

Protocol analysis via probabilistic model checking with PRISM

Marta Kwiatkowska
School of Computer Science



UNIVERSITY OF
BIRMINGHAM

www.cs.bham.ac.uk/~mzk

UC Berkeley, 7th December 2006

Overview

- Ubiquitous computing: the trends and challenges
 - Why probability?
 - Role of verification techniques
- Probabilistic/quantitative model checking
 - What does it involve?
 - Models and specifications
 - Quantitative analysis
- Lessons learnt
 - Bluetooth device discovery
- Future directions
 - and challenges that remain!

Ubiquitous computing: the trends...

- **Devices, ever smaller**
 - Laptops, phones, PDAs, ...
 - Sensors, motes, Smartdust, ...
- **Networking, wireless, wired & global**
 - Mobile ad hoc
 - Wireless everywhere
 - Internet everywhere
 - Global connectivity
- **Systems/software**
 - Decentralised
 - Self-*
 - Autonomous
 - Mobile
 - Adaptive
 - Context-aware



Ubiquitous computing: the challenge

- **Grand Challenge in Computing**
 - Initiated in the UK but global
- **How to design & engineer**
 - Complex
 - Reliable
 - Reconfigurable
 - Fast & efficient
 - Large-scale networked systems?
- **How to ensure correctness, dependability, and performance?**
 - **The science**, or rigorous foundation
 - **Verification**, techniques and tools
- **How to embed within the natural environment**
 - Human/computer interfaces
 - Social & ethical implications



This talk: emphasis on verification

- **Automated** techniques for the assurance of properties below, e.g. using model checking,
 - safety
 - security, privacy & trust
 - performance
 - dependability
- **NB, quantitative**, as well as qualitative requirements:
 - How **quickly** will I receive a reply in my car network?
 - Is my mobile phone **energy efficient**?
 - How **trustworthy** is the supplier?
- **Our focus, probabilistic model checking**
 - **Probabilistic** models, also **costs/rewards**
 - Range of **quantitative** analyses
 - **Automated** verification techniques



Why probability?

- **Randomisation** used in distributed coordination algorithms
 - As a **symmetry breaker**, in **gossip routing** to reduce flooding
- **To model uncertainty and performance**
 - To **quantify** rate of **failures**, express **Quality of Service**
- **For quantitative analysis of software and systems**
 - To **quantify resource usage** given a policy
 - "the **minimum battery capacity** for a given scenario is .."
- **In evidence-based, statistical analysis of behaviours**
 - To **quantify** trust, anonymity, etc
- **In modelling of biological processes**
 - "the **probability** that there are 50 Na molecules at time 10 is ... "
 - "the **expected long-run percentage** of Na molecules is ... "

Real-world protocol examples

- Protocols featuring **randomisation**
 - Randomised back-off schemes
 - E.g. IEEE 802.11 (WiFi) Wireless LAN MAC protocol
 - Random choice of waiting time
 - Bluetooth, device discovery phase
 - Random choice of routes to destination
 - Crowds, anonymity protocol for internet routing
 - Random choice over a set of possible addresses
 - IPv4 dynamic configuration (link-local addressing)
 - and more
- **Continuous** probability distribution needed to model network traffic, node mobility, random delays...

Verification techniques


- Probabilistic model checking, e.g. this lecture
 - Finite-state models only, but **automatic**
 - **Quantitative** analysis, including probabilistic reachability and expectations
 - Broad range of models
 - Emphasis on **experimentation** and **bug-hunting** rather than **verification** for all scenarios
- Program correctness proofs
 - Not limited to finite-state models
 - Mechanised using **theorem provers**
 - **Not** automatic and involved
 - More limited range of models and specifications
 - Emphasis on finding **correctness proofs** that apply **in all** scenarios

Probabilistic model checking inputs

- **Models : variants of Markov chains**
 - Discrete-Time Markov Chains (DTMCs)
 - Markov Decision Processes (MDPs)
 - Continuous-Time Markov Chains (CTMCs)
 - Probabilistic Time Automata (PTAs)
- **Specifications (informally)**
 - "probability of shutdown occurring is at most..."
 - "probability of delivery within time deadline is ..."
 - "expected time to message delivery is ..."
 - "expected power consumption is ..."
- **Specifications (formally)**
 - Extensions of temporal logic with probabilistic and expectation operators (PCTL, CSL, PTCTL)
 - Probability, time, cost/rewards

Markov Decision Processes (MDPs)

- Characteristics:
 - Discrete probability, nondeterminism
 - **No real-time**
- Formally, $(S, s_0, Steps, L, Act, c)$:
 - S finite set of **states**
 - s_0 **initial** state
 - **Steps** maps states s to **sets of probability distributions** μ over S
 - **Act** labelling of steps with **actions**
 - $L: S \rightarrow 2^{AP}$ **atomic propositions**
 - $c: S \times Act \rightarrow \mathbb{R}_{\geq 0}$ **cost function**
- Unfold into **infinite paths** $s_0 a_0 \mu_0 s_1 a_1 \mu_1 s_2 \dots$ s.t. $\mu_i(s_i, s_{i+1}) > 0$, all i
- Probability space induced on $Path_s$ by **adversary** (policy) A mapping finite path $s_0 a_0 \mu_0 s_1 a_1 \mu_1 \dots s_n$ to a distribution from s_n




Probability and cost

- Intuitively, for a **fixed adversary** $A \in \text{Adv}$:

- **Sample space** = infinite paths Path_s^A from s
- **Event** = set of paths
- **Basic event** = cone

$s_0 \dots s_k$



- Formally, probability space $(\text{Path}_s^A, \Sigma, \text{Pr}^A)$

- Assign probability $\text{P}(\cdot)$ to finite paths $\omega = s_0 a_0 \mu_0 s_1 a_1 \mu_1 s_2 a_2 \dots s_n$
- Define $\text{Pr}^A(C(\omega)) = \text{P}(\omega)$, for cones $C(\omega)$:

$$C(\omega) = \{ \pi \in \text{Path}_s^A \mid \omega \text{ is prefix of } \pi \}$$

- Then **cost** for a finite path ω and target set $F \subseteq S$

$$\text{cost}(F)(\omega) = \begin{cases} \sum_{i=0}^{\min\{i : \omega(i) \in F\}} c(a) & \text{if } \exists i. \omega(i) \in F, \\ & (a, \mu) \in \text{Steps}(\omega(i-1)) \\ \infty & \text{otherwise} \end{cases}$$

The logic PCTL: syntax

- Probabilistic Computation Tree Logic [HJ94,BdA95,BK98]
 - Based on CTL, for DTMCs/MDPs
 - Add **probabilistic operator**, e.g. $\text{send} \rightarrow \mathcal{P}_{\geq 0.9}(\diamond \text{deliver})$
 - and **expected cost operator**, e.g. $\mathcal{E}_{> 1}(\text{heads})$
- The syntax of **state** and **path** formulas of PCTL is:

$$\phi ::= \text{true} \mid a \mid \phi \ \mathcal{A} \ \phi \mid \neg \phi \mid \mathcal{P}_{\sim p}(\alpha) \mid \mathcal{E}_{\sim d}(\phi)$$

$$\alpha ::= X \phi \mid \phi \ U \ \phi$$

where $p \in [0,1]$ is a **probability threshold**, $d \in \mathbb{R}_{\geq 0}$ is a **cost bound** and $\sim \in \{<, >, \dots\}$

The logic PCTL: semantics

- Semantics is parameterised by a class of adversaries Adv
 - “under any scheduling, the probability/cost is ... at state s ”
 - reasoning about worst-case/best-case scenario

- The probabilistic operator:

$$s \models_{\text{Adv}} \mathcal{P}_{\sim p}(\alpha) \quad \Leftrightarrow \quad \Pr_s^A \{ \pi \in \text{Path}_s^A \mid \pi \models_{\text{Adv}} \alpha \} \sim p \\ \text{for all } A \in \text{Adv}$$

- The expectation operator:

$$s \models_{\text{Adv}} \mathcal{E}_{\sim d}(\phi) \quad \Leftrightarrow \quad E_s^A (\text{cost}\{ s' \in S \mid s' \models_{\text{Adv}} \phi \}) \sim d \\ \text{wrt } \Pr_s^A, \text{ for all } A \in \text{Adv}$$

- Semantics of remaining formulas standard

PCTL Until model checking for MDPs

- Reduces to **minimum/maximum** probability computation
 - $p^{\max}_s(\alpha), p^{\min}_s(\alpha)$, defined as sup/inf over all adversaries
- Obtain **linear optimisation** problems solvable iteratively, e.g.

$$\text{Sat}(\mathcal{P}_{\geq p}(\phi_1 \text{ U } \phi_2)) \Leftrightarrow \{s \in S \mid x_s \geq p\}$$

Maximise $\sum_{s \in S^?} x_s$ subject to the constraints:

$$x_s \leq \sum_{s' \in S^?} \mu(s') \cdot x_{s'} + \sum_{s' \in S^{\text{yes}}} \mu(s')$$

for all $s \in S^?, (a, \mu) \in \text{Steps}(s)$

- Note
 - $S^{\text{yes}} (S^{\text{no}})$ satisfy until with probability 1 (0) resp.
 - $S^? = S \setminus (S^{\text{no}} \cup S^{\text{yes}})$

PCTL Until: iterative solution

- S^{yes} , S^{no} can also be precomputed by **graph traversal** (BDD fixed point) [dAKNP97]
- The linear equation generalises to **linear optimisation** problems solvable iteratively, e.g.

$$\text{Sat}(\mathcal{P}_{\geq p}(\phi_1 \cup \phi_2)) \quad \Leftrightarrow \quad \{s \in S \mid \mathbf{x}_s \geq \mathbf{p}\}$$
$$\mathbf{x}_s^{n+1} = \begin{cases} 0 & \text{if } s \in S^{no}, \text{ all } n \\ 1 & \text{if } s \in S^{yes}, \text{ all } n \\ \min_{\mu \in \text{Steps}(s)} \sum_{s' \in S} \mu(s') \cdot \mathbf{x}_{s'}^n & \text{if } s \in S \setminus (S^{no} \cup S^{yes}) \end{cases}$$

- Need to combine
 - Conventional **graph-theoretic traversal**
 - **Linear optimisation** (simplified value iteration)

Expected cost model checking

- Reduces to **minimum/maximum** expected cost computation
 - $e_s^{\max}(\phi)$, $e_s^{\min}(\phi)$, defined as sup/inf over all adversaries
- Again, obtain **linear optimisation** problems solvable iteratively, e.g.

$$\text{Sat}(\mathcal{E}_{\geq d}(\phi)) \quad \Leftrightarrow \quad \{s \in S \mid x_s \geq d\}$$

Maximise $\sum_{s \in S^?} x_s$ subject to the constraints:

$$x_s \leq c(a) + \sum_{s' \in S^{\text{yes}}} \mu(s') \cdot x_{s'}$$

for all $s \in S^?$, $(a, \mu) \in \text{Steps}(s)$

- Note
 - S^{yes} , S^{no} , $S^? = S \setminus (S^{\text{no}} \cup S^{\text{yes}})$ can be precomputed as before
 - Unique solution, cost possibly ∞
 - Algorithm of [dAlf'97]

Probabilistic model checking involves...

- Construction of **models**
 - from a **high-level modelling language**, such as a probabilistic process algebra or probabilistic automaton
 - e.g. Markov decision processes (MDPs)
- Implementation of **probabilistic model checking algorithms**
 - **graph-theoretical algorithms**, combined with
 - (probabilistic) reachability
 - qualitative model checking (for 0/1 probability)
 - **numerical computation - iterative methods**
 - quantitative model checking (plot **probability values**, **expectations**, **rewards**, steady-state, etc, for a **range** of parameters)
 - typically, **linear equation** or **linear optimisation**
 - exhaustive, unlike simulation
 - also **sampling-based** (statistical) for approximate analysis
 - e.g. hypothesis testing based on simulation runs

The PRISM probabilistic model checker

- **Functionality**

- Modelling language: probabilistic reactive modules
- Construction of **models**: DTMCs/CTMCs, MDPs
- Probabilistic temporal logics: **PCTL** and **CSL**
- Extension with **costs/rewards**, **expectation** operator

- **Underlying computation involves**

- Reachability, qualitative model checking, BDD-based
- Linear equation system solution - Jacobi, Gauss-Seidel, ...
- Uniformisation (CTMCs)
- Dynamic programming (MDPs)
- Explicit and symbolic (MTBDDs, etc.)

- **Similar tools**

- ETMCC/MRMC, PROBMELA, Vesta, Rapture, Ymer, APMC, APNN-Toolbox, SMART, Mobius

PRISM Screenshots

The screenshot displays the PRISM 3.0 beta1 software interface. The main window shows the PRISM Model File: cluster.sm. The left-hand pane displays the model structure, including Modules (Left, Right, Repairman, Line, ToLeft, ToRight) and Constants (N, left_mx, right_mx, k, line_rate, Toleft_rate, Toright_rate, OPERATIONAL, MINIMUM, REPAIR). The main text editor contains the following code:

```
// workstation cluster [HHK00]
// dxp/gxn 11/01/00

stochastic

const int N; // number of workstations in each cluster
const int left_mx = N; // number of work stations in left cluster
const int right_mx = N; // number of work stations in right cluster

// minimum 005 requires 3/4*N connected workstations operational
const int k=floor(0.75*N);
formula minimum = (left_n>=k & Toleft_n) |
                  (right_n>=k & Toright_n) |
                  ((left_n+right_n)>=k & Toleft_n & Toright_n);

// rates
const double line_rate = 0.0002;
const double Toleft_rate = 0.00025;
const double Toright_rate = 0.00025;

// left cluster
module Left

    left_n : [0..left_mx] init left_mx; // number of workstations operational
    left : bool; // being repaired?

    [startLeft] !left & (left_n<left_mx) -> 1 : (left'=true);
    [repairLeft] left & (left_n<left_mx) -> 1 : (left'=false) & (left_n'=left_n+1);
    [] (left_n>0) -> 0.002*left_n : (left_n'=left_n-1);

endmodule

// right cluster
module Right = Left[left_n=right_n,
                  left=right,
                  left_mx=right_mx,
                  startLeft=startRight,
                  repairLeft=repairRight ]
```

The bottom status bar indicates "Building model... done." and the bottom-left corner shows "Built Model: No of states: 4180, No of transitions: 19552".

PRISM Screenshots

PRISM 3.0.beta1

File Edit Model Properties Options

Properties list: /data/private/luser/prism-examples/cluster/cluster.csl

Properties

- S=? ["premium"]
- S=? [!"minimum"]
- P>=1 [true U "premium"]
- P=? [true U<=T !"minimum"]
- P=? [true U[T,T] !"minimum" {"!"minimum"}{max}]
- P=? [true U<=T "premium" {"minimum"}{min}]
- P=? ["minimum" U<=T "premium" {"minimum"}{min}]
- P=? [!"minimum" U>=T "minimum" {"!"minimum"}{max}]
- R=? [I=T {"!"minimum"}{min}]
- R=? [C<=T]
- R=? [C<=T]

e that QOS drops below minimum quality within T time units (from the initial state)

Constants

Name	Type	Value
T	double	

Labels

Name	Definition
minimum	(left_n>=k&Toleft_n)(right_n>=k&Tori...
premium	(left_n>=left_mx&Toleft_n)(right_n>=r...

Experiments

Property	Defined Const...	Progress	Status	Method
P=? [true U[T... T=0.0:1.0E-...		660/660 (100%)	Done	Verification
P=? [true U[T... N=3,T=0.0:1...		101/101 (100%)	Done	Simulation
P=? [true U[T... N=3,T=0.0:1...		44/101 (43%)	Stopped	Verification
P=? [true U<... N=3,T=0.0:1...		21/21 (100%)	Done	Verification
P=? [true U<... N=3:1:5,T=0...		63/63 (100%)	Done	Verification

Graph1 Graph2 Graph3 Graph4 Graph5

New Graph

Probability

T

N=3
N=4
N=5

Model Properties Simulator Log

Running experiment... done.

PRISM Screenshots

PRISM 3.0.beta1

File Edit Model Properties Options

✂️ 📄 🗑️ 🗑️

Exploration

🔄 Auto Update

No. Steps:

↻ Do Update

State time: Auto

Action	Rate	Update
Left	0.016	left_n'=7
Right	0.02	right_n'=9
Line	2.0E-4	line_n'=false
ToRight	2.5E-4	Toright_n'=fal
[startLeft]	10.0	left'=true, r'=
[startToLeft]	10.0	r'=true, Toleft

Path Modification

🔙 Backtrack Remove

To Step: Before:

Formulae

Path formulae:

- 🔍 true U<=T !"minimum"
- 🔍 true U[T,T] !"minimum"
- ✅ true U<=T "premium"
- ✅ !minimum U<=T !"minimum"

State labels:

- ❌ init
- ❌ deadlock
- ✅ minimum
- ✅ premium

Simulation Path

New Path

Reset Path

Export Path

Model Type: Stochastic (CTMC)

State Rewards: 4373.594264201196

Defined Constants: N=10,T=100.0

Path Length: 18

Transition Rewards: 0.0

Total Time: 44.030215385327985

Total Reward: 4373.594264201196

Step	left_n	left	right_n	right	r	line	line_n	Toleft
0	10	false	10	false	false	false	true	false
1	10	false	9	false	false	false	true	false
2	10	false	9	true	true	false	true	false
3	10	false	10	false	false	false	true	false
4	9	false	10	false	false	false	true	false
5	9	true	10	false	true	false	true	false
6	10	false	10	false	false	false	true	false
7	10	false	9	false	false	false	true	false
8	10	false	10	true	true	false	true	false
9	10	false	10	false	false	false	true	false
10	9	false	10	false	false	false	true	false
11	9	false	10	false	false	false	true	false
12	9	false	9	false	false	false	true	false
13	8	false	9	false	false	false	true	false
14	7	false	9	false	false	false	true	false
15	8	true	9	false	true	false	true	false
16	8	false	9	false	false	false	true	false
17	8	false	10	true	true	false	true	false
18	8	false	10	false	false	false	true	false

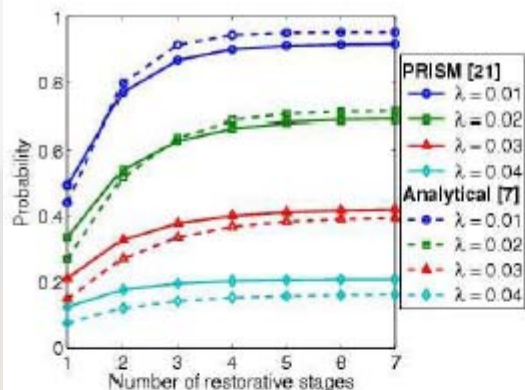
Model Properties Simulator Log

Loading properties... done.


Probabilistic specifications

- As in probabilistic extensions of temporal logic
 - incorporates PCTL for DTMCs/MDPs, CSL for CTMCs
- Examples:
 - $P < 0.001 [F \text{ shutdown }]$ - "shutdown eventually occurs with probability at most 0.001"
 - $P < 0.2 [F[t,t] (\text{deliv_rate} < \text{min})]$ "the probability that the current packet delivery rate has dropped below minimum at time t is less than 0.2"
 - $P \geq 0.95 [!\text{repair } U_{\leq 200} \text{ done }]$ - "with probability 0.95 or greater, the process will successfully complete within 200 hours and without requiring any repairs"
- No counterexamples (yet) in prob. model checking


Results: probabilistic analysis



Probability that 10% of gate outputs are erroneous for varying gate failure rates and numbers of stages



Probability that parties gain unfair advantage for varying numbers of secret packets sent



Optimum probability of leader election by time T for various coin biases


Quantitative specifications

- Focus on quantitative properties, compute actual values
 - $P=? [F \leq T \text{ "shutdown" }]$ - "what is the probability of shutdown occurring within T hours?"
- Best/worst-case scenarios
 - $P=? [F \text{ "error" } \{ \text{"init"} \} \{ \text{max} \}]$ - "what is the worst-case error probability over all possible initial configurations?"
 - $P_{\min}=? [! \text{end2 } U \text{ end1 }]$ - "what is the minimum probability of process 1 finishing before process 2, over all possible schedulings of the processes?"
- Experiments - ranges of model/property parameters
 - $P=? [F \leq T \text{ error }]$ for $N=1..5, T=1..100$
 - identify patterns, trends, anomalies in results


Cost- and reward-based properties

- **Costs and rewards**
 - real-valued quantities assigned to states/transitions
- **Instantaneous** - state-based measures
 - current queue size, number of operational channels, ...
 - "what is the **expected size** of the message queue at time t ?"
 - "what is the **long-run expected size** of the queue?"
- **Cumulative** - state or transition (impulse) costs/rewards
 - time, power consumption, messages lost, ...
 - "what is the **expected power consumption** during the first 2 hours of operation?"
 - "what is the **worst-case expected time** taken for the protocol to terminate?"


Results: quantitative analysis



Worst-case expected number of steps to stabilise for initial configurations with K tokens amongst N processes



Expected reactant concentrations over the first 12 hours



Maximum expected time for leader election for various coin biases

Efficiency: Symbolic techniques

- **State space explosion**
 - models of real-life systems typically **huge**
- **Symbolic probabilistic model checking**
 - data structures based on **binary decision diagrams (BDDs)**
 - compact storage: exploit model structure and regularity
 - efficient implementation of graph traversal fixed point algorithms
- **PRISM: multiple computation engines**
 - **MTBDDs** (BDD extension): storage/analysis of very large models (given structure/regularity), numerical computation can blow up
 - **sparse matrices**: fastest solution for smaller models ($<10^6$ states), prohibitive memory consumption for larger models
 - **hybrid**: combine MTBDD storage with explicit storage, ten-fold increase in analysable model size ($\sim 10^7$ states)

Efficiency: Other strategies

- **Approximate model checking** (see also APMC [LHP06])
 - sampling using Monte Carlo discrete-event simulation
 - performed at modelling language level - better scalability
 - potentially huge number of samples for accurate answers
 - also: statistical hypothesis testing, see e.g. [YS02]
- **Parallelisation of model checking**
 - distribution of storage/computation across multi-processor machines [KPZM04], networked clusters [ZPK05], grids
 - potentially promising for symbolic approaches - reduced I/O
 - simulation-based computations much easier to distribute

Ongoing research areas

- **Abstraction and refinement**, see e.g. [DJJL01,KNP06a]
 - construct smaller, abstract model by removing information/variables not relevant to property being checked, iteratively refine abstraction if analysis fails
- **Symmetry reduction** [DM06, KNP06b]
 - exploit replication of identical components
- **Partial order reduction**, see e.g. [BGC04], [DN04]
 - exploit commutativity of concurrently executed transitions
- **Compositionality**, see e.g. [dAHJ01,Che06]
 - analyse full model based on analysis of sub-components

PRISM real-world case studies

- Network and coordination protocols
 - Bluetooth device discovery [ISOLA'04, STTT'06]
 - IEEE 802.15.4 (Zigbee) [ISOLA'06]
 - IPv4 Zeroconf dynamic configuration [FORMATS'03, FMSD'06]
 - Randomised consensus [CAV'01]
 - Power Management (with Shukla & Gupta) [HLDVT'02, FAC'05]
- Security protocols
 - Crowds anonymity protocol (by Shmatikov) [JCS'04]
 - Probabilistic contract signing (with Shmatikov) [JCS'04]
- Reliability
 - NAND multiplexing for nanotech (with Shukla) [VLSI'04, TCAD'05]
 - Dependability of embedded controller [INCOM'04]
- Biology
 - FGF signalling pathway [CMSB'06]




Bluetooth device discovery

- **Bluetooth**: short-range low-power wireless protocol
 - widely available in phones, PDAs, laptops, ...
 - personal area networks (PANs)
 - open standard, specification freely available
- **Uses frequency hopping scheme**
 - to avoid interference (uses unregulated 2.4GHz band)
 - pseudo-random selection over 32 of 79 frequencies
- **Network formation**
 - piconets (1 master, up to 7 slaves)
 - self-configuring: devices discover themselves
- **Device discovery**
 - mandatory first step before any communication possible
 - **performance crucial** to reduce power consumption

Master (sender) behaviour

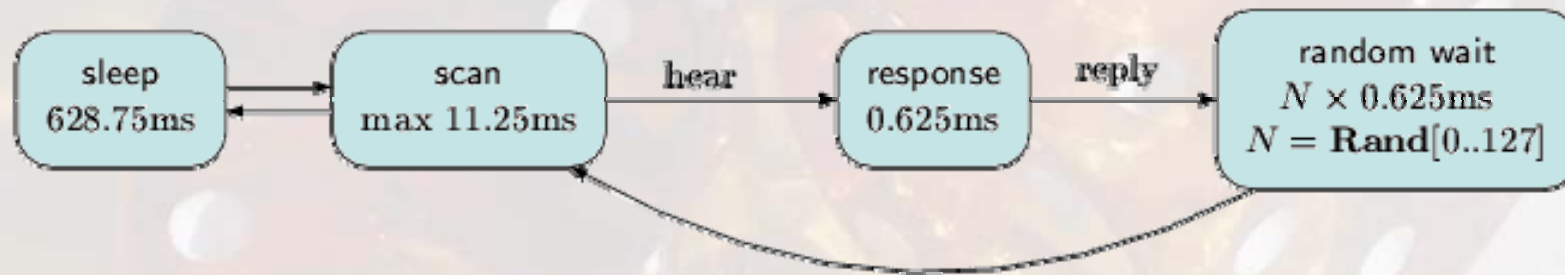
- 28 bit free-running clock **CLK**, ticks every **312.5μs**
- Frequency hopping sequence determined by clock:
 - $\text{freq} = [\text{CLK}_{16-12} + k + (\text{CLK}_{4-2,0} - \text{CLK}_{16-12}) \bmod 16] \bmod 32$
 - 2 trains of 16 frequencies (determined by offset k), 128 times each, swap between every 2.56s
- Broadcasts inquiry packets on two consecutive frequencies, then listens on the same two



1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
17	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
1	2	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
17	18	19	20	21	6	7	8	9	10	11	12	13	14	15	16	
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
1	2	3	4	5	6	23	24	25	26	27	28	29	30	31	32	
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
17	18	19	20	21	22	23	24	9	10	11	12	13	14	15	16	
17	18	19	20	21	22	23	24	25	10	11	12	13	14	15	16	
1	2	3	4	5	6	7	8	9	10	27	28	29	30	31	32	
1	2	3	4	5	6	7	8	9	10	11	28	29	30	31	32	
17	18	19	20	21	22	23	24	25	26	27	28	13	14	15	16	
17	18	19	20	21	22	23	24	25	26	27	28	29	14	15	16	
1	2	3	4	5	6	7	8	9	10	11	12	13	14	31	32	
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	32	
17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
1	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
17	18	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
17	18	19	4	5	6	7	8	9	10	11	12	13	14	15	16	
1	2	3	4	5	6	7	8	9	10	27	28	29	30	31	32	
1	2	3	4	5	22	23	24	25	26	27	28	29	30	31	32	
17	18	19	20	21	22	7	8	9	10	11	12	13	14	15	16	
17	18	19	20	21	22	23	8	9	10	11	12	13	14	15	16	
1	2	3	4	5	6	7	8	9	25	26	27	28	29	30	31	32
1	2	3	4	5	6	7	8	9	26	27	28	29	30	31	32	
17	18	19	20	21	22	23	24	25	26	11	12	13	14	15	16	
17	18	19	20	21	22	23	24	25	26	27	12	13	14	15	16	
1	2	3	4	5	6	7	8	9	10	11	12	29	30	31	32	
1	2	3	4	5	6	7	8	9	10	11	12	13	30	31	32	
17	18	19	20	21	22	23	24	25	26	27	28	29	30	15	16	
17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	16	

Slave (receiver) behaviour

- **Listens** (scans) on frequencies for **inquiry packets**
 - must listen on right frequency at right time
 - cycles through frequency sequence at much slower speed (every 1.28s)



- On hearing packet, **pause**, send **reply** and then wait for a **random delay** before listening for subsequent packets
 - avoid repeated collisions with other slaves


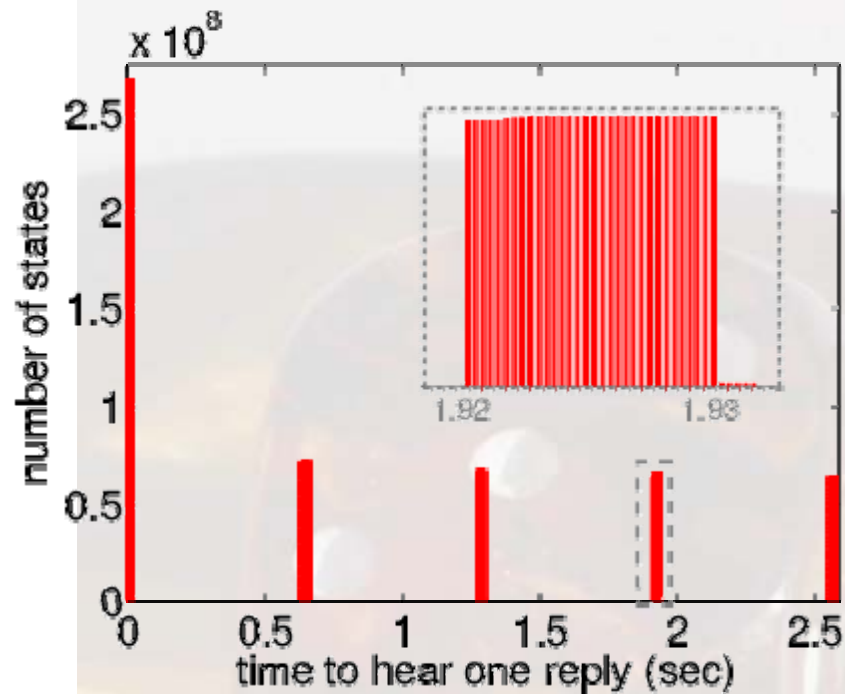
Bluetooth - PRISM model

- Modelling in PRISM [DKNP06]
 - model **one sender and one receiver**
 - **synchronous** (clock speed defined by Bluetooth spec)
 - **randomised** behaviour - use **DTMC**
 - model at lowest-level (one clock-tick = one transition)
 - use **real values** for delays, etc. from Bluetooth spec
- Modelling challenges
 - **complex interaction** between sender/receiver
 - combination of short/long time-scales - cannot scale down
 - sender/receiver not initially synchronised, huge number of possible initial configurations (**17,179,869,184**)

Bluetooth - Results



- **Huge DTMC** - initially, model checking infeasible
 - partition into 32 scenarios, i.e. 32 separate DTMCs
 - on average, approx. 3.4×10^9 states, **536,870,912** initial
 - can be built/analysed with PRISM's MTBDD engine
- **Compute for the first time:**
 - **$R=? [F \text{ replies}=K \{ \text{"init"} \} \{ \text{max} \}]$**
 - "worst-case expected time to hear K replies over all possible initial configurations"
 - also look at:
 - how many initial states for each possible expected time
 - cumulative distribution function assuming equal probability for each initial state

Bluetooth - Time to hear 1 reply



- worst-case expected time = 2.5716 sec
- in 921,600 possible initial states
- best-case = 635 μ s

Bluetooth - Time to hear 2 replies



- worst-case expected time = 5.177 sec
- in 444 possible initial states
- compare actual CDF with derived version which assumes times to reply to first/second messages are independent

Related projects

- **Modelling and verification of mobile ad hoc network protocols**
 - Modelling language with mobility and randomisation
 - Model checking algorithms & techniques
 - Formal methods for power management in PANs
- **Ubival (forthcoming)**
 - Verification via model checking, testing and simulation of ubiquitous software
 - Focus on mobility, context-awareness and adaptiveness
 - Apply probabilistic model checking to quantify the likelihood of context
- **Probabilistic model checking of biological processes**
 - Using PRISM and CTMC models to analyse signalling networks


Ongoing work

- Exploiting **structure**, model reductions
 - **Abstraction**, employing games-based approaches, [QEST'06]
 - **Symmetry reduction**, [CAV'06]
 - **Partial order reduction**, (with Baier et al) [FST&TCS'06]
- Extension for **mobility**
 - Extension for ambients and spatial logic (with Vigliotti)
 - Combine with techniques for pi-calculus (with Peng Wu)
 - Applications in systems biology
- **Multi-criteria** optimisation
 - For MDPs (with Etessami et al)
- **Real** software, not models! **E.g. verify source code containing**
`x := random(1..65024); {ZeroConf}`

Challenges for future

- **Scalability!**
 - Abstraction/refinement
 - Compositionality, e.g. assume-guarantee
- **Proof assistant** for probabilistic verification
 - Verify **for all** scenarios, **parametric** analysis
- **Efficiency**
 - Approximation methods
 - Statistical analysis
 - Parallelisation
- Extension for **hybrid** systems
 - E.g. **context-awareness** for geographical coordinates

For more information...



J. Rutten, M. Kwiatkowska, G. Norman and
D. Parker

Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems

P. Panangaden and F. van Breugel (editors),
CRM Monograph Series, vol. 23, AMS
March 2004



www.cs.bham.ac.uk/~dyp/prism/

- Case studies, statistics, group publications
- Download, version 3.1 (4700+ downloads)
- Unix/Linux, Windows, Apple platforms
- Publications by others and courses that feature PRISM...

PRISM Contributors

