Polynomial-Time Verification of PCTL Properties of MDPs with Convex Uncertainties and its Application to Cyber-Physical Systems

Alberto Puggelli

DREAM Seminar - November 26, 2013

Collaborators and PIs:
Wenchao Li
Dorsa Sadigh
Katherine Driggs Campbell
A. L. Sangiovanni-Vincentelli
S. A. Seshia

TerraSwarm

University of California Berkeley

Goal of this talk

- Spur collaborations with other researchers in the department
  - Developed theoretical framework
  - Developed (prototype) tool implementation
  - Now it is time to apply the framework to relevant case studies
- Success stories
  - Verification of human driver behavior (D. Sadigh, K. Driggs Campbell)
  - On-going integration of the algorithms within PRISM (state-of-the-art tool developed at the University of Birmingham and Oxford University, UK)

Verify a Hybrid World with Uncertainties

Sensor Networks

SoC Power Management

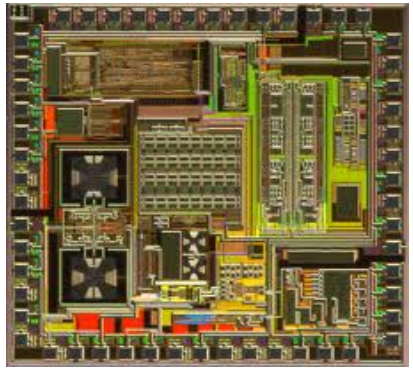

Biochemical Synthesis

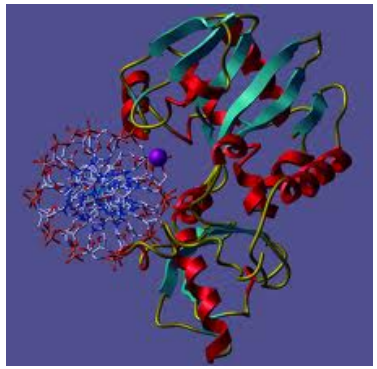

Need to **formally** verify and **quantitatively** analyze system performances in the presence of **uncertainties** (unmodeled dynamics, errors in parameter estimation, faulty and malicious behaviors)

Stock Market Exchange

Renewables Scheduling

Robot Path Planning

Behavior of a Human Driver

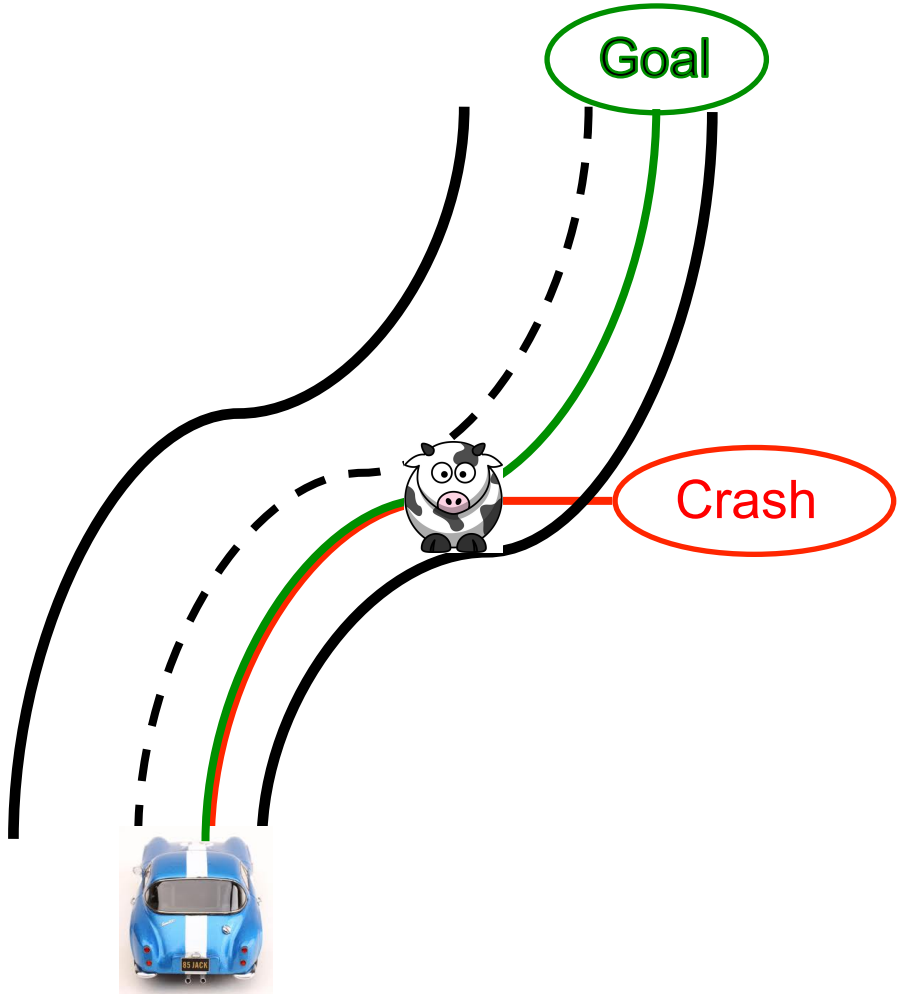

"The driver will always eventually perform the maneuver correctly" – FALSE
"The driver will perform the maneuver correctly with probability higher than 90%" – TRUE

Setup for Model Training

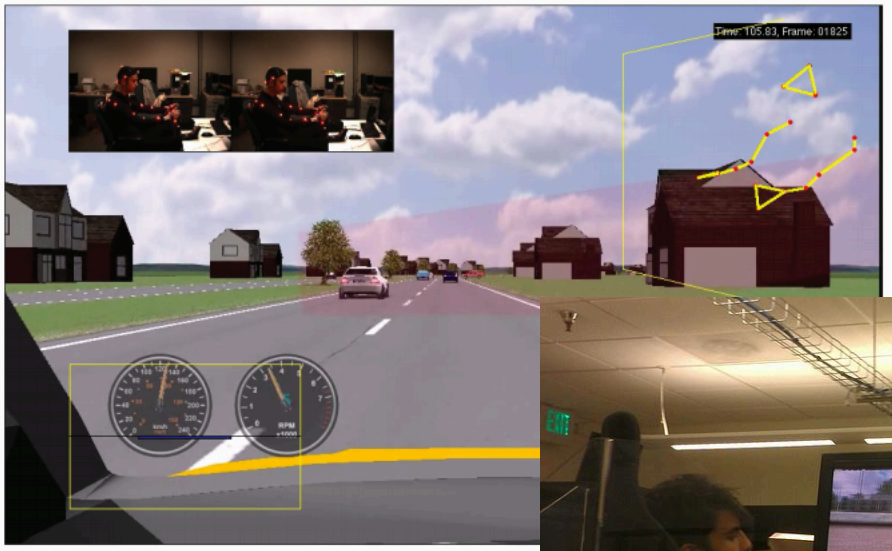

courtesy of V. Vasudevan,
K. Driggs Campbell,
G. Juniwal

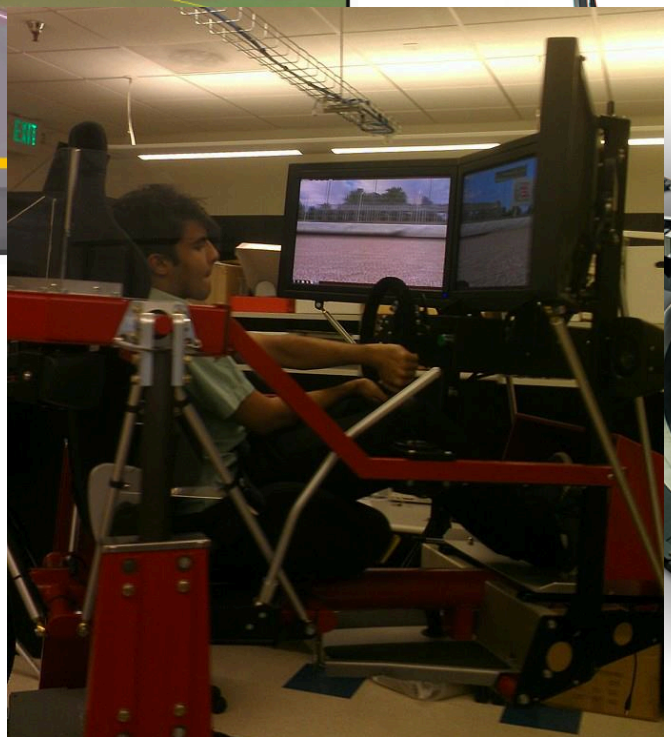

- Intrinsic uncertainties in modeling the human behavior!
- How can we account for this at verification time?

Two More Steps Towards the Goal

[Puggelli et al. '13]
Polynomial-time algorithm for PCTL verification of Convex-MDPs

[Chatterjee et al. '08]
PCTL verification for Interval-MDPs is at most in co-NP

[Kozine et al. '02]
Interval-MDP: Interval Uncertainties in transition probabilities of MDPs

*[*Kwiatkowska *et al. '00]*
PRISM: Algorithms and Tool for PCTL verification of MDPs

[Hansson et al. '94]
Probabilistic Computation Tree Logic (PCTL)

[Bianco'95-Courcoubetis'95]
Verification algorithms for Markov Decision Processes (MDPs)

Outline

- Background
  - Convex-MDP: MDP with Convex Uncertainty Sets
  - Probabilistic Computation Tree Logic (PCTL)
- Polynomial-Time Verification Algorithm[1]
- Case Studies
  - Randomized Consensus Protocol
  - ZeroConf Protocol
  - Behavior of a Human Driver[2]

1. A. Puggelli *et al.*, Proceedings of CAV2013
2. D. Sadigh *et al.*, submitted to AAAI 2014 Symposium

Outline

- Background
  - Convex-MDP: MDP with Convex Uncertainty Sets
  - Probabilistic Computation Tree Logic (PCTL)
- Polynomial-Time Verification Algorithm[1]
- Case Studies
  - Randomized Consensus Protocol
  - ZeroConf Protocol
  - Behavior of a Human Driver[2]

1. A. Puggelli *et al.*, Proceedings of CAV2013
2. D. Sadigh *et al.*, submitted to AAAI 2014 Symposium

Convex-MDP

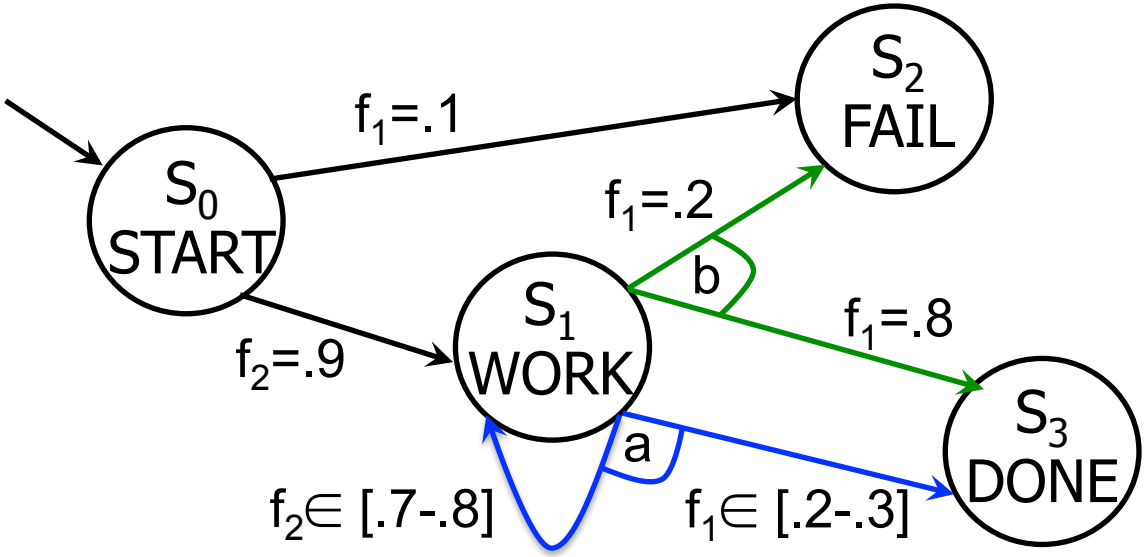

- Action chosen by an *Adversary*
- Transition probability distribution chosen by *Nature*
- Transition probabilistically executed

Interval

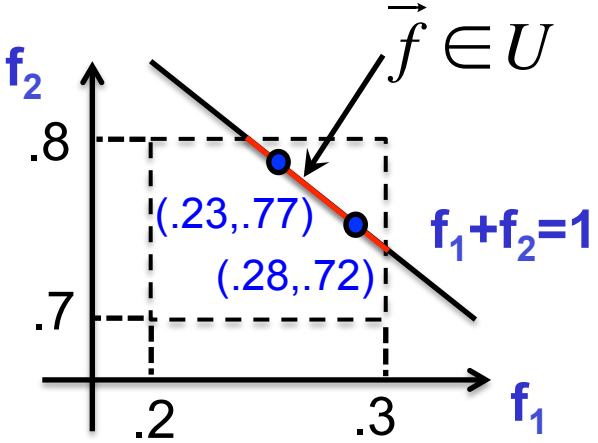

Ellipsoidal

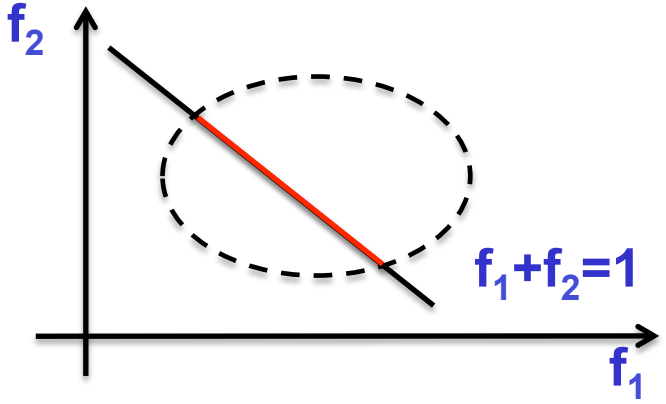

Likelihood

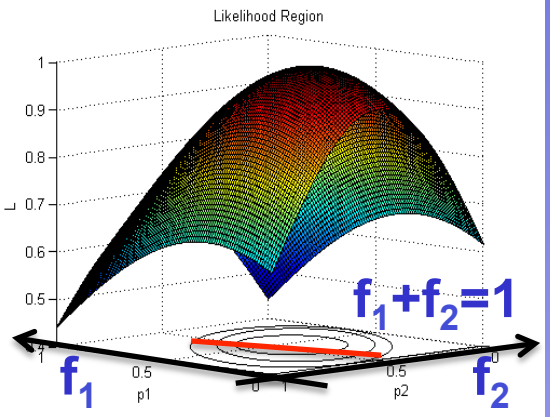

A. Nilim, "Robust Control of Markov Decision Processes with Uncertain Transition Matrices", 2005

Semantics of Convex-MDPs (1)

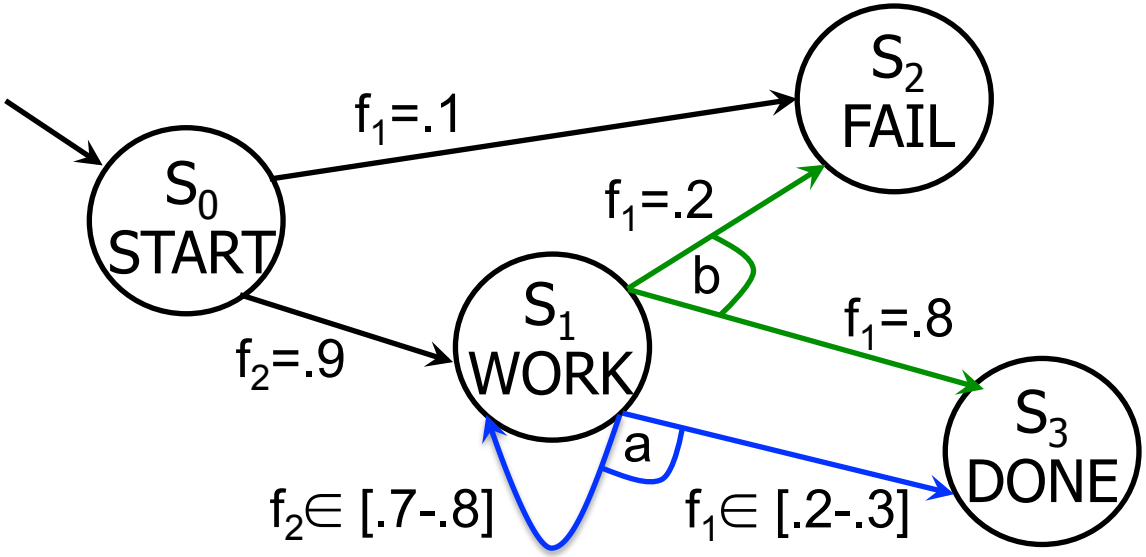

- Action chosen by an *Adversary*
- Transition probability distribution chosen **once** by *Nature*
- Transition probabilistically executed

Path

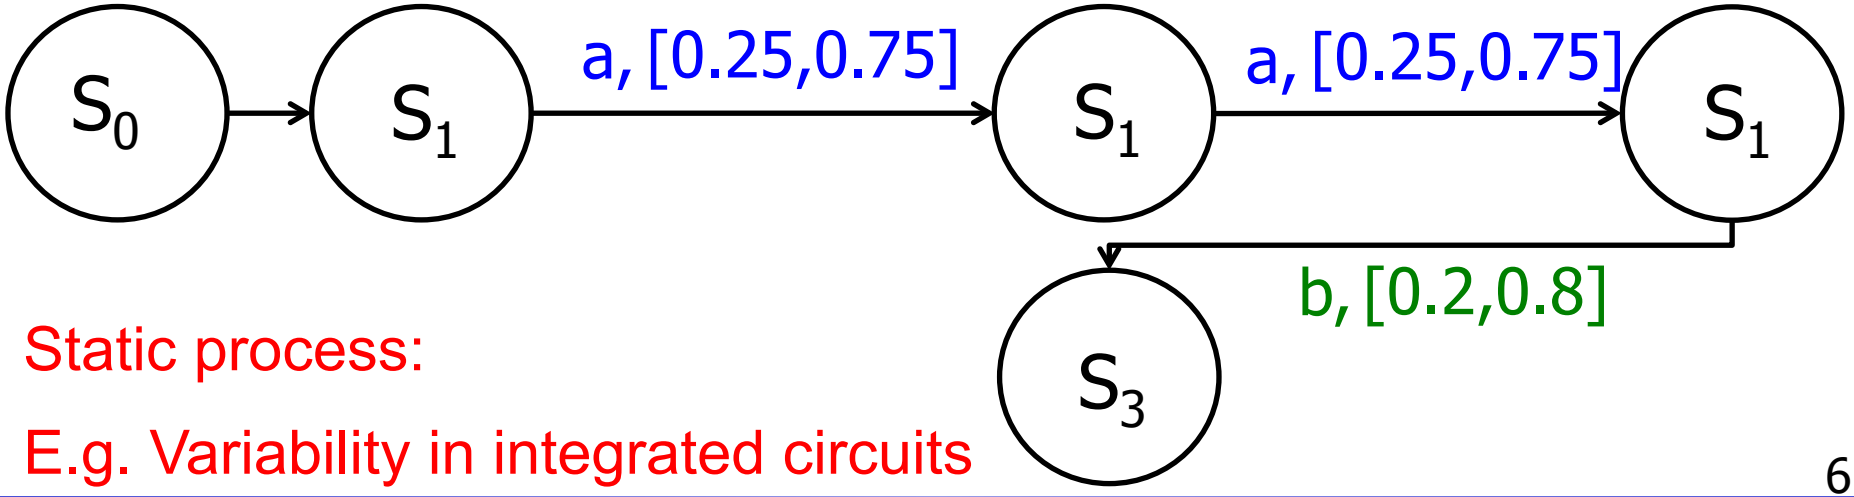

Static process:

E.g. Variability in integrated circuits

Semantics of Convex-MDPs (2)

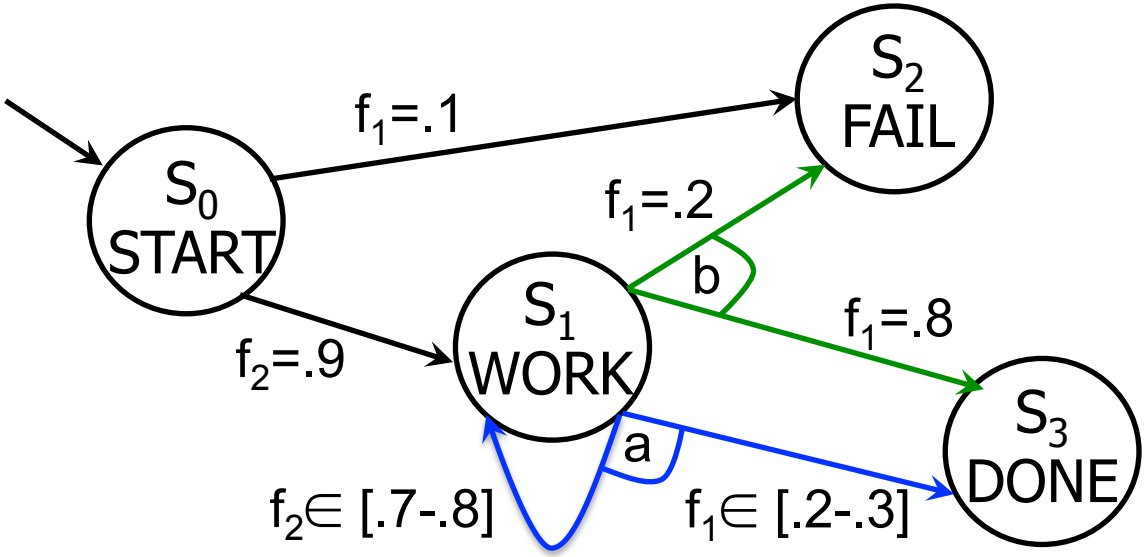

- Action chosen by an *Adversary*
- Transition probability distribution chosen **at each step** by *Nature*
- Transition probabilistically executed

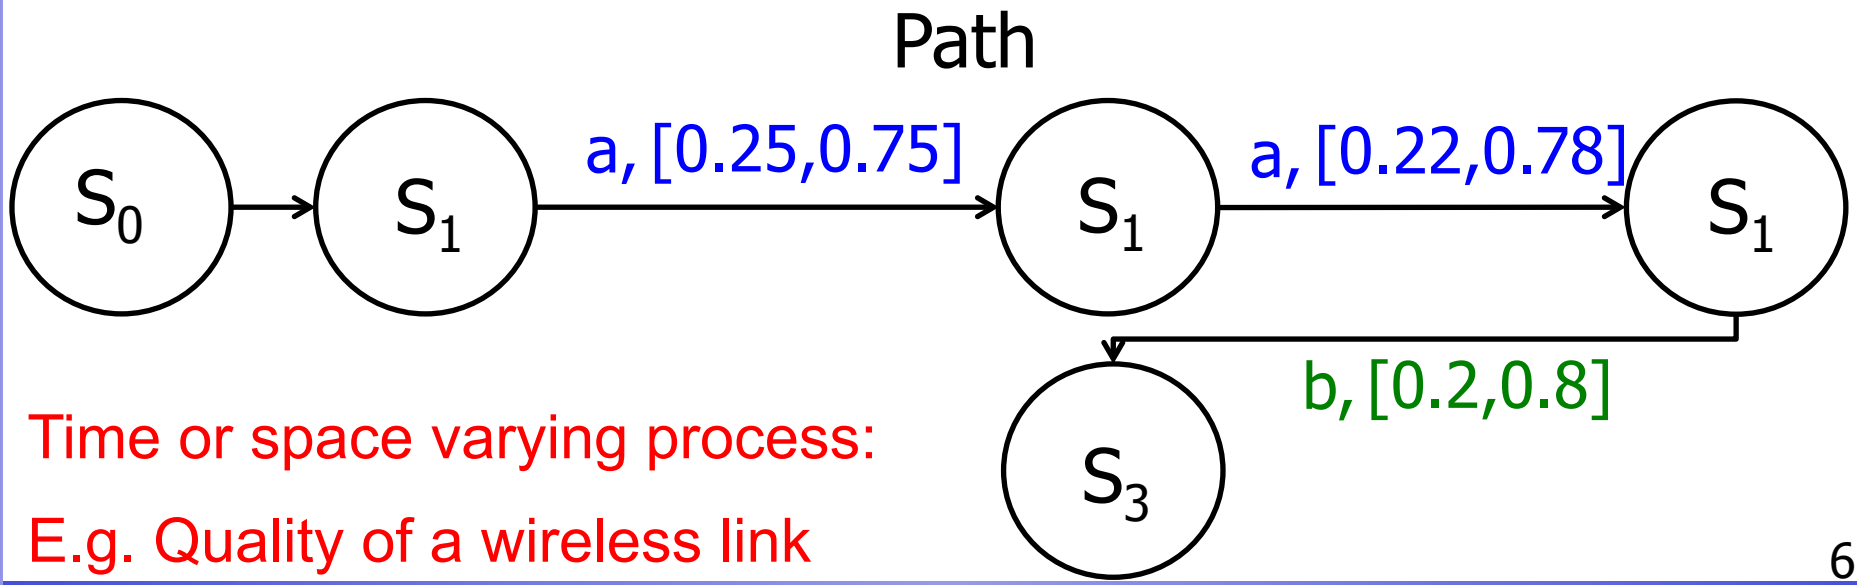

Time or space varying process:

E.g. Quality of a wireless link

Probabilistic Computation Tree Logic

- Logic syntax

$$\phi ::= True \mid \omega \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid P_{\bowtie p}[\psi] \qquad \text{state formulas}$$

$$\psi ::= \mathcal{X}\phi \mid \phi_1 \, \mathcal{U}^{\leq k} \phi_2 \mid \phi_1 \, \mathcal{U} \phi_2 \qquad \text{path formulas}$$

Next — Bounded Until — Unbounded Until

- Logic semantics

$s \models \text{True}$

$s \models \omega$ iff $\omega \in L(s)$

$s \models \neg\phi$ iff $s \not\models \phi$

$s \models \phi_1 \wedge \phi_2$ iff $s \models \phi_1 \wedge s \models \phi_2$

$s \models P_{\bowtie p}[\psi]$ iff $Prob(\{\pi \in \Pi_s(\alpha, \eta^a) \mid \pi \models \psi\}) \bowtie p$

$\forall \alpha \in Adv$ and $\eta^a \in Nat$

- Verification algorithm: solve the optimization problem

$$P_s^{max}[\psi] = \max_{a \in \mathcal{A}(s)} \max_{\mathbf{f}_s^a \in \mathcal{F}_s^a} P_s(a, \mathbf{f}_s^a)[\psi] \overset{?}{\leq} p \qquad P_s^{min}[\psi] = \min_{a \in \mathcal{A}(s)} \min_{\mathbf{f}_s^a \in \mathcal{F}_s^a} P_s(a, \mathbf{f}_s^a)[\psi] \overset{?}{\geq} p$$

Which Logic to Use?

- Qualitative logics (LTL, CTL):
  - Pros: efficient algorithms,
  - Cons: only give “yes/no” answers

- Quantitative logics:
  - PCTL
    - Pros: efficient algorithms, enables quantitative analysis
    - Cons: can’t express arbitrary liveness and fairness properties
  - ω-PCTL[1]
    - Pros: quantitative analysis, express safety, liveness, fairness
    - Cons: no efficient algorithm

1. K. Chatterjee *et al.*, “Model-Checking **ω**-Regular Properties of Interval Markov Chains”, TACAS 2008

Outline

- Background
  - Convex-MDP: MDP with Convex Uncertainty Sets
  - Probabilistic Computation Tree Logic (PCTL)
- Polynomial-Time Verification Algorithm[1]
- Case Studies
  - Randomized Consensus Protocol
  - ZeroConf Protocol
  - Behavior of a Human Driver[2]

1. A. Puggelli *et al.*, Proceedings of CAV2013
2. D. Sadigh *et al.*, submitted to AAAI 2014 Symposium

New Results in Theoretical Complexity

- Size of Convex-MDP
  - $\mathcal{R}$ = O(#States x #Transitions x #Actions)
- Size of PCTL formula
  - $\mathcal{Q}$ = O(#Operators(excluding $\mathcal{U}^{\leq k}$) + #($\mathcal{U}^{\leq k}$) x k_{max})

PCTL Operator	Verification Complexity			
	Puggelli'13		Chatterjee'08*	
	In $\mathcal{R}$	In $\mathcal{Q}$	In $\mathcal{R}$	In $\mathcal{Q}$
Qualitative	P	P	P	P
Next ($\mathcal{X}$)	P	P	co-NP	P
Bounded Until ($\mathcal{U}^{\leq k}$)	P	Pseudo-P in k_{max}	-	-
Unbounded Until ($\mathcal{U}$)	P	P	co-NP	P

*Only interval uncertainties

Unbounded Until in Markov Chains

$$\phi = P_{\geq 0.7}\,[\neg \mathrm{FAIL}\; \mathcal{U}\; \mathrm{DONE}]$$

1. x_i = Probability of satisfying ϕ for i-th state
2. Set up and solve the **system of equations**

$$\begin{cases} x_2 = 0 \\ x_3 = 1 \\ x_0 = 0.9x_1 + 0.1x_2 \\ x_1 = 0.2x_1 + 0.7x_3 + 0.1x_2 \end{cases}$$

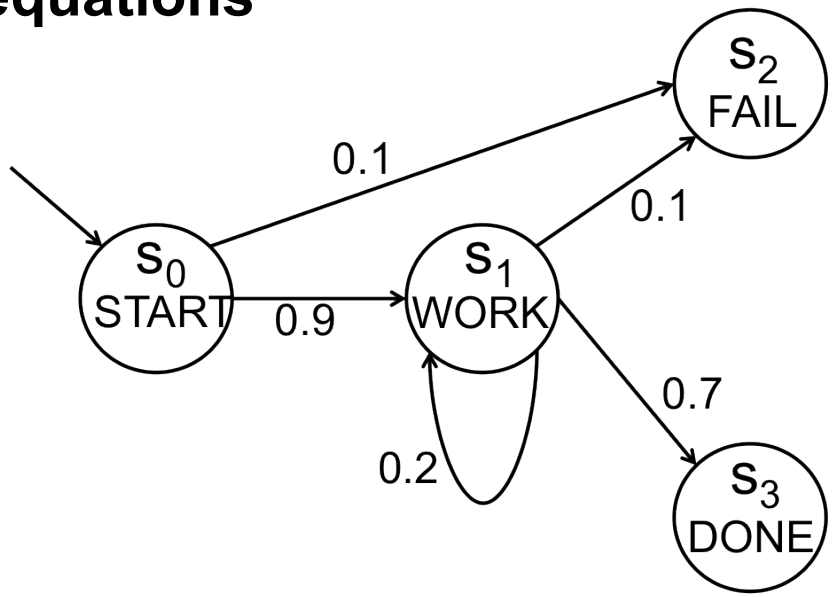

3. #Equations = #States = N
4. Algorithmic complexity $O(N^3)$ ➔ Polynomial in $\mathcal{R}$

Unbounded Until in MDP

$$\phi = P_{\geq 0.7}\,[\neg \text{FAIL}\; \mathcal{U}\; \text{DONE}]$$

1. Need to consider the worst-case adversary (historyless-deterministic enough)
2. Set up and solve the **linear program**

$$\max_x \sum x_i$$

$$s.t.\quad x_2 = 0$$

$$x_3 = 1$$

$$\boxed{\begin{array}{l} x_1 \leq 0.3x_1 + 0.7x_3 \\ x_1 \leq 0.2x_2 + 0.8x_3 \end{array}}$$

$$x_0 = 0.9x_1 + 0.1x_2$$

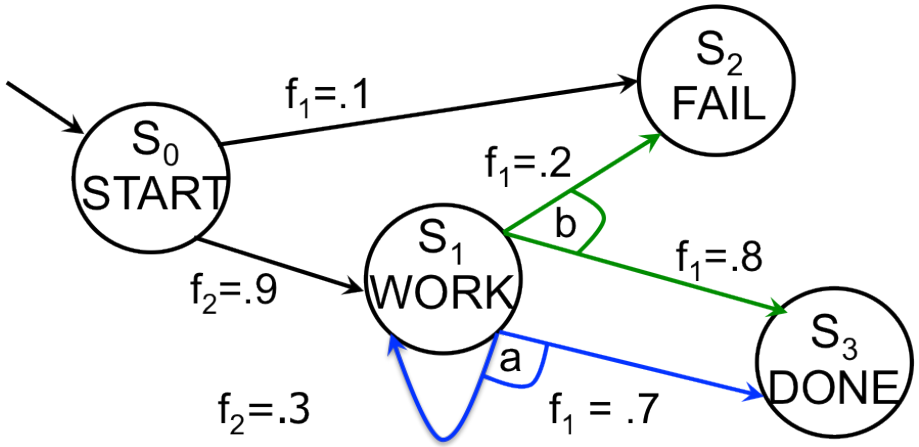

3. #Constraints = O(#States x #Actions)
4. Interior Point ➔ Algorithmic complexity polynomial in $\mathcal{R}$

Unbounded Until in Convex-MDP

$$\phi = P_{\geq 0.7}\left[\neg \text{FAIL}\, \mathcal{U}\, \text{DONE}\right]$$

1. Need to consider the worst-case adversary and nature
2. Set up the **optimization problem**

$$\max_x \sum x_i$$

$$s.t. \quad x_2 = 0$$

$$x_3 = 1$$

$$x_0 = 0.9x_1 + 0.1x_2$$

$$x_1 \leq 0.2x_2 + 0.8x_3$$

$$\boxed{x_1 \leq \min_{\vec{f} \in U} f_1 x_3 + f_2 x_1}$$

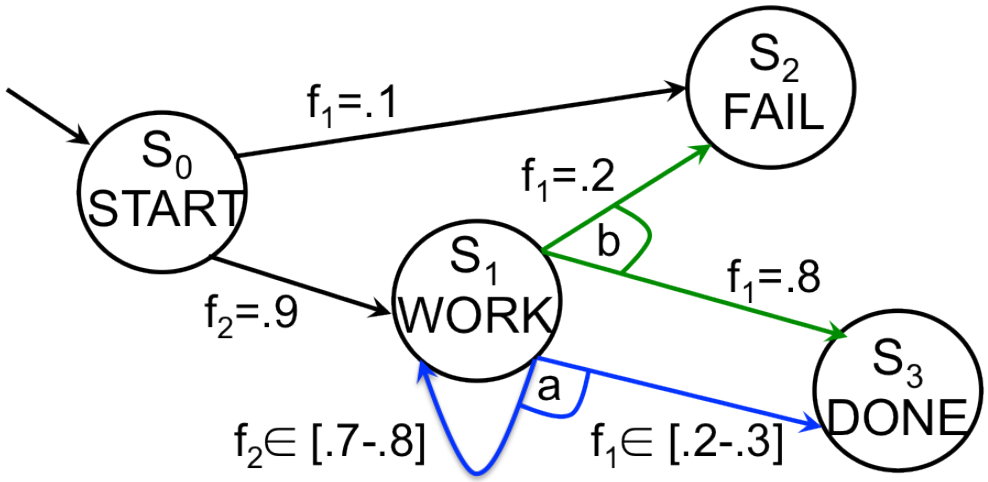

3. The **adversarial nature minimizes the upper bound** on x_i
4. To maintain convexity, need to add one constraint $\forall \vec{f} \in U$
5. Uncountably infinite number of constraints: cannot solve

Unbounded Until in Convex-MDP

- Try all Probability Distributions?

$$\max_x \sum x_i$$

s.t.

$$x_2 = 0$$

$$x_3 = 1$$

$$x_0 = 0.9x_1 + 0.1x_2$$

$$x_1 \le 0.2x_1 + 0.7x_3 + 0.1x_2$$

$$x_1 \le 0.2x_1 + 0.8x_3$$

$$x_1 \le 0.25x_1 + 0.75x_3$$

$$x_1 \le 0.21x_1 + 0.79x_3$$

...

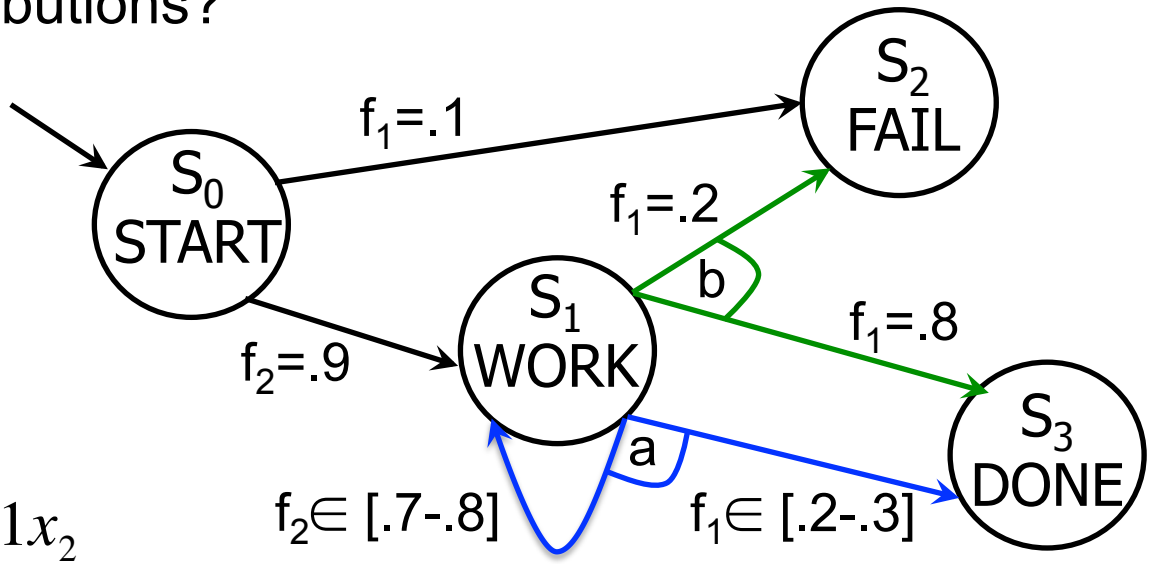

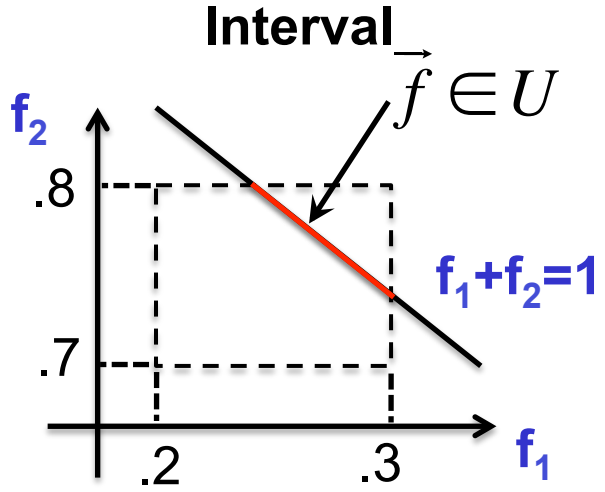

- NO: Uncountably infinite number of distributions

Dual Transformation for the Inner Problem

- Primal Problem $\sigma(\vec{x}) = \min_{\vec{f} \in U} f_1 x_3 + f_2 x_1$
- Dual Problem $d(\vec{x}) = \max_{\vec{\lambda} \in D} g(\vec{\lambda}, \vec{x})$
  - Convex
  - Number of dual variables and constraints is polynomial in $\mathcal{R}$
  - $\sigma(\vec{x}) \geq g(\vec{\lambda}, \vec{x}) \quad \forall \vec{\lambda} \in D$
  - Strong duality holds: $d(\vec{x}) = \sigma(\vec{x})$

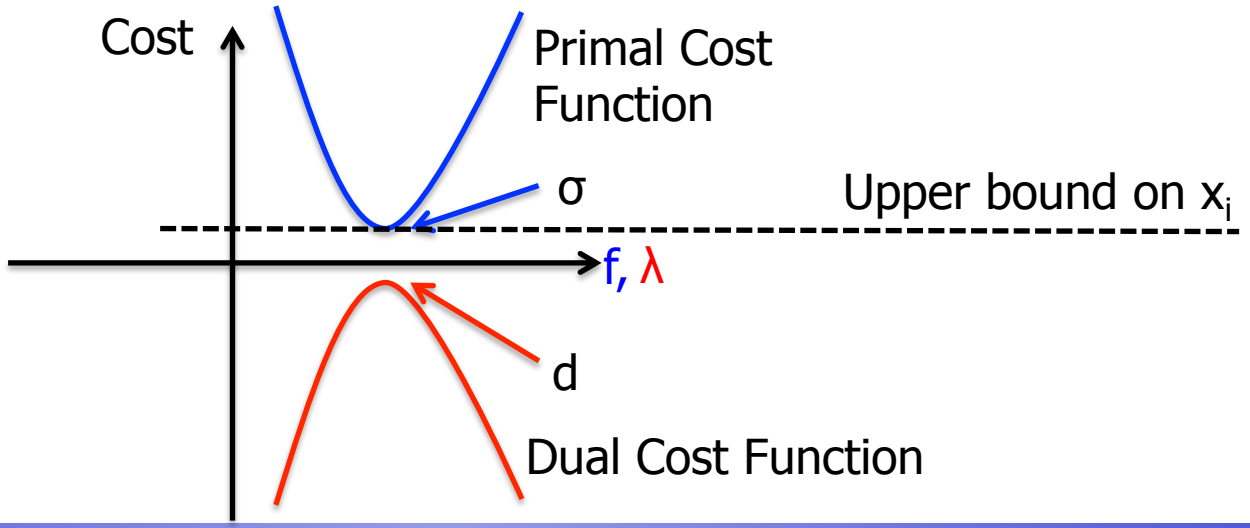

New Formulation

Original formulation

$$\max_x \sum x_i$$
$$s.t.\ x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le 0.2x_2 + 0.8x_3$$
$$x_1 \le \min_{\vec{f} \in U} f_1 x_3 + f_2 x_1$$

Dual transformation of the inner problems

$$\max_x \sum x_i$$
$$s.t.\quad x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le 0.2x_2 + 0.8x_3$$
$$x_1 \le \max_{\lambda \in D} g(\lambda, x)$$

New formulation (drop all inner problems)

$$\max_{x,\lambda} \sum x_i$$
$$s.t.\quad x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le 0.2x_2 + 0.8x_3$$
$$x_1 \le g(\lambda, x)$$
$$\lambda \in D$$

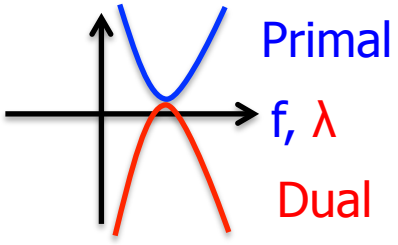

Unbounded Until can be verified by solving one convex problem with a number of variables and constraints polynomial in $\mathcal{R}$.

Outline

- Background
  - Convex-MDP: MDP with Convex Uncertainty Sets
  - Probabilistic Computation Tree Logic (PCTL)
- Polynomial-Time Verification Algorithm[1]
- Case Studies
  - Randomized Consensus Protocol
  - ZeroConf Protocol
  - Behavior of a Human Driver[2]

1. A. Puggelli *et al.*, Proceedings of CAV2013
2. D. Sadigh *et al.*, submitted to AAAI 2014 Symposium

Randomized Consensus Protocol [Aspnes'90]

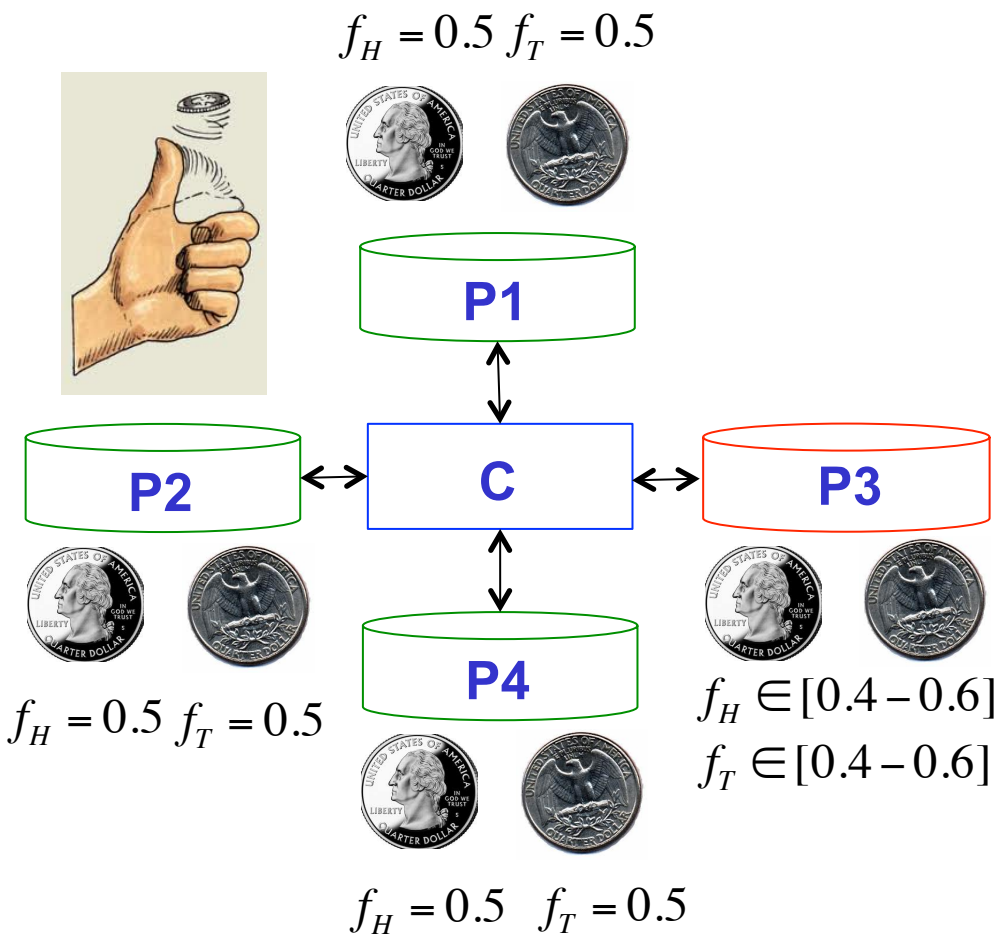

- Study the probability of agreement in a network of asynchronous processes
- Uncertainty models a **faulty/compromised process** which tosses a biased coin

$$P_{s_0}^{min}[\psi] := P_{s_0}^{min}(\mathbf{F}(\{finished\} \wedge \{all\_coins\_equal\_1\}))$$

Randomized Consensus Protocol

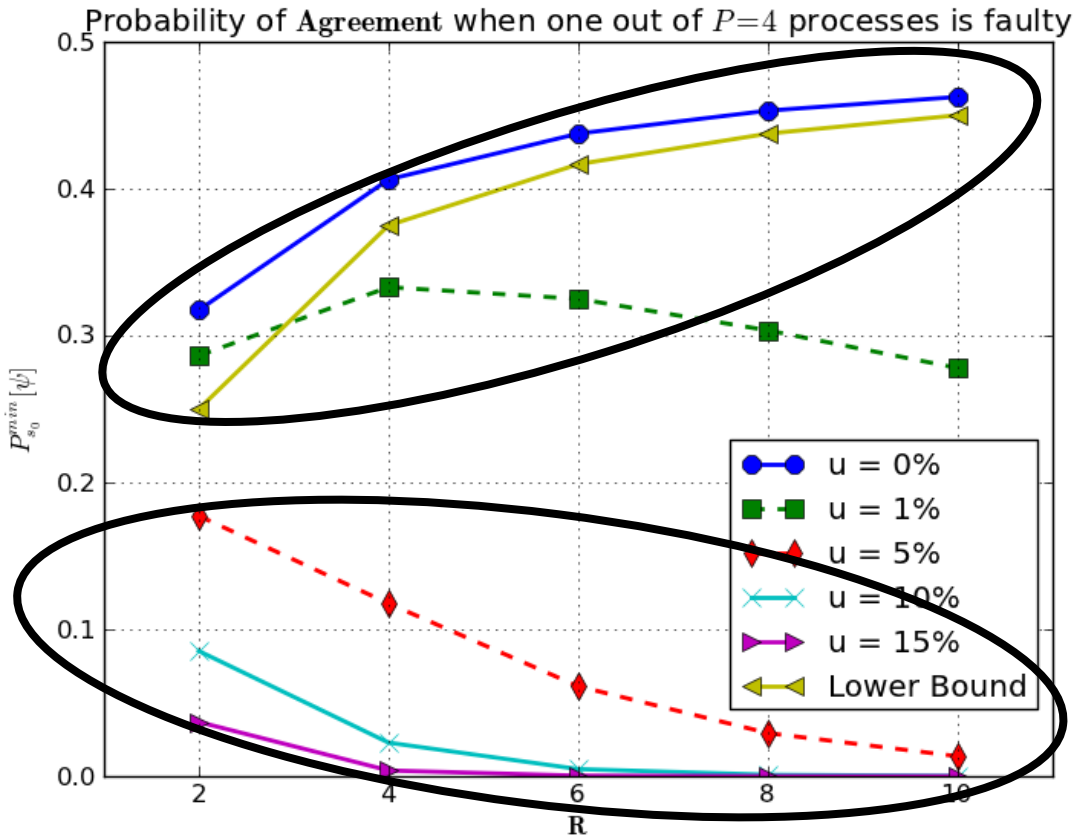

- With **fair** coins, the probability of agreement **increases** for increasing protocol rounds
- In the presence of **uncertainty,** increasing the protocol rounds instead **decreases** the probability of agreement

The proposed analysis allows a better tuning of protocol parameters to accommodate for faulty/compromised processes

Runtime Analysis

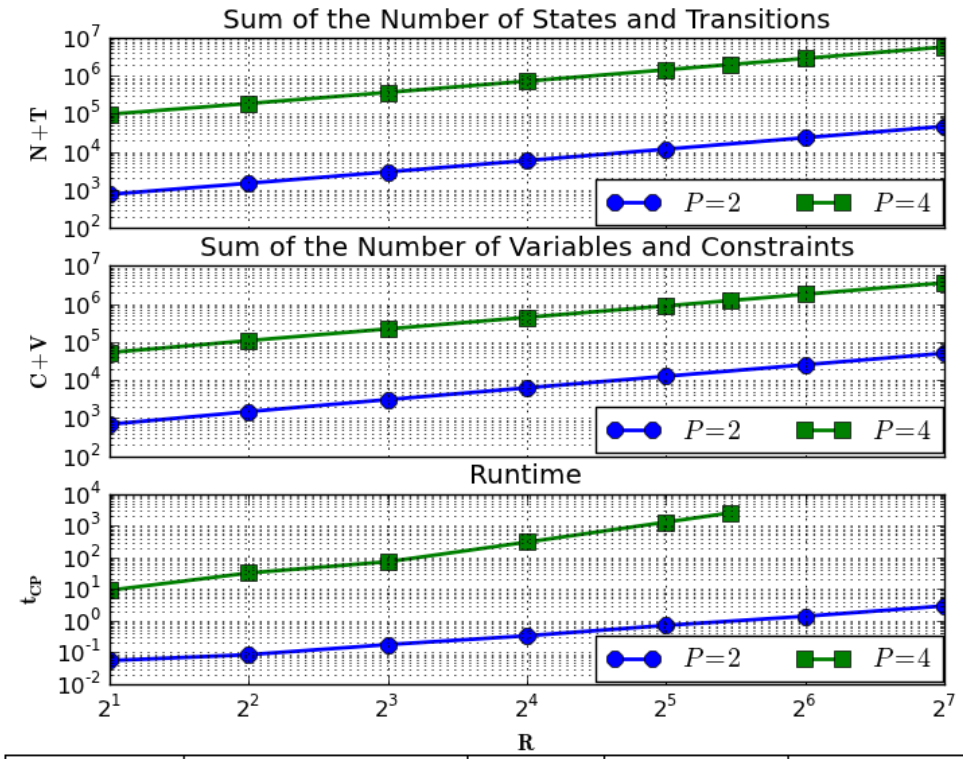

- Use MOSEK as background LP solver
- Size of the convex problem and **runtime scale polynomially**
- Comparable with PRISM[2] and 1000x faster than PARAM[3]

Tool	$P=2, R=2$ $N+T=764$	$R=7$ 2,604	$R=128$ 47,132	$P=4, R=2$ 97,888	$R=32$ 1,262,688	$R=44$ 1,979,488	$P=6, R=4$ 14,211,904
CP	$0.02s$	$0.1s$	$2.1s$	$8.3s$	$1,341s$	2,689	TO
PRISM	$0.01s$	$0.09s$	$196s$	$1s$	$2,047s$	TO	$1860s$
PARAM	$22.8s$	$657s$	TO	TO	TO	TO	TO

1. www.mosek.com
2. Kwiatkowska et al., "PRISM 4.0: Verification of Probabilistic Real-time Systems"
3. Hahn et al., "Synthesis for PCTL in Parametric Markov Decision Processes"

ZeroConf Protocol [Cheshire'05]

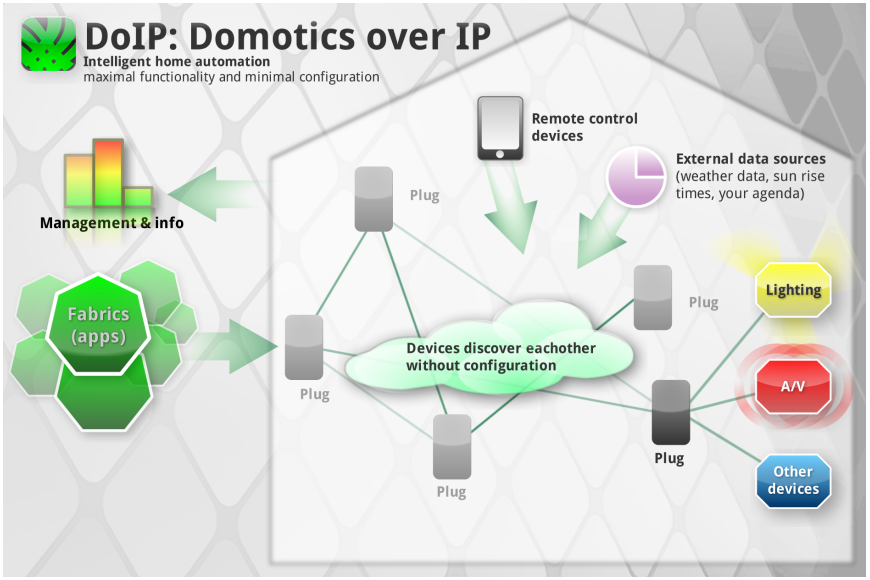

(source: doip.org)

- Study the **QoS of a network** configuration **protocol** for domotic applications
- Model the network as a Timed Automata

- **Maximum likelihood estimator** to model the losses in the **(physical) wireless channel**

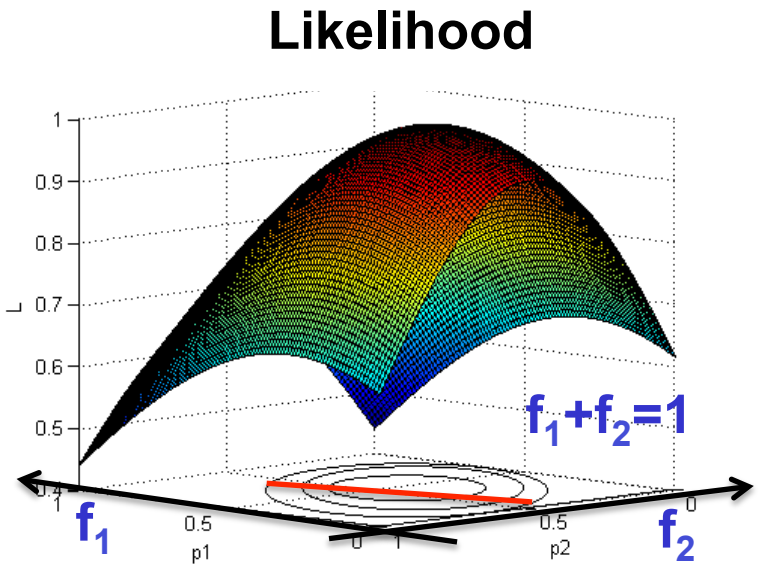

ZeroConf Protocol

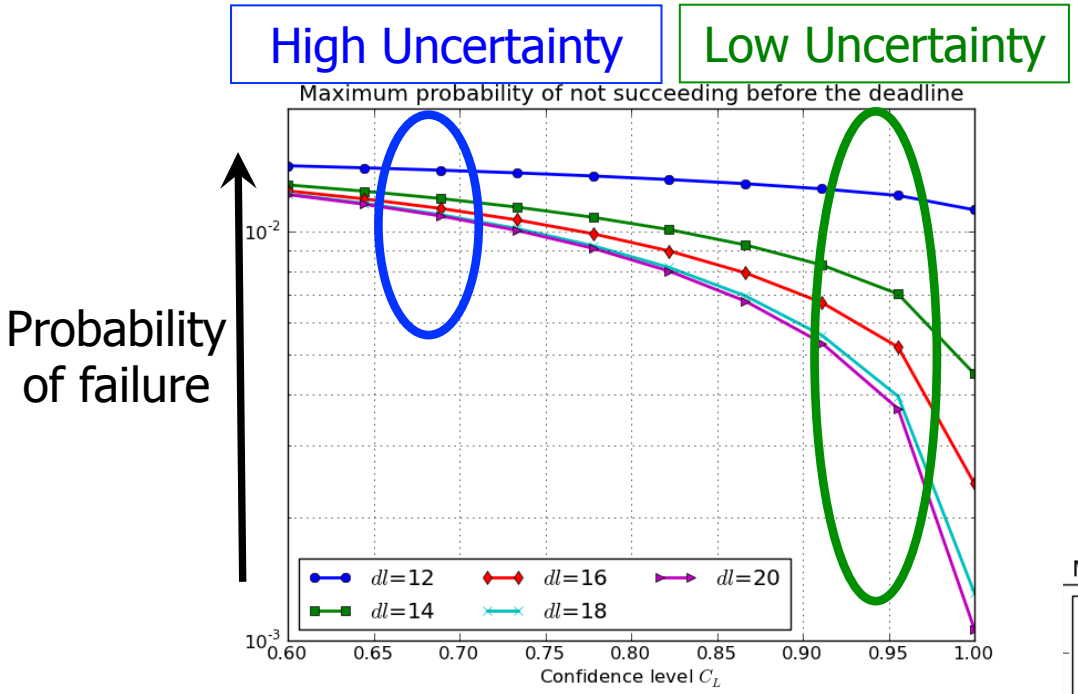

- Probability of failing to register to the network within a preset deadline
- Analysis **with no uncertainties largely underestimates** the probability of failure

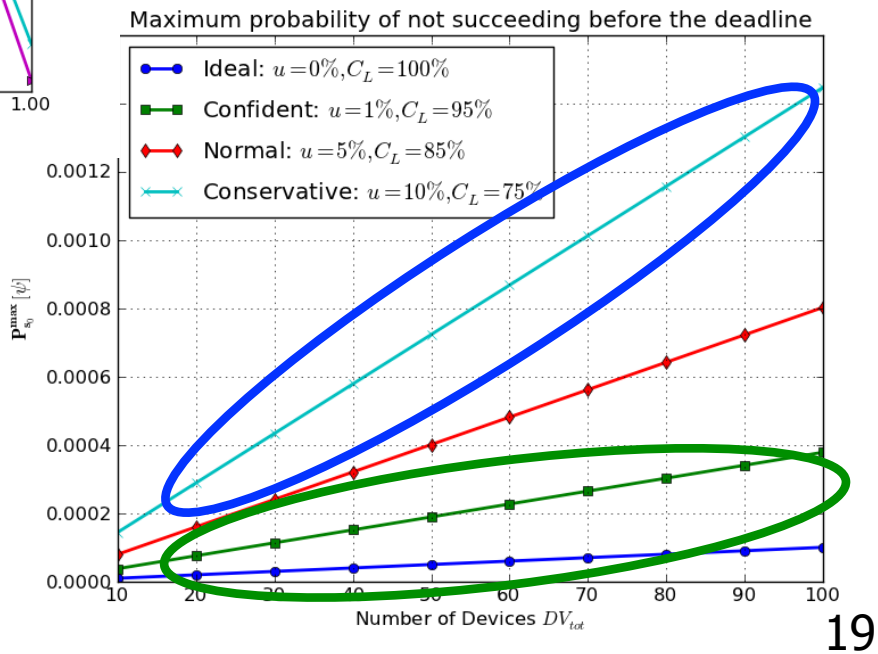

Our analysis enables a robust configuration of protocol parameters to fit variable conditions of operation

Why Modeling the Driver Behavior?

More effective teaching strategies

Collision avoidance

Lane changing

Assisted maneuvers

(Semi)-Autonomous Driving

Driving regulations and insurance terms

Data Collection

- Focus on modeling differences between attentive and distracted driving[1]

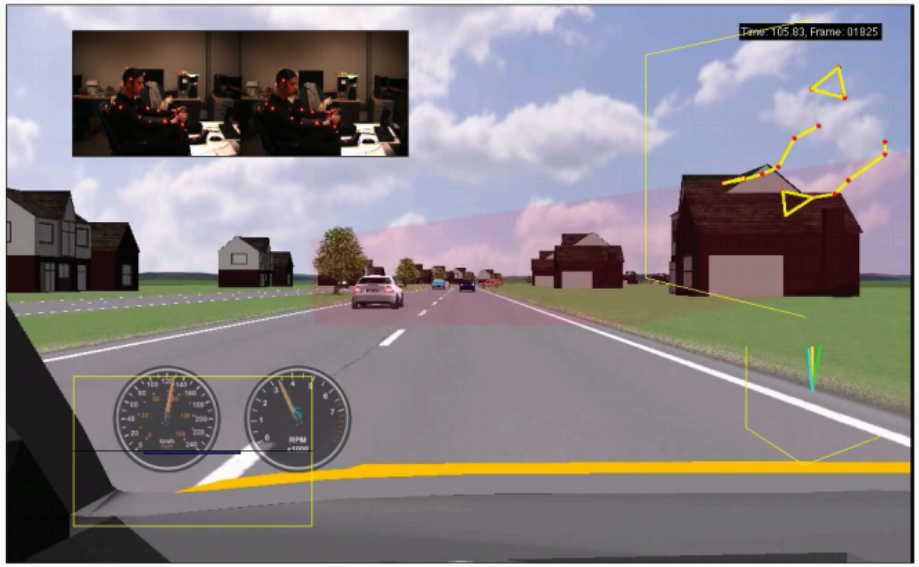

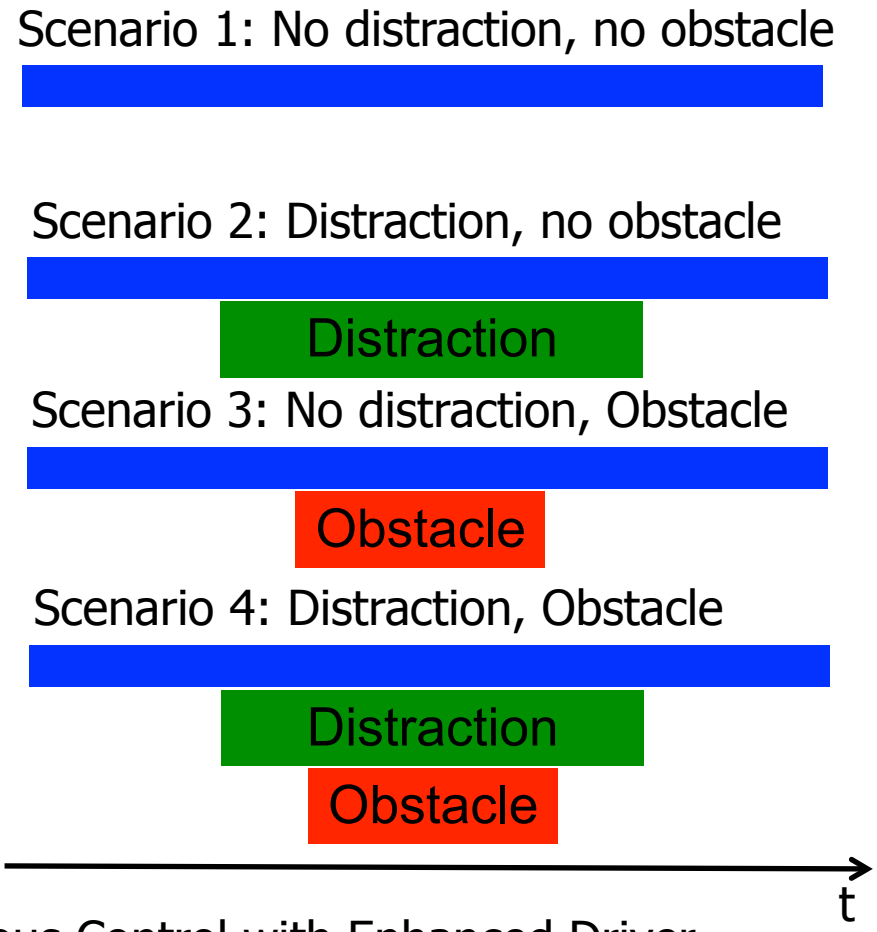

1. V. Vasudevan *et al.*, "Safe Semi-Autonomous Control with Enhanced Driver Modeling", ACC 2012

Library of Atomic Behaviors

- Library of atomic labels L = {distracted, attentive, swerving, braking, accelerating, right lane, left lane…}
- Modes $\subseteq 2^L$ E.g. m_1=(distracted, right lane)
- **Goal:** Predict vehicle trajectories for each mode
- Measured inputs:
  - Driver steering angle (every 30ms)
  - Driver pose -> proxy for attention level
- Cluster measured inputs into the available atomic modes
- For each mode, use a model of vehicle dynamics to predict possible trajectories for 1.2s

Example:
Mode = (right lane, straight, distracted)

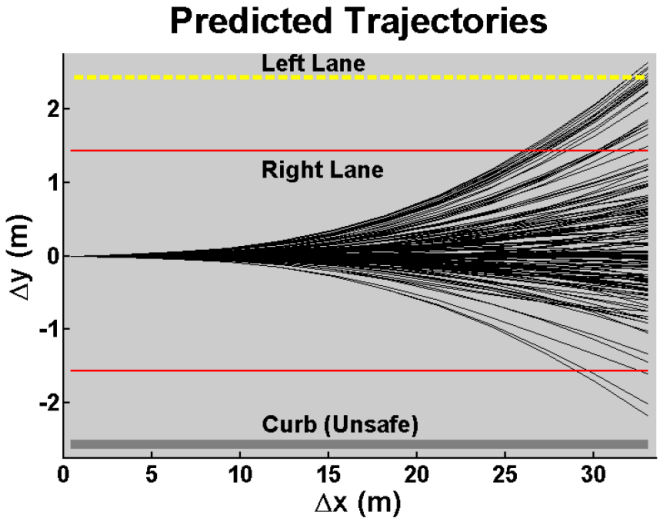

Model Creation

- **Modes are** interpreted as **states** of the Convex-MDP
- Transition **probabilities are** computed based on **empirical frequencies** of trajectory end-points.

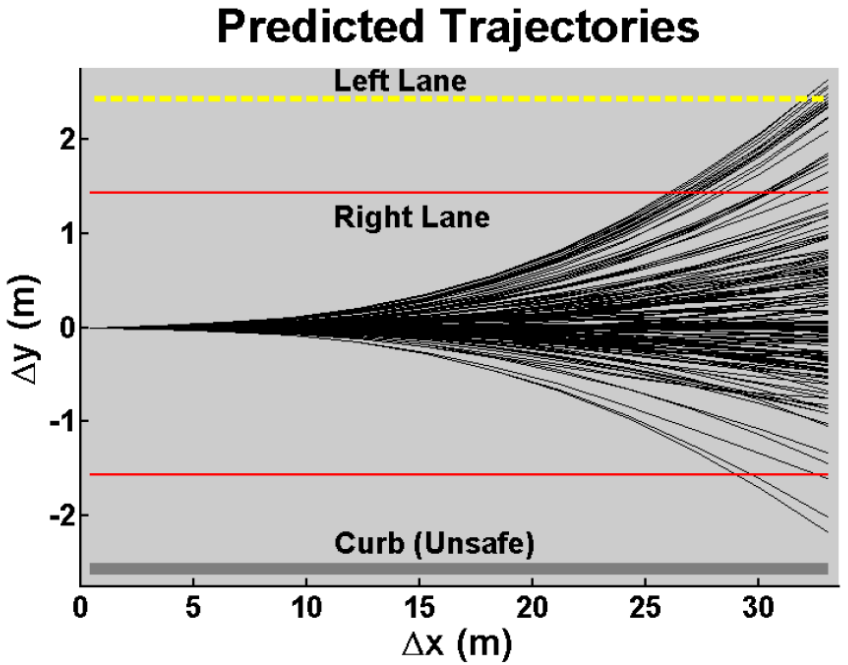

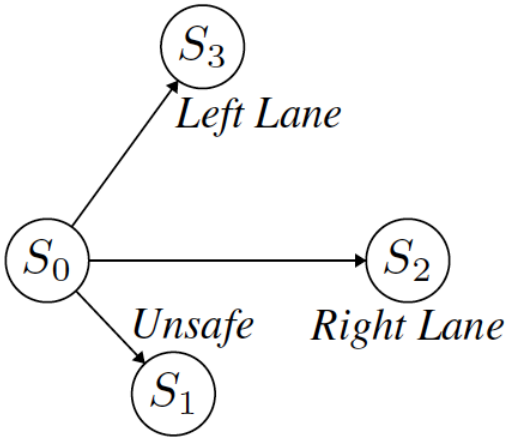

Transition	Transition Probability Interval
$S_0 \rightarrow S_1$	[0.019,0.021]
$S_0 \rightarrow S_2$	[0.890,0.980]
$S_0 \rightarrow S_1$	[0.048,0.053]

Analysis of a Complex Maneuver

- Repeat the process to build the model of a complex maneuver

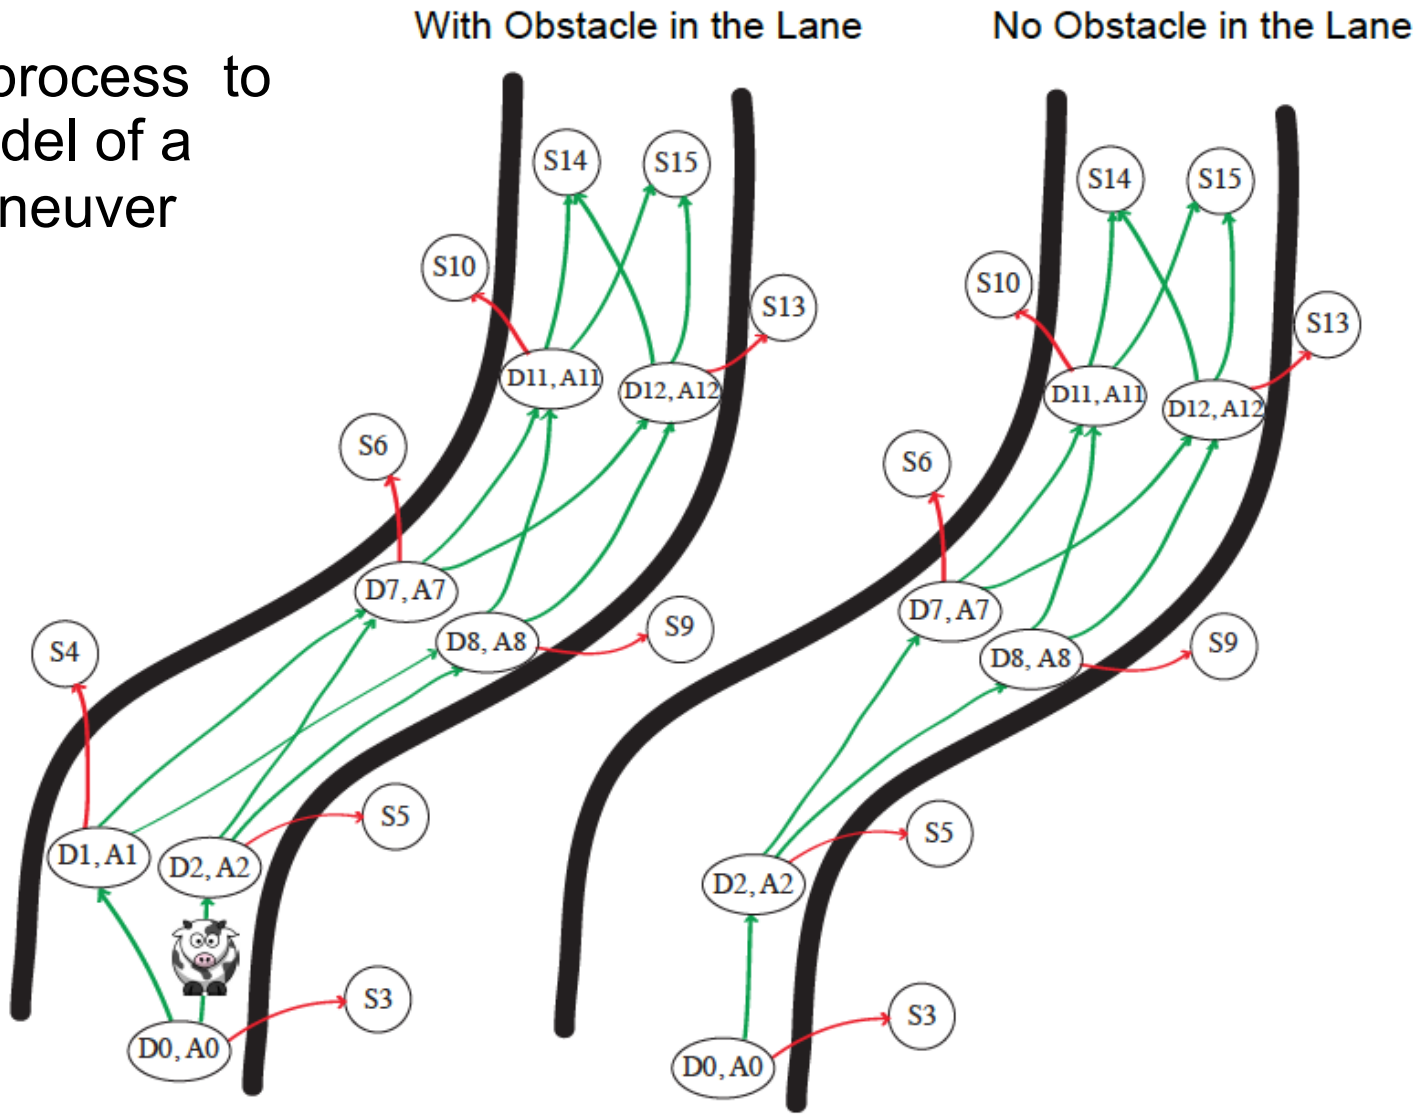

Verified Properties

Table 1: Verified Properties

P1	P_{max}/P_{min} [Attention $\mathcal{U}$ $Unsafe$]
P2	P_{max}/P_{min} [(Attention $\wedge$ $\neg Swerving$) $\mathcal{U}$ $Final$]
P3	P_{max}/P_{min} [(Attention $\wedge$ $Right\ Lane$) $\mathcal{U}$ $Final$]
P4	P_{max}/P_{min} [(Attention $\wedge$ $\neg Braking$) $\mathcal{U}$ $Final$]
Attention is a placeholder for either *Attentive* or *Distracted*	

- Evaluating different driving styles
- Estimating probability of threats

Comparison among Uncertainty Models

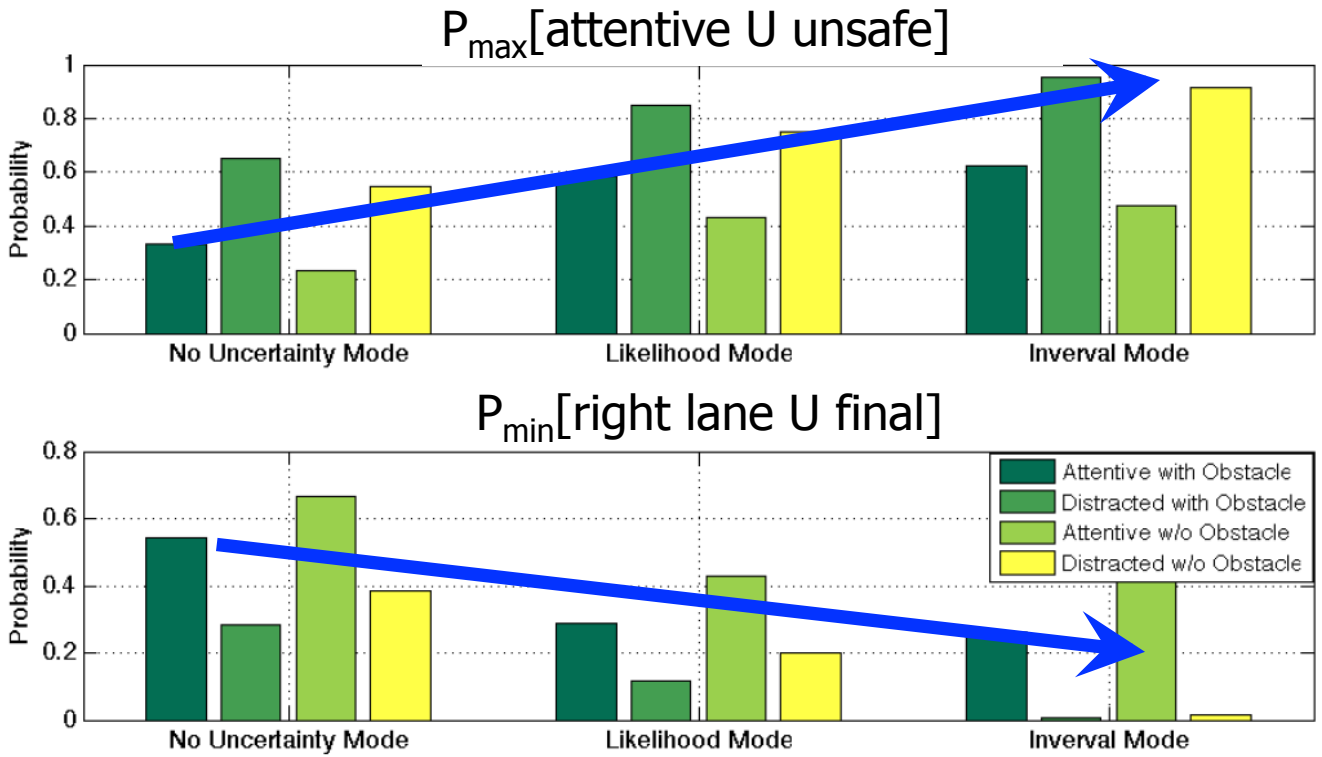

- With no uncertainty, results might be overly optimistic
- Both uncertainty models trained with 95% confidence
  - Interval model might be overly pessimistic
  - Likelihood model is a statistically-valid compromise

Sensitivity to the Uncertainty Level

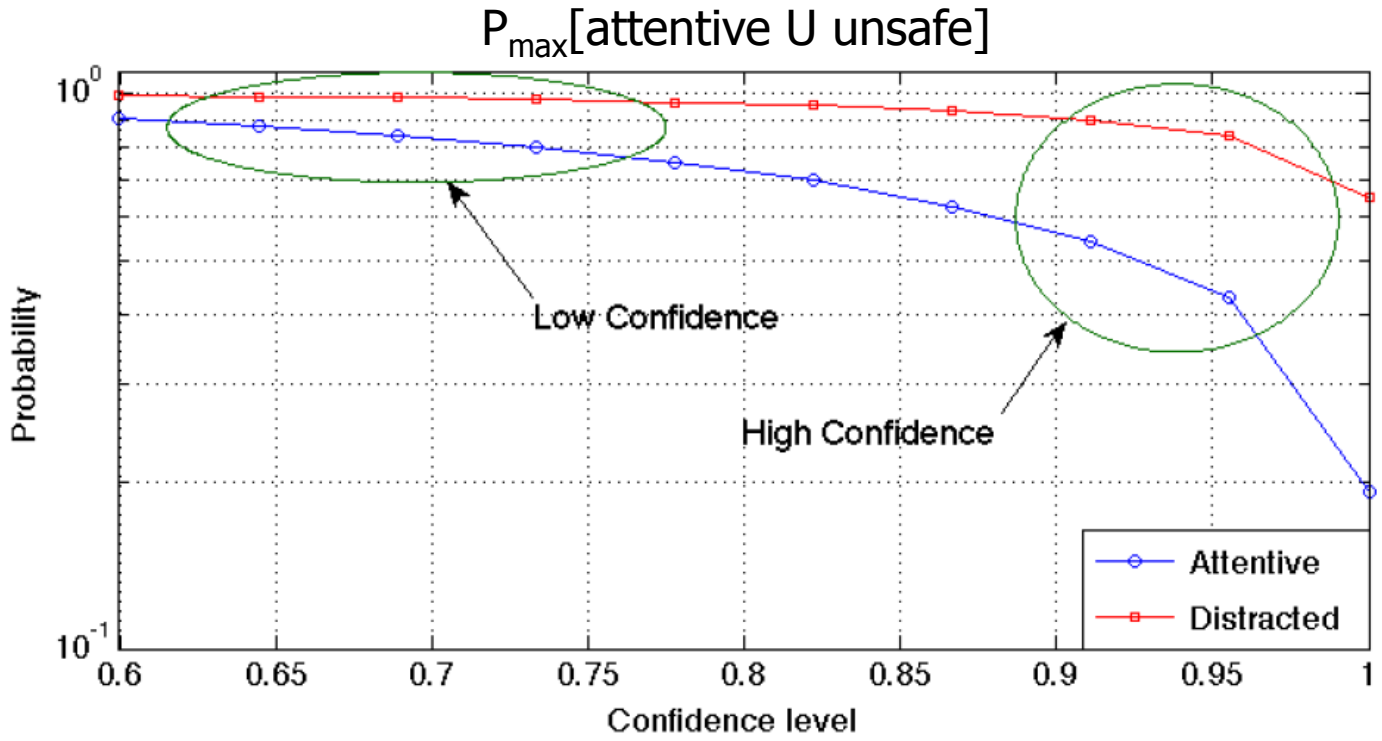

- Attentive driver always perform better (gap varies among individuals!)
- Depending on the specification, a different level of confidence is required -> guide on how to train the model!

Characterization of Individual Driving Styles

Table 1: Verified Properties

P1	$P_{max}/P_{min}\ [\text{Attention}\ \mathcal{U}\ Unsafe]$
P2	$P_{max}/P_{min}\ [(\text{Attention}\ \wedge\ \neg Swerving)\ \mathcal{U}\ Final]$
P3	$P_{max}/P_{min}\ [(\text{Attention}\ \wedge\ Right\ Lane)\ \mathcal{U}\ Final]$
P4	$P_{max}/P_{min}\ [(\text{Attention}\ \wedge\ \neg Braking)\ \mathcal{U}\ Final]$
Attention is a placeholder for either *Attentive* or *Distracted*	

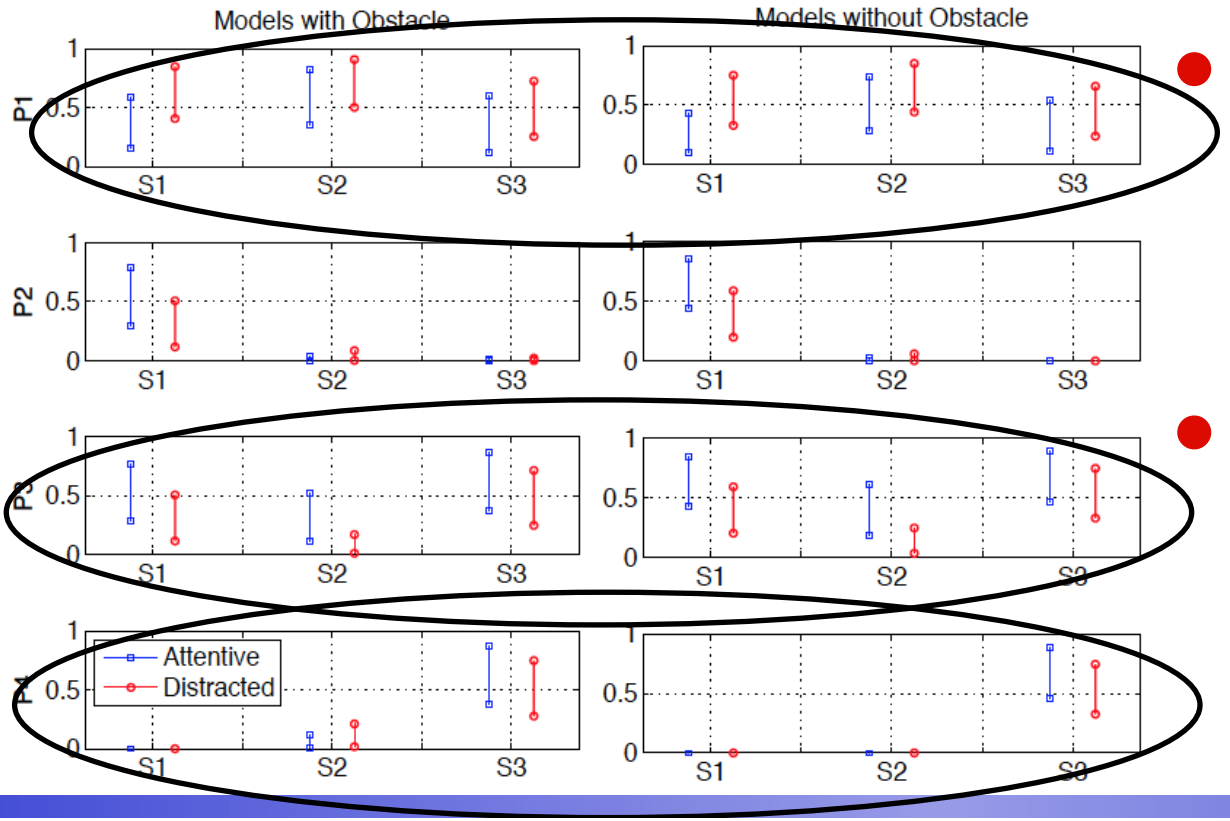

- Compare driving styles
  - S2 worst on keeping the right lane
  - S3 brakes less often
- The presence of an obstacle always increases the probability of threats

Conclusions and Future Work

- Proposed a **polynomial time algorithm** for the verification of PCTL properties of MDPs
- Lowered theoretical complexity for Interval-MDPs from co-NP to P and extended to a large class of **non-linear convex models of uncertainty**
- Applied to the verification of the behavior of a human driver
- Application to further case studies (e.g. pricing of renewable energy)
- Theory extensions:
  - Continuous-Time Markov Chains
  - Compositional methods (assume-guarantee)
  - Stochastic control

Source code available at:
http://www.eecs.berkeley.edu/~puggelli/

Runtime Analysis

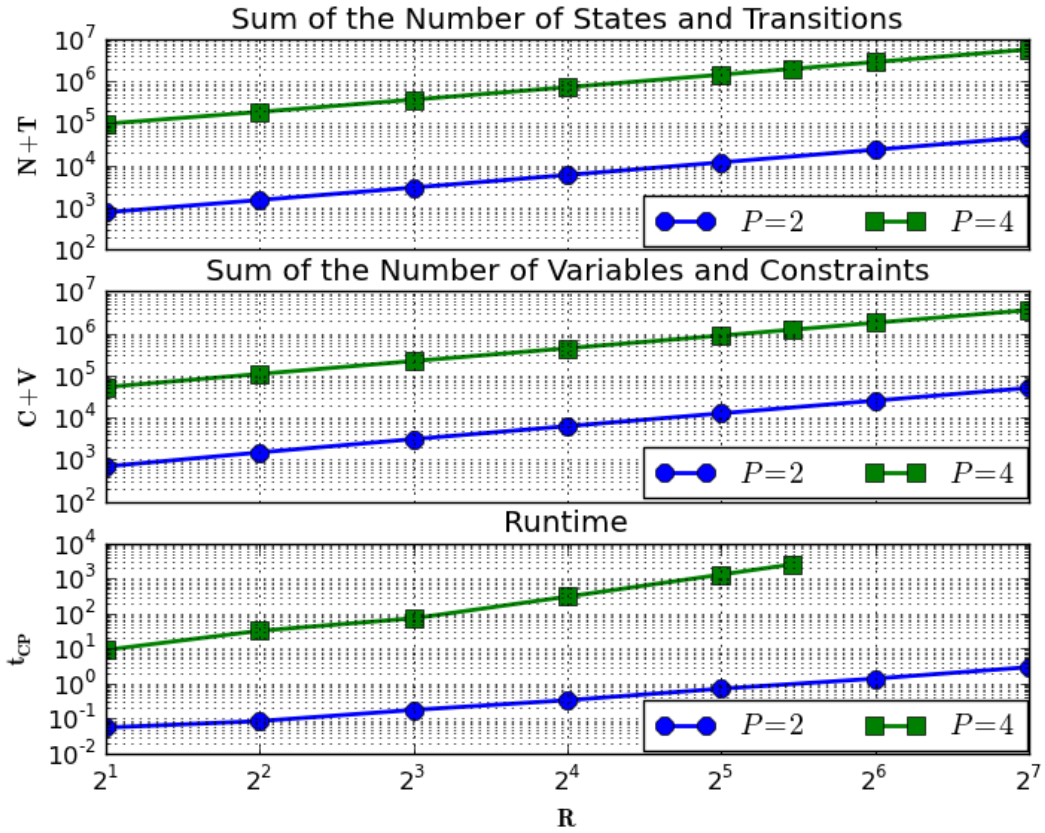

- Use MOSEK as background convex solver
- Size of the convex problem and **runtime scale polynomially**
- Comparable with PRISM[2] and 1000x faster than PARAM[3]

1. www.mosek.com
2. Kwiatkowska et al., "PRISM 4.0: Verification of Probabilistic Real-time Systems"
3. Hahn et al., "Synthesis for PCTL in Parametric Markov Decision Processes"

Unbounded Until in Convex-MDP

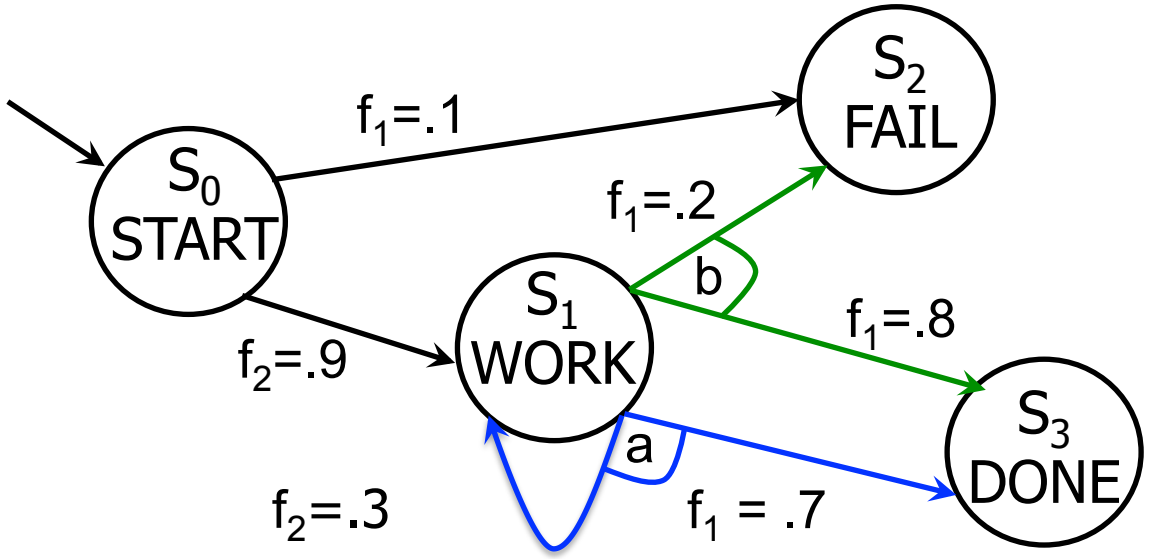

Unbounded Until in Convex-MDP

- Try all Probability Distributions?

$$\max_x \sum x_i$$

s.t.

$$x_2 = 0$$

$$x_3 = 1$$

$$x_0 = 0.9x_1 + 0.1x_2$$

$$x_1 \le 0.2x_1 + 0.7x_3 + 0.1x_2$$

$$x_1 \le 0.2x_1 + 0.8x_3$$

$$x_1 \le 0.25x_1 + 0.75x_3$$

$$x_1 \le 0.21x_1 + 0.79x_3$$

...

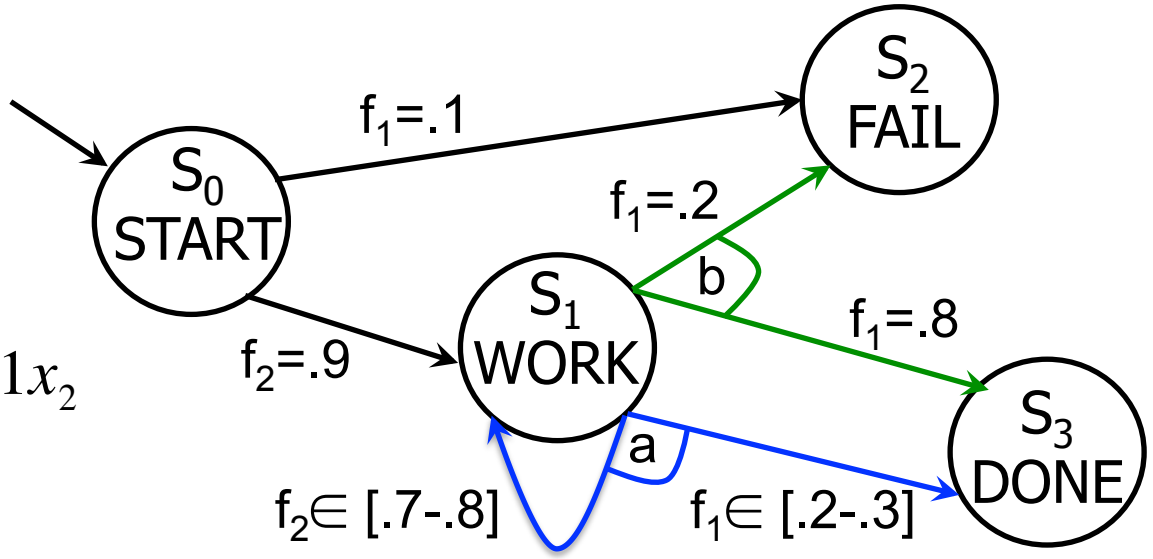

- NO: Uncountably infinite number of distributions

Until Operator in CMDPs: Duality

$$\max_x \sum x_i$$
$$\text{s.t. } x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le 0.2x_1 + 0.7x_3 + 0.1x_2$$
$$x_1 \le 0.2x_1 + 0.8x_3$$
$$x_1 \le 0.25x_1 + 0.75x_3$$
$$x_1 \le 0.21x_1 + 0.79x_3$$
$$\ldots$$

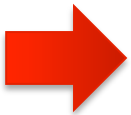

$$\max_x \sum x_i$$
$$\text{s.t. } x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le \min_{p \in U_1} p_{1,1}x_1 + p_{1,2}x_2 + p_{1,3}x_3$$
$$x_1 \le \min_{p \in U_2} p_{2,1}x_1 + p_{2,3}x_3$$

- Worst-case:
  - Minimize the upper bound
- $\min_{p \in U_1} p_{1,1}x_1 + p_{1,2}x_2 + p_{1,3}x_3$
  - Primal problem

Until Operator in CMDPs: Duality-Theory Approach

$$\max_x \sum x_i$$
$$\text{s.t. } x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le \min_{p \in U_1} p_{1,1}x_1 + p_{1,2}x_2 + p_{1,3}x_3$$
$$x_1 \le \min_{p \in U_2} p_{2,1}x_1 + p_{2,3}x_3$$

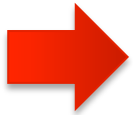

$$\max_x \sum x_i$$
$$\text{s.t. } x_2 = 0$$
$$x_3 = 1$$
$$x_0 = 0.9x_1 + 0.1x_2$$
$$x_1 \le \max_{\lambda_1 \in D_1} g_1(\lambda_1, x)$$
$$x_1 \le \max_{\lambda_2 \in D_2} g_2(\lambda_2, x)$$

- Substitute each primal problem with the corresponding dual problem