Performance Analysis of Probabilistic Timed Automata Using Digital Clocks

Marta Kwiatkowska, University of Birmingham
Gethin Norman, University of Birmingham
Dave Parker, University of Birmingham
Jeremy Sproston, Università di Torino

Overview

- Probabilistic timed automata (PTAs)
- Expected time/cost properties
- Digital clocks
- Case study: IPv4 ZeroConf protocol

Motivation

In real-life systems, timing behaviour often coexists with probabilistic behaviour

- Randomized algorithms:
  - IEEE 1394 (FireWire) root contention protocol
  - Backoff strategies in communication protocols (Ethernet, IEEE 802.11)
  - Bluetooth wireless protocol
- Unpredictable environment:
  - Message loss in communication protocols
  - Failures/faults

Probabilistic Models

- Formalisms for probabilistic timed systems
  - Discrete-time Markov chains (DTMCs)
    - discrete time/probabilities
  - Continuous-time Markov chains (CTMCs)
    - exponential distributions
  - Markov decision processes (MDPs)
    - discrete time/probabilities + nondeterminism
  - Probabilistic timed automata (PTAs)
    - dense real time, discrete probabilities

Probabilistic Timed Automata

Timed automata with probabilistic branching over the edges

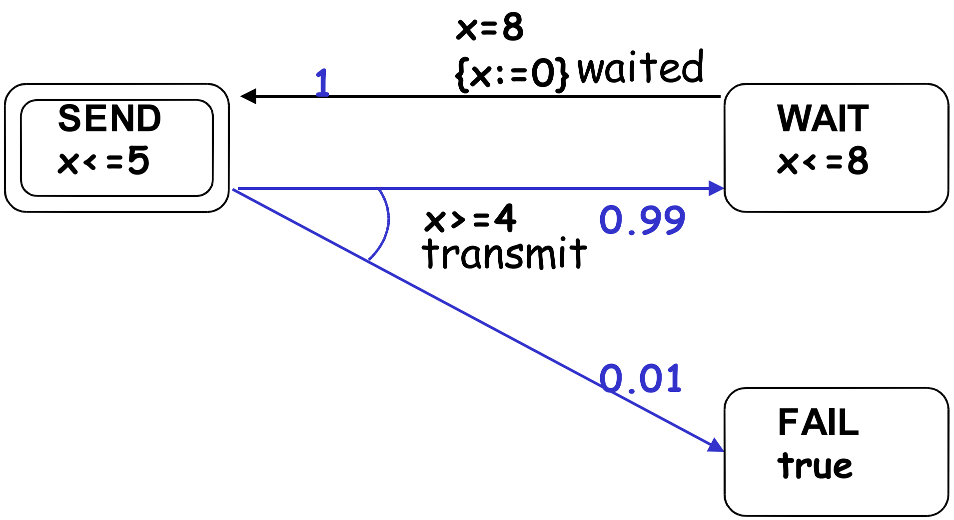

Traditionally, clocks take values in $R_{>=0}$

Probabilistic Timed Automata

Formalism	Semantics
Timed automata	Transition systems
Probabilistic timed automata	Markov decision processes

- States: location, clock valuation pairs (l,v) (v is in $(R_{>=0})^{|clocks|}$)
  - Real-valued clocks give infinitely many states
- Transitions: 2 classes

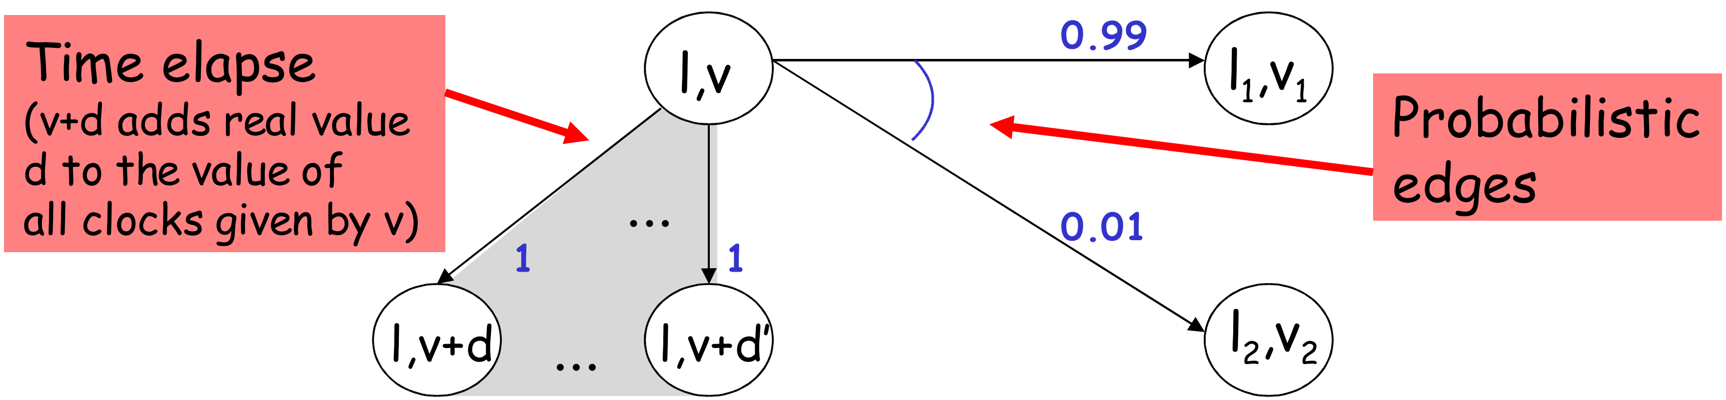

Properties

- Probabilistic timed reachability

Example :	"With probability 0.05 or less, the system aborts within 30 seconds"

  - PTA context: [KNS01, KNSS02, KNS03]

- Expected reachability

Example :	"The expected time elapsed before the first data packet is delivered is at most 0.1 seconds"

Example :	"The expected cost accumulated before a host chooses an IP address is at most 40"

  - PTA context: this talk

Costs

- At the level of probabilistic timed automata
  - Cost pair: (r,e)
    - r is in $R_{>=0}$: rate at which cost is accumulated as time passes
    - e maps from events to $R_{>=0}$: event-cost function assigning a cost with each event
  - Special case: time=cost, with r=1 and e(.)=0
- Example:

r=1
e(transmit)=2
e(waited)=0

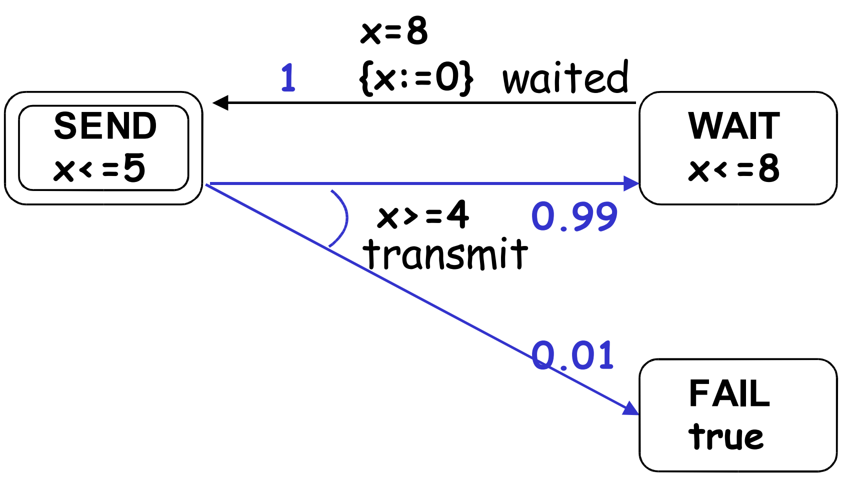

Expected Costs

- The coexistence of nondeterministic and probabilistic choice means that there is no unique probability space over paths

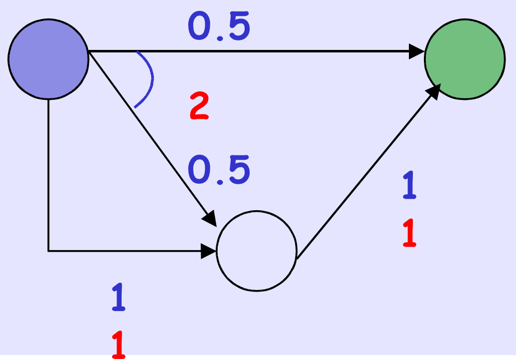

Probabilities in blue, costs in red

Accumulated cost from to
Minimum expected cost = 2
Maximum expected cost = 2.5

- Therefore, we obtain the minimum and maximum expected accumulated costs before reaching a set of locations

Computing Expected Costs

- Compute: minimum/maximum expected cost accumulated before reaching a set of locations
  - Assume that probability of reaching this set is 1
- Challenge: infinite state space
  - We know how to compute expected accumulated costs for finite-state Markov decision processes [de Alfaro97]
- Solution: use digital clocks (digitization)
  - Digital clocks: take values in $\mathbb{N}$, rather than $\mathbb{R}_{>=0}$
  - Results in a finite state space

Digital Clocks

- We consider time-elapse transitions of duration 1 only

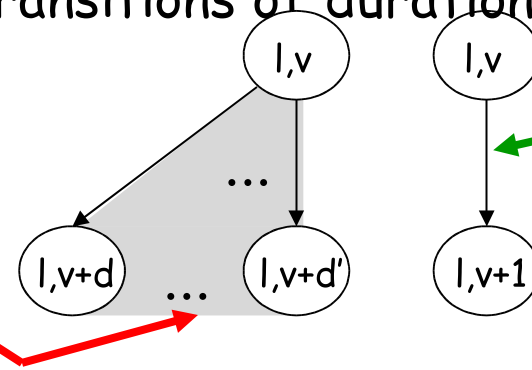

- Can also assume clocks do not exceed max+1
  - where max is the maximal constant used in clock constraints

- Correctness based on [HMP92]: "digitization" maps between digitally-clocked and real-clocked behaviours
- Correctness requires: closed, diagonal-free (P)TA

  For example: x<5 NO x-y>=3 NO x>=4 YES

Expected Costs and Digital Clocks

- Let PTA be a closed, diagonal-free probabilistic timed automaton, L′ be a set of its locations, (r,e) be a cost pair
- Central result - using digitization, we prove that:
  - Minimum expected costs w.r.t (r,e) accumulated before reaching L′ in real-clocked PTA and digitally-clocked PTA agree
  - Same for the maximum expected costs
- Proof idea:
  - for each scheduler of nondeterminism in real-clocked PTA, we can construct a discrete-clocked scheduler with a lower expected cost
    - How? Digitize real-clocked paths of the scheduler such that total duration along the path is always rounded down; then total time-elapse cost is also always rounded down
  - for maximum: symmetric (round durations up)

Case Study: IPv4 ZeroConf Protocol

- IPv4 ZeroConf protocol [Cheshire,Adoba,Guttman'02]
  - New IETF standard for dynamic network self-configuration
  - Link-local (no routers within the interface)
  - No need for an active DHCP server
  - Aimed at home networks, wireless ad-hoc networks, hand-held devices
  - "Plug and play"

- Self-configuration
  - Performs assignment of IP addresses
  - Symmetric, distributed protocol
  - Uses random choice and timing delays

IPv4 ZeroConf Protocol

- Select an IP address out of 65024 at random
- Send a probe querying if address in use, and listen for 2 seconds
  - If positive reply received, restart
  - Otherwise, continue sending probes and listening (2 seconds)
- If K probes sent with no reply, start using the IP number
  - Send 2 packets, at 2 second intervals, asserting IP address is being used
  - If a conflicting assertion received, either:
    - defend (send another asserting packet)
    - defer (stop using the IP address and restart)

IPv4 ZeroConf Protocol…

- Possible problem…
  - IP number chosen may be already in use, but:
    - Probes or replies may get lost or delayed (host too busy)
- Issues:
  - Self-configuration delays may become unacceptable
    - Would you wait 8 seconds to self-configure your PDA?
  - No justification for parameters
    - for example K=4 in the standard

PTA Model of the Protocol

- Different models studied:
  - Discrete-time Markov chain and Markov reward models (analytical)
    - [Bohnenkamp-van der Stok-Hermanns-Vaandrager03] and [Andova-Katoen03]
  - Timed automata model using UPPAAL [Zhang-Vaandrager03]
  - PTA model with digital clocks using PRISM, this talk

- Parallel composition of two PTAs:
  - one (joining) host, modelled in detail
  - environment (communication medium + other hosts)

- Variables:
  - K (number of probes sent before the IP address is used)
  - the probability of message loss
  - the number of other hosts already in the network

Expected Costs

- Compute minimum/maximum expected cost accumulated before obtaining a valid IP address?
- Costs:
  - Time should be costly: the host should obtain a valid IP address as soon as possible
  - Using an IP address that is already in use should be very costly: minimise probability of error
- Cost pair: (r,e)
  - r=1 (t time units elapsing corresponds to a cost of t)
  - $e=10^{12}$ for the event corresponding to using an address which is already in use
  - e=0 for all other events

Performance Analysis

- Use the probabilistic model checker PRISM
  - prototype extension for cost-based properties
- PTAs with digital clocks can be encoded directly in the PRISM modelling language
  - as can PTA costs
- Implemented algorithms of [de Alfaro97]
  - stochastic shortest path problem for finite-state MDPs
  - similar to existing PCTL model checking algorithms

Results

Prob. of message loss = 0.001

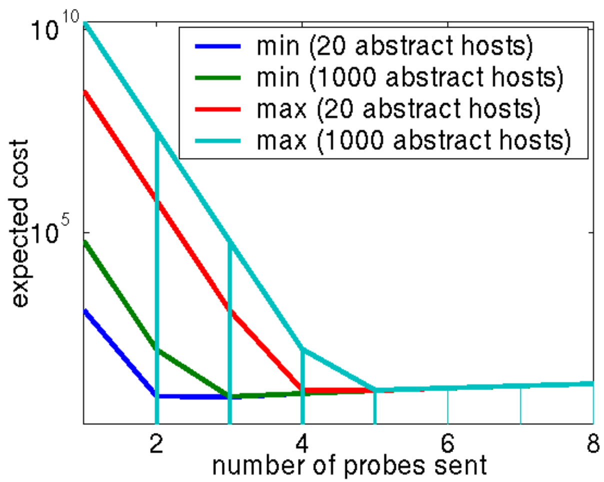

Prob. of message loss = 0.01

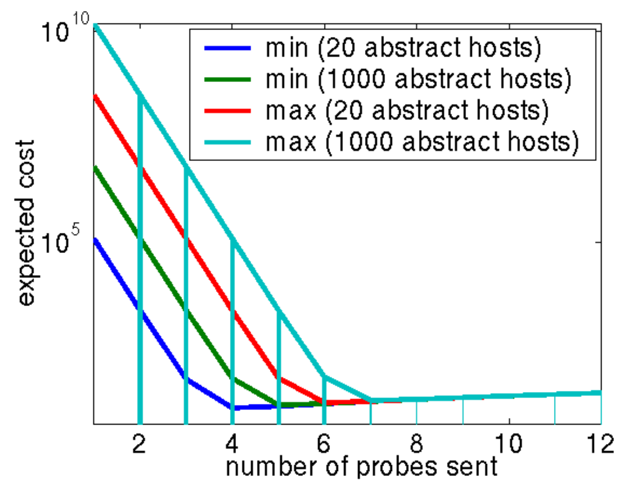

- Sending a high number of probes increases the cost
  - increases delay before a fresh IP address can be used
- Sending a low number of probes increases the cost
  - increases probability of using an IP address already in use
- Similar results to the simpler model of [BvdSHV03]

Conclusions

- Computed expected-cost properties of probabilistic timed automata
- Employed digitally-clocked models to obtain results which also hold for real-clocked models
- More results available at the PRISM web-page

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

- Extensions:
  - Lift the restriction on constant time-elapse cost rates
  - Try other solution methods: regions, zones