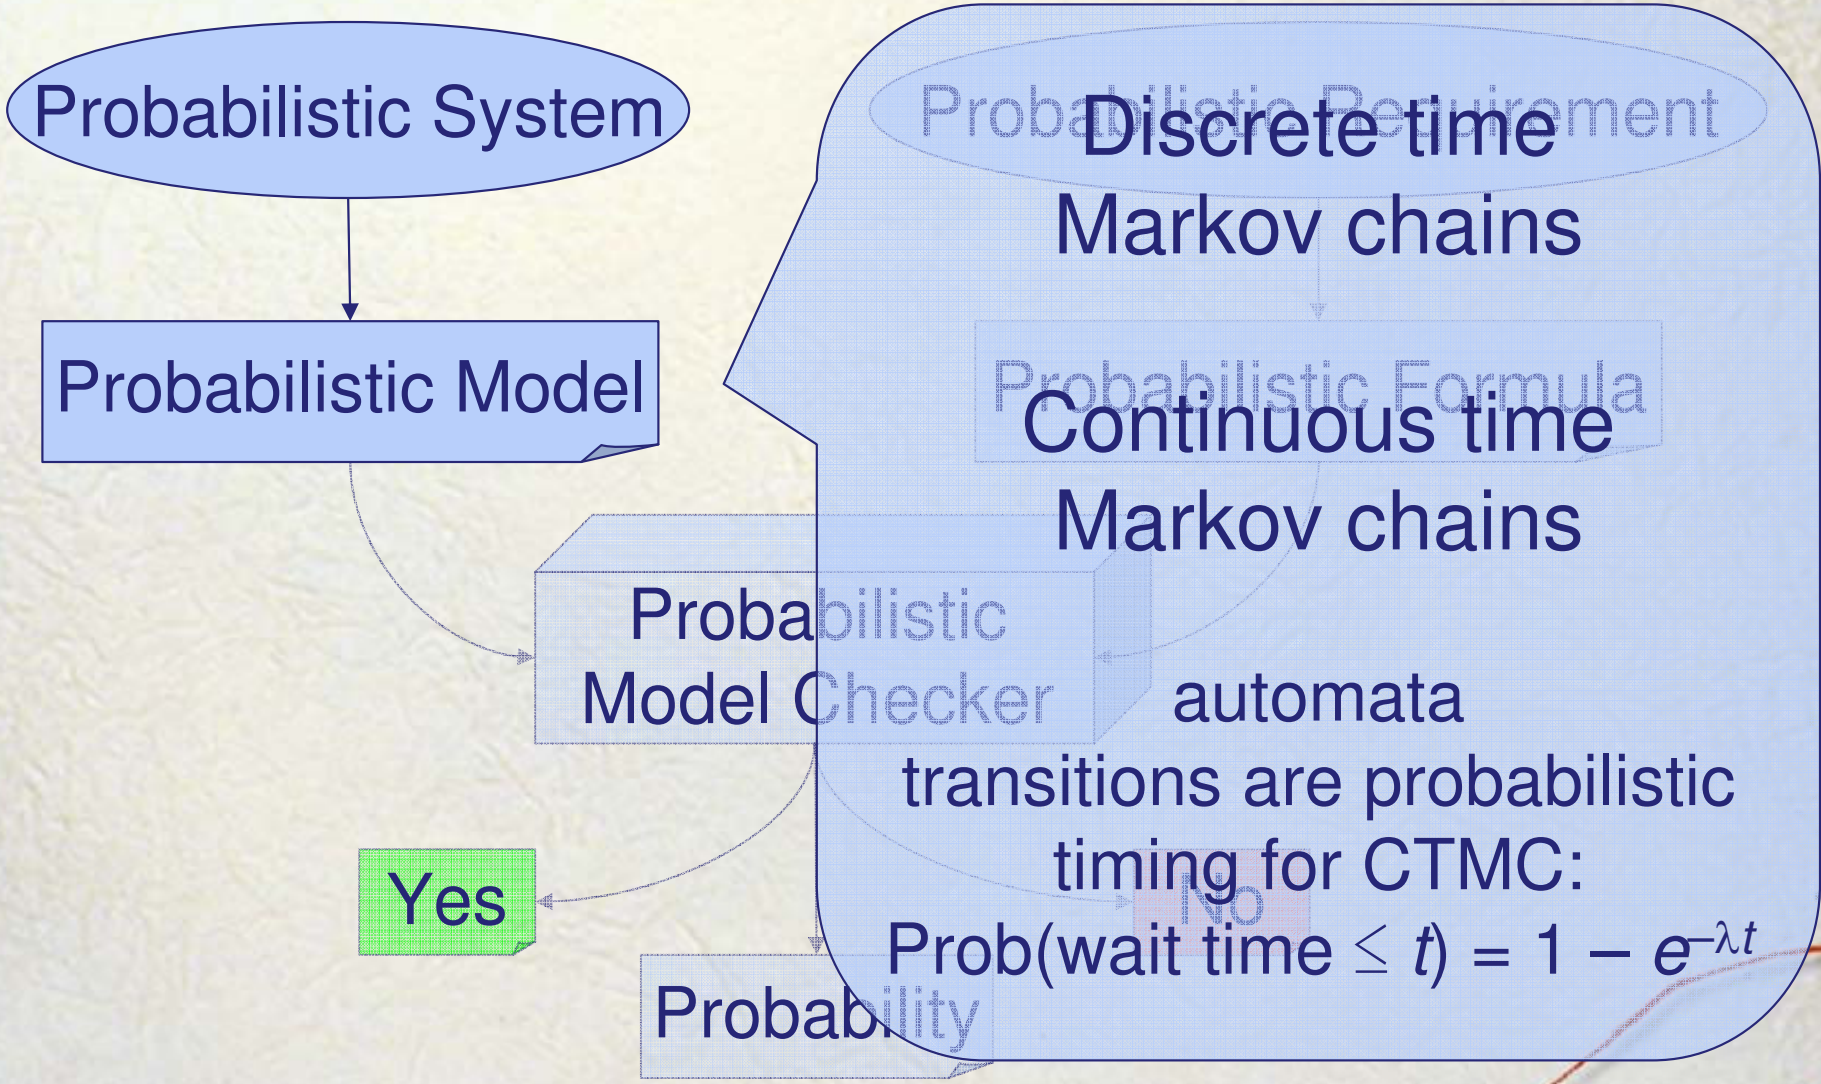

Synchronous Leader Election

- nodes in a ring elect a leader
  - each node selects random number as id
  - passes it around the ring (synchronously)
  - if $\exists$ unique id,
    node with maximum unique id is leader

- [Itai & Rodeh 1990]

Synchronous Leader Election

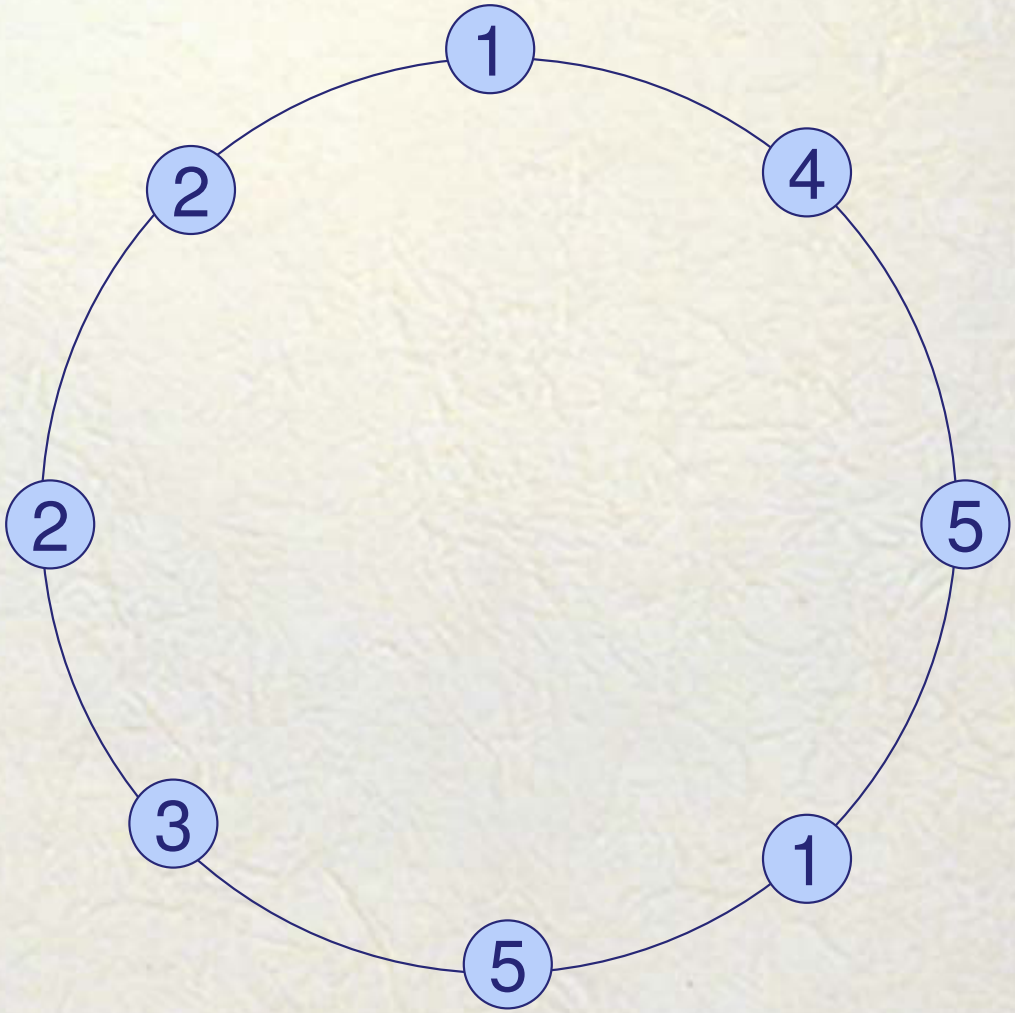

Synchronous Leader Election

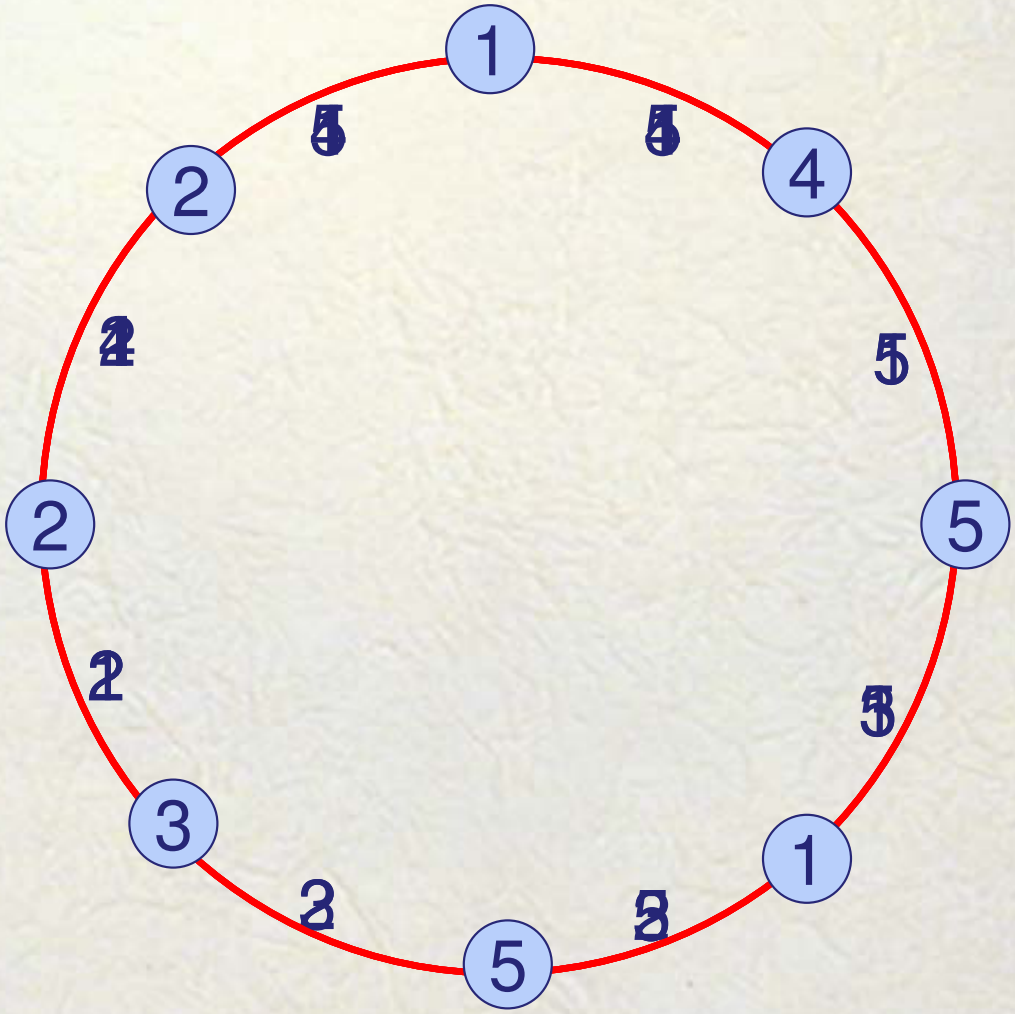

Models

- Discrete time Markov chain
  - transitions are fully probabilistic
  - timing is irrelevant
- Continuous time Markov chain
  - transitions are fully probabilistic
  - and timing also:
    Prob(wait time $\leq t$) = $1 - e^{-\lambda t}$

Probabilistic Formulas

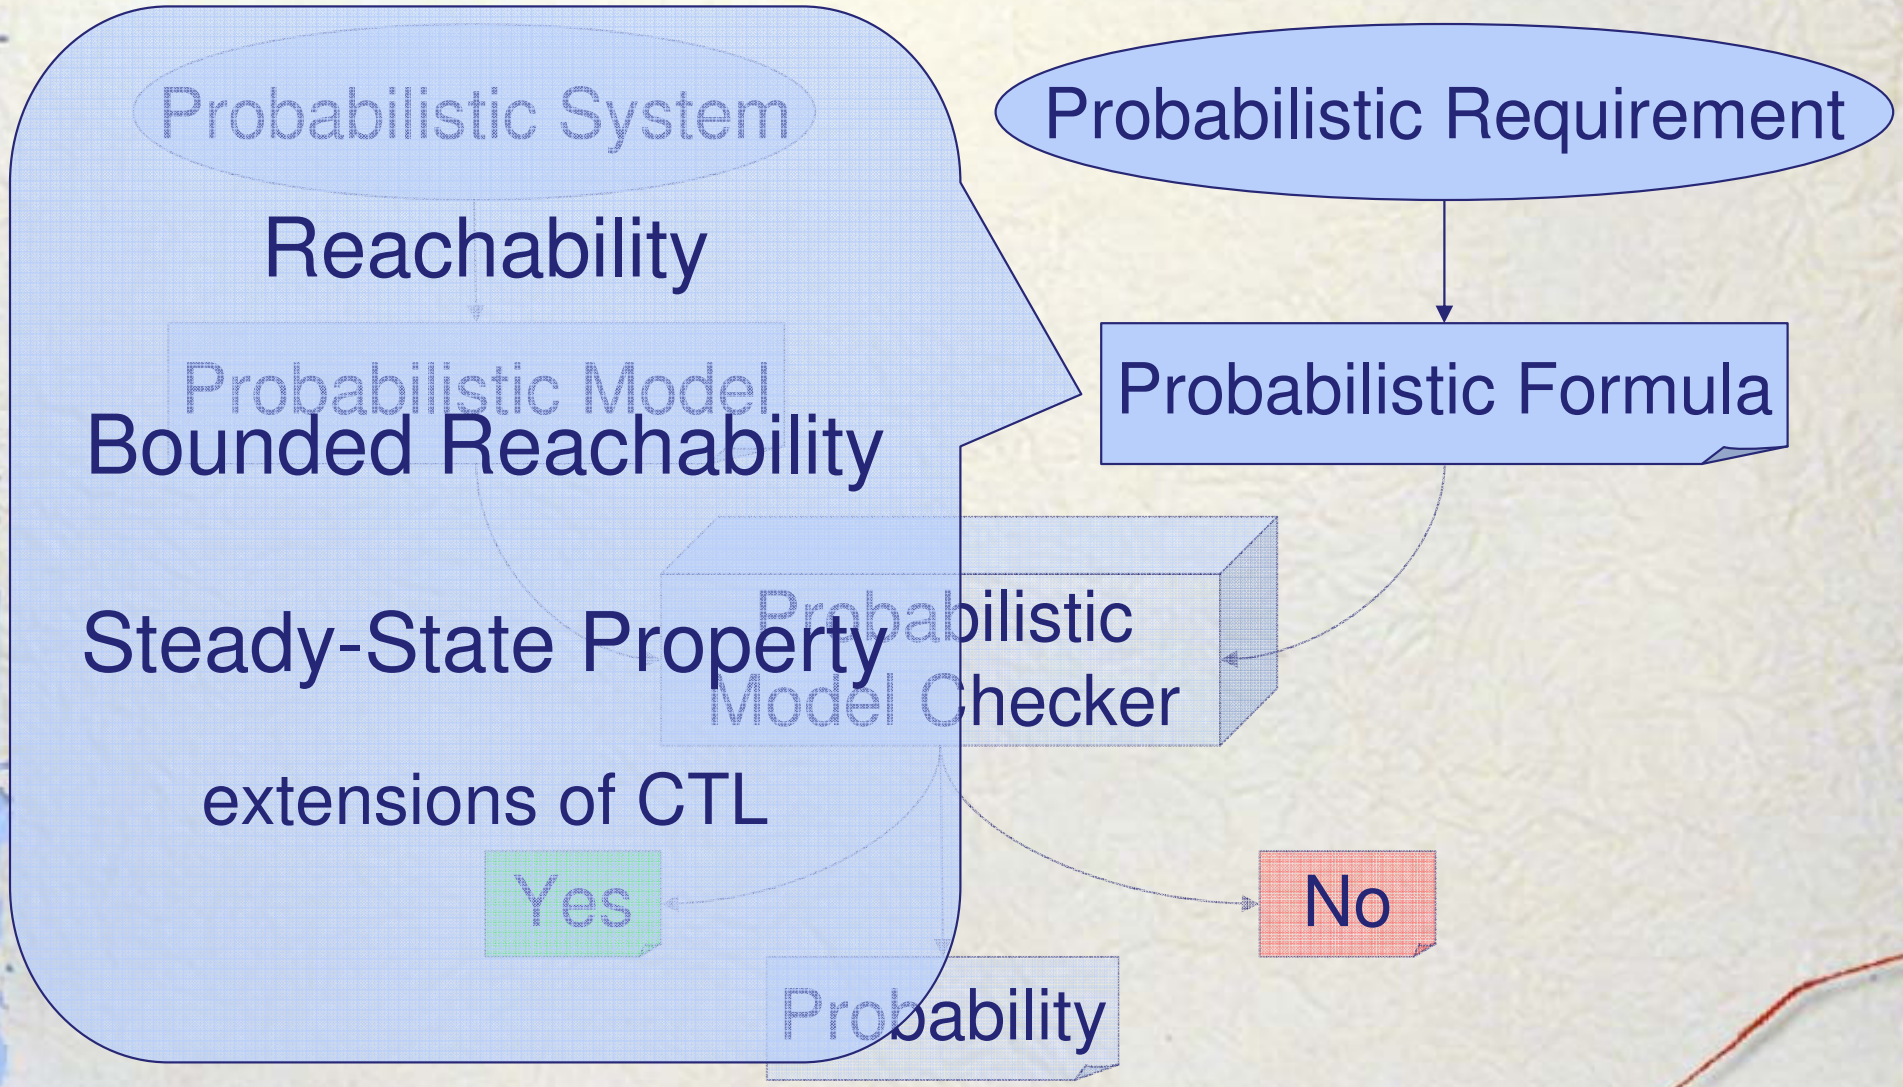

Probabilistic Model Checkers

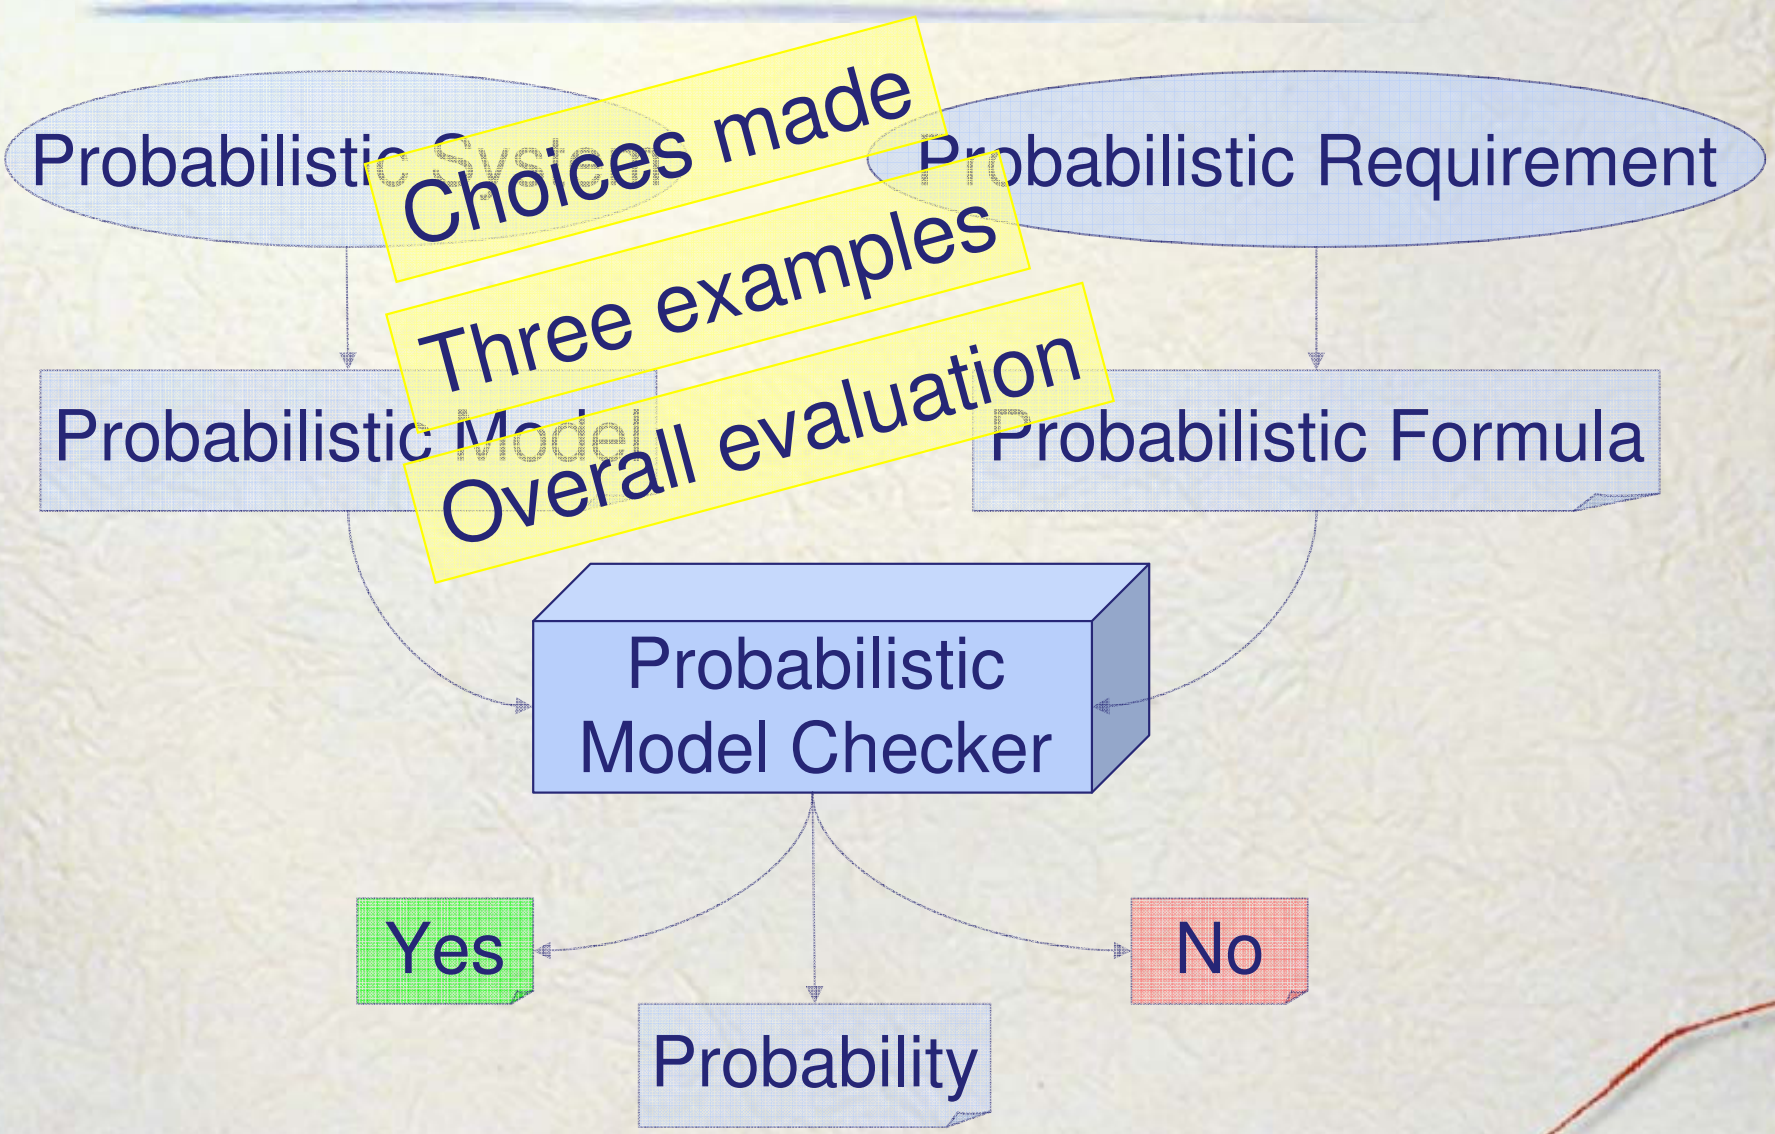

Tools

ETMC	numerical	CTMC	Java
MRMC		DTMC + CTMC	C
PRISM hybrid		DTMC + CTMC MTBDD for transition matrix	C(++) and Java
PRISM sparse		DTMC + CTMC sparse transition matrix	
VESTA	statist	CTMC, reachability	Java
YMER		CTMC, bounded reach	C(++)

Modelling

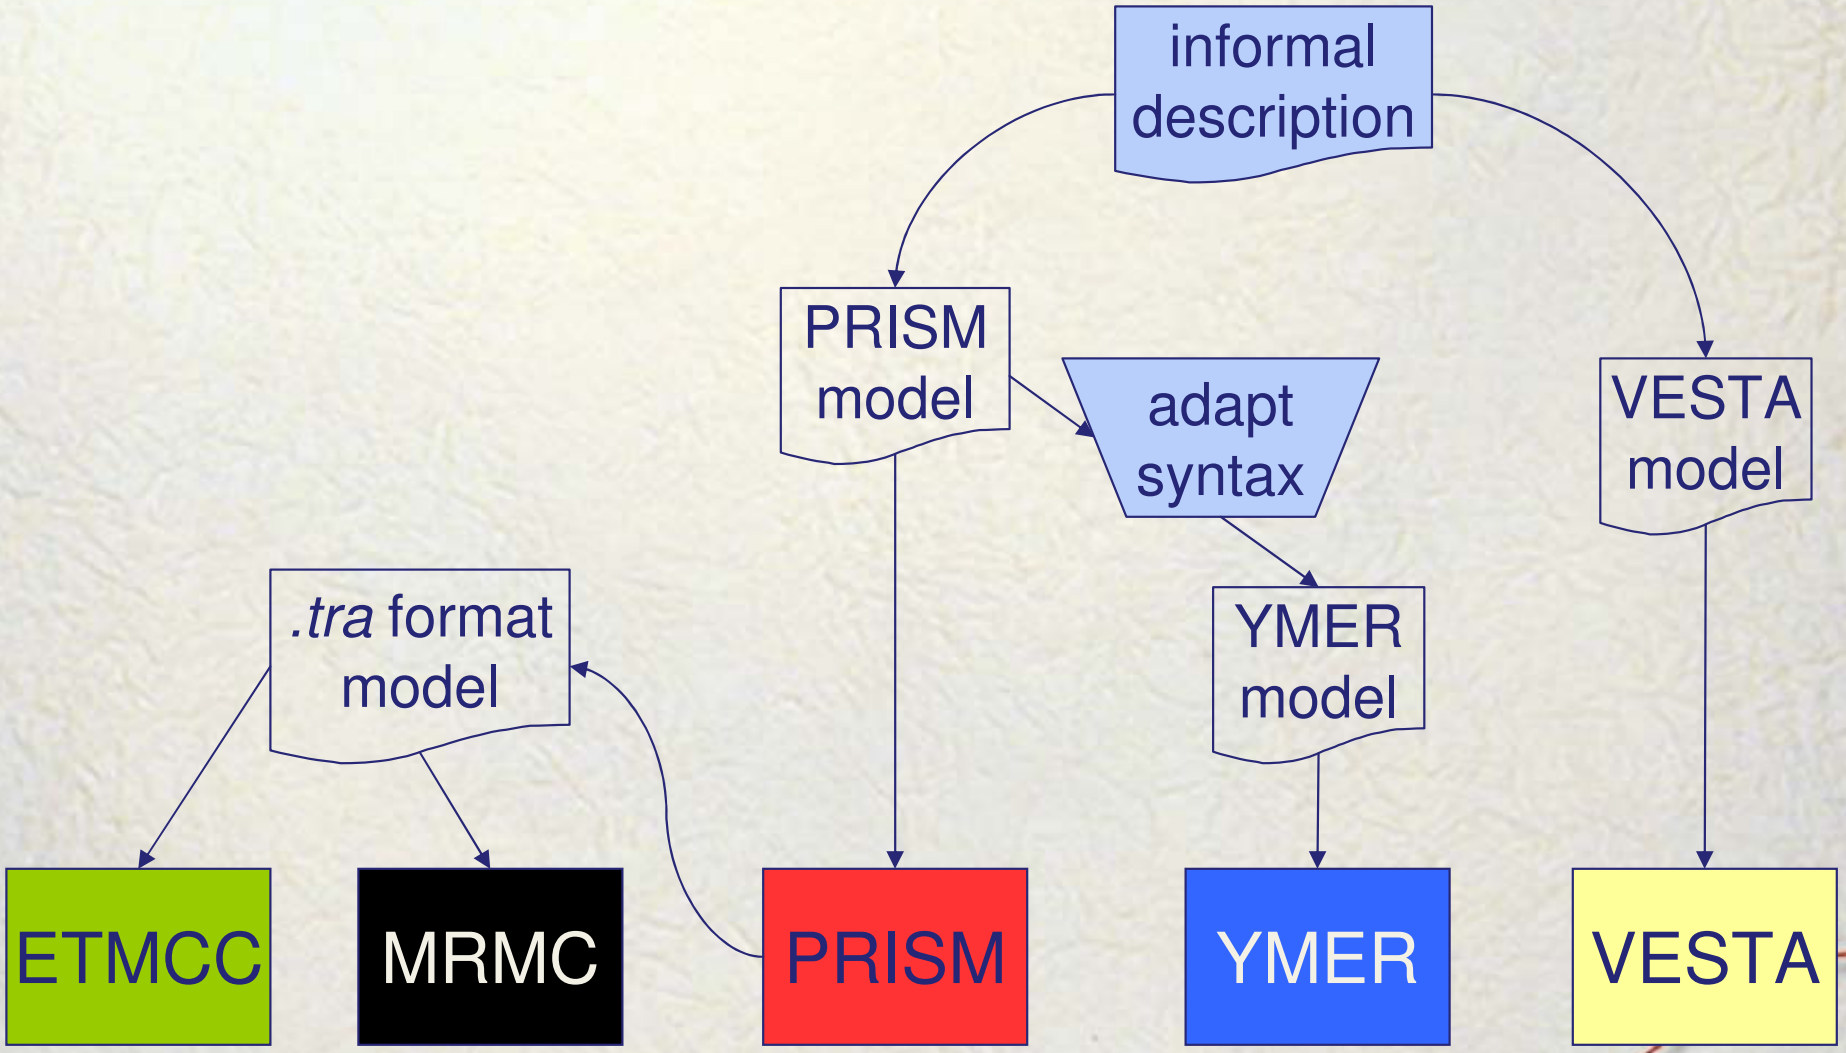
informal description
PRISM model
adapt syntax
VESTA model
.tra format model
YMER model
ETMCC
MRMC
PRISM
YMER
VESTA

Selected Benchmarks

Synchronous Leader Election	discrete time
Randomized Dining Philosophers	discrete time
Birth–Death Process	discrete time
Tandem Queuing Network	continuous time
Cyclic Server Polling	continuous time

Experiment Relevance

- Repeatable
- Verifiable
- Significant
- Encapsulated

Experiment 1
Reachability

- **Cyclic Polling Server:**
  server cycles over *n* stations
  and serves each one in turn
  - e.g. teacher walks through class,
    each pupil may ask a question

- $busy_1 \rightarrow P_{\geq 1}(true \ \mathrm{U} \ poll_1)$
  If station 1 is busy,
  the server will poll it eventually

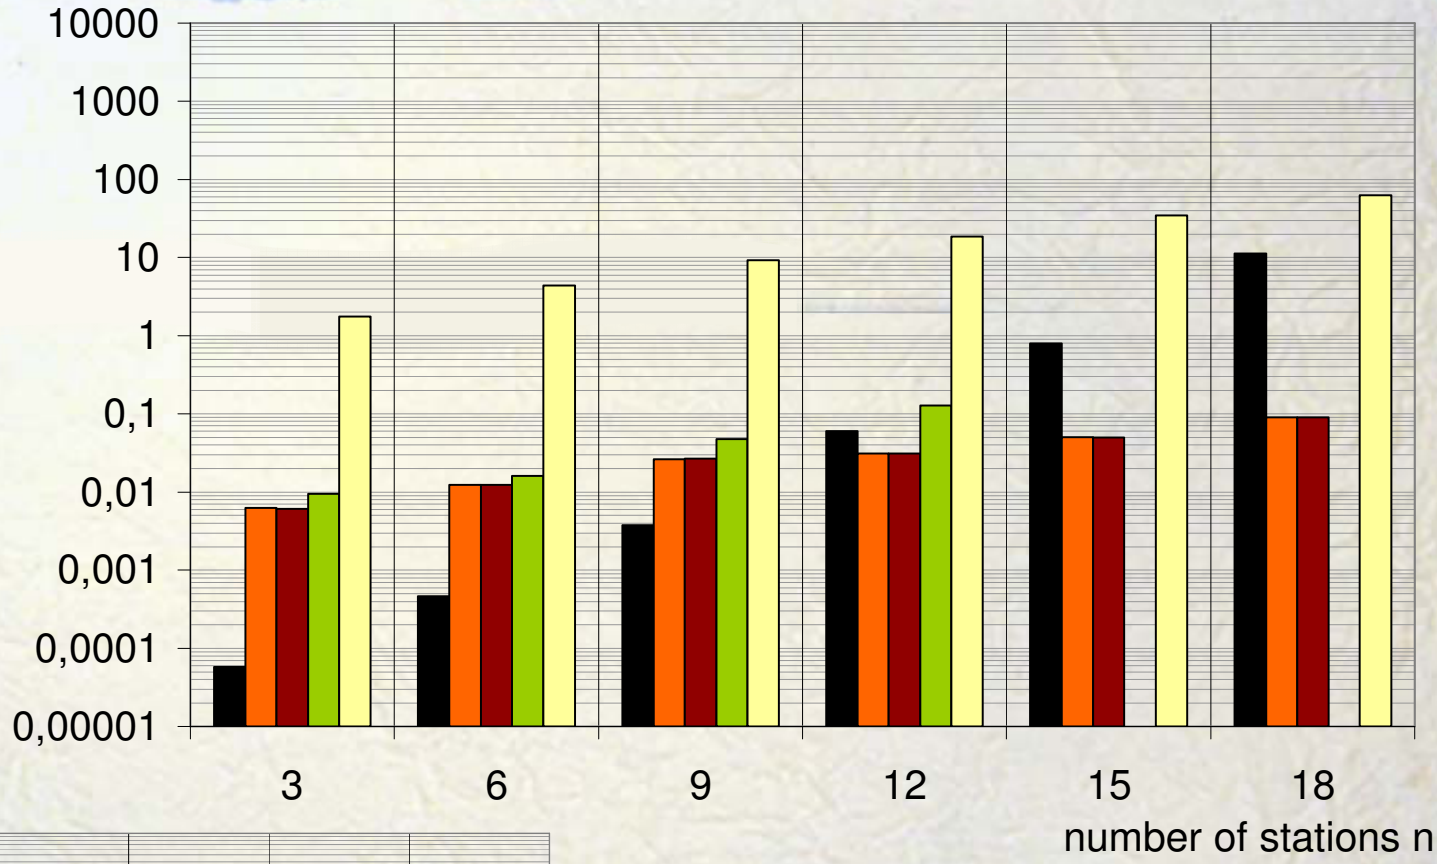
10000
1000
100
10
1
0,1
0,01
0,001
0,0001
0,00001
3
6
9
12
15
18
number of stations n

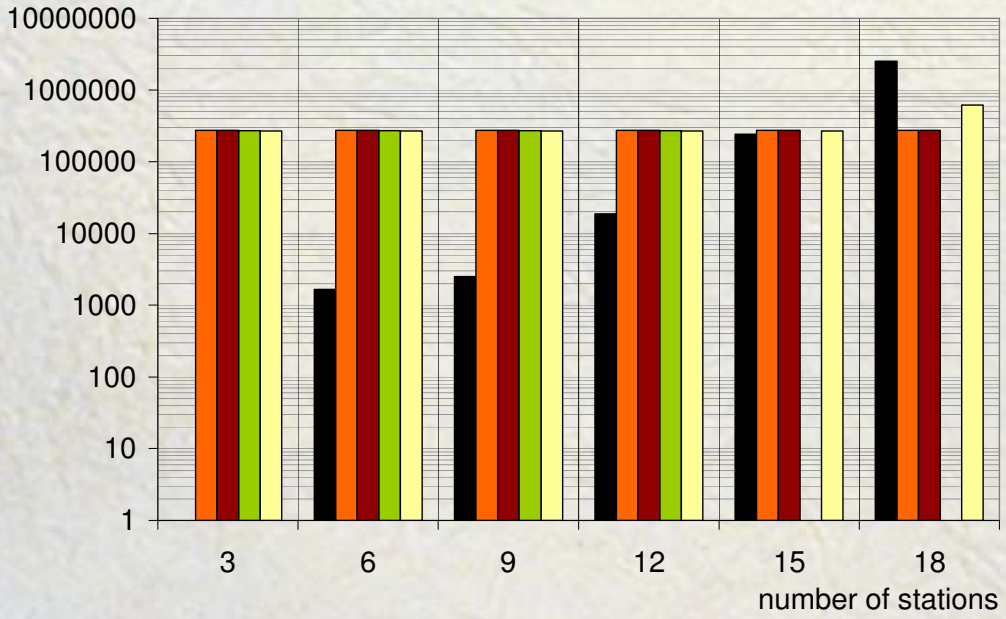
10000000
1000000
100000
10000
1000
100
10
1
3
6
9
12
15
18
number of stations n

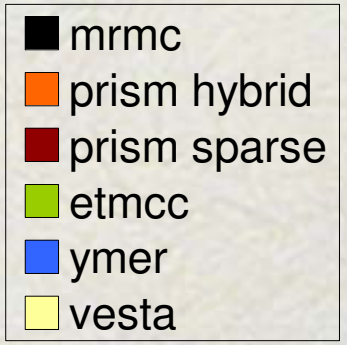
mrmc
prism hybrid
prism sparse
etmcc
ymer
vesta

Analysis

ETMCC	slow, out of memory
PRISM	only symbolic → sparse=hybrid
MRMC	fastest tool for small models
VESTA	excessive number of samples, slow
YMER	not implemented

PRISM: MTBDD Size

- Multi-Terminal BDD = data structure for transition matrix
- size heavily depends on model
- large MTBDD → slow

Model	# states	# MTBDD nodes
Synchronous leader election	500.000	1.000.000 .
Cyclic polling	7.000.000	< 3.000 .

CPS versus SLE runtime

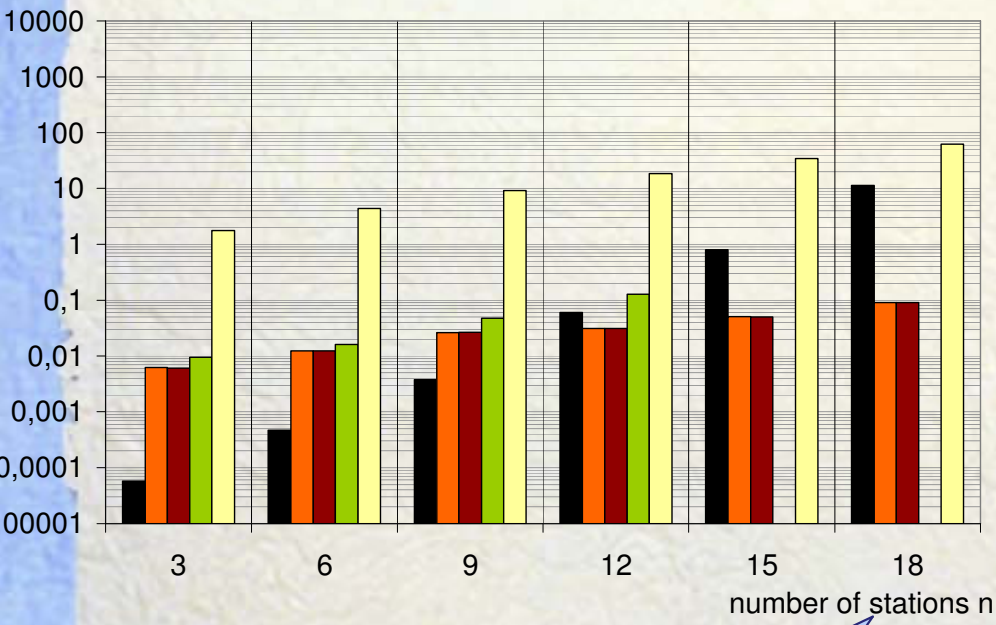

7.077.888 states
2.745 MTBDD nodes

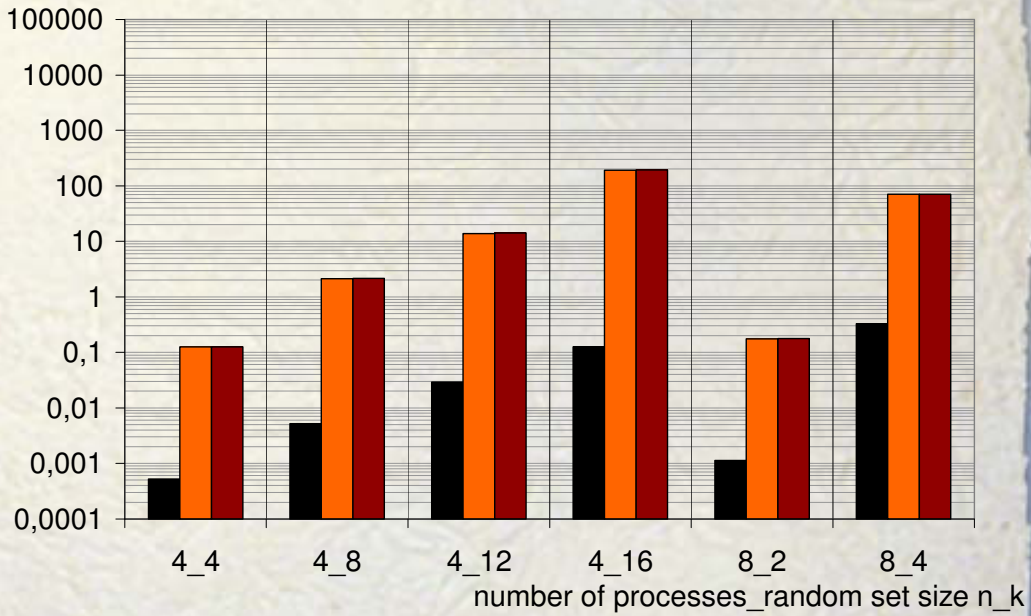

458.847 states
1.131.806
MTBDD nodes

- mrmc
- prism hybrid
- prism sparse
- etmcc
- ymer
- vesta

VESTA: simulation problem

- actual probability close to bound $P_{\geq p}(...)$
- estimate is almost always in $[p-\varepsilon, p+\varepsilon]$
- some irregularity stops the simulation
- $0.95 = \mathit{Prob}(\text{yes} \mid \text{actual } \mathit{Prob} \geq p)$
  $\neq \mathit{Prob}(\text{actual } \mathit{Prob} \geq p \mid \text{yes})$

Result Overview: Timing

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	–/0	N/A

depends heavily on MTBDD size

Result Overview: Memory

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	0/+	N/A

MTBDD size varies heavily

almost independent from model size

Experiment 2
Bounded Reachability

- Tandem Queueing Network
  - two queues after each other

- $P_{<0.01}(\mathit{true}\ \mathrm{U}^{\leq 2}\ \mathit{full})$

  Is the probability
  that the system gets full in 2 time units
  small?

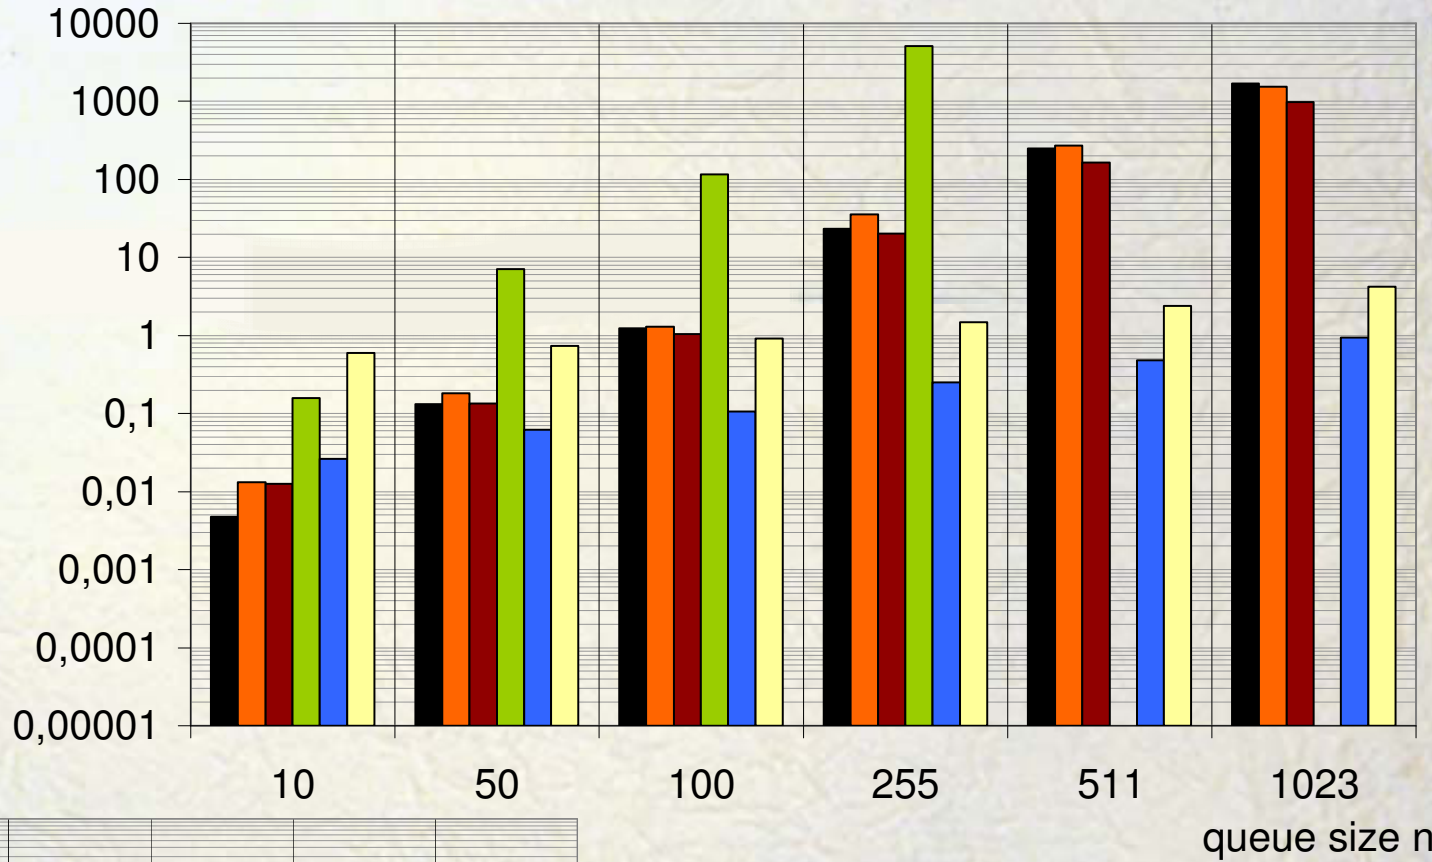
10000
1000
100
10
1
0,1
0,01
0,001
0,0001
0,00001
10
50
100
255
511
1023
queue size n

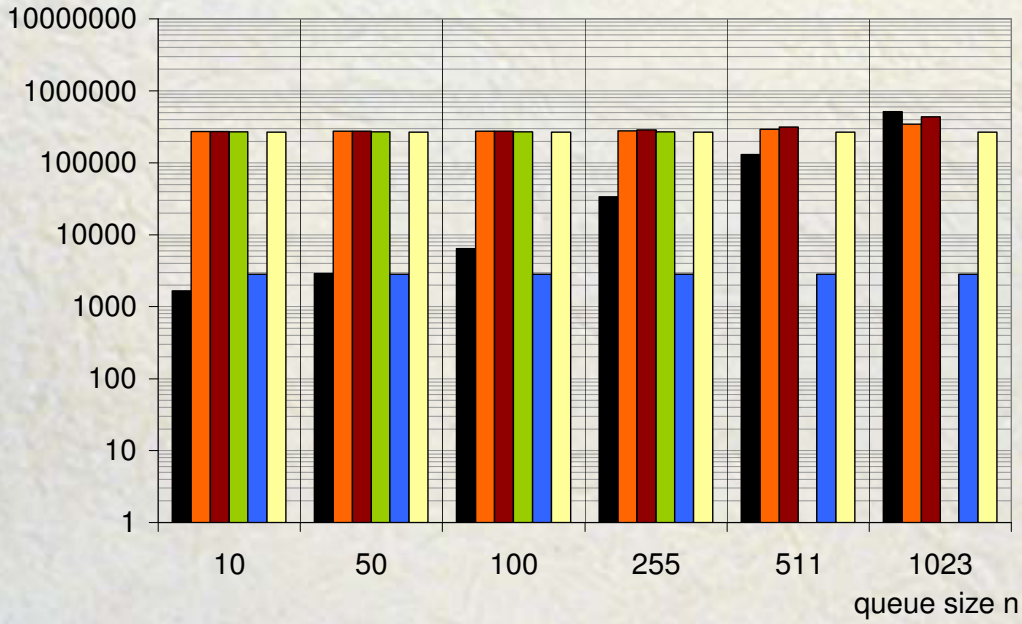
10000000
1000000
100000
10000
1000
100
10
1
10
50
100
255
511
1023
queue size n

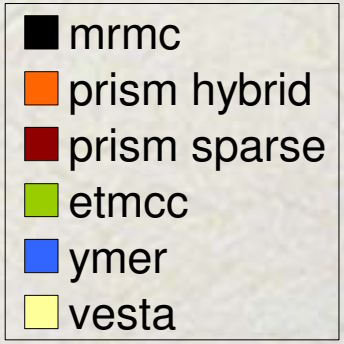
mrmc
prism hybrid
prism sparse
etmcc
ymer
vesta

Analysis

ETMCC	slow, out of memory
PRISM	sparse=faster, hybrid=smaller
MRMC	fast for small models
VESTA	ok if you can afford statistical errors
YMER	best choice if you can afford statistical errors

Result Overview: Timing

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	–/0	N/A
bounded U	–	+	0/+	+/++	+	++

Result Overview: Memory

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	0/+	N/A
bounded U	–	+	+/++	+	+	++

Experiment 3
Steady State Property

- Tandem Queuing Network

- $S_{>0.2}(\ P_{>0.1}(\mathbb{X}\ \text{2nd queue full})\ )$

  In equilibrium,
  the probability to satisfy $P_{>0.1}(\mathbb{X}\ \ldots)$
  is > 0.2

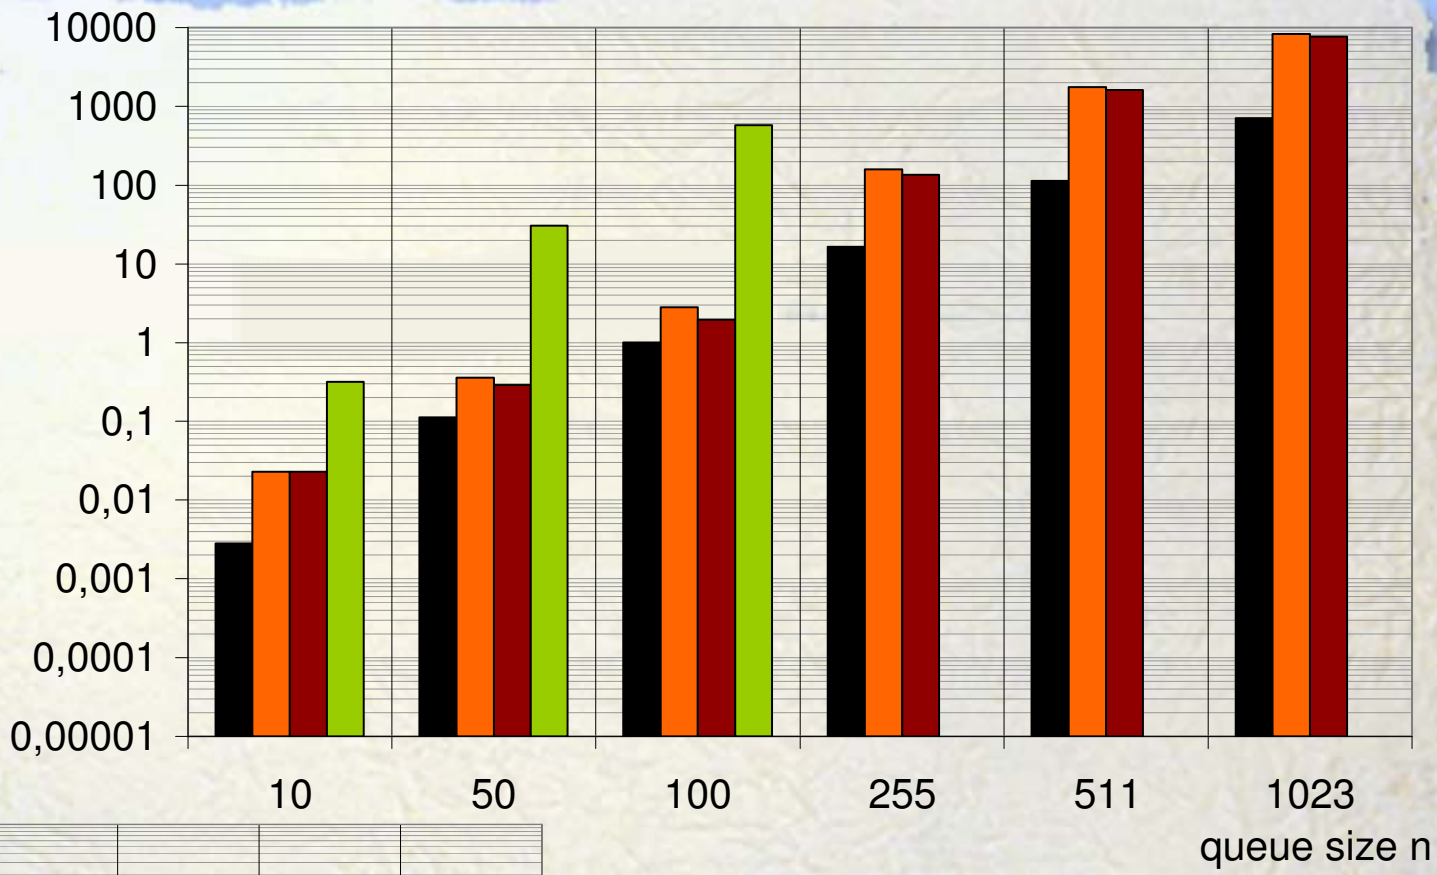
10000
1000
100
10
1
0,1
0,01
0,001
0,0001
0,00001
10
50
100
255
511
1023
queue size n

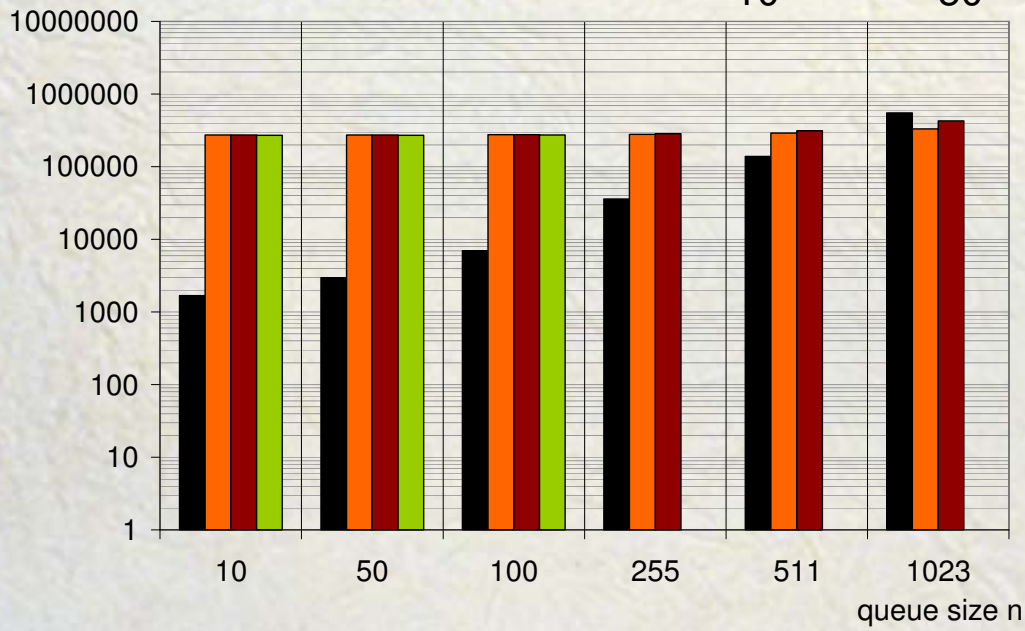
10000000
1000000
100000
10000
1000
100
10
1
10
50
100
255
511
1023
queue size n

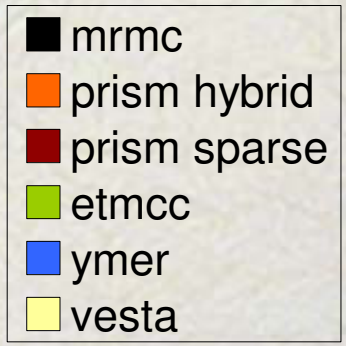
mrmc
prism hybrid
prism sparse
etmcc
ymer
vesta

Analysis

ETMCC	slow, out of memory
PRISM	sparse=faster hybrid=slightly smaller
MRMC	fastest
VESTA	not implemented
YMER	not implemented

Simulating Steady State?

- simulation of bounded reachability has clear stopping criterion
- simulation of unbounded reachability ≈ reachability with very large bound

- simulation of steady state?
  → never stops

Result Overview: Timing

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	–/0	N/A
bounded U	–	+	0/+	+/++	+	++
steady state	–	++	0/+	+	N/A	N/A

Result Overview: Memory

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	0/+	N/A
bounded U	–	+	+/++	+	+	++
steady state	–	+	+/++	+	N/A	N/A

Nested Formulas

- we also checked nested properties

$$P_{\geq 0.8}(P_{\geq 0.9}(\mathit{true}\ \mathrm{U}^{\leq 100}\ n_{70})\ \mathrm{U}\ n_{50})$$

- not detailed here

Result Overview: Timing

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	–/0	N/A
bounded U	–	+	0/+	+/++	+	++
steady state	–	++	0/+	+	N/A	N/A
nested	–	++	0/+	+	––	N/A

based on a single property only: did not terminate

Result Overview: Memory

	ETMCC	MRMC	PRISM hybrid	PRISM sparse	VESTA	YMER
unb'nded U	–	+	+/++	+/++	0/+	N/A
bounded U	–	+	+/++	+	+	++
steady state	–	+	+/++	+	N/A	N/A
nested	–	+	+/++	+	N/A	N/A

Conclusions

ETMCC	worst, only small models
MRMC	fastest for small models
PRISM hybrid	fast if MTBDD is small
PRISM sparse	fast
VESTA	rather slow, statistical errors
YMER	slim & very fast, only bounded reach, few statistical errors