Checking & Distributing Statistical Model Checking [*]

Peter Bulychev[1], Alexandre David[1], Kim G. Larsen[1], Axel Legay[2],
Marius Mikučionis[1], Danny Bøgsted Poulsen[1],

[1] Computer Science, Aalborg University, Denmark
[2] INRIA/IRISA, Rennes Cedex, France

Abstract. In this paper we propose a general framework for distributed
statistical model checking of networks of priced timed automata. The
first contribution is a new algorithm to distribute sequential hypothesis
testing without introducing bias in the results. The second contribution is
an implementation of this algorithm in UPPAAL. The major contribution
is an experimental and analytical evaluation of the approach through
case studies, including an analysis of the SMC algorithm itself.

1 Introduction

Statistical Model Checking techniques (SMC) [8,12,17], can be seen as a trade-
off between testing and formal verification. The core idea of the approach is to
conduct some simulations of the system and verify if they satisfy some given
property. The results are then used together with statistical algorithms in order
to decide whether the system satisfies the property with some probability. Statis-
tical model checking techniques can also be used to estimate the probability that
a system satisfies a given property [8,6]. Of course, in contrast to an exhaustive
approach, a simulation-based solution does not guarantee a correct result with
100% confidence. However, it is possible to bound the probability of making an
error. SMC gets widely accepted in various research areas and applied to prob-
lems that are beyond the scope of classical formal techniques [1,2,10,11,13,19,20].

Unfortunately, extremely huge sized problems and a demand of extremely
high confidence may require generation of a large number of simulation runs, each
of which may itself be extremely time consuming. To address this *confidence-
explosion* problem, we suggest in this paper to take advantage of PC-clusters
and GRID computers. In fact, it is well-known that statistical solutions methods
that use samples of independent observations are often trivially parallelizable.
As observed in [18], SMC algorithms can be distributed through the help of a
master/slave architecture where multiple computers are used to generate the
simulations. The idea is as follows: one or more slave processes register their
ability to generate simulation with a single master process that is used to collect
those simulations and perform the statistical test.

---

[*] Work partially supported by VKR Centre of Excellence – MT-LAB, an "Action de
Recherche Collaborative" ARC (TP)I

However, this process may become complex when considering sequential hypothesis testing (when the number of simulations is not known in advance). The problem is that there might exist a correlation between a time needed to generate a random simulation and the fact that a property is satisfied by this simulation. Thus it is important to guarantee that the technique will not introduce a bias towards the results that are generated by shorter simulations.

In a series of recent works [4], we have extended UPPAAL with SMC algorithms applied to Networks of Priced Timed Automata – hence leading to the first implementation of SMC for real-timed stochastic systems. The objective of this paper is to go one step further and propose the first complete study of distributed SMC, in general, and in the framework of UPPAAL in particular. Our contributions are:

1. A *distributed implementation* of the estimation algorithm proposed in [8]. Building on classical Monte Carlo techniques [7], an estimation algorithm precomputes the number of simulations needed to estimate the probability to satisfy a property with a given confidence. Such an algorithm which is trivially parallelizable amounts to equally distribute the number of simulations to perform between the slave computers.

2. A *new distributed algorithm* for sequential hypothesis testing where simulations are computed on the fly until a threshold is passed and a decision is taken. Here, it is important to avoid introducing bias in the results, which may be potentially complex and eventually decrease the benefit of using several processors. To counter this, [18] proposed a round-Robin solution where the runs are counted in rounds. We generalise the solution in [18] by introducing *batches* and *buffers*. The batch is used to reduce communication by sending an aggregate result of predefined size (instead of individual results). The buffer is used to improve concurrency since the nodes are more loosely synchronized.

3. A *thorough evaluation* of our implementation through new applications of SMC algorithms. In particular, we apply the distributed SMC engine to an analysis of an instance of the LMAC protocol of unprecedented size. Additionally, a thorough evaluation of the distributed SMC framework itself is made aiming at identifying optimal settings of the parameters for the framework. The evaluation is carried out both experimentally (using the implementation) as well as analytically (using SMC) based on a model of the distributed SMC algorithm itself, and with high consistency in identifications made by the two approaches.

2   Statistical Model-Checking in Uppaal

This section introduces the formalisms used in UPPAAL for modeling systems and and specifying properties. Then, we briefly survey existing Statistical Model Checking (SMC) algorithms. Finally, a novel application of SMC is presented.

2.1   Networks of Priced Timed Automata

The new SMC engine of UPPAAL [3] supports the analysis of Priced Timed Automata (PTAs) that are timed automata whose clocks can evolve with dif-

ferent rates, and with no restrictions in guards and invariants. Additionally, we support other features of the UPPAAL model checker's input language such as integer variables, data structures and user-defined functions. We also assume PTAs are input-enabled, deterministic (with a probability measure defined on the sets of successors), and non-zeno. PTAs communicate via broadcast channels and shared variables to generate Networks of Price Timed Automata (NPTA).

Fig. 1 provides an NPTA with three components A, B, and T as specified using the UPPAAL GUI. One can easily see that the composite system $(A|B|T)$ has the transition sequence:

$$((A_0, B_o, T_0), [x = 0, y = 0, C = 0]) \xrightarrow{1} \xrightarrow{a!}$$
$$((A_1, B_0, T_1), [x = 1, y = 1, C = 4]) \xrightarrow{1} \xrightarrow{b!}$$
$$((A_1, B_1, T_2), [x = 2, y = 2, C = 6]),$$

demonstrating that the final location T_3 of T is reachable. In fact, location T_3 is reachable

Fig. 1: An NPTA, $(A|B|T)$.

within cost 0 to 6 and within total time 0 and 2 in $(A|B|T)$ depending on when (and in which order) A and B choose to perform the output actions a! and b!. Assuming that the choice of these time-delays is governed by probability distributions, a measure on sets of runs of NPTAs is induced, according to which quantitative properties such as *"the probability of T_3 being reached within a total cost-bound of 4.3"* become well-defined.


Fig. 2: Cumulative probabilities for `time` and `Cost`-bounded reachability of T_3.

In our early works [4], the stochastic semantic of PTA components associates probability distributions on both the delays one can spend in a given state as well as on the transition between states. In UPPAAL uniform distributions are applied for bounded delays and exponential distributions for the case where a component can remain indefinitely in a state. In a network of PTAs the components repeatedly race against each other, i.e. they independently and stochastically decide on their own how much to delay before outputting, with the "winner" being the component that chooses the minimum delay. For instance, in the NPTA of Fig. 1, A wins the initial race over B with probability 0.75.

Properties For specifying properties of NPTAs, we use cost-constraint temporal properties over runs of the form $\psi = \Diamond_{C \leq c} \varphi$. Here C is an observer clock (that is never reset and should grow to infinity on any infinite run), $c \in \mathbb{R}_{\geq 0}$ and φ is a state-predicate. We say that a run π satisfies $\psi = \Diamond_{C \leq c} \varphi$ if there exists a state (ℓ, v) in π satisfying φ and with $v(C) \leq c$. For an NPTA M we define $\mathbb{P}_M(\psi)$ to be the probability that a random run of M satisfies ψ.

Reconsider the example of Fig. 1, we can evaluate the the probabilities `Pr[time<=2](◇ T.T3)` and `Pr[C<=6](◇ T.T3)` in UPPAAL, obtaining as ex-

pected 0.75 for the composition $(A|B|T)$ for both of these probabilities. In fact Fig. 2 gives the time- and cost-bounded reachability probabilities for $(A|B|T)$ for a range of bounds.

2.2 Statistical Model Checking Algorithms

We briefly recall statistical algorithms allowing to answer the following two types of questions : (1) *Qualitative :* Is the probability that a random run of a given NPTA $\mathcal{A}$ satisfies a property $\Diamond_{C \leq c} \varphi$ greater than a certain threshold θ? and (2) *Quantitative :* What is the probability that a random run of $\mathcal{A}$ satisfies $\Diamond_{C \leq c} \varphi$? For both question a run of the system is encoded as a Bernoulli random variable that is true if the run satisfies the property and false otherwise.

Qualitative Question This reduces to test the hypothesis H: $p = \mathbb{P}_{\mathcal{A}}(\Diamond_{C \leq c} \varphi) \geq \theta$ against $K : p < \theta$. To bound the probability of making errors, we use strength parameters α and β and we test the hypothesis $H_0 : p \geq p_0$ and $H_1 : p \leq p_1$ with $p_0 = \theta + \delta_0$ and $p_1 = \theta - \delta_1$ (δ_0 and δ_1 are parameters of the algorithm). The interval $p_0 - p_1$ defines an indifference region, and p_0 and p_1 are used as thresholds in the algorithm. The parameter α is the probability of accepting H_0 when H_1 holds and the parameter β is the probability of accepting H_1 when H_0 holds. The above test can be solved by using Wald's sequential hypothesis testing [16]. This test, which is presented in Algorithm 1, computes a proportion r among those runs that satisfy the property. With probability 1, the value of the proportion will eventually cross $\log(\beta/(1 - \alpha))$ or $\log((1 - \beta)/\alpha)$ and one of the two hypothesis will be selected.

Algorithm 1: Hypothesis testing

   function hypothesis(S:model , ψ: property)

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

   end

Quantitative Question This reduces to a Monte Carlo approach that computes the number N of runs needed in order to produce an approximation interval $[p - \epsilon, p + \epsilon]$ for $p = Pr(\psi)$ with a confidence $1 - \alpha$. The values of ϵ and α are chosen by the user and N relies on the Chernoff-Hoeffding bound.

2.3 Analysing SMC in Uppaal

In this section we will use the SMC engine of UPPAAL to our first non-trivial task, namely to analyse itself! More precisely, by suitably modeling the sequentual

testing algorithm as well as a sample model M, we will be able to use the SMC engine of UPPAAL to analyse the performance of SMC on M. Later, in Section 4, this will allow us to evaluate various naive (and even faulty) proposals for distributing SMC.

The sample model M given in Fig. 3a[3] makes an initial probabilistic choice between the two branches, each having a looping transition taken repeatedly with a delay choosen uniformly from $]0, 2]$. Performing sequential testing of the hypothesis H_0: `Pr[<=100](◇ OK)≥ 0.5` some 10 times with $\alpha = 0.05$ as level of significance and with an indifference region of ± 0.01, we consistently (and correctly) dismiss the hypothesis with an average of 408.6 runs and with standard deviation 127.5.


(a) M. (b) Master. (c) Generator.

Fig. 3: Sample model M (a) satisfying `Pr[<=100](◇ OK)= 0.42` and modeling SMC of M (b,c) with respect to H_0: `Pr[<=100](◇ OK)≥ 0.5` with 0.05 as level of significance and $[0.49, 0.51]$ as indifference region.

Now, aiming at obtaining a better understanding of sequential testing[4] we may simply model the sequential testing algorithm of M directly and analyse its (expected) performance using UPPAAL SMC. The resulting model is given in Fig. 3 and consists of an extension of the sample model M into the component Generator that will repeatedly generate random runs of M (of time-duration 100) and report the outcome to a Master using the channels `add` (when 100 time-units has elepased without `OK` having been observed) and `sub` (used as soon as it is observed that the `OK` branch has been taken, note the absence of the `time>=100` guard on the right side of the Generator model). The Master has the obligation of adjusting appropriately the ratio-variable r according to Alg. 1, and conclude on H_0 or H_1 as soon as the value of r exceeds the given threshold. Given the indifference region $[0.49, 0.51]$ and level of significance 0.05, we find that the approximate values to be used[5] in Alg. 1 are: $-\log(p_1/p_0) = \log(1 - p_0/1 - p_1) = 0.01715$ and $\log((1 - \beta)/\alpha) = -\log(\beta/(1 - \alpha)) = 1.2787$ ($\approx 1.262 + 0.017$). In the model of Fig. 3 we are using scaled integer constants for these values. Now, looking at the estimation of `Pr[#<=20000](◇ Master.H₁)` in Fig. 4, we find – as expected – that the probability of accepting H_1 (H_0) tends

---

[3] M is a timed variant of the model proposed in [17] and used to demonstrate bias in a naive distributed approach to SMC.

[4] The performance of sequential testing has been subject to significant studies and is well-understood [15]. The aim here is to demonstrate that our UPPAAL SMC engine is a useful tool for obtaining such an insight.

[5] Those values are obtained by observing Wald's ratio on several application of the SMC algorithm to the same problem, and then take the average of the observations.

to 1 (0) as the number of steps increases. We also see that the average number of runs is estimated to 481.4. The "mismatch" with the experimentally found average 408.6 is due to early termination when the threshold for H_0 is exceeded.

(a) $\texttt{Pr}[\# \leq 80000](\diamond \ \texttt{Master.H}_1)$

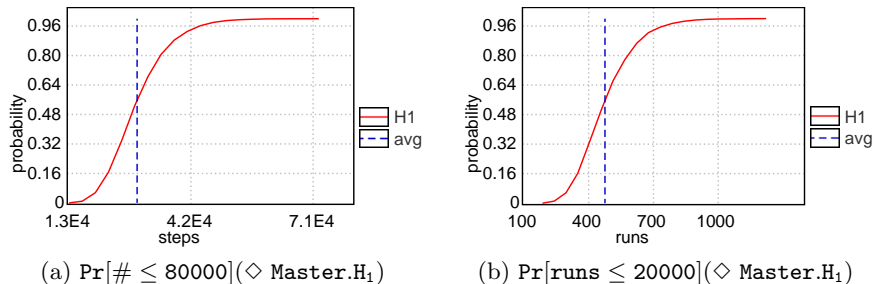(b) $\texttt{Pr}[\texttt{runs} \leq 20000](\diamond \ \texttt{Master.H}_1)$

Fig. 4: Cumulative probability plots over number of steps and runs.

3 Distributed Statistical Model-Checking in Uppaal

SMC suffer from the fact that high confidence required by an answer may demand a large number of simulation runs, each of which may itself be time consuming. As an example, the first hypothesis test shown later in this section can generate between 14,000 and 2,390,000 runs if the parameters α, β, δ range between 0.01 and 0.001. A possible way to leverage this problem is to use several computers working in parallel using a master/slaves architecture: one or more slave processes register their ability to generate simulation with a single master process that is used to collect those simulations and perform the statistical test. When working with an estimation algorithm, this collection is trivially performed as the number of simulations to perform is known in advance and can be equally distributed between the slaves.

When working with sequential algorithms, the situation gets more complicated. Indeed, we need to avoid introducing bias when collecting the results produced by the slave computers. This means that results should not be collected arbitrarily as illustrated by considering the model of Section 2.3 with several instances of the `Generator` template. Checking the property `Pr[runs<=20000](`$\diamond$ `Master.H1)` Fig. 5 shows that different distributions can be obtained with different numbers of generator nodes,

Fig. 5: Probability distributions obtained with 1 (top), 5, 10, and 20 (bottom) generator nodes.

hence revealing a bias in the results. In fact the probability of accepting H_1 tends (incorrectly) to 0 when the number of `Generator` components increases.

A solution, which was proposed in [17], consists in observing that Wald's ratio r is updated as a function of the Bernouilli random variable x as $r+ = x*r_{acc}+(1-x)*r_{rej}$ with r_{acc} and r_{rej} being constants depending on the tested hypothesis. To reduce blocking and still update r, the non-biased algorithm updates two safe approximations for r (r_1 and r_2). If x is unknown then it updates with $r_1+ = r_{rej}$ and $r_2+ = r_{acc}$, and then testing if $r_1 \leq I$ to accept H_0 or if $r_2 \geq S$ to accept H_1[6]. When all outcomes of a round are known then we can reset $r_1 := r_2 := r$. This allows us to accept H_0 even if some accepting outcomes are missing or conversely to accept H_1 if some rejecting outcomes are missing.

We generalize [17] by aggregating the outcomes x by *batches* (of size B) and also by implementing a *buffer* (of size K) of incoming results. The batch is used to reduce communication by sending B aggregate results. The buffer is used to improve concurrency since the nodes are more loosely synchronized and they can be K runs ahead of the slowest node. Fig. 6 illustrates our algorithm at the master node that receives asynchronous messages from all other nodes in a buffer. A message is an aggregate result containing the outcome of B runs. The master may take a decision as soon as $r_1 \leq I$ or $r_2 \geq S$. When all outcomes at the bottom line of the buffer are known we reset $r_1 := r_2 := r$ with the


Fig. 6: Buffer of results at the master node.

exact updated value of r with those outcomes, and free the bottom line of the buffer. In practice, our algorithm is calibrated to count the runs up to a certain depth in the buffer. Indeed, the outcomes are weighted by B so few missing aggregated outcomes can prevent the algorithm from deciding. We have implemented this algorithm with asynchronous communications (using OpenMPI). There can be at most K pending messages due to the size of the buffer. If a slave tries to send more messages, then the communication will block waiting for a "slot" to be free. The experiment performed in the remainder of the paper has been carried out on varying numbers of nodes on a cluster with dual Xeons 5650 (hexa-cores at 2.66GHz) interconnected with infiniband.

We first make two types of experiments to exhibit the performance characteristics of our algorithm. The experiments are carried out using the train-gate example available as a demo of UPPAAL. This model comprises a number of trains crossing a bridge with only one track. A gate controller stops and restart the trains to ensure mutual exclusion on the bridge and absence of starvation for the trains. Our first experiment concerns 6 trains and the property of being in a state where train 5 is crossing while all the other trains are stopped.

```
Pr[<=100](<> Train(5).Cross and
```

---

[6] $I = \log(\beta/(1 - \alpha))$ and $S = \log((1 - \beta)/\alpha)$ as stated in Alg. 1.

```
(forall (i : id_t) i != 5 imply Train(i).Stop)) >= 0.46188
```

The runs are relatively short with few components so they will be cheap to compute and we expect the throughput of messages to be high. In addition, the hypothesis we are testing is not deterministic, which means that the outcomes and computation times of the runs will vary. The property is checked with high confidence (99.999%) and small indifference regions (+/- 0.00001) to have a precise and reliable result – and to stress our distributed algorithm.

Our second experiment considers a "large" instance with 20 trains, where we check if the model satisfies mutual exclusion on the bridge, expressed by the property

```
Pr[<=1000]([] forall (i : id_t) forall (j : id_t)
         Train(i).Cross and Train(j).Cross imply i == j) >= 0.9999
```

Here, the runs are random but bounded by the same large bound and since the inner property `[]forall(i : id_t)forall(j : id_t)...` holds by model-checking, all the runs will all reach their bounds. In addition, we have 20 trains and the runs are long (1000 time units) so they are relatively expensive to generate. This means that all the runs are implicitly synchronized and small deviations are amortized by the long runs. The throughput of messages will be low, which means a low overhead compared to the actual useful work of generating the runs.
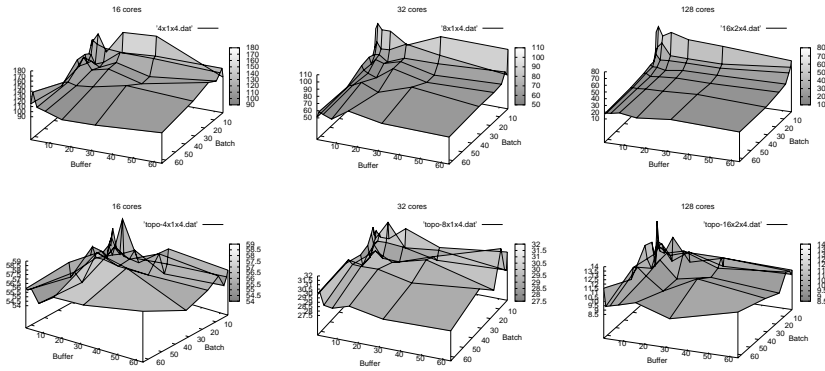

Fig. 7: Verification times on 16, 32, and 128 cores in function of B and K for the "small" model (first row) and the "large" model (second row).

Figure 7 shows our results for different number of cores. The solution in [18] corresponds to the particular case with K and B are equal to one, exhibiting in all the experiments the worst verification time, and with performance deteriorating with increasing number of cores (i.e. for 128 cores performance loss is a factor of 4). Though the impact of the buffer size is less, the experiments indicate that a buffer of size 2-4 will suffice. The results also demonstrate linear scalability of

8

(a) Master.

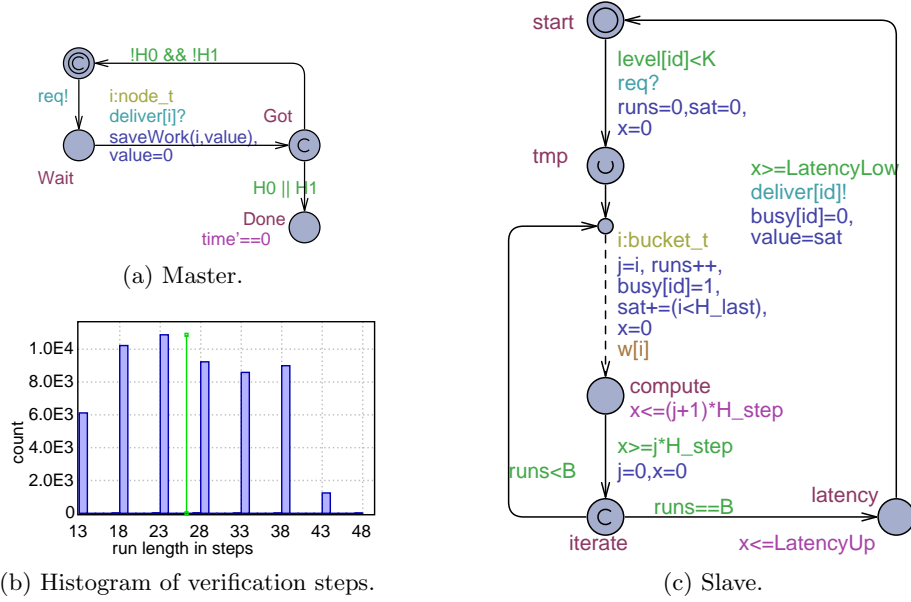(b) Histogram of verification steps.

(c) Slave.

Fig. 8: Timed automata model of a statistical model checking process.

our distributed implementation: for $B = 32$ and $K = 2$ the verification times for 16, 32 and 128 cores are 108, 56 and 19 seconds (respectively).

4 Analyzing Distributed SMC in Uppaal

In this section we model the implemented distributed algorithm of sequential hypothesis testing and we check it using the SMC engine of Uppaal. The goal is to estimate the verification time and processor utilization, check for bias in the distributed algorithm, and explore the parameters of our distributed SMC algorithm in an analytical manner.

Modeling. We model the master and slave processes described in Section 3 as shown in Figure 8. The master sends a broadcast request `req!` to verify batches of runs (of size B). We use a standard modeling pattern to synchronize on the corresponding `req?` as soon as possible. The master gathers the results with its `saveWork` function and loops again if neither H_0 nor H_1 is accepted. Listing 1.1 shows this `saveWork` function that implements the distributed hypothesis testing algorithm of Section 3. Uppaal uses floating point numbers that are not available in the modeling language. Instead we encode fixed point arithmetics with integers and we use precomputed tables for logarithm values (shown in Listing 1.3 in the appendix). Once the master accepts H_0 or H_1, it moves to the location Done and stops the clock `time`.

Slave processes proceed to compute their batches if their communication buffers are not full (`level[id] < K`) or wait for the condition to hold. The

9

Listing 1.1: Master code.

```
1  // buffer  portion  for  early  termination:
2  const int  P = (K<=4)?K : ((K<=8)?5 : ((K<=16)?8 : ((K<=32)?10 : 12)));
3  bool H0 = false,  H1 = false;  // for  hypothesis  H0 and H1
4  int  batch[N][K];  // buffer  of  batches  (K batches for  N nodes)
5  long satisfied  =0, unsatisfied  =0; // information  about filled   lines
6  long sat=0, unsat=0, unknown=N*P*B; // early results in unfilled   lines
7  long logRatio  = 0, ratioLow = 0, ratioUp  = 0; // scaled  by p.scale
8  void  saveWork(const node_t node, const int  value) {
9      if  (level [node]<=P) { // entered the early results   portion
10         sat  += value; unsat += B−value; unknown −= B;
11     }
12     batch[node][level [node]]  = value; level [node]++; // store
13     if  (level [node]==1) { // entered at the lowest  level
14         bool filled     = forall (i : node_t) level [i]>0;
15         if  (filled  ) { // line  at the lowest  level   has been filled
16             int  L;
17             for  (i : node_t) { // shift   all   queues one by one
18                 satisfied    += batch[i][0];  // count as firm  results
19                 unsatisfied   += B−batch[i][0];
20                 sat  −= batch[i][0];   // discount  from early   results
21                 unsat −= B−batch[i][0];   unknown += B;
22                 level [i]−−; // remove from buffer
23                 for  (L=0; L<level[i];  ++L) {
24                     batch[i][ L] = batch[i][ L+1]; // shift
25                     if  (L==P) { // entered the early results    portion
26                         sat  += batch[i][L+1]; unsat += B−batch[i][L+1];
27                     }
28                 }
29                 batch[i][ level [i]]=0;  // cleanup
30             }
31             logRatio  = p.valAcc*satisfied   + unsatisfied *p.valRef;
32             if  (logRatio  <= p.logInf) H0 = true;
33             if  (logRatio  >= p.logSup) H1 = true;
34         }
35     }
36     ratioLow = p.valAcc*(satisfied      +sat+unknown) +
37                p.valRef*(unsatisfied   +unsat);
38     ratioUp  = p.valAcc*(satisfied      +sat) +
39                p.valRef*(unsatisfied   +unsat+unknown);
40     if  (ratioUp   <= p.logInf) H0 = true;
41     if  (ratioLow >= p.logSup) H1 = true;
42 }
```

compute location models the computation time of a run, chosen according to the distribution shown in Figure 8b. This is encoded using probabilistic edges with

weights matching the distribution. The distribution comes from a real verification of the property in Section 3:
```
Pr[<=100](<> Train(5).Cross and
          (forall (i : id_t) i != 5 imply Train(i).Stop)) >= 0.46188
```
The last weighted edge (case `i=H`) is reserved for the runs that did not satisfy the property.

Verification. In the hypothesis we test, the actual probability is very close to 0.46188. Since the real probability falls in the indifference region of our test, we would expect that a non-biased implementation would accept H_0 or H_1 equally often. Estimating the probability of confirming the hypothesis H_0 with the query `Pr[<=10000000](<> master.H0)` gives the probability 0.503 ± 0.005 with 99.9% confidence, confirming that our algorithm is not biased as well as the validity of our model.

Similarly, we obtain the distribution of the verification time by the query `Pr[<=10000000](<> master.Done)` for a model with number of nodes $N = 128$, batch size $B = 64$, and buffer size $K = 4$. The result is 9557.6 time units in average and the distribution histogram is depicted in Figure 9a. To estimate the processor usage time, we add another process with a single location with the invariant `usage'==sum(i:node_t)busy[i]`. Here, `usage` is a clock that grows with a rate equal to the number of busy nodes.

(a) Estimated time distribution.   (b) Processor usage time distribution.
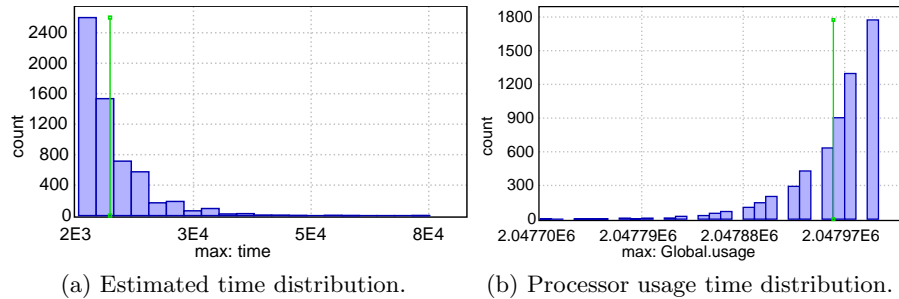
Fig. 9: Time estimation from 6000 runs of DSMC model.

The question is now to find a good settings for the parameters of our algorithm (B and K). We perform parameter sweep to estimate the verification time for values of B and K taking values in $1, 2, 4, 8, 16, 32, 64$ for three topologies with the number of processing nodes $N = 16, 32$, or 128. The results are depicted in Fig. 10, where it is visible that extremely small batch size requires more time. Large batch sizes can also be detrimental in a large cluster setting (Fig. 10c where too many runs are requested in bulk than actually needed to establish the result). Buffer size of one has a huge penalty of blocking with small blocks, but it is barely noticeable otherwise. This confirms the experimental findings of Section 3.

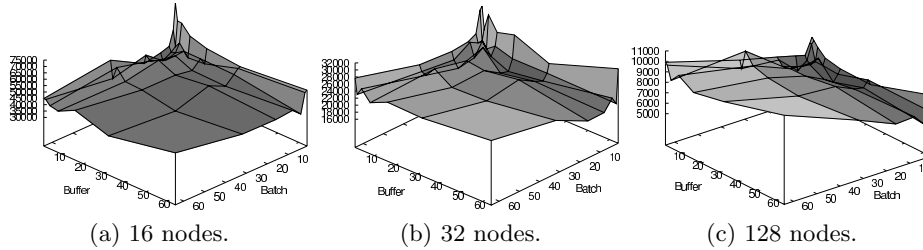(a) 16 nodes.   (b) 32 nodes.   (c) 128 nodes.

Fig. 10: Estimated verification times in model time units.

5   Lightweight Media Access Control

LMAC is a Lightweight Media Access protocol (studied in [4,5]) used for scheduling communication in wireless sensor networks where the topology is determined by physical location and radio connectivity of the individual nodes. One of the goals of the LMAC protocol is to minimize the number of collisions in the network and to reconfigure the network to avoid further collisions.

The original model has been developed in [5] where topologies of 4-5 nodes are studied exhaustively using classical UPPAAL and a number of topologies are identified as problematic, containing perpetual collisions. In this paper we provide new insight as to the likelihood of perpetual collisions in different topologies. This insight could not be delivered by the use of classical UPPAAL and the experiment conducted is of unprecedented size.

In LMAC communication media access time is discretized into time frames and each time frame is divided into time slots. The goal of the protocol is to allocate the time slots to each node efficiently. The challenge is that there is no central node distributing and as-

Fig. 11: LMAC protocol phases.

signing slots and nodes cannot themselves listen while transmitting, hence neighbours are responsible for detecting and informing each other about collisions.

After waiting phase, the node moves to a discovery phase and listens for an entire time frame and notes which time slots are used by its neighbours. The collision counting expression `collisions=++cc;` is added on the edge from `rec_one0` to `done0` in Fig. 12b. After one time frame of discovery phase, the node chooses seemingly unused time slot and moves to an active phase. The node falls back to waiting phase if there are no neighbours (no signal received) or all slots are occupied. During active and discovery phases the node listens and notes any collisions (several receptions during the same slot). During active phase the node transmits information about collisions it has detected during its time slot and listens for collisions and information about collisions during other time slots. From the active phase the node may fall back to discovery phase if
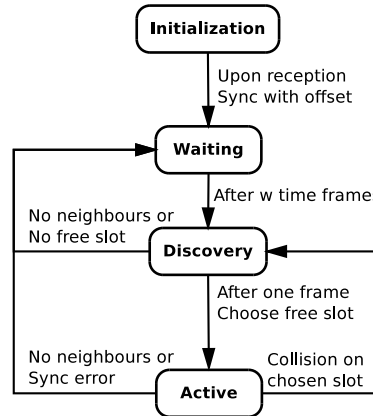
12

it is notified about the collisions on its time slot and falls back to the waiting phase if it detects that neighbours are gone.

Figure 11 shows the four phases of the protocol. Initially all nodes except the gateway are listening and waiting for a radio signal from its neighbourhood during the initialization phase. The communication is triggered by a dedicated gateway node. Upon reception of signal, the node notes the relative time offset of the signal and moves to waiting phase, during which it chooses to wait for a random amount of time frames. The random delay is modeled using probabilistic branching (see Fig. 12a) with geometrical weights (`weight` array).
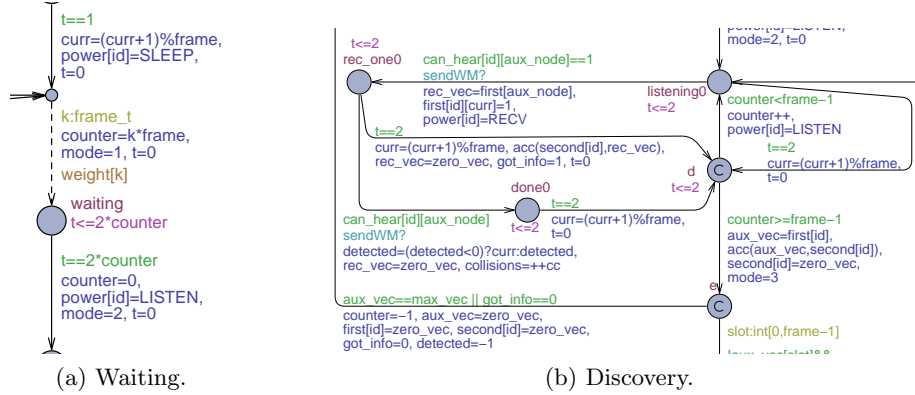

(a) Waiting.   (b) Discovery.

Fig. 12: LMAC phases in the model.

Starting from the model[7] of [5], we removed the verification optimizations constraining the parallelism, annotated it with power consumption and collision counting (as cost variables). The model contains twice as many slots as nodes, whereas one slot per node is enough to schedule flawless communication in any topology if nodes were aware of each others choices.

First we examine the distribution of the first collisions over time. The first row of Fig. 13 is a result of a query `Pr[<=1000](◇ collision>0)` and it shows that most collisions happen early in time and in a ring topology some collisions may be discovered later (possibly when the first signal propagation meets at the opposite of the ring). In the second row of Fig. 13 the distribution of possible number of collisions is examined using a query `Pr[collisions<=100](◇ time>=1000)`: in a chain and a ring topologies the collisions are unlikely to occur ($> 90\%$ probability of 0 collisions), but in a star it is almost guaranteed to occur (only 8% probability of 0 collisions). The third row of Fig. 13 shows the probability distribution of collision counts after twice as long period of time (using query `Pr[collisions<=100](◇ time>=2000)`). Notice that the shape of distributions has not changed, but the small bumps have shifted to the right at exactly twice the number of collisions and almost identical probability density, which implies that those particular collisions are accumulating proportionally to the progress of time, and in other words it means that collisions are reoccurring perpetually
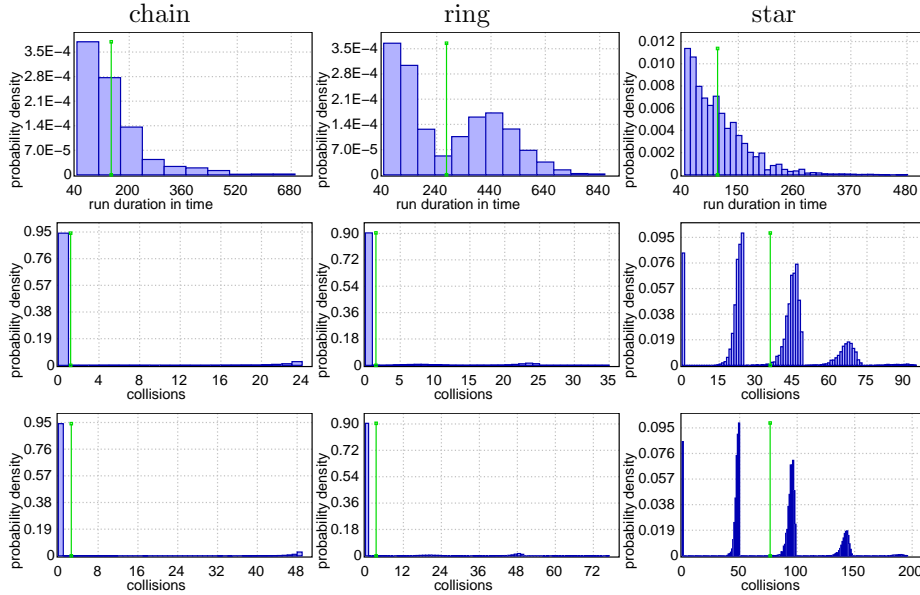

Fig. 13: Collision statistics in three different topologies, in rows: probability of a collision over time, probabilities of a number of collisions up to 1000 and up to 2000 time units.

without recovery. We checked these three properties on a 128 cores cluster with high precision (with $\alpha = \beta = 0.0001$ and $\varepsilon = 0.0005$) in about 30 minutes, which generated around 19 million runs.

We have demonstrated how UPPAAL SMC can be used to identify problematic topologies and distributed implementation can provide a high degree of accuracy in spotting the reoccurring collisions.

6 Conclusion and Future work

In this paper we have developed, implemented, applied and evaluated a general and scalable framework for distributed statistical model checking. We have thoroughly investigated the distribution of sequential algorithms where bias can be introduced when collecting the samples produced by slave computers. In particular, we have identified best choices of batch and buffer sizes both experimentally and analytically, with agreement in the findings of the two approaches. In the future, we plan to implement and distribute other SMC algorithms, principaly the Bayesian algorithms introduced in [20,9].

Finally, it is worth mentioning that we have tried to use other distributed SMC model checkers such as Ymer [18] or PVesta [14]. Aside from the fact that the Gui of those two tools is quite restricted, we observed that Ymer does not work anymore and that PVesta only distributes those algorithms where the number of simulations is precomputed in advance. A comparison between our toolset and PVesta is given in the Appendix.

References

1. J. Bogdoll, L.-M. Fiorti, A. Hartmanns, and H. Hermanns. Partial order methods for statistical model checking and simulation. In *FORTE*, LNCS. Springer, 2011. to appear.

2. E. M. Clarke, J. R. Faeder, C. J. Langmead, L. A. Harris, S. K. Jha, and A. Legay. Statistical model checking in biolab: Applications to the automated analysis of t-cell receptor signaling pathway. In *CMSB*, LNCS, pages 231–250, 2008.

3. A. David, K. Larsen, A. Legay, Z.Wang, and M. Mikucionis. Time for real statistical model-checking: Statistical model-checking for real-time systems. In *CAV*, LNCS. Springer, 2011.

4. A. David, K. G. Larsen, A. Legay, M. Mikučionis, D. B. Poulsen, J. V. Vliet, and Z. Wang. Statistical model checking for networks of priced timed automata. In *FORMATS*, LNCS, pages 80–96. Springer, 2011.

5. A. Fehnker, L. van Hoesel, and A. Mader. Modelling and verification of the lmac protocol for wireless sensor networks. In J. Davies and J. Gibbons, editors, *Integrated Formal Methods*, volume 4591 of *LNCS*, pages 253–272. Springer Berlin / Heidelberg, 2007.

6. R. Grosu and S. A. Smolka. Monte carlo model checking. In *Proc. of 11th Int. Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 3440 of *LNCS*, pages 271–286. Springer, 2005.

7. J. M. Hammersley and D. C. Handscomb. *Monte Carlo Methods*. Methuen, 1975.

8. T. Hérault, R. Lassaigne, F. Magniette, and S. Peyronnet. Approximate probabilistic model checking. In *VMCAI*, LNCS, pages 73–84, 2004.

9. S. K. Jha, E. M. Clarke, C. J. Langmead, A. Legay, A. Platzer, and P. Zuliani. A bayesian approach to model checking biological systems. In *CMSB*, volume 5688 of *LNCS*, pages 218–234. Springer, 2009.

10. A. Legay and B. Delahaye. Statistical model checking : An overview. *CoRR*, abs/1005.1327, 2010.

11. D. E. Rabih and N. Pekergin. Statistical model checking using perfect simulation. In *ATVA*, volume 5799 of *LNCS*, pages 120–134. Springer, 2009.

12. K. Sen, M. Viswanathan, and G. Agha. Statistical model checking of black-box probabilistic systems. In *CAV*, LNCS 3114, pages 202–215. Springer, 2004.

13. K. Sen, M. Viswanathan, and G. Agha. On statistical model checking of stochastic systems. In *CAV*, LNCS 3576, pages 266–280, 2005.

14. K. Sen, M. Viswanathan, and G. A. Agha. Vesta: A statistical model-checker and analyzer for probabilistic systems. In *QEST*, pages 251–252. IEEE Computer Society, 2005.

15. A. Wald. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics*, 16(2):117–186, 1945.

16. R. Wald. *Sequential Analysis*. Dove Publisher, 2004.

17. H. L. S. Younes. *Verification and Planning for Stochastic Processes with Asynchronous Events*. PhD thesis, Carnegie Mellon, 2005.

18. H. L. S. Younes. Ymer: A statistical model checker. In *CAV*, volume 3576 of *LNCS*, pages 429–433. Springer, 2005.

19. H. L. S. Younes, M. Z. Kwiatkowska, G. Norman, and D. Parker. Numerical vs. statistical probabilistic model checking. *STTT*, 8(3):216–228, 2006.

20. P. Zuliani, A. Platzer, and E. M. Clarke. Bayesian statistical model checking with application to simulink/stateflow verification. In *HSCC*, pages 243–252. ACM ACM, 2010.
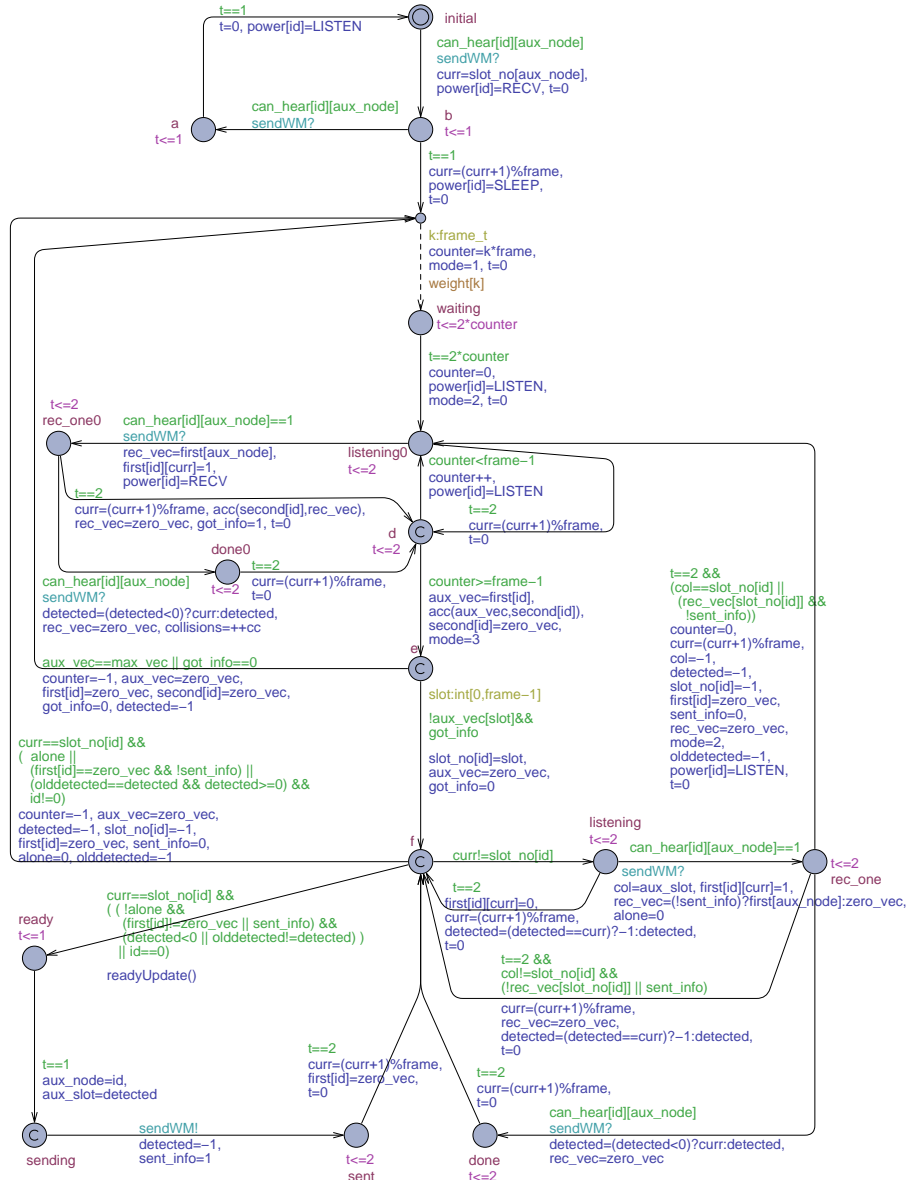
A   LMAC Model

Fig. 14: LMAC node model.

B DSMC Model Listings

Listing 1.2: DSMC model global declarations.

```
1  typedef int [-1024*1024*1024, 1024*1024*1024] long; // high precision
2  const int  N = 128; // number of processors
3  const int  B = 32; // number of runs in a batch
4  const int  K = 4; // maximum difference of batches among processors
5  const int  LatencyLower = 19; // latency lower bound
6  const int  LatencyUpper = 20; // latency higher bound
7  typedef int [0,  N-1] node_t; // node id type
8  typedef struct  { // SMC verification  parameters:
9       long scale ; // denominator for fixed  point operations
10      long alpha; // probability  of false  positive  (multiplied  by scale )
11      long beta;   // probability  of false  negative  (multiplied  by scale )
12      long theta ; // the probability   for hypothesis  testing
13      long deltaM;// lower probability    deviation
14      long deltaP; // upper probability    deviation
15      long p1;     // lower probability   bound [theta-deltaM]
16      long p0;     // upper probability   bound [theta+deltaP]
17      long valAcc; // value accumulation [log(p1/p0)]
18      long valRef; // value reference    [log((1-p1)/(1-p0))]
19      long logInf ; // log infimum (lower value  bound) [log(beta/(1-alpha))]
20      long logSup;// log supremum (upper value bound) [log((1-beta)/alpha)]
21  } SMC_params_t; // the structure  is  passed to Master upon instantiation
22  const int  H_first  = 13;// the start  of  the histogram
23  const int  H_last  = 48; // the end  of  the histogram
24  const int  H_step = 1;  // time step  of  one bucket
25  typedef int [H_first ,  H_last ] bucket_t; // type of integer  with specific   range
26  const int  w[bucket_t]  = { 6207,0,0,0,0,10463,0,0,0,5,10903,0,0,0,4,
27      9133,0,0,0,3,8569,0,0,0,2,8837,0,0,0,1,1233,0,0,0,1,64469      };
28  broadcast chan req;  // master requests
29  broadcast chan deliver [ node_t]; // slave  delivers
30  int  level [ node_t];  // level   of  the batch queue for  each node
31  bool busy[node_t];  // encodes whether the node is  computing
```

Listing 1.3: DSMC instantiation and system declaration.

```
1  int  value ; // shared variable  for  transfering   results
2  const SMC_params_t p[3] = {
3  //scale   alpha   beta theta +delta -delta   p1    p0    valAcc valRef logInf  logSup
4   {100000, 100, 100, 46188, 100, 100, 46088, 46288, -433, 372,-690675,690675},
5   {100000, 1000, 1000, 46188, 1000, 1000, 45188, 47188, -4331, 3717,-459512,459512},
6   {100000, 5000, 5000, 46188, 5000, 5000, 41188, 51188,-21736,18637,-294444,294444}
7  };
8  master = Master(p[1], value );
9  slave (const node_t id ) = Slave(id,  value );
10 system master, slave ,  Global;
```

(a) 16 nodes.   (b) 32 nodes.   (c) 128 nodes.

Fig. 15: Estimated processor usage times: `E[≤16000; 150](max: usage)`

B.1 Distributing Estimation

Distributing the estimation algorithm is much simpler than distributing sequential hypothesis testing. We need a *fixed* number of runs to compute an estimate of the probability value with given confidence level. This is an embarrassingly parallel problem since we can simply divide the work equally and gather the result at the end. The speed of generating the runs has no influence here since every computer node is allocated the same number of runs and if this number is high enough (and the computers identical), the running time on every node will be closely the same. This latter claim is indeed confirmed by our experiments since we observe that the performance scales almost linearly with the number of nodes. Interestingly, the loss in efficiency in the later cases exhibits the overhead of starting up all the processors (around 3-4 seconds), which would be compensated for much longer runs. The jobs here are too small and fast, which is an extreme case. Figure 1 shows running time and relative efficiency for estimating a few probabilities on the Firewire and LMAC protocol[8] with confidence 99.9% and uncertainty interval 0.005. Time is not very sensitive to the physical placement of the cores here because there is essentially one communication round when all the nodes have finished their jobs.

C Comparison with other toolsets

To the best of our knowledge, there are only two tools that implement distributed statistical model checking algorithms, namely Ymer and PVESTA. We could not conduct a direct comparison with distributed Ymer has the implementation is currently broken (Personnal communication with H. Youness). However, we could compare the batch algorithm implemented in Ymer with ours in the setting of UPPAAL (see Section 3).

PVESTA is a distributed statistical model checker based on vesta [14]. It performs hypothesis testing but instead of using a sequential approach, it relies on a single sampling where the number, N, of samples needed to solve the qualitative question is pre-computed. As for the estimation algorithm, sampling plan is

---

[8] The model and properties are available on http://people.cs.aau.dk/~adavid/smc/.

PxC/N	Firewire					LMAC				
	1	2	4	8	16	1	2	4	8	16
1x1	621.7s	316.7s	160.2s	81.1s	44.7s	279.3s	140.7s	73.0s	37.0s	19.5s
	1.00	*0.98*	*0.97*	*0.96*	*0.87*	*1.00*	*0.99*	*0.96*	*0.94*	*0.90*
1x2	300.9s	162.2s	80.5s	47.6s	24.3s	144.3s	71.0	37.5s	19.2s	10.4s
	1.03	*0.96*	*0.97*	*0.82*	*0.80*	*0.97*	*0.98*	*0.93*	*0.91*	*0.84*
1x4	161.2s	84.0s	44.8s	24.1s	16.0s	74.2s	36.1s	19.3s	9.6s	8.1s
	0.96	*0.93*	*0.87*	*0.81*	*0.61*	*0.94*	*0.97*	*0.90*	*0.91*	*0.54*
2x4	85.1s	46.5s	23.1s	14.1s	8.5	35.5s	19.6s	10.1s	10.2s	6.4s
	0.91	*0.84*	*0.84*	*0.69*	*0.57*	*0.98*	*0.89*	*0.86*	*0.43*	*0.34*

Table 1: Time in seconds and efficiency (italic) to estimate probabilities on the Firewire and LMAC model in function of the number of nodes (N), processors per node (P) and cores per processor (C).

thus trivially parallelizable. To compare PVesta with our implementation we created a Uppaal model of the cyclic polling server example of the PVesta distribtution.

	Time (900 st.)			Samples (900 st.)			Time (9000 st.)	
$Nodes$	UPPAAL_h	PVesta	UPPAAL_e	UPPAAL_h	PVesta	UPPAAL_e	UPPAAL_h	UPPAAL_e
1	≤ 4	46.0	12.15	115	16906	18448	≤ 4	84.0
2	≤ 4	23.7	7.5	190	16906	18448	≤ 4	44.4
4	≤ 4	12.7	3.91	557	16906	18448	≤ 4	23.8
8	≤ 4	7.2	5.5	2340	16906	18448	≤ 4	12.5

Table 2: Verification time for Uppaal and PVesta. The nodes column refers only to the nodes used for sample generation (PVesta actually used $nodes + 1$ processing units for the verification). The UPPAAL_h column shows the verification time in seconds for hypothesis testing and UPPAAL_e for estimation using Uppaal.

We carry our experiments on the *polling* example distributed with PVesta that we translate into an equivalent Uppaal model. We test for the hypothesis provided with the example, namely P >= 0.5 [<> < 20.0 @0] and its equivalent in Uppaal and we estimate this probability in Uppaal. The results are given in table 2. We compare the tools with a configuration of 900 stations in the model. We note that for 900, the computation time of Uppaal is so short that we can only measure the overhead of starting and distributing the computation that takes less than 4 seconds, which is why we show ≤ 4 in the table. To see the scalability of Uppaal, we experiment with 9000 stations in the model. We can only show the results for Uppaal since PVesta reached its time limit of one hour before it could give a result. The experiments show that Uppaal and

PVESTA scale almost linearly and UPPAAL is at least two orders of magnitude faster than PVESTA on for hypothesis testing.