

Probabilistic verification and synthesis

Marta Kwiatkowska

Department of Computer Science, University of Oxford

KTH, Stockholm, August 2015



Lecture plan

- Course slides and lab session
 - <u>http://www.prismmodelchecker.org/courses/kth15/</u>
 - 5 sessions: lectures 9-12noon, labs 2.30-5pm
 - 1 Introduction
 - 2 Discrete time Markov chains (DTMCs)
 - 3 Markov decision processes (MDPs)
 - 4 LTL model checking for DTMCs/MDPs
 - 5 Probabilistic timed automata (PTAs)
- For extended versions of this material
 - and an accompanying list of references
 - see: <u>http://www.prismmodelchecker.org/lectures/</u>

Probabilistic models

	Fully probabilistic	Nondeterministic
Discrete time	Discrete-time Markov chains (DTMCs)	Markov decision processes (MDPs)
		Simple stochastic games (<mark>SMGs</mark>)
Continuous time	Continuous-time Markov chains (<mark>CTMCs</mark>)	Probabilistic timed automata (PTAs)
		Interactive Markov chains (IMCs)

Part 5

Probabilistic Timed Automata

Recall – MDPs

- Markov decision processes (MDPs)
 - mix probability and nondeterminism
 - in a state, there is a nondeterministic choice between multiple probability distributions over successor states



- Adversaries
 - resolve nondeterministic choices based on history so far
 - properties quantify over all possible adversaries
 - e.g. $P_{<0.1}$ [\diamond err] maximum probability of an error is < 0.1

Real-world protocol examples

- Systems with probability, nondeterminism and real-time
 - e.g. communication protocols, randomised security protocols
- Randomised back-off schemes
 - Ethernet, WiFi (802.11), Zigbee (802.15.4)
- Random choice of waiting time
 - Bluetooth device discovery phase
 - Root contention in IEEE 1394 FireWire
- Random choice over a set of possible addresses
 - IPv4 dynamic configuration (link-local addressing)
- Random choice of a destination
 - Crowds anonymity, gossip-based routing

Overview (Part 5)

- Time, clocks and zones
- Probabilistic timed automata (PTAs)
 - definition, examples, semantics, time divergence
- PTCTL: A temporal logic for PTAs
 - syntax, examples, semantics
- Model checking for PTAs
 - the region graph
 - digital clocks
 - zone-based approaches:
 - (i) forwards reachability
 - (ii) backwards reachability
 - (iii) game-based abstraction refinement
- Costs and rewards
- Parameter synthesis

Real-time systems verification

- Classical model checking
 - labelled transition systems as models
 - CTL as specification notation

• Many systems feature real-time aspects

- embedded systems
- in-car and in-flight systems
- communication protocols
- controllers
- etc

• Real-time model checking (e.g. UPPAAL)

- timed automata as models
- TCTL as specification notation

Modelling...



Spec:

Modelling...



Spec:

Modelling...



Spec:

Modelling...



Spec:

Modelling...



Spec:

Modelling...



Spec:

Modelling with time...



Spec:

Timed automata: basics

A timed automaton is a finite graph:



- Finite set of locations
- Finitely many labelled transitions between locations
- Transitions take no time (are instantaneous)
- Automaton can remain in a location for a period of time

Counting time: clocks



Clocks, here t

- real-valued variables
- increase at the same rate as time
- initially t=0
- after a period in Loc₁, it is reset to zero

Guards

Guards enable progress



Transitions in timed automata

- can be guarded
- a guard, e.g. t<3, is a constraint on the value of clock t
- specifies when the transition is enabled
- i.e. t=4 means precisely

Invariants

Invariants enforce progress



Locations in timed automata

- can have invariants
- i.e. a constraint for remaining in the location

Time, clocks and clock valuations

- Dense time domain: non-negative reals $\mathbb{R}_{\geq 0}$
 - from this point on, we will abbreviate $\mathbb{R}_{\geq 0}$ to \mathbb{R}
- Finite set of clocks $x \in X$
 - variables taking values from time domain $\ensuremath{\mathbb{R}}$
 - increase at the same rate as real time
- A clock valuation is a tuple $v \in \mathbb{R}^{X}$. Some notation:
 - v(x): value of clock x in v
 - v+t: time increment of t for v

 $\cdot \ (v+t)(x) = v(x)+t \ \forall x \in X$

-v[Y:=0]: clock reset of clocks $Y \subseteq X$ in v

· v[Y:=0](x) = 0 if $x \in Y$ and v(x) otherwise

Zones (clock constraints)

• Zones (clock constraints) over clocks X, denoted Zones(X):

$$\zeta ::= \mathbf{x} \le d \ | \ \mathbf{c} \le \mathbf{x} \ | \ \mathbf{x} + \mathbf{c} \le \mathbf{y} + d \ | \ \neg \zeta \ | \ \zeta \lor \zeta$$

- where $\textbf{x},\textbf{y} \in \textbf{X}$ and $\textbf{c},\textbf{d} \in \mathbb{N}$
- used for both syntax and algorithms

• Can derive:

- logical connectives, e.g. $\zeta_1\wedge\zeta_2\equiv\neg(\neg\zeta_1\vee\neg\zeta_2)$
- strict inequalities, through negation, e.g. x>5 $\equiv \neg(x \le 5)...$

• Some useful classes of zones:

- closed: no strict inequalities (e.g. x > 5)
- convex: define a convex set, efficient to manipulate

Zones and clock valuations

- A clock valuation v satisfies a zone ζ , written v $\triangleright \zeta$ if - ζ resolves to true after substituting each clock x with v(x)
- The semantics of a zone ζ ∈ Zones(X) is the set of clock valuations which satisfy it (i.e. a subset of ℝ^X)
 - NB: multiple zones may have the same semantics
 - e.g. $(x \le 2) \land (y \le 1) \land (x \le y+2)$ and $(x \le 2) \land (y \le 1) \land (x \le y+3)$
- We consider only canonical zones
 - i.e. zones for which the constraints are as 'tight' as possible
 - $O(|X|^3)$ algorithm to compute (unique) canonical zone [Dil89]
 - allows us to use syntax for zones interchangeably with semantic, set-theoretic operations
 - c-closure ,close(ζ ,c), ignores all constraints which are greater than c



Operations on zones - Set theoretic

• Intersection of two zones: $\zeta_1 \cap \zeta_2$



Operations on zones - Set theoretic

• Union of two zones: $\zeta_1 \cup \zeta_2$



Operations on zones - Set theoretic

• Difference of two zones: $\zeta_1 \setminus \zeta_2$



Operations on zones – Clock resets

- $\zeta[Y:=0] = \{ v[Y:=0] \mid v \triangleright \zeta \}$
 - clock valuations obtained from $\boldsymbol{\zeta}$ by resetting the clocks in Y





Operations on zones: c-closure

- c-closure: close(ζ,c)
 - ignores all constraints which are greater than c



Overview (Part 5)

- Time, clocks and zones
- Probabilistic timed automata (PTAs)
 - definition, examples, semantics, time divergence
- PTCTL: A temporal logic for for PTAs
 - syntax, examples, semantics
- Model checking for PTAs
 - the region graph
 - digital clocks
 - zone-based approaches:
 - (i) forwards reachability
 - (ii) backwards reachability
 - (iii) game-based abstraction refinement
- Costs and rewards
- Parameter synthesis

Probabilistic timed automaton (PTA)

- Models a probabilistic real-time communication protocol
 - starts in location di; after between 1 and 2 time units, the protocol attempts to send the data:
 - with probability 0.9 data is sent correctly, move to location sr
 - \cdot with probability 0.1 data is lost, move to location si
 - in location si, after 2 to 3 time units, attempts to resend
 - · correctly sent with probability 0.95 and lost with probability 0.05



Probabilistic timed automata (PTAs)

- Probabilistic timed automata (PTAs)
 - Markov decision processes (MDPs) + real-valued clocks
 - or: timed automata + discrete probabilistic choice
 - model probabilistic, nondeterministic and timed behaviour
- Syntax: A PTA is a tuple (Loc, I_{init}, Act, X, inv, prob, L)
 - Loc is a finite set of locations
 - $-I_{init} \in Loc$ is the initial location
 - Act is a finite set of actions
 - X is a finite set of clocks
 - inv : Loc \rightarrow Zones(X) is the invariant condition
 - prob \subseteq Loc×Zones(X)×Dist(Loc×2^X) is the probabilistic transition relation
 - L : Loc \rightarrow AP is a labelling function



Probabilistic transition relation

- Probabilistic edge relation
 - − prob ⊆ Loc×Zones(X)×Act×Dist(Loc×2^x)
- Probabilistic transition $(I,g,a,p) \in prob$
 - I is the source location
 - g is the guard
 - a is the action
 - p target distribution

• Edge (I,g,a,p,I',Y)

- from probabilistic edge (l,g,a,p) where p(l',Y)>0
- l' is the target location
- Y is the set of clocks to be reset



PTAs – Behaviour

- A state of a PTA is a pair (I,v) \in Loc $\times \mathbb{R}^{X}$ such that v \triangleright inv(I)
- A PTAs start in the initial location with all clocks set to zero
 let <u>0</u> denote the clock valuation where all clocks have value 0
- For any state (I,v), there is nondeterministic choice between making a discrete transition and letting time pass
 - discrete transition (l,g,a,p) enabled if $v \triangleright g$ and probability of moving to location l' and resetting the clocks Y equals p(l',Y)
 - time transition available only if invariant inv(l) is continuously satisfied while time elapses

PTA – Example



PTAs – Formal semantics

- Formally, the semantics of a PTA P is an infinite-state MDP $M_P = (S_P, s_{init}, Steps, L_P)$ with:
- States: $S_P = \{ (I,v) \in Loc \times \mathbb{R}^X \text{ such that } v \triangleright inv(I) \}$
- Initial state: $s_{init} = (I_{init}, \underline{0})$

actions of MDP M_P are the actions of PTA P or real time delays

- Steps: $S_P \rightarrow 2^{(Act \cup \mathbb{R}) \times Dist(S)}$ such that $(\alpha, \mu) \in Steps(I,v)$ iff:
 - (time transition) $\alpha = t \in \mathbb{R}$, $\mu(I, v+t) = 1$ and $v+t' \triangleright inv(I)$ for all $t' \leq t$
 - (discrete transition) $\alpha = a \in Act$ and there exists (l,g,a,p) \in prob

such that $v \triangleright g$ and, for any $(l',v') \in S_P$: $\mu(l',v') =$ $\sum p(l', Y)$ Y⊆X∧v[Y:=0]=v' Labelling: $L_P(I,v) = L(I)$ multiple resets may give same clock valuation

35

Time divergence

- We restrict our attention to time divergent behaviour
 - a common restriction imposed in real-time systems
 - unrealisable behaviour (i.e. corresponding to time not advancing beyond a time bound) is disregarded
 - also called non-zeno behaviour
- For a path $\omega = s_0(\alpha_0, \mu_0)s_1(\alpha_1, \mu_1)s_2(\alpha_2, \mu_2)...$ in the MDP M_P
 - $D_{\omega}(n)$ denotes the duration up to state s_n
 - i.e. $D_{\omega}(n) = \Sigma \{ \mid \alpha_i \mid 0 \leq i < n \land \alpha_i \in \mathbb{R} \mid \}$
- A path ω is time divergent if, for any $t \in \mathbb{R}_{\geq 0}$:
 - there exists $j \in \mathbb{N}$ such that $D_{\omega}(j){>}t$
- Example of non-divergent path:
 - $s_0(1,\mu_0)s_0(0.5,\mu_0)s_0(0.25,\mu_0)s_0(0.125,\mu_0)s_0...$
PTCTL – Syntax

- PTCTL: Probabilistic timed computation tree logic
 - derived from PCTL [BdA95] and TCTL [AD94]



- where:
 - where Z is a set of formula clocks, $\zeta\in \text{Zones}(X\cup Z),\,z\in Z,$
 - a is an atomic proposition, $p \in [0,1]$ and $\textbf{\sim} \in \{<,>,\leq,\geq\}$

PTCTL – Examples

- + z . $P_{>0.99}$ [packet2unsent U packet1delivered \wedge (z<5)]
 - "with probability greater than 0.99, the system delivers packet 1 within 5 time units and does not try to send packet 2 in the meantime"
- z . $P_{>0.95}$ [(x \leq 3) U (z=8)]
 - "with probability at least 0.95, the system clock x does not exceed 3 before 8 time units elapse"
- + z . $P_{\leq 0.1}$ [G (failure \lor (z \leq 60))]
 - "the system fails after the first 60 time units have elapsed with probability at most 0.01"

PTCTL – Semantics

- Let $(I,v)\in S_P$ and $E\in \mathbb{R}^Z$ be a formula clock valuation



PTCTL – Semantics of until

- Let ω be a path in M_P and \mathcal{E} be a formula clock valuation - $\omega, \mathcal{E} \models \psi$ satisfaction of ψ by ω , assuming \mathcal{E} initially
- $\omega, \mathcal{E} \models \varphi_1 \cup \varphi_2$ if and only if there exists $i \in \mathbb{N}$ and $t \in D_{\omega}(i+1)-D_{\omega}(i)$ such that
 - $\omega(i)+t, \mathcal{E}+(D_{\omega}(i)+t) \vDash \varphi_2$
 - \forall t' \leq t . $\omega(i)+t'$, $\mathcal{E}+(D_{\omega}(i)+t') \vDash \varphi_1 \lor \varphi_2$
 - $\hspace{0.1cm} \forall \hspace{0.1cm} j {<} i \hspace{0.1cm} . \hspace{0.1cm} \forall \hspace{0.1cm} t' {\leq} \hspace{0.1cm} D_{\omega}(j{+}1) {-} D_{\omega}(j) \hspace{0.1cm} . \hspace{0.1cm} \omega(j) {+} t', \\ \mathcal{E} {+} (D_{\omega}(j) {+} t') \vDash \varphi_1 \hspace{0.1cm} \lor \hspace{0.1cm} \varphi_2$
- + Condition " $\varphi_1 \lor \varphi_2$ " different from PCTL and CSL
 - usually φ_2 becomes true and φ_1 is true until this point
 - difference due to the density of the time domain
 - to allow for open intervals use disjunction $\varphi_1 \vee \varphi_2$
 - for example consider $x{\le}5$ U $x{>}5$ and $x{<}5$ U $x{\ge}5$

Probabilistic reachability in PTAs

- For simplicity, in some cases, we just consider probabilistic reachability, rather than full PTCTL model checking
 - i.e. min/max probability of reaching a set of target locations
 - can also encode time-bounded reachability (with extra clock)
- Still captures a wide range of properties
 - probabilistic reachability: "with probability at least 0.999, a data packet is correctly delivered"
 - probabilistic invariance: "with probability 0.875 or greater, the system never aborts"
 - probabilistic time-bounded reachability: "with probability 0.01 or less, a data packet is lost within 5 time units"
 - bounded response: "with probability 0.99 or greater, a data packet will always be delivered within 5 time units"

Overview (Part 5)

- Time, clocks and zones
- Probabilistic timed automata (PTAs)
 - definition, examples, semantics, time divergence
- PTCTL: A temporal logic for for PTAs
 - syntax, examples, semantics
- Model checking for PTAs
 - the region graph
 - digital clocks
 - zone-based approaches:
 - (i) forwards reachability
 - (ii) backwards reachability
 - (iii) game-based abstraction refinement
- Costs and rewards
- Parameter synthesis

PTA model checking – Summary

- Several different approaches developed
 - basic idea: reduce to the analysis of a finite-state model
 - in most cases, this is a Markov decision process (MDP)
- Region graph construction [KNSS02]
 - shows decidability, but gives exponential complexity
- Digital clocks approach [KNPS06]
 - (slightly) restricted classes of PTAs
 - works well in practice, still some scalability limitations
- Zone-based approaches:
 - (preferred approach for non-probabilistic timed automata)
 - forwards reachability [KNSS02]
 - backwards reachability [KNSW07]
 - game-based abstraction refinement [KNP09c]

The region graph

- Region graph construction for PTAs [KNSS02]
 - adapts region graph construction for timed automata [ACD93]
 - partitions PTA states into a finite set of regions
 - based on notion of clock equivalence
 - construction is also dependent on PTCTL formula
- + For a PTA P and PTCTL formula φ
 - construct a time-abstract, finite-state MDP $R(\phi)$
 - translate PTCTL formula φ to PCTL formula φ'
 - $-\phi$ is preserved by region quivalence
 - i.e. φ holds in a state of M_P if and only if φ' holds in the corresponding state of $R(\varphi)$
 - model check $R(\phi)$ using standard methods for MDPs

The region graph – Clock equivalence

- Regions are sets of clock equivalent clock valuations
- Some notation:
 - let c be largest constant appearing in PTA or PTCTL formula
 - let [t] denotes the integral part of t
 - t and t' agree on their integral parts if and only if
 - $(1) \lfloor t \rfloor = \lfloor t' \rfloor$
 - (2) t and t' are both integers or neither is an integer
- The clock valuations v and v' are clock equivalent (v \cong v') if:
 - for all clocks $x \in X$, either:
 - \cdot v(x) and v'(x) agree on their integral parts
 - v(x) > c and v'(x) > c
 - for all clock pairs $x, x' \in X$, either:
 - \cdot v(x) v(x') and v'(x) v'(x') agree on their integral parts
 - v(x) v(x') > c and v'(x) v'(x') > c

Region graph – Clock equivalence

• Example regions (for 2 clocks x and y)



Region graph – Clock equivalence

- Fundamental result: if $v\cong v'$, then $v \vartriangleright \zeta \iff v' \vartriangleright \zeta$
 - it follows that $r \vartriangleright \zeta$ is well defined for a region r
- r' is the successor region of r, written succ(r) = r', if
 - for each v \in r, there exists t>0 such that v+t \in r' and v+t' \in r \cup r' for all t'< t



The region graph

- The region graph MDP is (S_R,s_{init},Steps_R,L_R) where...
 - the set of states S_R comprises pairs (I,r) such that I is a location and r is a region over $X\,\cup\,Z$
 - the initial state is $(I_{init}, 0)$
 - the set of actions is {succ} \cup Act
 - succ is a unique action denoting passage of time
 - the probabilistic transition function Steps_R is defined as:
 - $S_R \times 2^{(\{succ\} \cup Act) \times Dist(S_R)}$
 - $\ (succ, \mu) \in \textbf{Steps}_R(l, r) \ iff \ \mu(l, succ(r)) \!=\! 1$
 - (a,µ) $\in \textbf{Steps}_R(I,r)$ if and only if \exists (I,g,a,p) \in prob such that

 $r \vartriangleright g \text{ and, for any (l',r')} \in S_{R:} \quad \mu(l',r') = \sum_{Y \subseteq X \land r[Y:=0]=r'} p(l',Y)$

- the labelling is given by: $L_R(I,r) = L(I)$

Region graph – Example

PTCTL formula: z.P_{~p} [true U (sr<4)]





50

Region graph construction

- Region graph
 - useful for establishing decidability of model checking
 - or proving complexity results for model checking algorithms

• But...

- the number of regions is exponential in the number of clocks and the size of largest constant
- so model checking based on this is extremely expensive
- and so not implemented (even for timed automata)
- Improved approaches based on:
 - digital clocks
 - zones (unions of regions)

Digital clocks

- Simple idea: Clocks can only take integer (digital) values
 - i.e. time domain is $\mathbb N$ as opposed to $\mathbb R$
 - based on notion of ϵ -digitisation [HMP92]
- Only applies to arestricted class of PTAs; zones must be:
 - closed no strict inequalities (e.g. x > 5)
 - Digital clocks semantics yields a finite-state MDP
 - state space is a subset of Loc $\times \ \mathbb{N}^X$, rather than Loc $\times \ \mathbb{R}^X$
 - clocks bounded by c_{max} (max constant in PTA and formula)
 - then use standard techniques for finite -state MDPs

Example – Digital clocks



Digital clocks

- Digital clocks approach preserves:
 - minimum/maximum reachability probabilities
 - a subset of PTCTL properties
 - (no nesting, only closed zones in formulae)
 - only works for the initial state of the PTA
 - (but can be extended to any state with integer clock values)

In practice:

- translation from PTA to MDP can often be done manually
- (by encoding the PTA directly into the PRISM language)
- automated translations exist
- many case studies, despite "closed" restriction
- Problem: can lead to very large MDPs
 - alleviated partially by efficient symbolic model checking

Zone-based approaches

- An alternative is to use **zones** to construct an MDP
- Conventional symbolic model checking relies on computing
 - post(S') the states that can be reached from a state in S' in a single step
 - pre(S') the states that can reach S' in a single step
- Extend these operators to include time passage
 - dpost[e](S') the states that can be reached from a state in S' by traversing the edge e
 - tpost(S') the states that can be reached from a state in S' by letting time elapse
 - pre[e](S') the states that can reach S' by traversing the edge e
 - tpre(S') the states that can reach S' by letting time elapse

Zone-based approaches

- Symbolic states (I, ζ) where
 - $I \in Loc$ (location)
 - $\boldsymbol{\zeta}$ is a zone over PTA clocks and formula clocks
 - generally fewer zones than regions
- tpost(I, ζ) = (I, $\land \zeta \land inv(I)$)
 - $\checkmark \zeta$ can be reached from ζ by letting time pass
 - $\checkmark \zeta \land inv(I)$ must satisfy the invariant of the location I
- tpre(I, ζ) = (I, $\checkmark \zeta \land inv(I)$)
 - $\checkmark \zeta$ can reach ζ by letting time pass
 - $\checkmark \zeta \land$ inv(l) must satisfy the invariant of the location l

Zone-based approaches

- For an edge e = (I,g,a,p,I',Y) where
 - I is the source
 - g is the guard
 - a is the action
 - l' is the target
 - Y is the clock reset
- dpost[e](I, ζ) = (I', ($\zeta \land g$)[Y:=0])
 - $\zeta \wedge g$ satisfy the guard of the edge
 - $(\zeta \land g)[Y:=0]$ reset the clocks Y
- dpre[e](l', ζ ') = (l, [Y:=0] ζ ' \land (g \land inv(l)))
 - $[Y=0]\zeta'$ the clocks Y were reset
 - [Y:=0] $\zeta' \land$ (g \land inv(l)) satisfied guard and invariant of l

Forwards reachability

- Based on the operation $post[e](I,\zeta) = tpost(dpost[e](I,\zeta))$
 - $(l',v') \in post[e](l,\zeta)$ if there exists $(l,v) \in (l,\zeta)$ such that after traversing edge e and letting time pass one can reach (l',v')

Forwards algorithm (part 1)

•

- start with initial state $S_F = \{tpost((I_{init}, \underline{0}))\}$ then iterate for each symbolic state $(I, \zeta) \in S_F$ and edge e add post[e](I, \zeta) to S_F
- until set of symbolic states $\rm S_{\rm F}$ does not change
- To ensure termination need to take c-closure of each zone encountered (c is the largest constant in the PTA)

Forwards reachability

- Forwards algorithm (part 2)
 - construct finite state MDP (S_F ,(I_{init} ,<u>0</u>),Steps_F,L_F)
 - states S_F (returned from first part of the algorithm)
 - $L_F(I,\zeta){=}L(I)$ for all $(I,\zeta)\in S_F$
 - $\mu \in \text{Steps}_F(I, \zeta)$ if and only if there exists a probabilistic edge (I,g,a,p) of PTA such that for any (I', ζ ') $\in Z$:

 $\mu(l',\zeta') = \sum \left\{ \left| p(l',X) \right| (l,g,\sigma,p,l',X) \in edges(p) \land post[e](l,\zeta) = (l',\zeta') \left| \right\} \right\}$

summation over all the edges of (I,g,a,p) such that applying **post** to (I, ζ) leads to the symbolic state (I', ζ ')

Forwards reachability – Example





Forwards reachability - Limitations

- Problem reduced to analysis of finite-state MDP, but...
- Only obtain upper bounds on maximum probabilities
 - caused by when edges are combined
- Suppose **post**[e₁](I, ζ)=(I_1, ζ_1) and **post**[e₂](I, ζ)=(I_2, ζ_2)
 - where e_1 and e_2 from the same probabilistic edge
- By definition of post
 - there exists $(I,v_i) \in (I,\zeta)$ such that a state in (I_i, ζ_i) can be reached by traversing the edge e_i and letting time pass
- Problem
 - we combine these transitions but are (I,v_1) and (I,v_2) the same?
 - may not exist states in (I, ζ) for which both edges are enabled

Forwards reachability – Example

- Maximum probability of reaching I_3 is 0.5 in the PTA
 - for the left branch need to take the first transition when x=1
 - for the right branch need to take the first transition when x=0
- · However, in the forwards reachability graph probability is 1

- can reach I_3 via either branch from ($I_0, x=y$)



Backwards reachability

- An alternative zone-based method: backwards reachability
 - state-space exploration in opposite direction, from target to initial states; uses pre rather than post operator
- Basic ideas: (see [KNSW07] for details)
 - construct a finite-state MDP comprising symbolic states
 - need to keep track of branching structure and take conjunctions of symbolic states if necessary
 - MDP yields maximum reachability probabilities for PTA
 - for min. probs, do graph-based analysis and convert to max.
- Advantages:
 - gives (exact) minimum/maximum reachability probabilities
 - extends to full PTCTL model checking
- Disadvantage:
 - operations to implement are expensive, limits applicability
 - (requires manipulation of non-convex zones)

Overview (Part 5)

- Time, clocks and zones
- Probabilistic timed automata (PTAs)
 - definition, examples, semantics, time divergence
- PTCTL: A temporal logic for for PTAs
 - syntax, examples, semantics
- Model checking for PTAs
 - the region graph
 - digital clocks
 - zone-based approaches:
 - (i) forwards reachability
 - (ii) backwards reachability
 - (iii) game-based abstraction refinement
- Costs and rewards
- Parameter synthesis

Abstraction

- Very successful in (non-probabilistic) formal methods
 - essential for verification of large/infinite-state systems
 - hide details irrelevant to the property of interest
 - yields smaller/finite model which is easier/feasible to verify
 - loss of precision: verification can return "don't know"
- Construct abstract model of a concrete system
 - e.g. based on a partition of the concrete state space
 - an abstract state represents a set of concrete states



Automatic generation of abstractions using refinement

- start with a simple coarse abstraction; iteratively refine

Abstraction of MDPs

- Abstraction increases degree of nondeterminism
 - i.e. minimum probabilities are lower and maximums higher



We construct abstractions of MDPs using stochastic games



- yields lower/upper bounds for min/max probabilities



Abstraction refinement

- Consider (max) difference between lower/upper bounds
 - gives a quantitative measure of the abstraction's precision



• If the difference ("error") is too great, refine the abstraction

- a finer partition yields a more precise abstraction
- lower/upper bounds can tell us where to refine (which states)
- (memoryless) strategies can tell us how to refine

Abstraction-refinement loop

Quantitative abstraction-refinement loop for MDPs



 Refinements yield strictly finer partition

• Guaranteed to converge for finite models

• Guaranteed to converge for infinite models with finite bisimulation

Abstraction refinement for PTAs

Model checking for PTAs using abstraction refinement



Abstraction refinement for PTAs

- Computes reachability probabilities in PTAs
 - minimum or maximum, exact values ("error" $\varepsilon{=}0)$
 - also time-bounded reachability, with extra clock
- Integrated in PRISM (development release)
 - PRISM modelling language extended with clocks
 - implemented using DBMs

In practice, performs very well

- faster than digital clocks or backwards on large example set
- (sometimes by several orders of magnitude)
- handles larger PTAs than the digital clocks approach



Costs and rewards

- Like other models, we can define a reward structure (ρ,ι) for a probabilistic timed automaton
- $\underline{\rho}$: Loc $\rightarrow \mathbb{R}_{\geq 0}$ location reward function
 - $\underline{\rho}(I)$ is the rate at which the reward is accumulated in location I
- $\iota : Act \to \mathbb{R}_{\geq 0}$ action reward function
 - $-\iota(a)$ is the reward associated with performing the action a
- Generalises notion for uniformly priced timed automata
- A useful special case is the elapsed time
 - $\underline{\rho}(I) = 1$ for all locations $I \in Loc$
 - $\iota(a)=0$ for all actions $a \in Act$

Expected reachability

• Expected reachability:

 min./max. expected cumulated reward until some set of states (locations) is reached

Example properties

- "the maximum expected time until a data packet is delivered"
- "the minimum expected number of retransmissions before the message is correctly delivered"
- "the maximum expected number of lost messages within the first 200 seconds"

Model checking

- digital clocks semantics preserves expected reachability
- so can use existing MDP reward model checking techniques
- zone-based approaches solved recently [FORMATS 2015]

Summary

- Probabilistic timed automata (PTAs)
 - combine probability, nondeterminism, real-time
 - well suited for e.g. for randomised communication protocols
 - MDPs + clocks (or timed automata + discrete probability)
 - extension with continuous distributions exists, but model checking only approximate
- PTCTL: Temporal logic for properties of PTAs
 - but many useful properties expressible with just reachability
- PTA model checking
 - region graph: decidability results, exponential complexity
 - digital clocks: simple and effective, some scalability issues
 - forwards reachability: only upper bounds on max. prob.s
 - backwards reachability: exact results but often expensive
 - abstraction refinement using stochastic games: performs well
PRISM: Recent & new developments

New features:

- 1. parametric model checking
- 2. parameter synthesis
- 3. strategy synthesis
- 4. stochastic multi-player games
- 5. real-time: probabilistic timed automata (PTAs) [CAV 2015]

Further new additions:

- enhanced statistical model checking (approximations + confidence intervals, acceptance sampling)
- efficient CTMC model checking (fast adaptive uniformisation)
- benchmark suite & testing functionality
- <u>www.prismmodelchecker.org</u>



- Probability + nondeterminism + real-time
 - timed automata + discrete probabilistic choice, or...
 - probabilistic automata + real-valued clocks
- PTA example: message transmission over faulty channel



PRISM modelling language

- textual language, based on guarded commands

pta

const int N;

module transmitter s : [0..3] init 0; tries : [0..N+1] init 0; x : clock; invariant (s=0 \Rightarrow x≤2) & (s=1 \Rightarrow x≤5) endinvariant [send] s=0 & tries \le N & x≥1 \rightarrow 0.9 : (s'=3) + 0.1 : (s'=1) & (tries'=tries+1) & (x'=0); [retry] s=1 & x≥3 \rightarrow (s' =0) & (x' =0); [quit] s=0 & tries>N \rightarrow (s' =2); endmodule rewards "energy" (s=0) : 2.5; endrewards

PRISM modelling language

- textual language, based on guarded commands



PRISM modelling language

- textual language, based on guarded commands



PRISM modelling language

- textual language, based on guarded commands



Model checking PTAs in PRISM

- Properties for PTAs:
 - min/max probability of reaching X (within time T)
 - min/max expected cost/reward to reach X
 (for "linearly-priced" PTAs, i.e. reward gain linear with time)
- PRISM has two different PTA model checking techniques...
- "Digital clocks" conversion to finite-state MDP
 - preserves min/max probability + expected cost/reward/price
 - (for PTAs with closed, diagonal-free constraints)
 - efficient, in combination with PRISM's symbolic engines
- Quantitative abstraction refinement
 - zone-based abstractions of PTAs using stochastic games
 - provide lower/upper bounds on quantitative properties
 - automatic iterative abstraction refinement

PRISM: Recent & new developments

New features:

- 1. parametric model checking
- 2. parameter synthesis
- 3. strategy synthesis
- 4. stochastic multi-player games
- 5. real-time: probabilistic timed automata (PTAs) [CAV 2015]

Further new additions:

- enhanced statistical model checking (approximations + confidence intervals, acceptance sampling)
- efficient CTMC model checking (fast adaptive uniformisation)
- benchmark suite & testing functionality
- <u>www.prismmodelchecker.org</u>



Case study: Cardiac pacemaker

- Develop model-based framework
 - timed automata model for pacemaker software [Jiang et al]
 - hybrid heart models in Simulink, adopt synthetic ECG model (non-linear ODE) [Clifford et al]
- Properties
 - (basic safety) maintain
 60-100 beats per minute
 - (advanced) detailed analysis
 energy usage, plotted against timing parameters of the pacemaker
 - parameter synthesis: find values for timing delays that optimise energy usage





Optimal timing delays problem

- Optimal timing delay synthesis for timed automata [EMSOFT2014][HSB 2015]
- The parameter synthesis problem solved is
 - given a parametric network of timed I/O automata, set of controllable and uncontrollable parameters, CMTL property φ and length of path n
 - find the optimal controllable parameter values, for any uncontrollable parameter values, with respect to an objective function O, such that the property ϕ is satisfied on paths of length n, if such values exist
- Consider family of objective functions
 - maximise volume, minimise energy
- Discretise parameters, assume bounded integer parameter space and path length
 - decidable but high complexity (high time constants)

Optimal probability of timing delays

- Previously, no nondeterminism and no probability in the model considered
- Consider parametric probabilistic timed automata (PPTA),
 - e.g. guards of the form $x \leq b$,
- Can we synthesise optimal timing parameters to optimise the reachability probability?
- Semi-algorithm [RP 2014]
 - exploration of parametric symbolic states, i.e. location, time zone and parameter valuations
 - forward exploration only gives upper bounds on maximum probability (resp. lower for minimum)
 - but stochastic game abstraction yields the precise solution...
- Implementation in progress

Quantitative verification - Trends

- Being 'younger', generally lags behind conventional verification
 - much smaller model capacity
 - compositional reasoning in infancy
 - automation of model extraction/adaptation very limited
- Tool usage on the increase, in academic/industrial contexts
 - real-time verification/synthesis in embedded systems
 - probabilistic verification in security, reliability, performance
- Shift towards greater automation
 - specification mining, model extraction, synthesis, verification, ...
- But many challenges remain!

Acknowledgements

- My group and collaborators in this work
- Project funding
 - ERC, EPSRC, Microsoft Research
 - Oxford Martin School, Institute for the Future of Computing
- See also
 - VERWARE <u>www.veriware.org</u>
 - PRISM www.prismmodelchecker.org

PhD Comics and Oxford...



PHD TALES FROM THE "2nd Desserts"

JORGE CHAM @ 2008







- You are welcome to visit Oxford!
- PhD scholarships, postdocs in verification and synthesis, and more

Thank you for your attention

More info here: www.prismmodelchecker.org

