

Lablet:

Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems

Geir E. Dullerud, Sayan Mitra (PI), Swarat Chaudhuri

NSA SoS Lablet, Bi-weekly Internal Meeting

February 26, 2015.

Project Goals

Hard problem addressed:

- predictive security metrics;
- scalability and composability.

Title: Static-Dynamic Analysis of Security Metrics for CPS

Goals:

- (a) Identify security **metrics** & **adversary models**;
- (b) develop **theory, algorithms** & **tools** for analyzing the metrics in the context of adversary models.

Impact and Approach

How is this project advancing SoS?

Automated metric-based methods for

- *Security analysis* of cyber-physical systems;
- *Synthesis of safe* cyber-physical systems.

Technical approach

- formal methods;
- control theory;
- programming languages (tools: z3, cvc4).

Important:: Emphasis on automation, scalability, and benchmarks.

Technical Accomplishments to Date

- Huang, Wang, Mitra and Dullerud, "Controller Synthesis for Linear Time-varying Systems with Adversaries", working paper, 2015.
- Huang, Fan, Mereacre, Mitra, Kwiatkowska, "Simulation-based Verification of Implantable Medical Devices with Guaranteed Coverage", submitted for review, 2014.
- Heemels, Dullerud, Teel, "L2-gain Analysis for a Class of Nonlinear Hybrid Systems with Applications to Reset and Event-triggered Control," submitted to IEEE Transactions on Automatic Control, 2014.
- R. Essick, J.-W. Lee, and G.E. Dullerud, "Control of Linear Switched Systems with Receding Horizon Modal Information". IEEE Transactions on Automatic Control, 2014.
- Path-By-Path Output Regulation of Switched Systems With a Receding Horizon of Modal Knowledge. R. Essick, J.-W. Lee, and G.E. Dullerud, In the proceedings of the American Control Conference (ACC), 2014.
- Invariant Verification of Nonlinear Hybrid Automata Networks of Cardiac Cells, Z. Huang, C. Fan, A. Mereacre, S. Mitra, and M. Kwiatkowska. To appear in Computer Aided Verification (CAV), LNCS, 2014.
- Proofs from Simulations and Modular Annotations, Zhenqi Huang and Sayan Mitra, in 17th International Conference on Hybrid Systems: Computation and Control (HSCC 2014), to be held as part of held as part of the seventh Cyber Physical Systems (CPSWeek 2014), Berlin, Germany.
- Entropy- minimizing Mechanism for Differential Privacy of Discrete-time Linear Feedback Systems by Yu Wang, Zhenqi Huang, Sayan Mitra, and Geir Dullerud, to appear in IEEE Conference on Decision and Control (CDC), 2014.
- Proving Abstractions of Dynamical Systems through Numerical Simulations. Mitra, Sayan, In Hot Topics in Science of Security (HOTSoS)
- Stabilization of Markovian Jump Linear Systems with Limited Information, Q. Xu, Zhang, C., and G. E. Dullerud, ASME Journal of SMC, 2014.
- Mishra, A., C. Langbort, and G.E. Dullerud, "Decentralized Control of Linear Switched Nested Systems with L-2 induced Norm Performance," to appear in IEEE Transactions on Control of Network Systems, 2015.
- On Price of Privacy in Distributed Control Systems, Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir Dullerud, in 3rd ACM International Conference on High Confidence Networked Systems (HiCoNS), April 15-17, 2014 in Berlin, Germany as part of Cyber Physical Systems Week 2014 (CPSWeek 2014).
- Differentially Private Iterative Synchronous Consensus, Zhenqi Huang, Sayan Mitra, and Geir Dullerud. In the proceedings of the Workshop on Privacy in the Electronic Society (WPES), collocated with of 19th ACM Conference on Computer and Communications Security (CCS), Raleigh, NC 2012. ACM press.

Secure Control Systems: Big Picture

State and Dynamics Attack:

*direct state modification,
alteration of dynamics*

“CPS Security is Security on Steroids”

Diverse attack mechanisms with complex performance metrics: both discrete and continuous.

Measurement Attack:



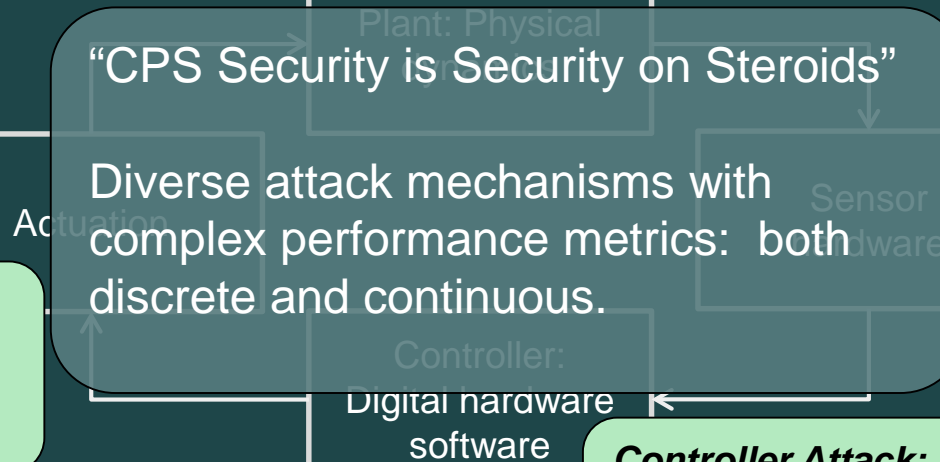
sensors or data links are compromised.

Actuation Attack:


*actuators degraded,
command latency.*

Controller Attack:

*malicious algorithm loaded,
or hardware interrupted.*



General application domain



Features:

- Information topology
- Communication constraints
- Complex hybrid Dynamics
- Sensor resolution
- Shared resources


Networks of interconnected physical and digital systems.

Distributed Robotics: Kiva Systems



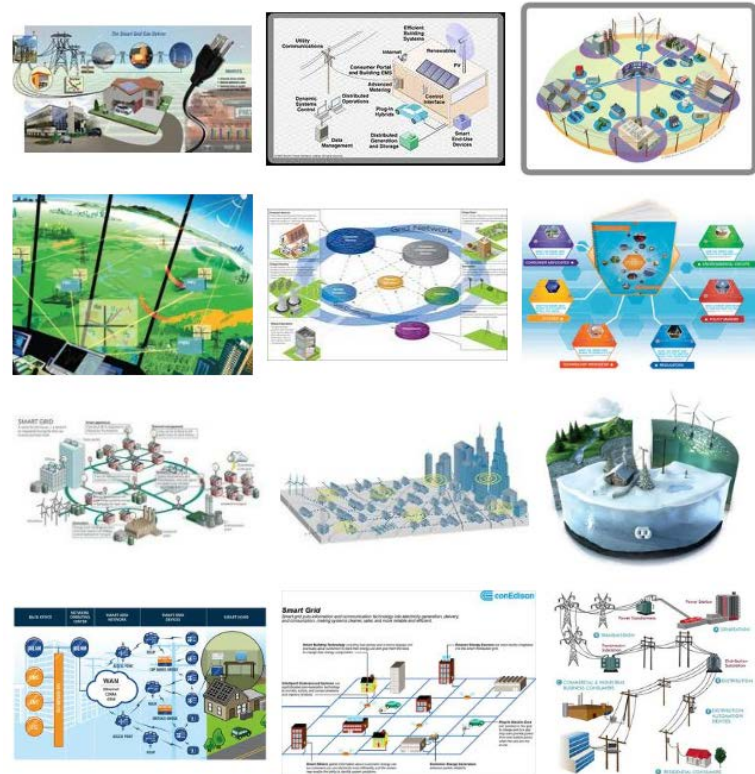
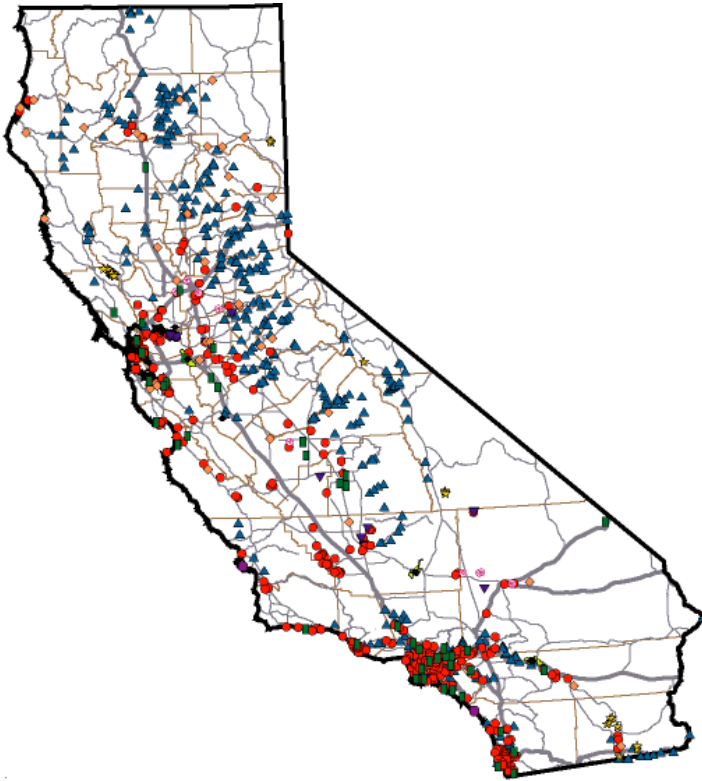
- networked dedicated robots
- light-weight general purpose robots

Drone and UAV Teams




- military applications
- civilian: construction, inspection

Power and Smart Grid



- add sensors, networking, algorithms
- infrastructure, renewables, PHEVs

Building Energy Systems




- interacting subsystems
- sensor network
- storage capability, variable load


Internet

CAIDA's IPv6 AS Core AS-level INTERNET GRAPH

Archipelago August 2010



copyright © 2010 UC Regents. all rights reserved.



- Building high capacity networks
- Big data and learning
- Internet of Things

Financial Markets

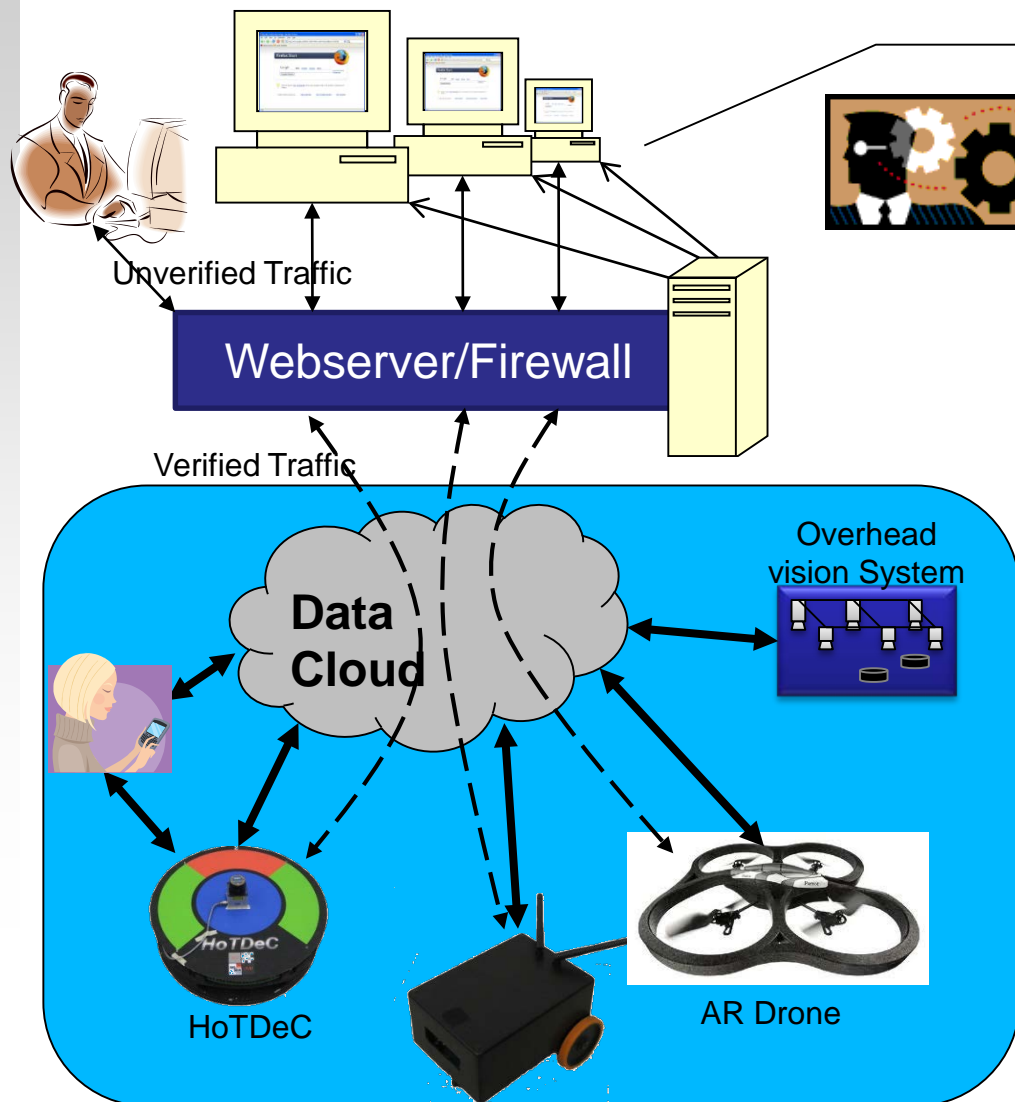


- large number of players
- very high speed dynamics
- complex interactions

Example: Cooperative Robotics



HoTDeC system



Web Scripting Interface

Programming Window

```


Process{
  Goto(x=10,y=50,speed=30);
  Goto(x=3, y=25,speed=10);
  Stop;
}

Process{
  if( LidarDetect(deg[-30,30])
    avoidRight;
}
    
```

Command

```

SetMotorSpeed
...
Stop
Turn(rad=<float
...
Sensors
LidarDetect( r...
EncoderPositi
...
    
```

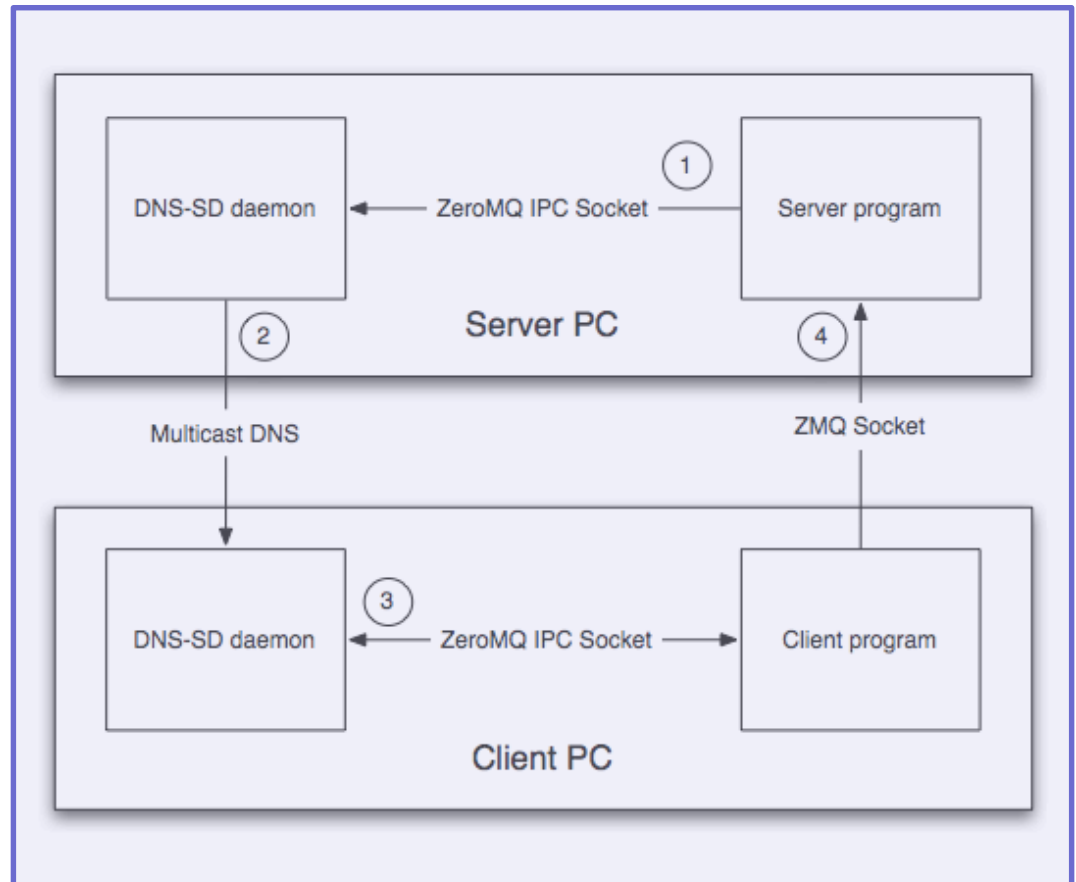


- Execute Program
- Simulate Program
- View Sensors
- Block Program
- Demos


- Agents of varying capability
- Human players in network
- Internet based
- Communication: wireline and wireless
- Multiresolution information
- Re-purposeable: many-testbed-in-one

Service Discovery

1. **DNS-SD server broadcasts availability of service on network**
2. **Client PC's DNS-SD server receives the broadcast**
3. **Client Program connects to server program based on data from the DNS-SD server**
4. **Server program registers availability with DNS Service Discovery server**



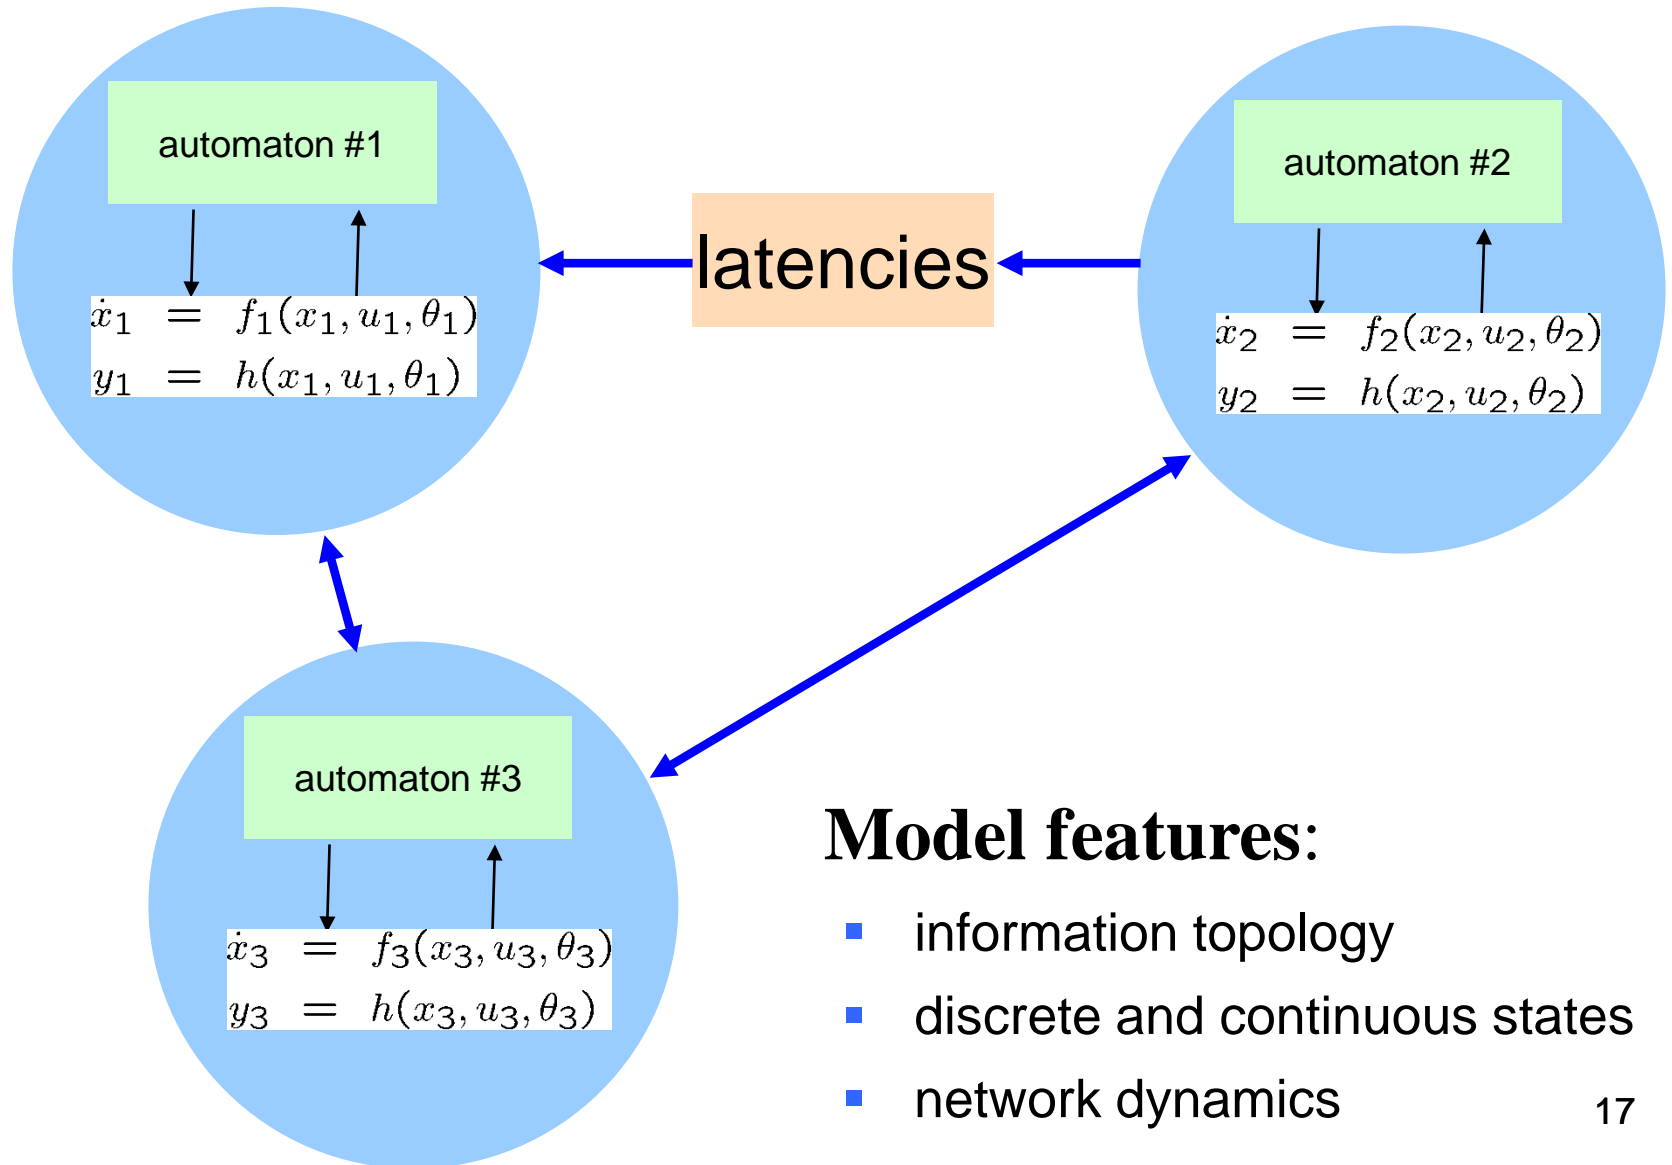
General application domain



Features:

- Information topology
- Communication constraints
- Complex hybrid Dynamics
- Sensor resolution
- Shared resources

Distributed Hybrid Models



Model features:

- information topology
- discrete and continuous states
- network dynamics

Nodal Dynamics

Discrete transition systems (countable states) FSM, PDA, TMs

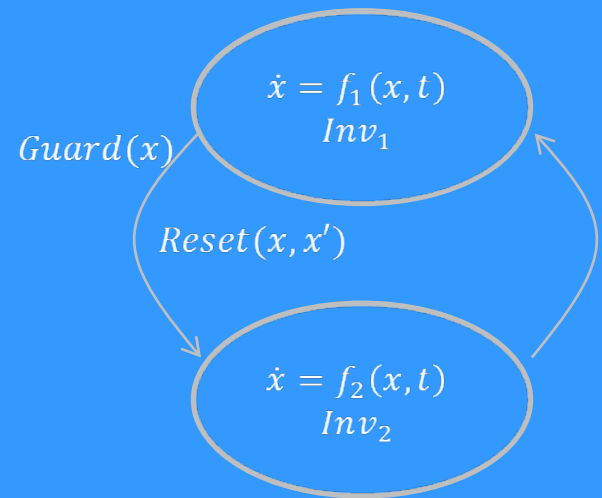
Communicating processes IO automata, process algebras

Dynamical Systems
 $\dot{x} = f(x, t, u)$

Nondeterministic transition systems

Switched Systems
 $\dot{x} = f_{\sigma(t)}(x, t, u)$

Networked Hybrid Automata*





CPS Metrics

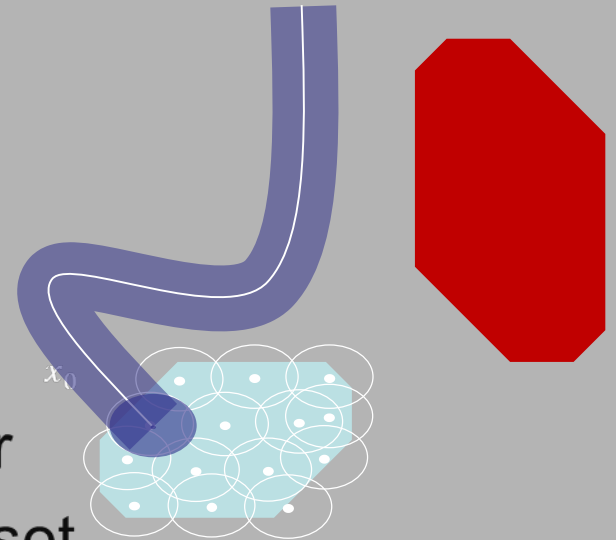
- **Quantized metrics:**
 - min-cut type on discrete-transition system (e.g., compromised sensors);
 - set inclusion or exclusion (e.g., unsafe set)
- **Continuum-valued metrics:**
 - Cost functions based on:
 - state and decisions;
 - noise and disturbances.
 - Degree of detectability and privacy
- **Composite metrics.**

Technical Issues Addressed

- Reachability
- Switching
- Adversarial noise and disturbances
- Model information and disinformation
- Distributed agents
- Safeness of state

Reachability Analysis

- Given start  and target 
- Compute finite cover of initial set
- Simulate from the center x_0 of each cover
- **Bloat** simulation so that bloated tube contains all trajectories from the cover
- Union = over-approximation of reach set
- Check intersection/containment with T
- Refine
- How much to bloat?



Games with Partial Modal Information

$$x_{t+1} = f(x_t; \theta_{1t}, \dots, \theta_{Mt}; u_{1t}, \dots, u_{Mt}; w_t)$$

$$y_{it} = g(x_t; \theta_{it}; v_t)$$

$$x_0 \sim \mathcal{N}(\bar{x}_0, X_0)$$

$$w_t \sim \mathcal{N}(0, W)$$

$$v_{it} \sim \mathcal{N}(0, V_i)$$

Stochastic
Player types

i : Player index

One time step delayed information sharing:

- Types θ_{it} are i -i.i.d. $\{\theta_{i,0:t}, y_{i,0:t}, v_{i,0:t-1}, y_{-i,0:t-1}\}$

- All players have knowledge of its distribution

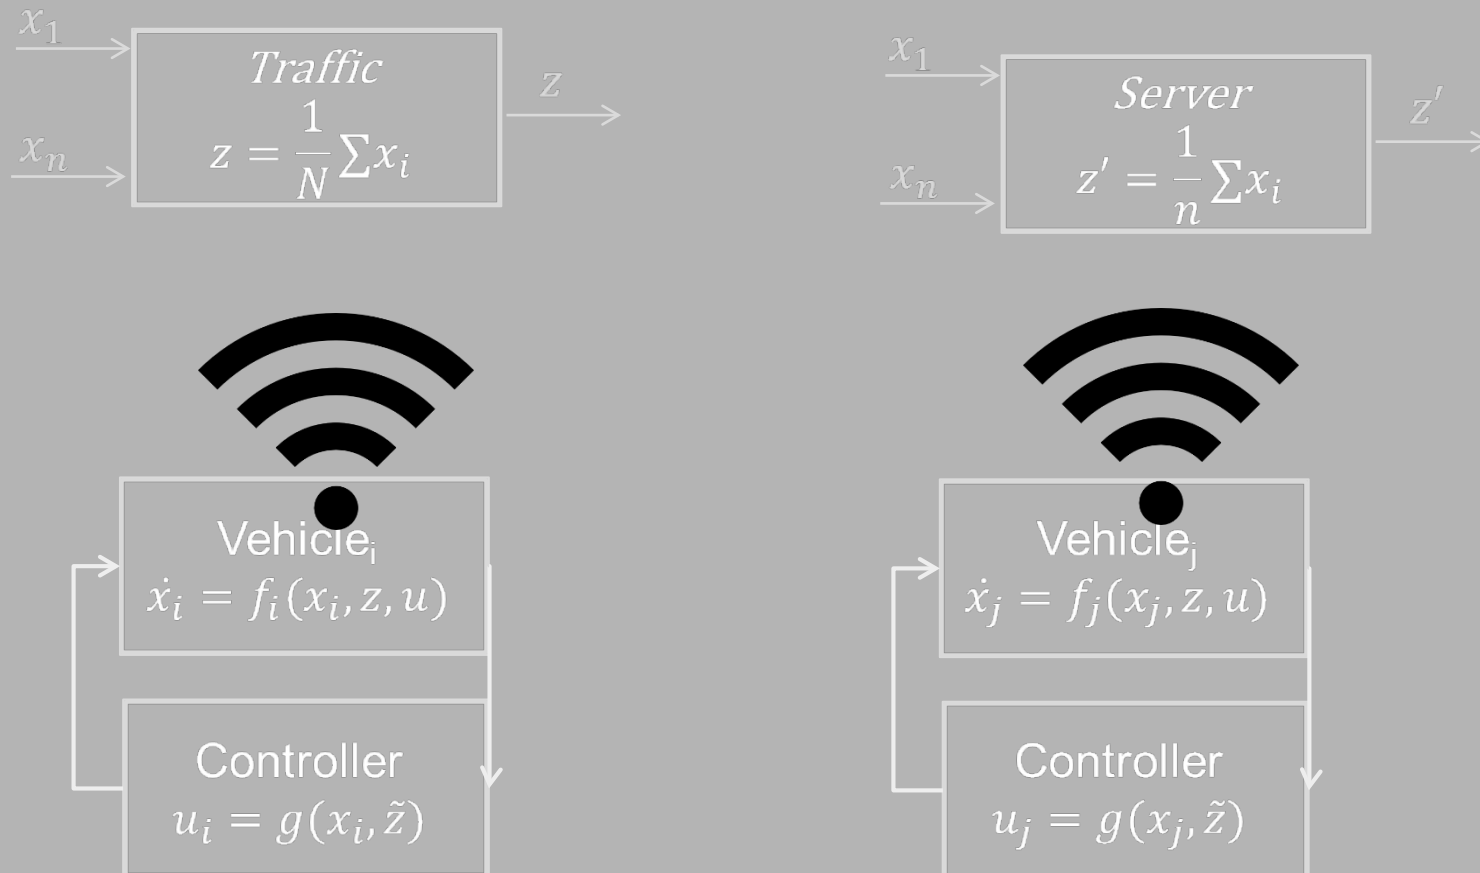
Goal: To find decentralized optimal strategies

$\gamma_{i,t}(I_{i,t}) = u_{i,t}$ that min/max a common cost $E[J(\gamma_1, \dots, \gamma_M)]$

$$J(\gamma_1, \dots, \gamma_M) = x_N^T Q x_N + \sum_{t=0}^{N-1} \{x_t^T Q x_t + u_{1,t}^T R_1 u_{1,t} + \dots + u_{M,t}^T R_M u_{M,t}\}$$

Controlling Agents in Shared Environment

Cost of Privacy in Control: Differential Privacy



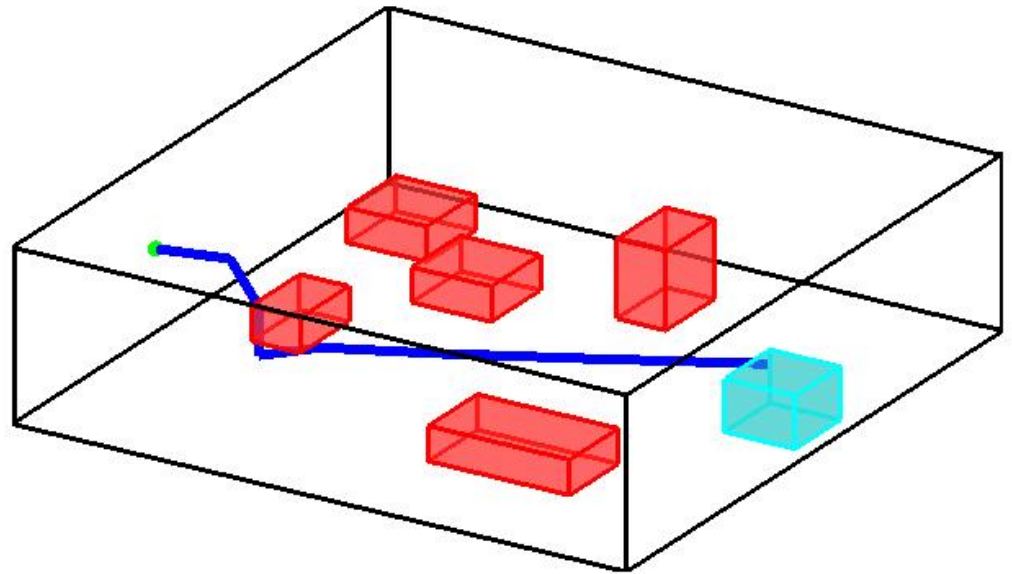
Section I: Adversary-resilient Controller Synthesis

A reach-avoid problem with adversary

- System Dynamics with control u and adversary input a
 - $x_{t+1} = f(x_t, u_t, a_t)$
 - constraints on control and adversarial input: $u \in Ctr, a \in Adv$
- Find a sequence of control u such that for any a :
 - Safe:
$$\forall t. x_t \in Safe$$
 - Winning:
$$x_T \in Goal$$

Case Study

- A linear helicopter model with adversary adding noise to the control



Constraint-based Synthesis

- The problem can be written in terms of logic formulas:
 - Find $u \in Ctr$ such that

$$\forall a \in Adv: \bigwedge_{t \in [T]} x(u, a, t) \in Safe \bigwedge x(u, a, T) \in Goal$$

- The current SMT solvers (e.g. z3, cvc4) can solve a linear, quantifier-free version of such formula


Approach: dynamic decoupling

- We represent the reach set with adversary as a direct sum of the adversary-free behavior and adversary leverage:
 - $Reach(u, Adv, t) = Reach(u, 0, t) \oplus A$
 - A is computed by statically analysis the dynamics of the system
- Then, strengthen the constraints:
 - $Safe' = Safe \ominus A, Goal' = Goal \ominus A$
 - Computed by solving optimization problem
- Then ship the quantifier-free formula to SMT solver
 - Find $u \in Ctr$ such that

$$\bigwedge_{t \in [T]} x(u, 0, t) \in Safe' \quad \bigwedge x(u, 0, T) \in Goal'$$

Preliminary results

Complete System	# x, u, a	T	$ \phi_{Safe} , \#Obs$	$ \phi_{Goal} , \phi_{Ctr} $	$ \phi $	Result	Run Time (s)
Vehicle	4,2,2	40	16, 3	4, 160	804	unsat	2.79
		80	20, 4	4, 320	1924	sat	16.49
		80	44, 10	4, 320	3844	sat	35.22
		80	84, 20	4, 320	7044	sat	53.8
		160	20, 5	4, 640	3844	sat	91.78
		320	24, 6	4, 1280	8964	sat	532.5
Helicopter	16,4,4	5	18, 3	6,40	136	sat	1.2
		5	24, 4	6,40	166	unsat	0.61
		7	24, 4	9, 56	213	sat	8.2
		9	36, 6	6, 72	402	sat	24.5
		12	24,4	6, 96,	338	sat	60.6
		15	24, 4	6, 96,	576	sat	158.8
		18	24, 4	10, 96,	640	-	-



More and Ongoing

- Synthesis of attacks
- Synthesis of State-dependent Control
- Synthesis for Nonlinear system
- Counter-example Guided Refinement

Section II:

On Differential Privacy of Distributed Control System



General Question

- For distributed control systems, how expensive is it to preserve privacy?
- Navigation
 - Routing delays vs location privacy
- Smart Grid
 - Peak demand vs schedule privacy


Differential Privacy (DP)

- Introduced in [1]: a private mechanism should not provide substantially different outputs if one users data changes
- In [2]: minimization of estimation error for open-loop dynamical systems with differential privacy
- [3] discuss cost of privacy for consensus algorithms

[1] C. Dwork et al. TCC2006

[2] JL. Ny and GJ Pappas. TAC2014

[3] Z. Huang et al. WPES2012




Differential Privacy (DP)

- **Def. DP:** M is a mechanism that gives ϵ -differential privacy with $\epsilon > 0$, if for all datasets D and D' that differ in one user's data, for all subset of observations $S \subseteq \text{Range}(M)$,

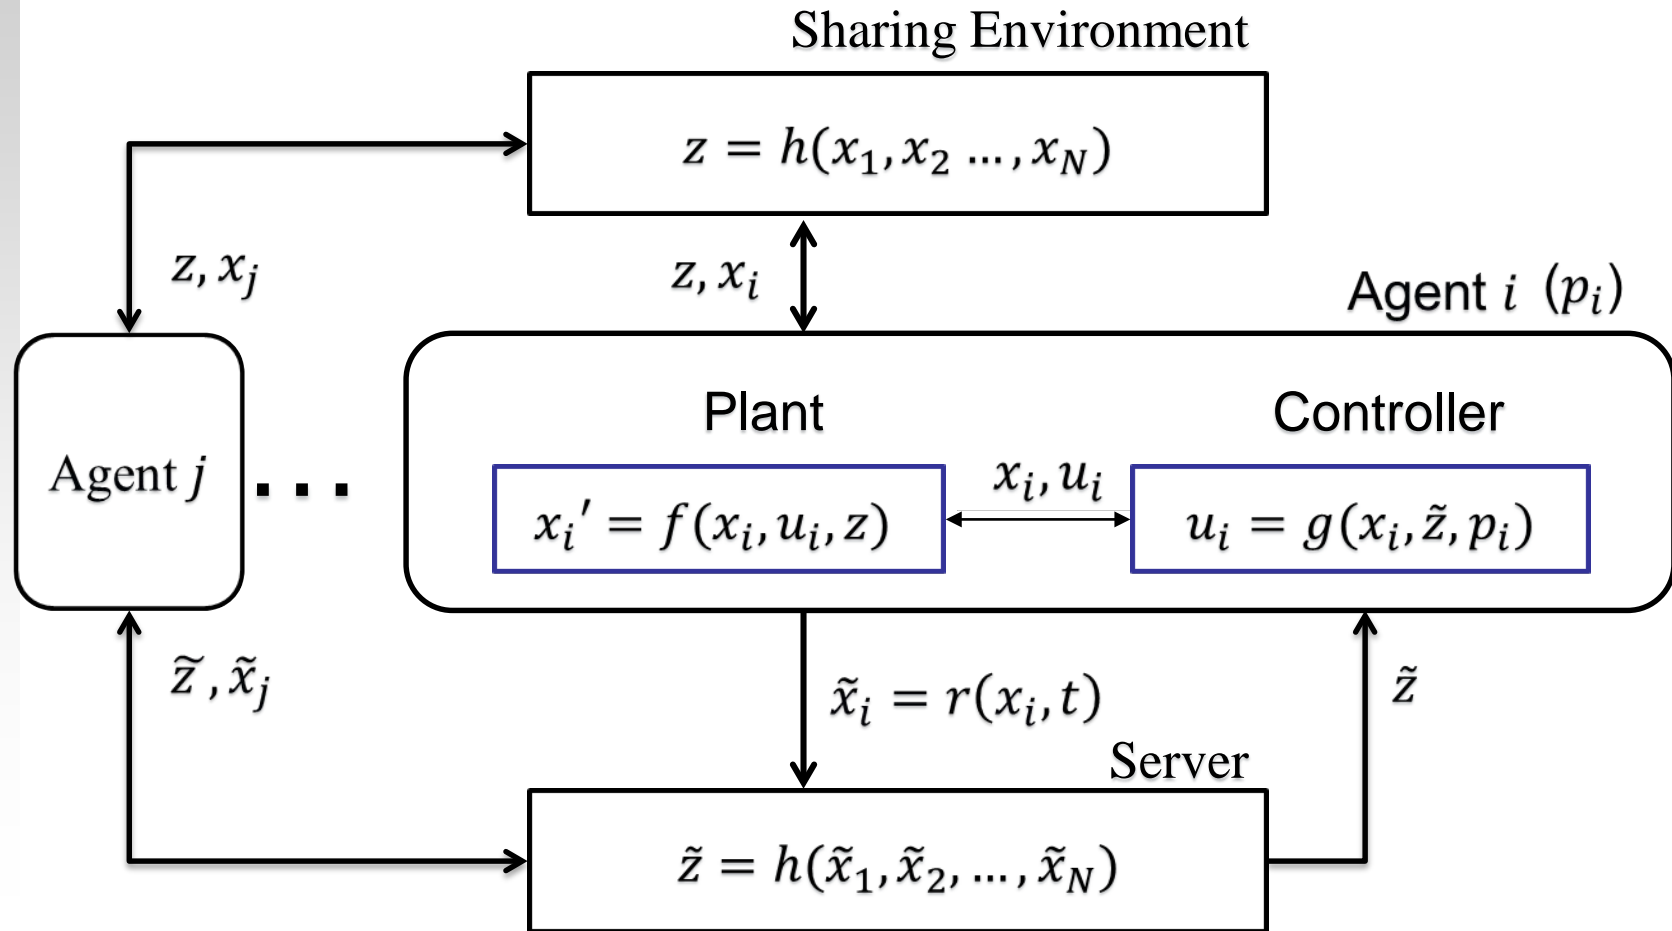
$$\Pr[M(D) \in S] \leq e^\epsilon \Pr[M(D') \in S]$$

A Simple Example of DP Algorithm

- **Weight watchers example**
 - Multiple people diet together. Each people participated has a weight value in the range $[a, b]$. They want to compute the average weight without reporting their exact weight.
 - Each participant will add a carefully chosen noise to his own weight and report it to the server.
 - The server can then publish the average without bleaching the individuals' privacy
 - Laplace noise $\sim \text{Lap}(\epsilon(b - a))$



Distributed Control System



Observables: \tilde{x}_i, \tilde{z} . **Valuable information:**

Example: Navigation

- Routing of N agents on a 2-D plan:
 - The agent i 's state $x_i \in \mathbb{R}^2$
 - Preference of agent i is a path with length T : $p_i \in \mathbb{R}^{2T}$
 - The environment state, center of mass: $z = \frac{1}{N} \sum_i x_i$
 - The update law of the individual agent's state at time t :
$$x_i(t+1) = x_i(t) + c(z(t) - x_i(t)) + u_i(t)$$
 - The control law:
$$u_i(t) = -cz_i(t) + 0.8(p_i(t) - x_i(t))$$
- To design: the report strategy:

DP for Distributed Control System

- The sensitive data of the system $p = \{p_1, p_2, \dots, p_N\}$

- p_i is the desired trajectory of agent i :

$$p_i = [p_i(0), p_i(1), \dots, p_i(T)]$$

Unbounded change in p_i results in unbounded change in system's behavior

- **Def. DP:** Let Obs be any observation stream of bounded time T and p, p' be different in agent i 's preferences. The DCS is ϵ -DP if:

$$\Pr[p \text{ leads to } Obs] \leq e^{-\epsilon} \Pr[p' \text{ leads to } Obs]$$

$$\uparrow \\ e^{\epsilon |p_i - p'_i|}$$

Cost of Privacy for Distributed Control System

- **Average Cost:** $\frac{1}{N} \sum_{t=0}^T \sum_i |x_i(t) - p_i(t)|^2$
 - Fixed a DCS, depends only on the preferences p

- **Baseline Mechanism M' :** $\tilde{x}_i(t) = x_i(t)$

- **The Cost of Privacy of a DP mechanism M is:**

$$CoP = \sup_p \mathbf{E}[\mathbf{Cost}_{M,p} - \mathbf{Cost}_{M',p}]$$

Sensitivity

- The sources of uncertainty of the system
 - The preferences of agents
 - The randomized report function
- Fixed a sequence of observation (Obs) and the agents' preferences (p), the trajectories of all the agents are fully specified: $x(Obs, p, t)$
- **Def. Sensitivity:** difference in the system's states resulting from change in individual's p_i

$$\Delta(t) = \sup_{Obs, adj(p, p')} \frac{|x(Obs, p, t) - x(Obs, p', t)|}{|p_i - p_i'|}$$

A DP Algorithm

- Theorem: The following distributed control system is ϵ -differentially private:

- At each time t each agent adds an vector of independent Laplace noise $Lap(\frac{\Delta(t)T}{\epsilon})$ to its actual state:

$$\tilde{x}_i(t) = x_i(t) + Lap\left(\frac{\Delta(t)T}{\epsilon}\right)$$

- Sensitivity and Cost of Privacy are properties of the dynamics of the system.

Linear Distributed Control Systems

- Linear distributed control system:

$$z_i(t) = \frac{1}{N} \sum_i x_i$$

$$x_i(t+1) = Ax_i(t) + cz(t) + u_i(t)$$

$$u_i(t) = -c\tilde{z}(t) + K(x_i(t) - p_i(t))$$

- We will design an ϵ -differentially private mechanism for this system and reason about cost of privacy.

Sensitivity of Linear Distributed Control Systems

- **Sensitivity:**

$$\Delta(t) = |(cI + K)^t| + |\sum_{s=0}^t (cI + K)^s (I - K)|$$

- Independent to the number of agents.
- Converges to a constant if the closed-loop dynamics is stable.
- Diverges exponentially otherwise.


- DP Mechanism:

$$\tilde{x}_i(t) = x_i(t) + Lap\left(\frac{\Delta(t)T}{\epsilon}\right)$$


CoP of Linear Distributed Control system

- **Cost of Privacy:** $O\left(\frac{T^3}{N^2\epsilon^2}\right)$
- The strategy works for system with many short-lived agents
- The cost of privacy is low for systems with large number of agents
- Improvement: protecting several waypoint instead of the whole desired trajectory $CoP \sim O\left(\frac{k^3}{N^2\epsilon^2}\right)$

Cost of Privacy



Cost v.s. (Decreasing) Privacy



Cost v.s. (Increasing) Number of agents

Conclusion

- A framework for studying the cost of differential privacy for distributed control systems.
- A communication strategy to guarantee differential privacy.
- A linear system with quadratic cost is specified
 - Cost of privacy has the order $O\left(\frac{T^3}{N^2\epsilon^2}\right)$ for stable dynamics, and grows exponentially otherwise.

Lablet:

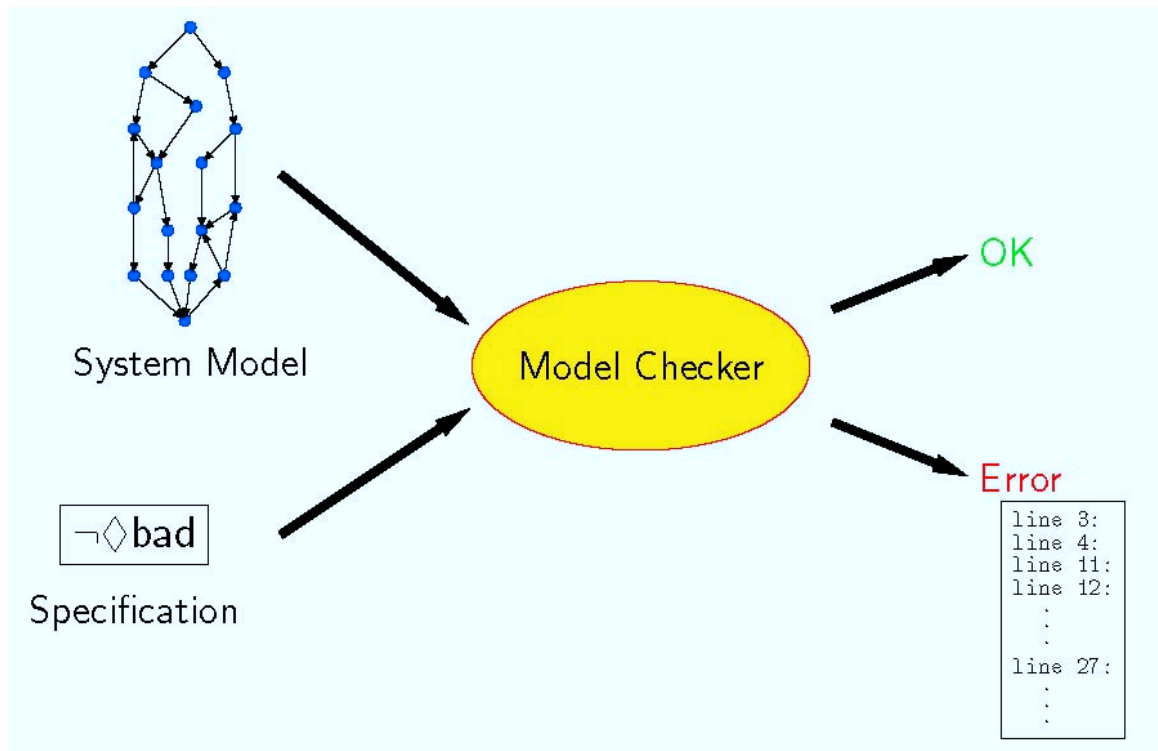
Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems

Geir E. Dullerud, Sayan Mitra (PI), Swarat Chaudhuri

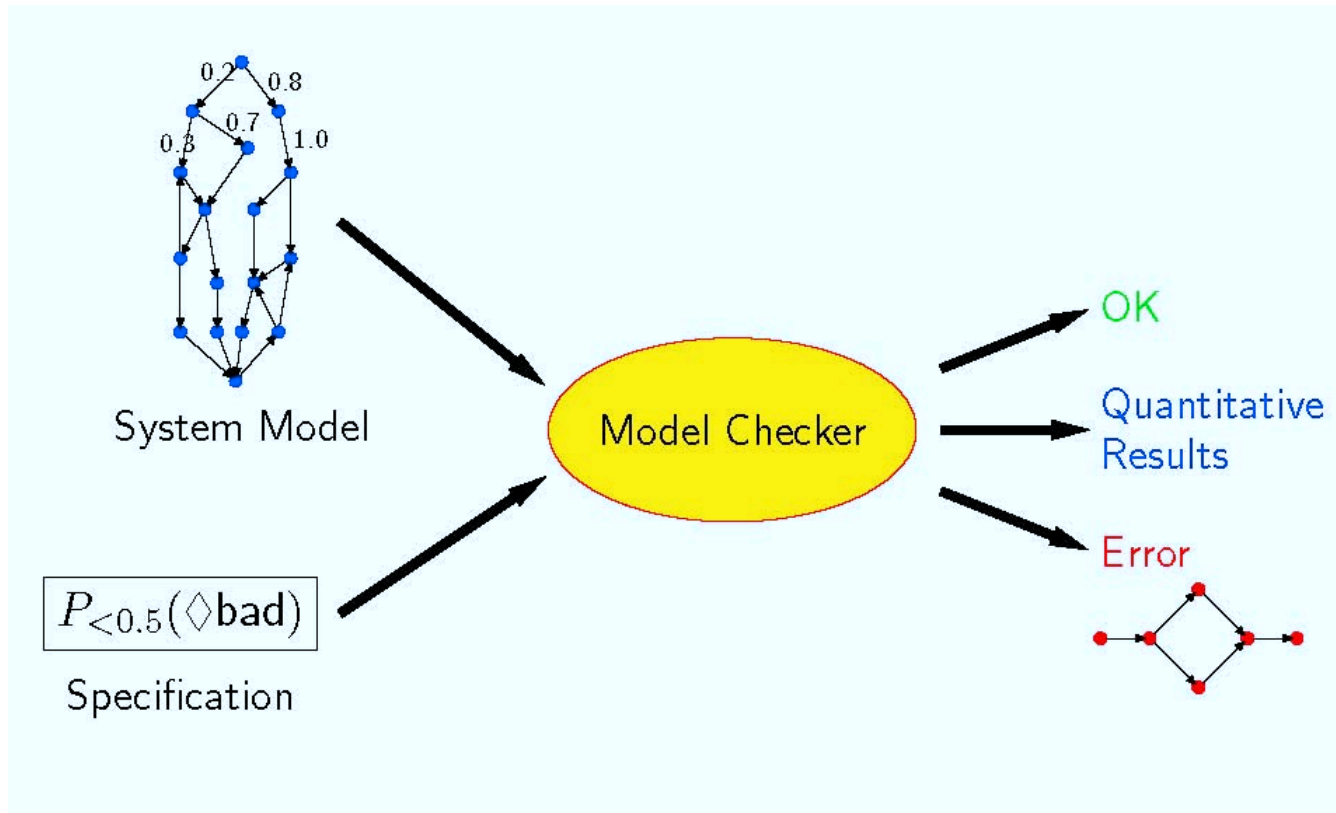
NSA SoS Lablet, Bi-weekly Internal Meeting

February 26, 2015.

Model Checking



Probabilistic Model Checking



What is Probabilistic Model Checking?

Overview



Probabilistic model checking is a *automatic, formal* verification technique for analysing systems that exhibit *stochastic* behavior.

- Systems modeled by (finite state) Markov Chains, Markov Decision Process, Continuous Time Markov Chains
- Properties reason over the measure space of *executions*. Allow one to quantify reliability, performance, security. Examples include “probability of shutdown is less than 0.02”, “the expected energy consumption is 15mW”


Algorithmic Approaches

- *Exact Methods*: Iterative algorithms that rely on techniques such as linear programming, and numerical integration.
- *Statistical*: Simulate the system, and statistically estimate the correctness of the system based on the sample executions drawn, using hypothesis testing

Probability Propagation in Physics-based Models




- IC distribution
- Uncertainty in physics



- Probability bounds
- Output distributions

Can frequently be
Converted to MC



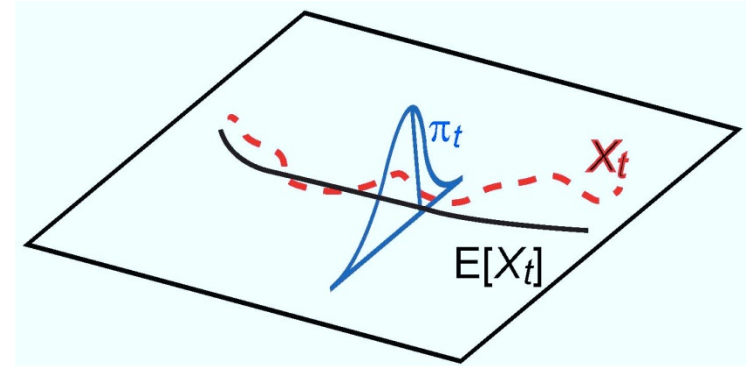

- Distributions
- Properties
- Design tradeoffs
- Provable bounds

Overview

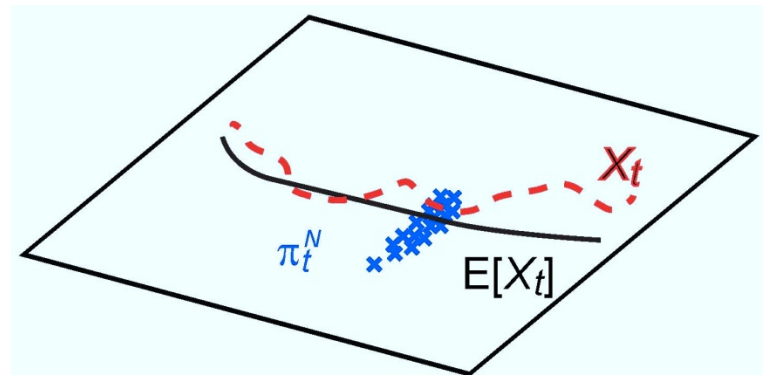

- Stochastic simulation and Poisson Variance Reduction.
- Finite channel Markov processes.
- Particle Filter.

Objective - Stochastic Simulation

Evolution of random processes through system dynamics

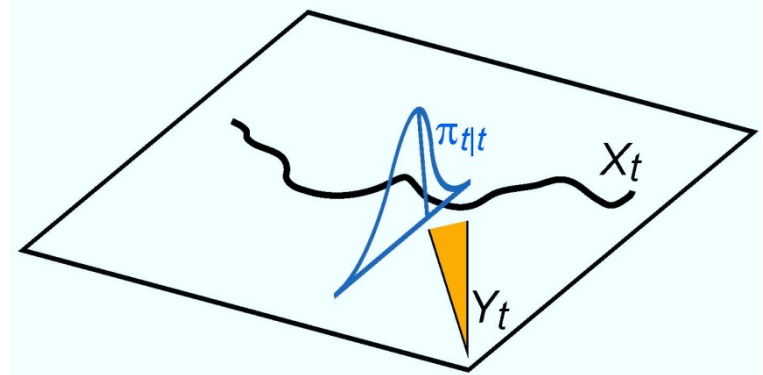
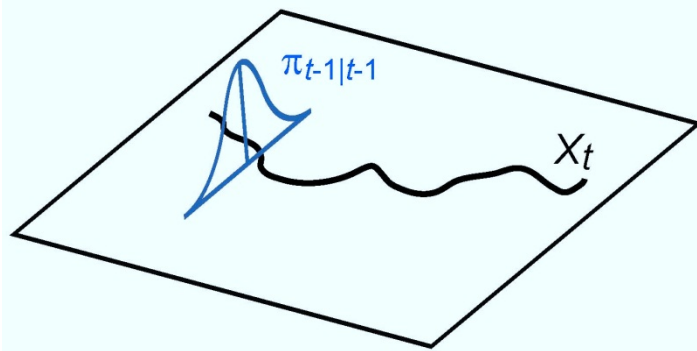


Numerical approximation via simulation


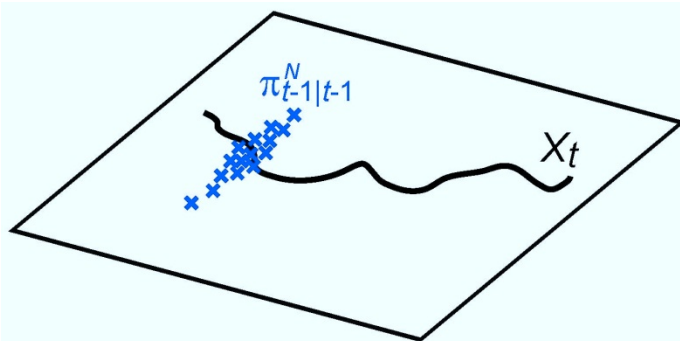


Objective - Filtering

Evolution of conditional probabilities through system dynamics



Numerical approximation via particle filtering



Overview

- Stochastic simulation and Poisson Variance Reduction.
- Finite channel Markov processes.
- Particle Filter.

Motivation

Consider the nonlinear evolution of a continuous-time stochastic process, given by:

$$X(t) = X(0) + \sum_{i=1}^I Y^i \left(\int_0^t \rho^i(s, X(s)) ds \right) \zeta^i,$$

where Y^i is a random process. Numerical integration \implies corresponding discrete-time stochastic process:

$$\tilde{X}_{\ell+1} = \tilde{X}_{\ell} + \sum_{i=1}^I S_{\ell}^i \left(\rho^i(t_{\ell}, \tilde{X}_{\ell}) \tau \right) \zeta^i,$$

where S_{ℓ}^i is also random.

- Mean pathwise behavior? Analytical solution is impossible.
- Estimation by simulation.

Motivation

- Stochastic simulation: Applicable to problems from physical modeling to control and estimation.
- Classic problem is estimation of the mean behavior based on random samples. Convergence of n averaged estimates is sure but slow $\mathcal{O}(\frac{1}{\sqrt{n}})$.
- How to reduce costs? Variance reduction: 2 orders of magnitude reduction in error of 1000 particle simulation.
- Property of method:
 - samples are fair draws;
 - ensemble members correlated.
- Tau-leaping: a fast, cheap algorithm to discretely approximate Markov processes.
- We apply variance reduction to the Poisson samples used to “tau-leap”.