SYSTEM SCIENCE OF SECURITY AND RESILIENCE FOR CYBER-PHYSICAL SYSTEMS (SURE)

XENOFON KOUTSOUKOS

VANDERBILT UNIVERSITY

VANDERBILT UNIVERSITY · Berkeley UNIVERSITY OF CALIFORNIA · Massachusetts Institute of Technology · INFORMATION & COMPUTER SCIENCES UNIVERSITY of HAWAI'I at MĀNOA

SURE

SYSTEM SCIENCE OF SECURITY AND RESILIENCE OF CPS

SURE

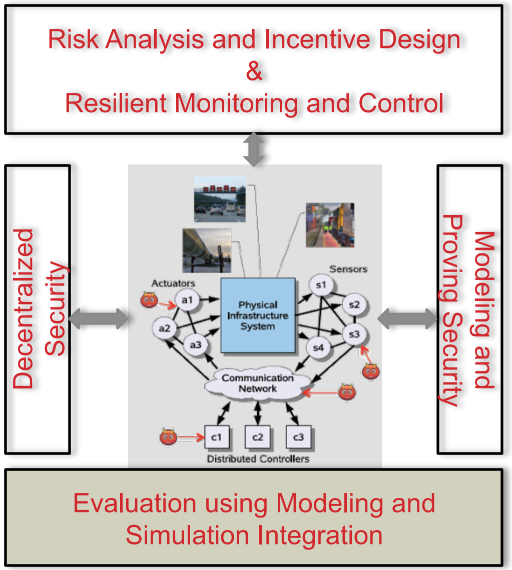

Key Ideas

1. Hierarchical Control and Coordination
   1. Risk analysis and incentive design that aim at developing regulations and strategies at the management level
   2. Resilient monitoring and control of the networked control system infrastructure
2. Science of decentralized security which aims to develop a framework that will enable reasoning about the security of all the integrated constituent CPS components
3. Reliable and practical reasoning about secure computation and communication in networks which aims to contribute a formal framework for reasoning about security in CPS
4. Evaluation and experimentation using modeling and simulation integration of cyber and physical platforms that directly interface with human decision makers.
5. Education and outreach

Impact

- Equip CPS designers and operators with foundations and theory-based comprehensive tools improve resilience against faults and intrusions
- Enable designers to take security decisions and allocate resources in a decentralized manner
- Enable experimentation, evaluation, and training using a modeling and simulation integration platform

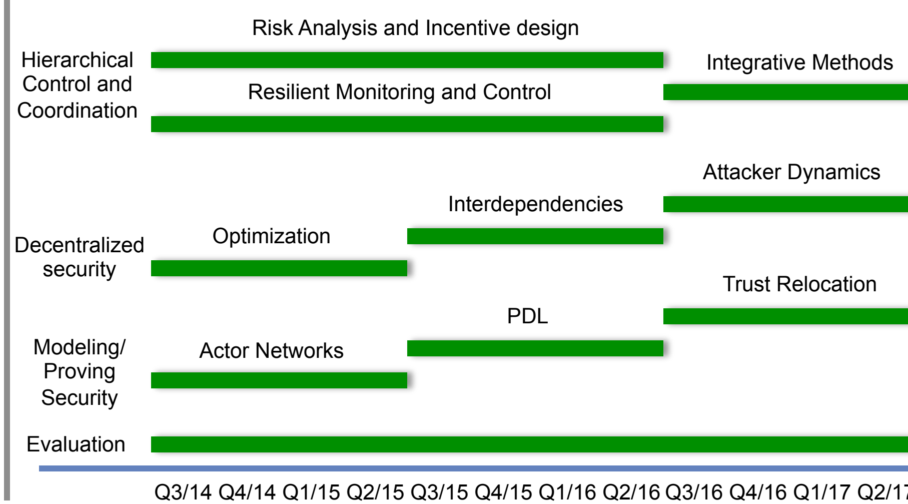

Berkeley UNIVERSITY OF CALIFORNIA — Massachusetts Institute of Technology — INFORMATION & COMPUTER SCIENCES UNIVERSITY of HAWAI'I at MĀNOA — VANDERBILT UNIVERSITY

OVERVIEW

- **Team**
- **Resilience of Cyber-Physical Systems**
- **Research Problems**
- **Project Thrusts**
  - Risk Analysis and Incentive Design
  - Resilient Monitoring and Control
  - Decentralized Security
  - Formal Reasoning about Security
  - Evaluation using Modeling and Simulation Integration
- **Resilient Monitoring**
- **Evaluation using Modeling and Simulation Integration**

TEAM

- Saurabh Amin (MIT)
- Katie Dey (Vanderbilt) – Outreach
- Anthony Joseph (UC Berkeley)
- Gabor Karsai (Vanderbilt)
- Xenofon Koutsoukos (Vanderbilt) – PI
- Dusko Pavlovic (U. of Hawaii)
- Larry Rohrbough (UC Berkeley)
- S. Shankar Sastry (UC Berkeley
- Janos Sztipanovits (Vanderbilt)
- Claire Tomlin (Vanderbilt)
- Peter Volgyesi (Vanderbilt) - Technology Integration and Evaluation
- Yevgeniy Vorobeychik (Vanderbilt)

- Team with interdisciplinary activities in multiple areas:
  - CPS, critical infrastructure, embedded software, mobile/distributed computing
  - Security and resilience, incentive design, game theory fault diagnosis, control theory, model-integrated computing, multi-agent systems, secure machine learning
- Successful collaborative projects
  - NSF Foundations of Hybrid and Embedded Systems ITR (2003- 2010)
  - Command and Control Wind Tunnel PRET (2006 - 2009)
  - High-Confidence Design of Networked Embedded Control Systems MURI (2006 – 2011)
  - NSF STC TRUST (2005 – 2014)
  - NSF CPS Frontier FORCES (2013 – 2018)

RESILIENCE OF CPS

Attributes of Resilience

- Functional correctness (by design)
- Robustness to *reliability* failures (faults)
- Survivability against *security* failures (attacks)

Challenges to Resilience

- Spatio-temporal dynamics
- Many strategic interactions with network interdependencies
- Inherent uncertainties
- Tightly coupled control and economic incentives

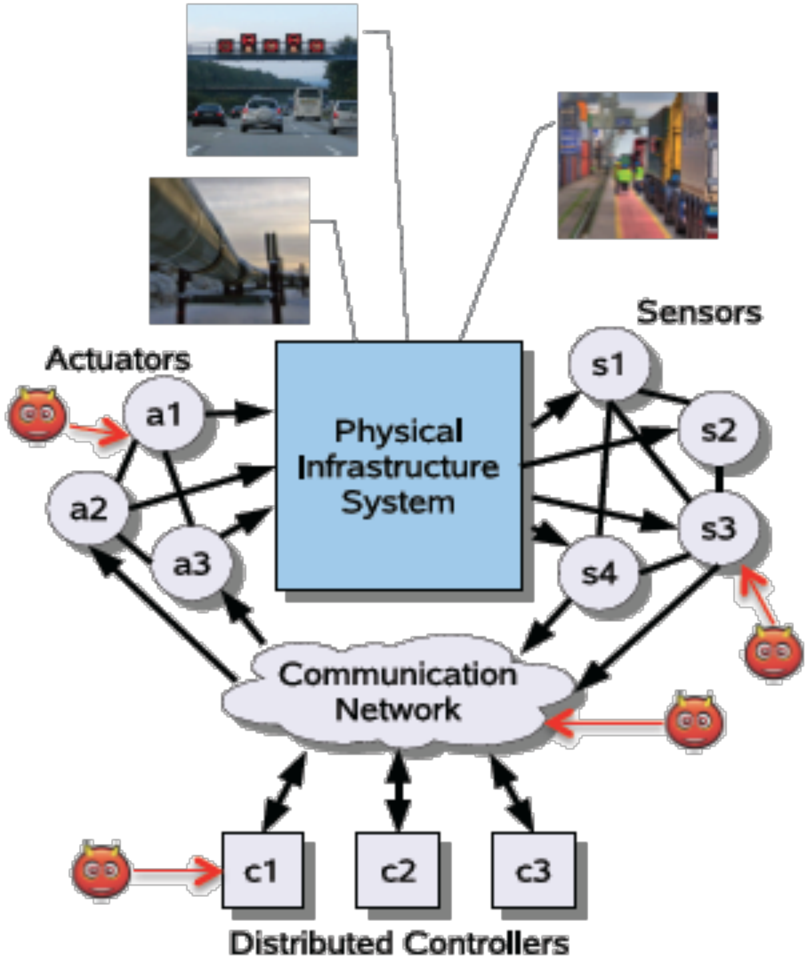

SCADA SYSTEMS FOR WATER DISTRIBUTION

Avencq cross-regulator

- **Regulatory control of canal pools**
  - Manipulate gate opening
  - Control upstream water level
  - Reject disturbances (offtake withdrawals)
- **SCADA components**
  - Level & velocity sensors
  - PLCs & gate actuators
  - Wireless communication

SCADA Interface

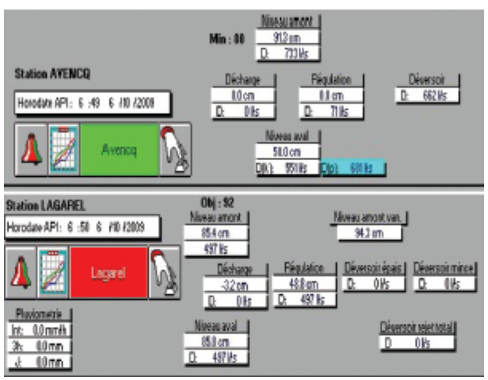

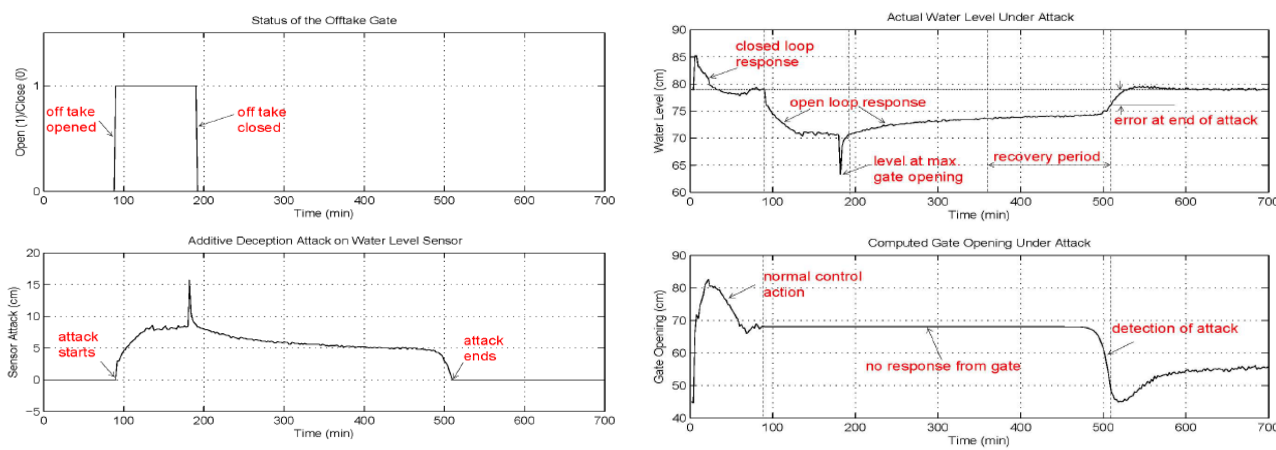

Successful attack: Field operation test (Oct. 12, 009)

TRAFFIC CONTROL SYSTEMS

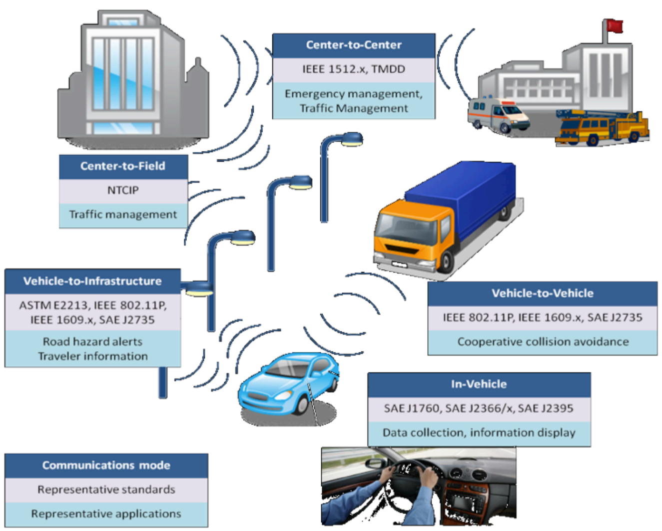

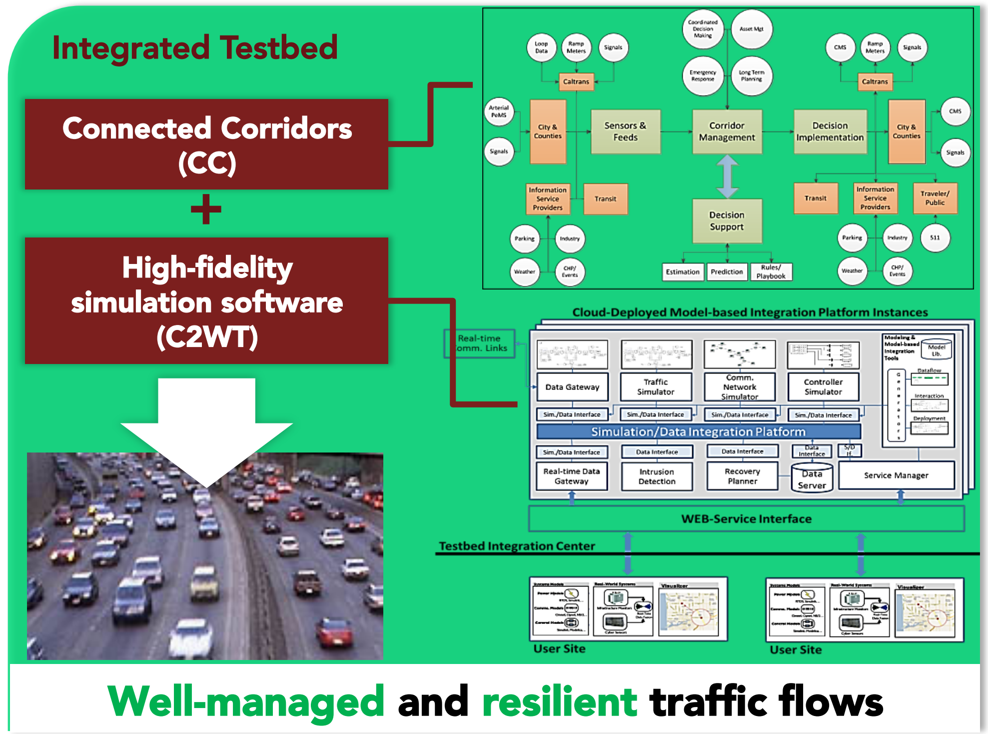

ACHIEVING RESILIENCE

A System Function *can be* allocated to various (combinations of) providers: Applications / Processes / Components

Processes / Components *can be* allocated to various (combinations of) platform Nodes

When a Node / Link / Process / Component fails (compromised), functionality can be restored by an

- alternative allocation of *functions* to *providers*, or
- alternative allocation of *providers* to *platform* nodes

RESEARCH PROBLEMS

Risk Analysis and Incentive Design

1. How the collection of agents in CPS can deal with strategic adversaries?
2. How strategic agents contribute to CPS efficiency and safety, while protecting their conflicting individual objectives?

Resilient Monitoring and Control

1. What are the control architectures that can improve resilience against intrusions and faults?
2. What types of dynamics can provide inherent robustness against impacts of faults and cyber attacks?
3. What are the physics-based invariants that can be used as “ground truth” in intrusion detection?

Decentralized Security

1. How can we design systems that are resilient event when there is significant decentralization of resources and decisions?

Formal Reasoning about Security in CPS

1. How do formally and practically reason about secure computation and communication?

Integrative Research and Evaluation

1. How to integrate and evaluate cyber & physical platforms and resilient monitoring & control architectures?
2. How to interface and support human decision makers?

PROJECT THRUSTS

1. **Hierarchical Coordination and Control**
   1. **Risk analysis and incentive design** that aim at developing regulations and strategies at the management level
   2. **Resilient monitoring and control** of the networked control system infrastructure
2. **Science of decentralized security** which aims to develop a framework that will enable reasoning about the security of all the integrated constituent CPS components
3. **Reliable and practical reasoning about secure computation and communication** in networks which aims to contribute a formal framework for reasoning about security in CPS
4. **Evaluation and experimentation** using modeling and simulation integration of cyber and physical platforms that directly interface with human decision makers.
5. **Education and outreach**

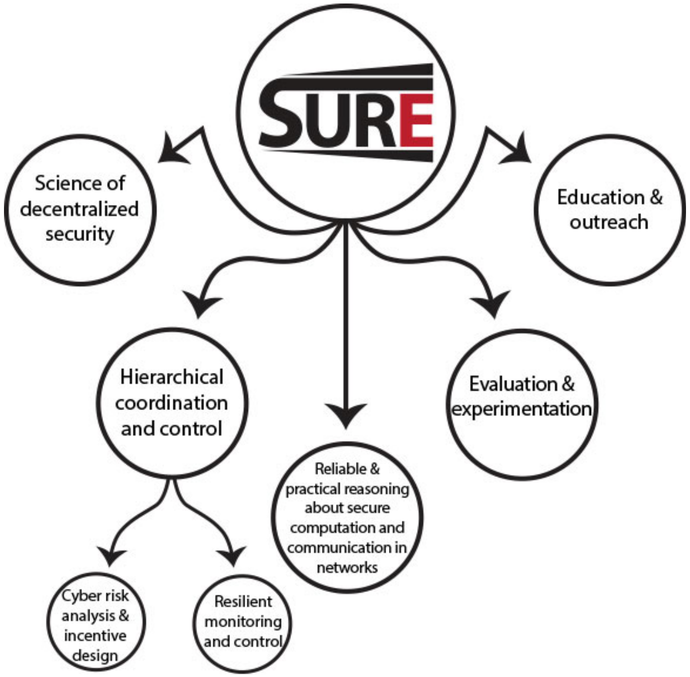

RISK ANALYSIS AND INCENTIVE DESIGN

SURE

1. *Game Theory*: How to model and solve large-scale network games that a) model both security (malicious attacks) and reliability (random faults) failures, b) account for the presence of dynamics and information incompleteness?
2. *Theory of incentives*: How to design and solve stochastic control and incentive-theoretic schemes, coupled with the outcome of the network games (mentioned above)?

A problem of incentives: Due to the presence of network-induced interdependencies, the individual optimal (Nash) security allocations are suboptimal

Goal: Develop mechanisms to reduce CPS incentive sub-optimality

Two-stage game of M plant-controller systems

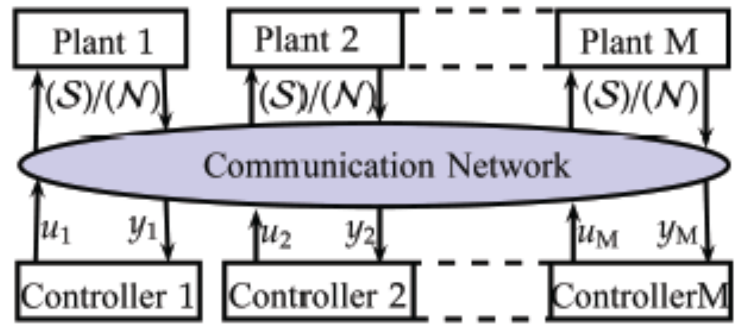

Theorem [Increasing incentive case]

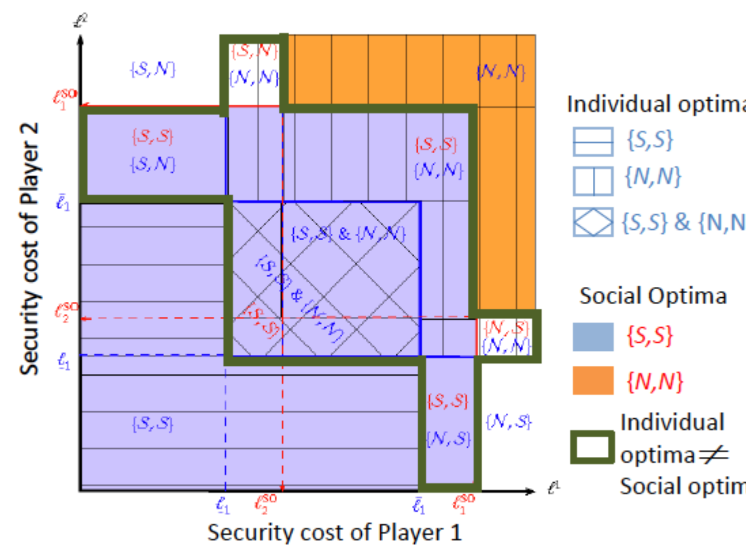

[Amin and Sastry]

VANDERBILT UNIVERSITY

RESILIENT MONITORING

1. How to to detect faults and attacks, which may degrade system performance, cause instability, and affect system operation and mission?
2. How to design resilient monitoring protocols that are robust to both random faults and adversarial attacks?
3. How to place and select sensors to improve resilience?

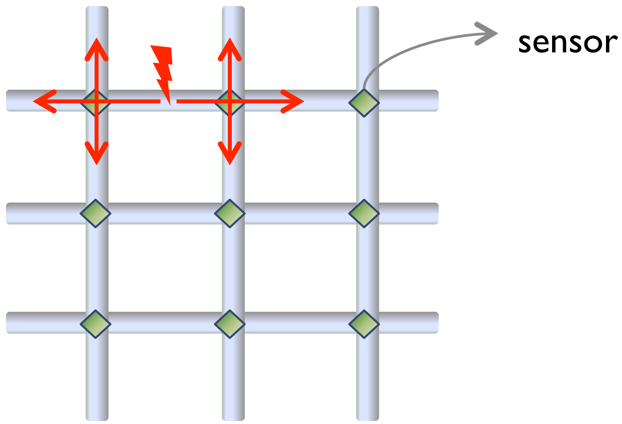

Resilient Fault Diagnosis for Flow Networks

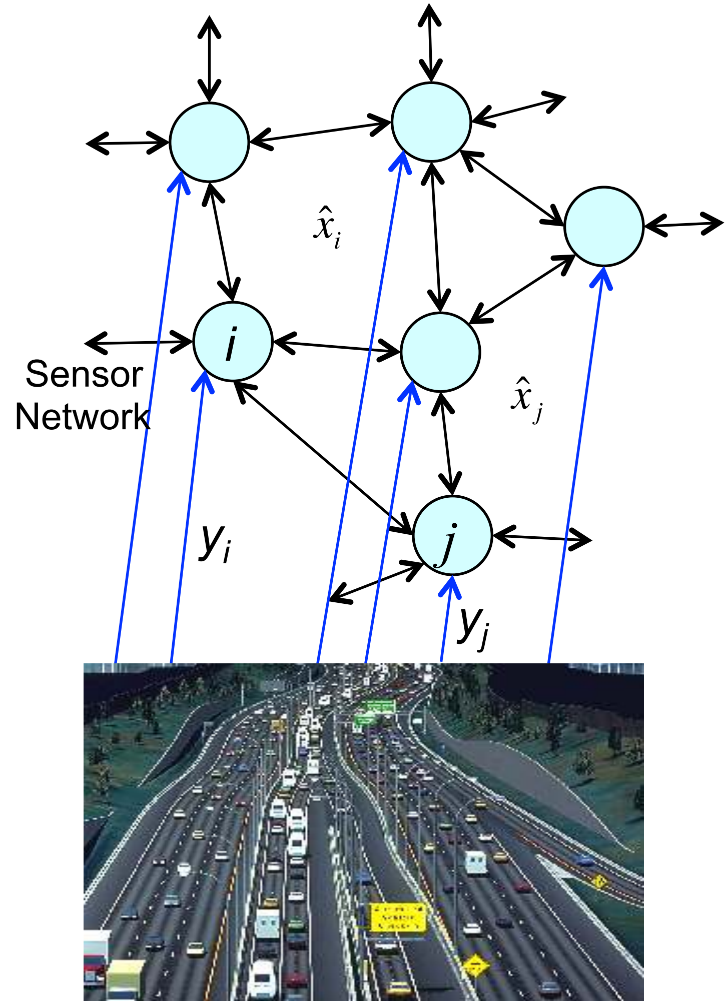

Resilient Distributed Consensus

[Amin and Koutsoukos]

ADVERSARIAL MACHINE LEARNING: RESOURCE AWARE LARGE-SCALE MALWARE CLASSIFICATION

- How to acquire labeled (ground truth) data for evaluation?
- How to achieve very high accuracy (low false positive and low false negative rates) and transparency?
- How to reduce human and machine workloads while retaining very high accuracy?
- How to explore these problems in a scientifically repeatable and valid environment?

SALT: Secure Active Learning Testbed

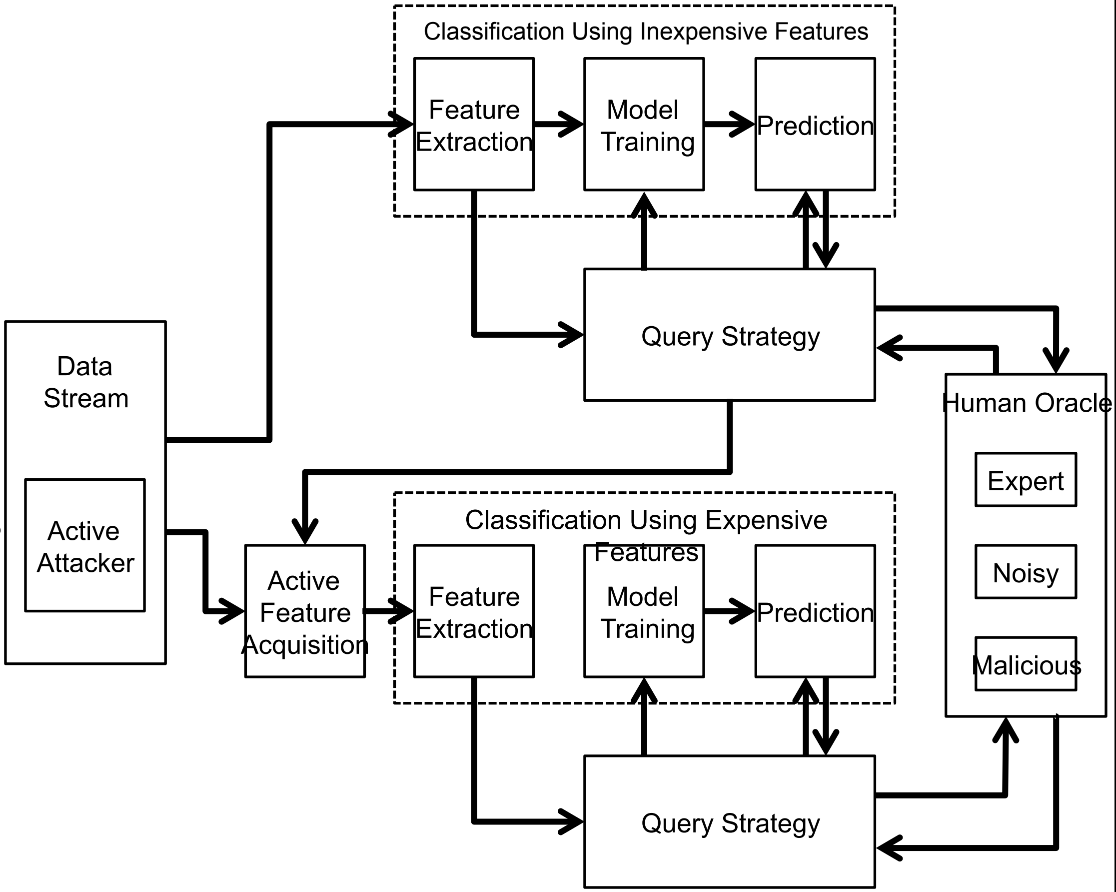

[Joseph]

VANDERBILT UNIVERSITY

RESILIENT CONTROL

Resilient network (supervisory) and local (regulatory) control: How to design practical control algorithms, which improve the survivability of CPS against network-level attacks and/or faults?

Manager's objective: Min social discomfort + inefficiencies

Zone's objective: Min individual discomfort + energy bill

Goal: Incentivize security via monitoring and control

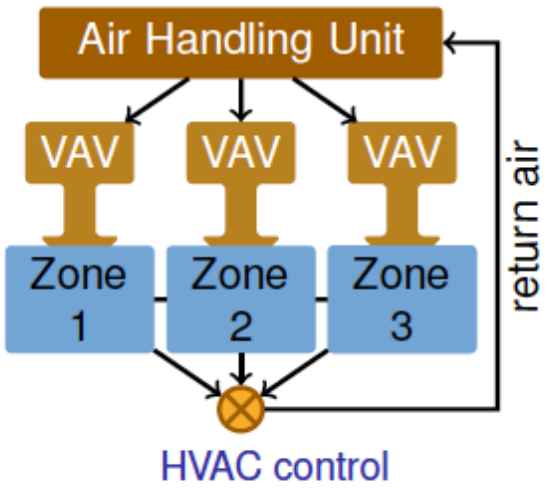

Resilient Control of Building Energy Systems

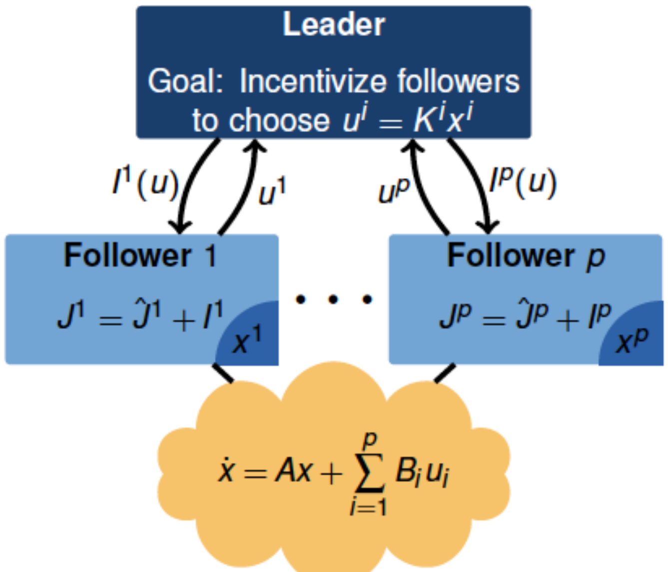

Stackelberg games for resilient control design

[Amin, Sastry, and Tomlin]

VANDERBILT UNIVERSITY

SENSOR/CONTROL NETWORK PLATFORM

SURE

Challenge: How to design and analyze system architectures that deliver required service in the face of compromised components?

Concept: Apply principles and techniques from run-time fault management to managing cyber effects

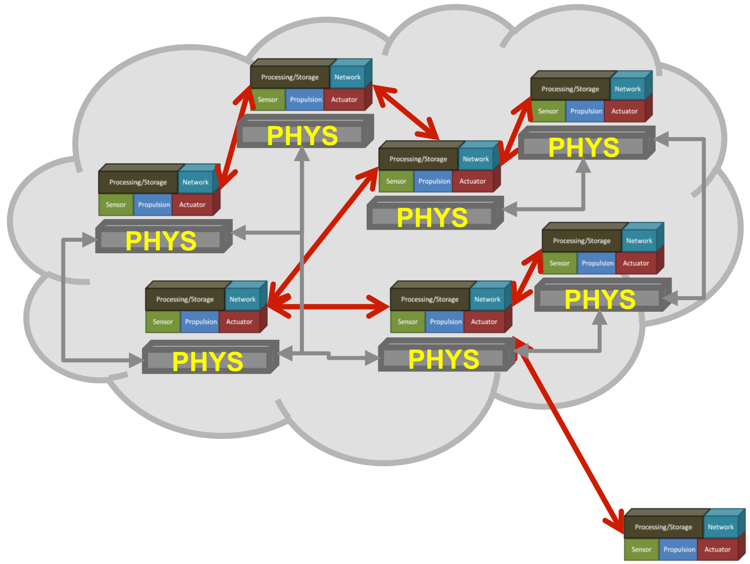

Resilience to faults:

- Detect anomaly
  - Locally or globally
- Isolate fault source
  - App, process, node, link, …
- Recover
  - Restart, replace, reconfigure

Platform provides:

- Overall architecture
- Reusable services for detection, diagnosis, mitigation

Application specific:

- Specific logic for detection, isolation, mitigation

Resilience to cyber *effects*:

- Detect anomaly
  - Locally or globally
- Isolate source of anomaly
  - App, process, node, link
- Recover
  - Restore, replace, reconfigure

Platform provides:

- Overall architecture
- Reusable services for detection, diagnosis, mitigation

Application specific:

- ???

[Karsai]

DECENTRALIZED SECURITY

How can we design systems that are resilient even when there is significant decentralization of resources and decisions?

- Defenders “jointly” own CPS (e.g., electric power grid; railway systems; transportation)
- Attacker chooses where to attack to cause the most damage (e.g., maximum disruption)
- Attacker responds to defensive measures (resilient control strategies; intrusion detection/prevention measures)

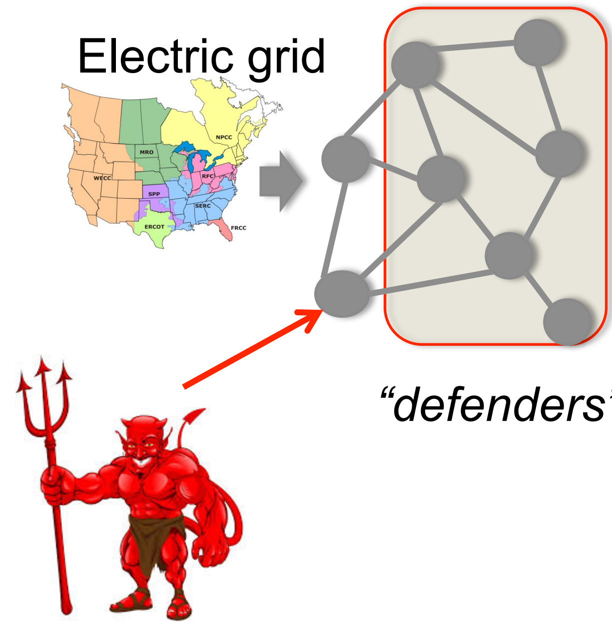

How do defenders who are primarily concerned about the portion of CPS they own choose their security measures?
Depends on the level of decentralization and the degree of system interdependence

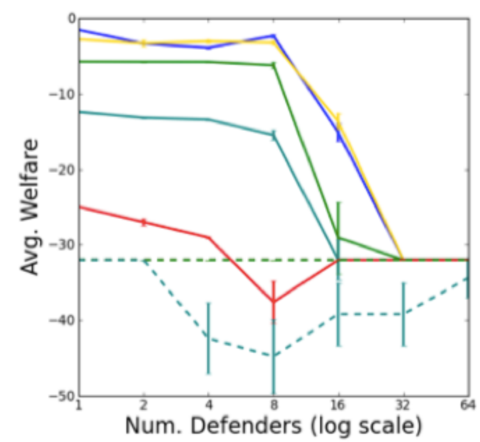

[Vorobeychik]

VANDERBILT UNIVERSITY

MODELING AND PROVING SECURITY IN NETWORKS

PROBLEM

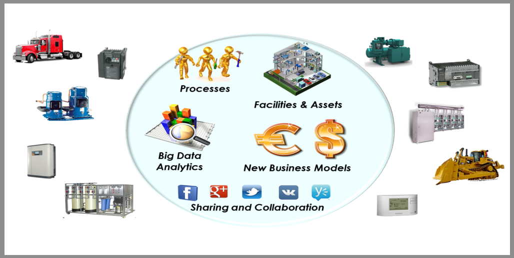

- *High assurance for Cyber Physical Systems*
- *Network computation with physical interface*

BACKGROUND

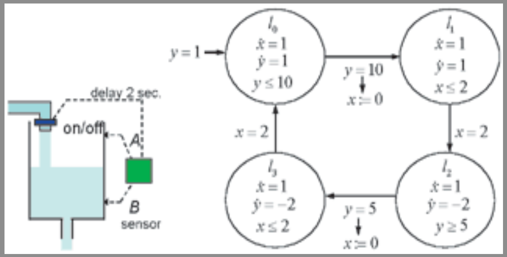

- *Hybrid systems, Petri nets*
- *Protocol Derivation Logic, Strand spaces*

APPROACH

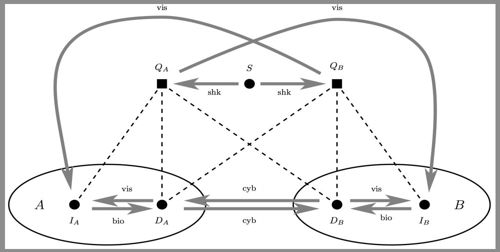

- ***Actor networks**: fibered state machines*
- ***Network computation**: partially ordered multisets (pomsets)*

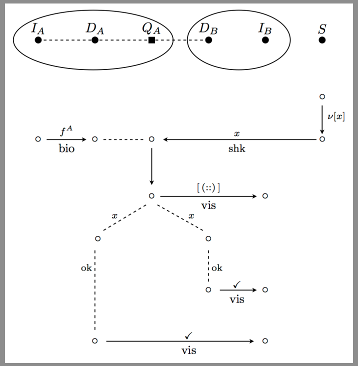

- ***Procedure Derivation Logic***
- ***Authentication templates** extended to capture physical and social channels*

VANDERBILT UNIVERSITY

[Pavlovic]

EVALUATION USING MODELING AND SIMULATION INTEGRATION

SURE

- **Validation of basic research**
  - Scenario-based experimentation
- **Collaboration**
  - Research thrusts and projects
  - Integration: Tools and languages
- **Motivation**
  - Red team vs Blue team scenarios and challenges
- **Outreach**
  - Accessible tools and technologies on the web
- **Model libraries and repositories**

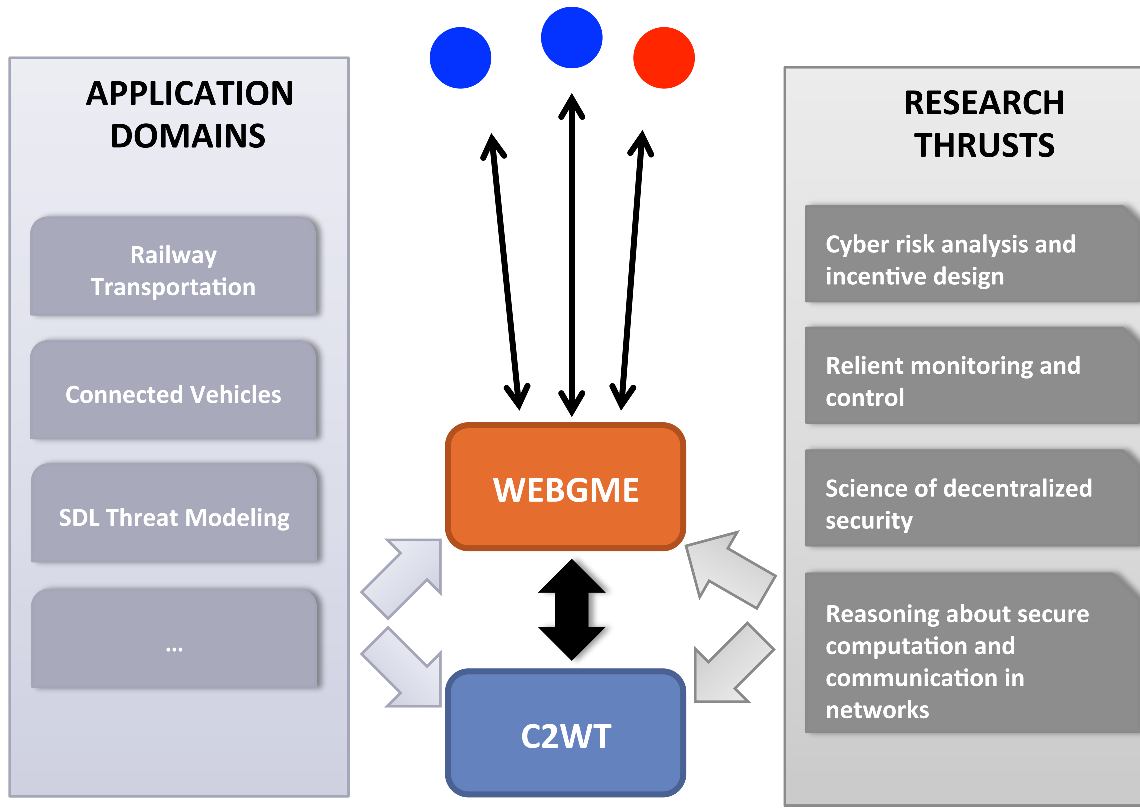

[Volgyesi and Sztipanovits]

EDUCATION AND OUTREACH

SURE

- **Classes**
  - S. Amin, 1.208 Resilient Infrastructure Networks, MIT, Fall 2014
  - X. Koutsoukos, CS 396 Security of CPS, Vanderbilt, Spring 2015.
- **Online Modules**
- **Workshops/Conferences**
  - How to Engineer Resilient Cyber-Physical Infrastructures, IEEE CDC 2014 [Amin]
  - Big Data Analytics for Societal Scale CPS: Energy Systems, IEEE CDC 2014 [Sastry]
  - Secure and Resilient Infrastructure CPS (HiCoNS) track, ICCPS 2015 [Koutsoukos]
- **Evaluation and Experimentation Testbed**
- **SOS-VO**

OVERVIEW

- Team
- Resilience of Cyber-Physical Systems
- Research Problems
- Project Thrusts
  - Risk Analysis and Incentive Design
  - Resilient Monitoring and Control
  - Decentralized Security
  - Formal Reasoning about Security
  - Evaluation using Modeling and Simulation Integration
- **Resilient Monitoring**
- Evaluation using Modeling and Simulation Integration

DISTRIBUTED PARAMETER ESTIMATION

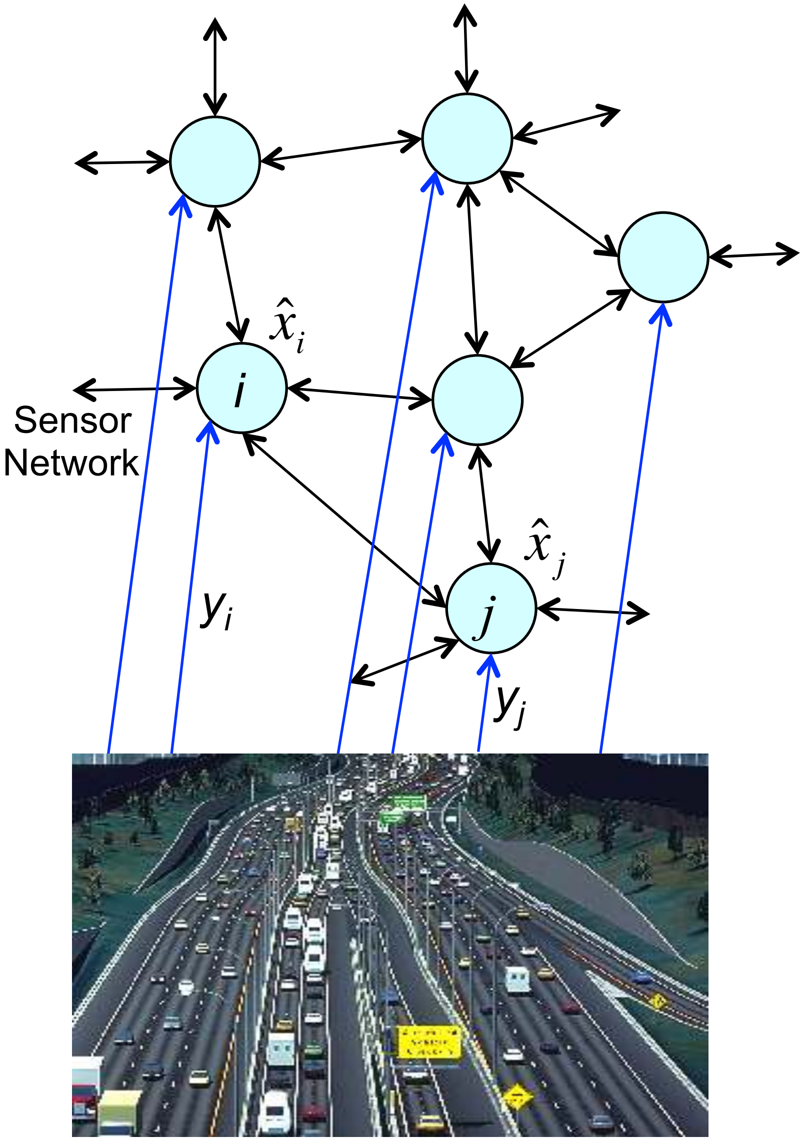

- All sensors measure independently some physical phenomenon with some error due to noise
  $y_i = \theta + v_i, v_i \sim N(0, \sigma_i^2), i = 1, 2, \ldots, n$
- The sensors improve their estimate by averaging the measurements
- Minimum variance estimate

$$\hat{\theta}_{MV} = \frac{\frac{1}{n}\sum_{i=1}^{n}\frac{1}{\sigma_i^2}y_i}{\frac{1}{n}\sum_{j=1}^{n}\frac{1}{\sigma_j^2}}$$

- It can be asymptotically computed in a distributed fashion using two average consensus algorithms in parallel

RESILIENT CONSENSUS IN THE PRESENCE OF ADVERSARIES

(3,2)-robust graph: resilient consensus in the presence of 1 adversary

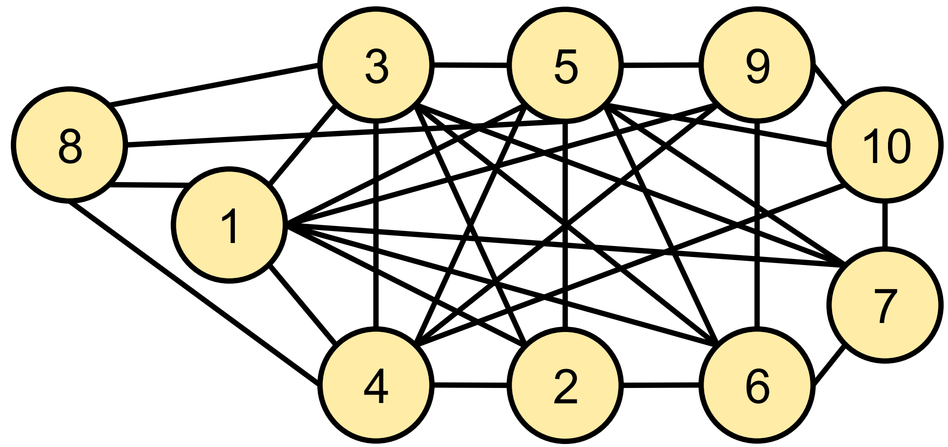

$$x_i(k+1) = \sum_{j \in \mathcal{R}_i(k)} w_{ij}\, x_j(k)$$

Adversarial Consensus Protocol

Adversary models

- Threat
- Scope

Robust network topologies

- Local redundancy

Resilience requires high degree of redundancy

Can we relax the redundancy requirements?

RESILIENT CONSENSUS WITH TRUSTED NODES (RCP-T)

Each normal node updates its value according to the following update rule

$$x_i(k+1) = \sum_{j \in \mathcal{R}_i(k)} w_{ij}\, x_j(k)$$

What is $\mathcal{R}_i(k)$?

- if node *i* is **not connected** to any trusted node

(F is the total number of attacks that can happen within the network)

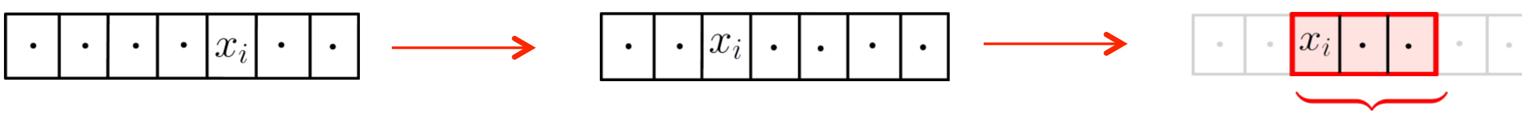

Collect values from neighbors

Sort

$\mathcal{R}_i(k)$

Remove F largest and F smallest values (Here, F=2)

- if node *i* is **connected** to at least one **trusted node**

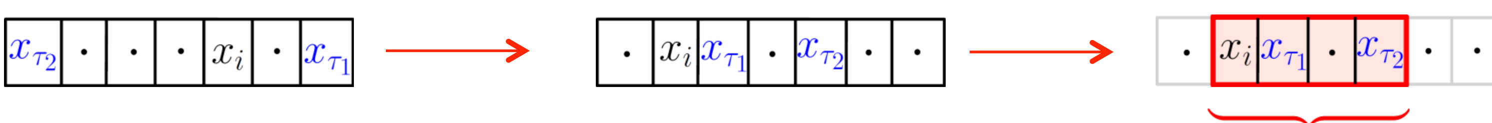

Collect values from neighbors

($x_{\mathcal{T}_1}, x_{\mathcal{T}_2}$ are trustworthy nodes' values)

Sort

$\mathcal{R}_i(k)$

RESILIENT CONSENSUS WITH TRUSTED NODES

Under RCP-T, consensus is always achieved in the presence of *arbitrary number of adversaries* iff there exists a set of trusted nodes that form a **connected dominating set**.

Under RCP-T

- *Any* number of attacks can be handled
- *Sparse* networks can be made resilient

Dominating Set:

$$D \subseteq V, \qquad \text{s.t.} \quad \bigcup_{v_i \in D} \mathcal{N}[v_i] = V$$

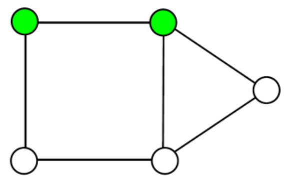

Connected Dominating Set:

Nodes in the dominating set induce a connected subgraph

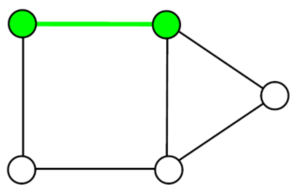

EXAMPLES

Example 1: (Tree – Sparse network)

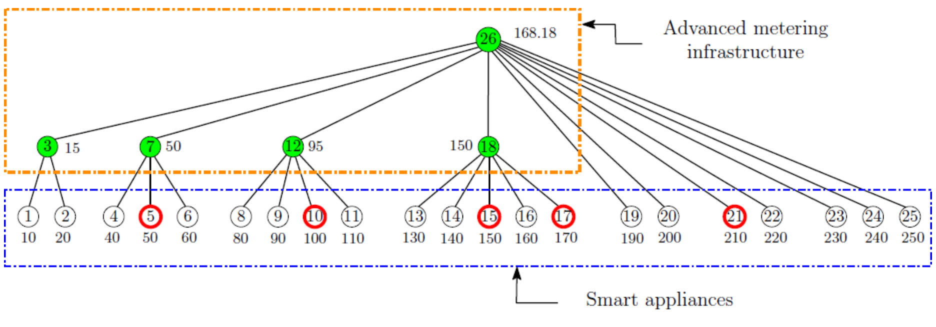

Example 2: (2,2) Robust graph

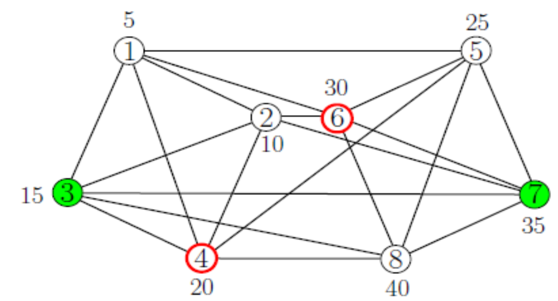

RCP-T

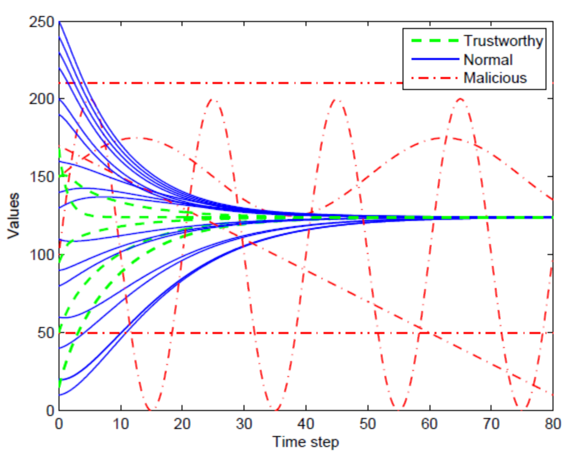

RCP-T

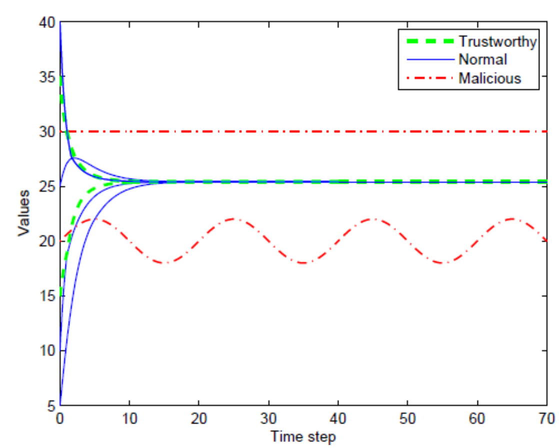

FAULT DIAGNOSIS IN FLOW NETWORKS

Objective: For a given flow network, the goal is to distribute the minimum number of sensors that can

1. Detect a link failure
2. Localize a link failure (uniquely identify a link failure)

Approach: Sensor network design for the detection and identification of faults

Methods: System (flow network, faults, sensor) model, combinatorial optimization

Performance evaluation: Resilience to random sensor faults and adversarial attacks

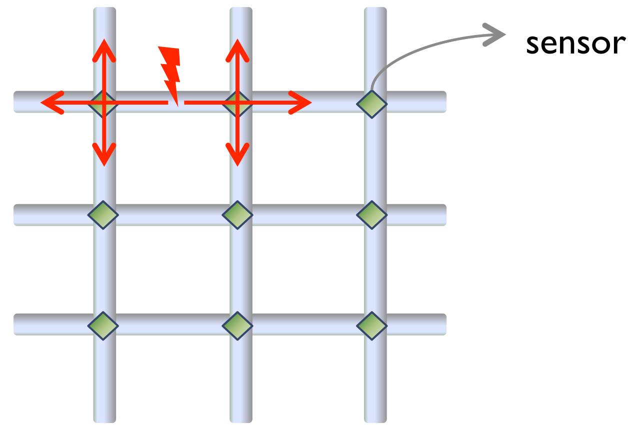

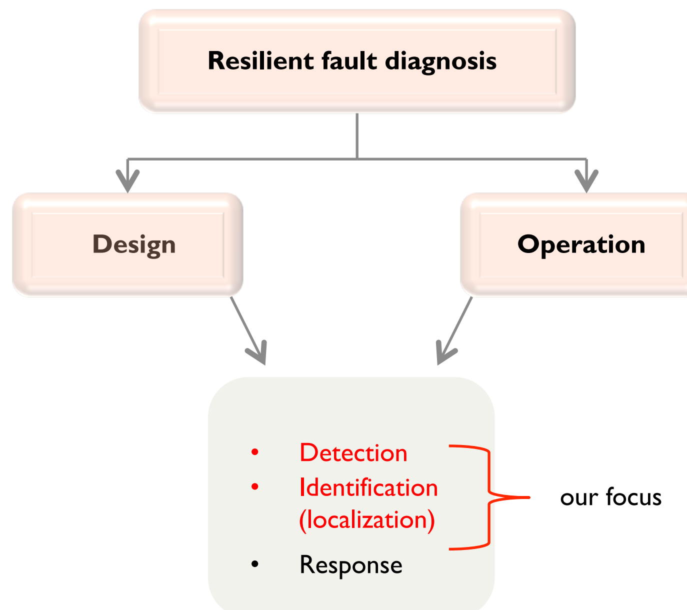

[Sela, Abbas, Koutsoukos, and Amin]

VANDERBILT UNIVERSITY

SYSTEM MODEL

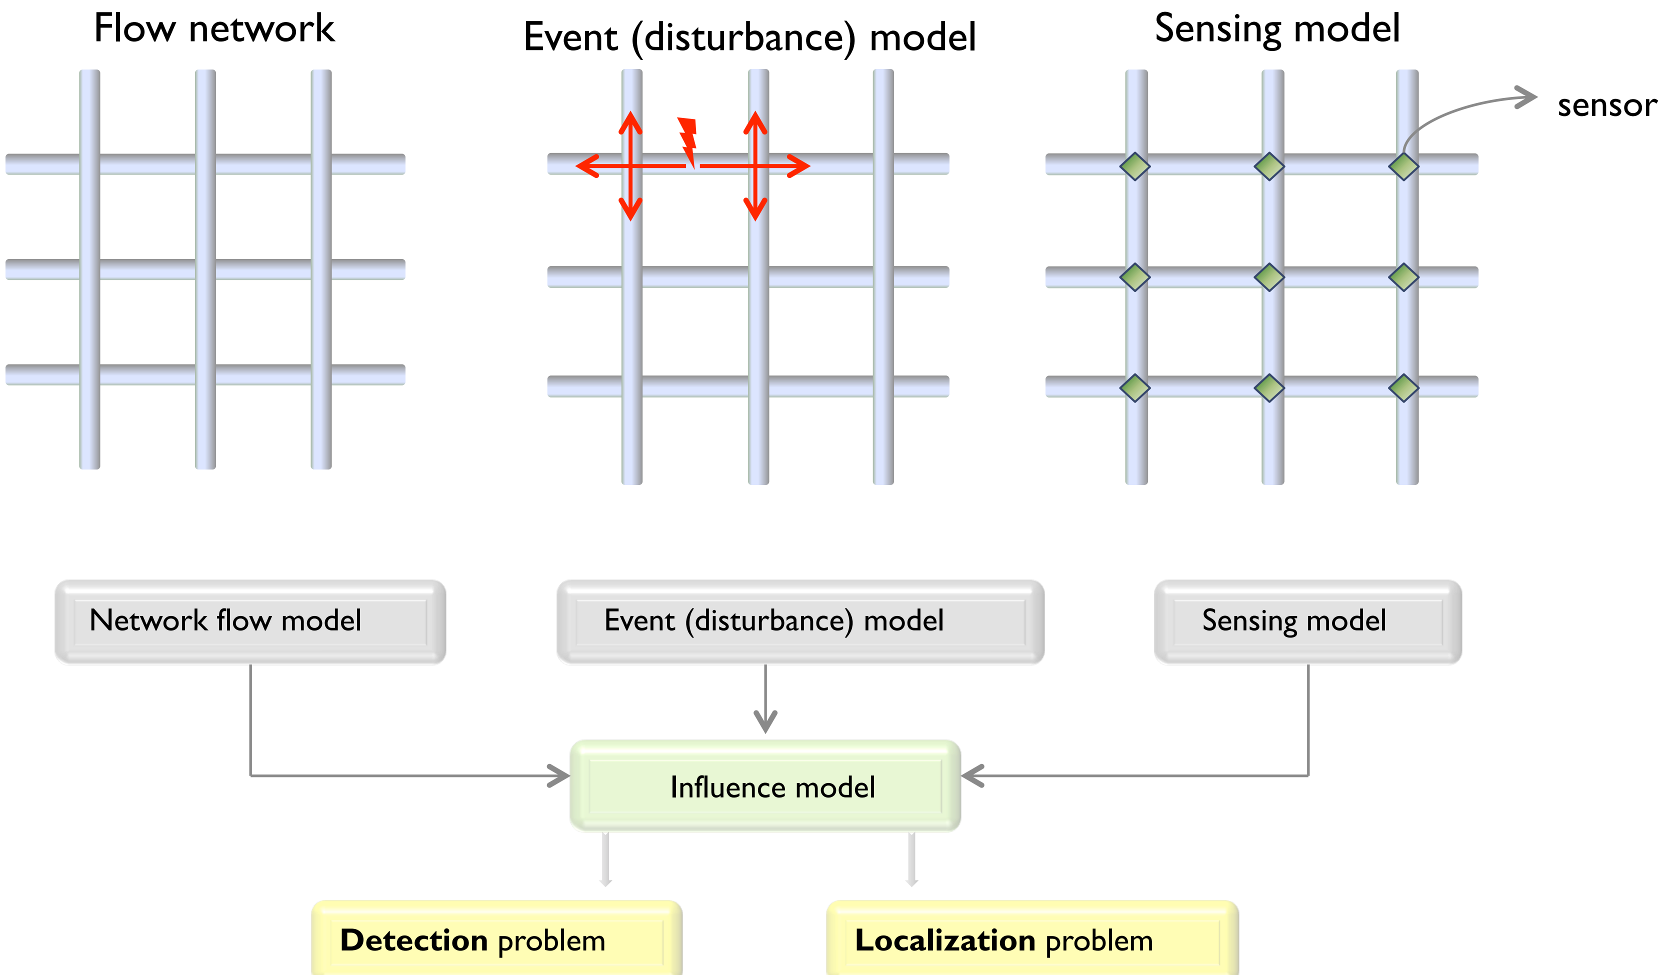

FLOW NETWORK MODEL

SURE

Flow network ⟷ Graph

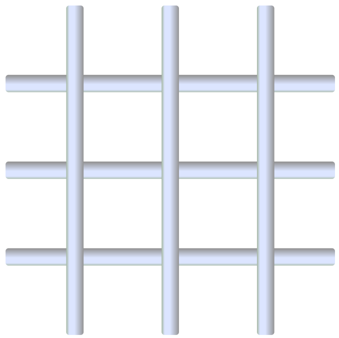

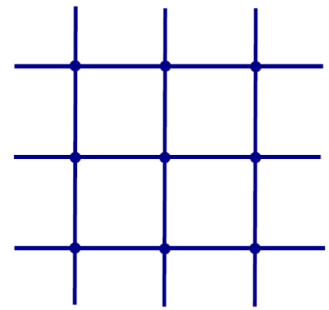

Links		Edges
Physical characteristics (e.g., length, material parameters etc.)	⟷	Edge weights
Physical characteristics (e.g., nodal demands etc.)		Node weights

The physical network is defined by

- Network topology (graph): G = G(V,E)
- Flow model over the graph G: f = f(Q,H,G)

Q = flows over network links
H = heads over network nodes
[H = p + z
p = pressure
z = elevation]

EVIDENCE (DISTURBANCE) MODEL

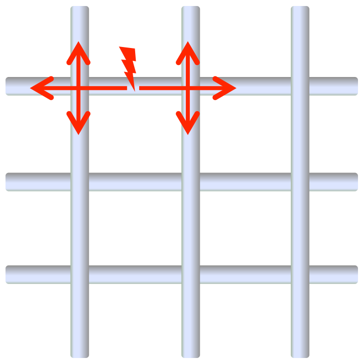

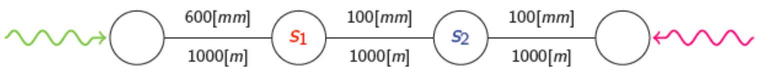

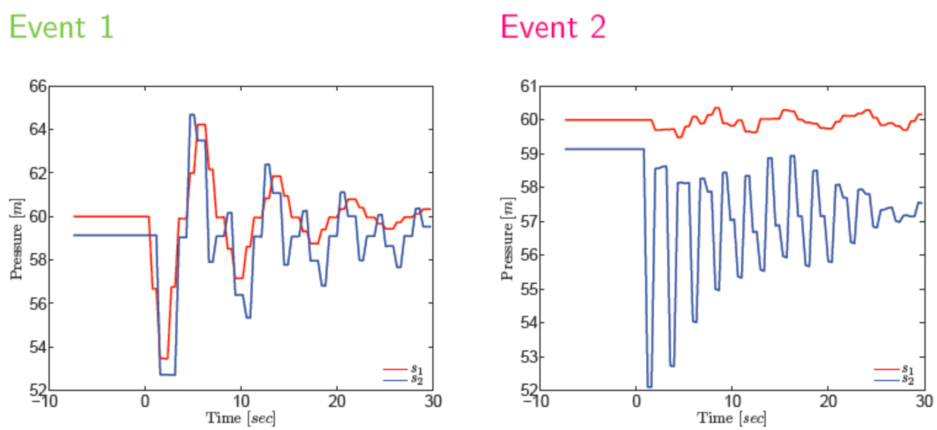

- Event model is comprised of a link failure and its impact
  - Pipe failure (random or induced) and the pressure transient generated.
  - Physically flushing an hydrant causing massive loss of water, increased load on the system and corresponding pressure losses.
  - Remotely closing or opening active elements (pumps, valves) that can cause severe transients in the systems.
- The signal propagates in all directions from the site of failure along the links of the network.

Along with the network topology, the **physical model** defined over the graph also affects the event model

SENSING MODEL

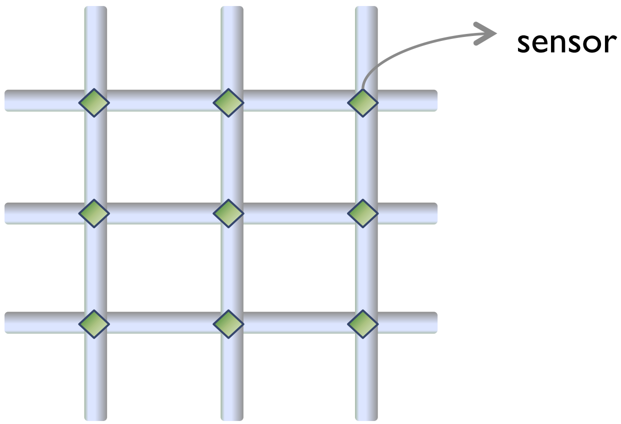

- Sensors are placed at the nodes.
- A sensor can detect the pressure signal from any direction.
- An alarm is raised when a sensor detects a signal.

First-order model:

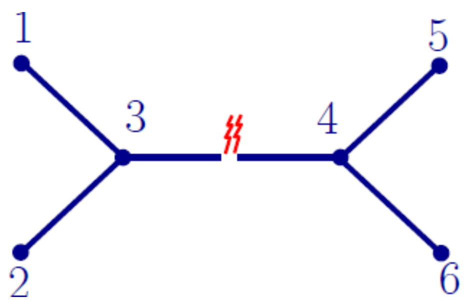

All sensors at the nodes adjacent to the end nodes of failed link will detect the fault.

R-disk model:

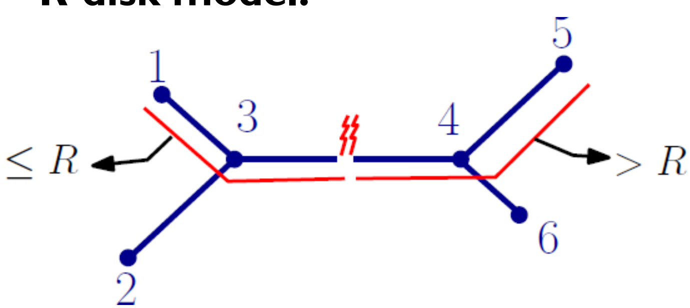

A sensor can detect a fault if and only if fault occurs at a link that lies within the distance R from the failure along the links.

INFLUENCE MODEL

- Network flow, event, and sensor model outputs are represented using an **influence matrix** M.
- $\boldsymbol{\ell}_i$ - i^{th} **row** corresponds to the **event** i.
- $\boldsymbol{\theta}_j$ - j^{th} **column** corresponds to the j^{th} **sensor**.
- $\boldsymbol{M_{ij}}$ - j^{th} sensor output in response to the event i.

Example: $\boldsymbol{M}$ is boolean matrix.

$$
\begin{array}{c}
\\ \ell_1 \\ \ell_2 \\ \vdots \\ \ell_i \\ \vdots \\ \ell_m
\end{array}
\begin{array}{c}
\begin{array}{cccccc} \theta_1 & \theta_2 & \dots & \theta_j & \dots & \theta_n \end{array} \\
\begin{pmatrix}
0 & 0 & \dots & 1 & 0 & 0 \\
1 & 1 & \dots & 0 & 1 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\
1 & 0 & \dots & 1 & 1 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\
0 & 0 & \dots & 1 & 0 & 1
\end{pmatrix}
\end{array}
$$

Sensor 1's output is 1 when event 2 occurs.

DETECTION AND LOCALIZATION

Detection

Find the minimum number of sensors and their locations so that every link failure can be detected by at least one sensor.

Event set: $\{\ell_1, \ell_2, \cdots, \ell_m\}$

Sensor set: $\{\theta_1, \theta_2, \cdots, \theta_n\}$

Detection set: C_i = Set of links whose failure is detected by the sensor i.

Minimum set cover problem

Localization

Find the minimum number of sensors and their locations so that every link failure can be uniquely identified and can be distinguished from any other link failure.

Event set: $\{\ell_1, \ell_2, \cdots, \ell_m\}$

Sensor set: $\{\theta_1, \theta_2, \cdots, \theta_n\}$

Identification set

Minimum test cover problem

EXAMPLE

- Consider a 10 by 10 grid network consisting of 100 nodes and 180 links.
- Influence matrix is obtained using the first order influence model.

Detection → (greedy) MSC

Localization → MTC → (transform) → (greedy) MSC

18 sensors are sufficient to **detect** any link failure.

46 sensors are sufficient to **localize** any link failure.

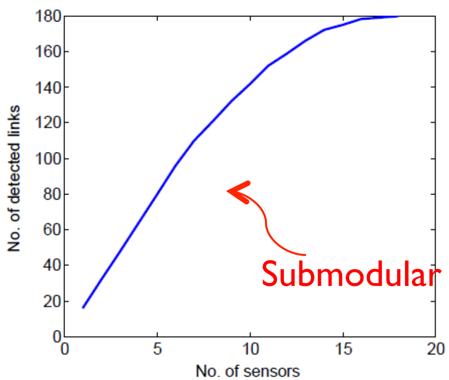

No. of detected links as a function of sensors deployed.

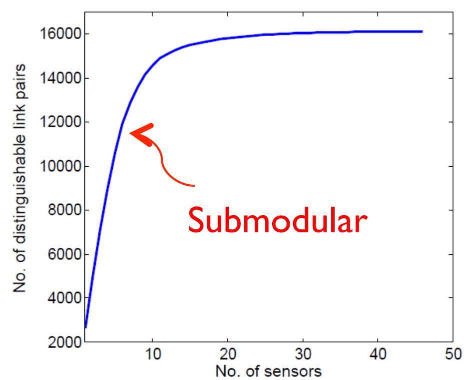

No. of distinguishable link pairs as a function of sensors deployed.

Total no. of link pairs

$$\binom{180}{2} = 16110$$

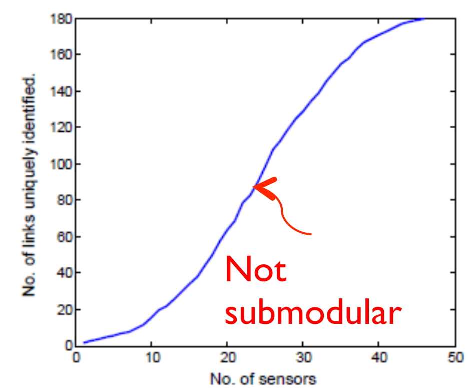

No. of links that can be uniquely identified as a function of sensors deployed.

FUTURE WORK

- Incorporate network topology and influence model to design efficient (scalable, improved approximation ratios) algorithms for detection and localization
- Characterize detection and localization of link failures as a function of number of sensors deployed (e.g., submodularity)
- To make the system resilient against these compromises, we might need to include redundancy (more sensors than required)
- How can we design a sensor networks for resilient localization and identification?
- How the detection & localization of link failures are dependent on the influence model and network topologies?
- For a given influence model, what are the (structural) constraints on the network topology such that *every* link failure can be detected, localized, in a resilient manner?
- Generalizations
  - Associating a probability distribution to the link failures.
  - Detecting (localizing) $k > 1$ simultaneous link failures.
  - Incorporating more generalized sensing and influence model.

OVERVIEW

- Team
- Resilience of Cyber-Physical Systems
- Research Problems
- Project Thrusts
  - Risk Analysis and Incentive Design
  - Resilient Monitoring and Control
  - Decentralized Security
  - Formal Reasoning about Security
  - Evaluation using Modeling and Simulation Integration
- Resilient Monitoring
- **Evaluation using Modeling and Simulation Integration**

GOALS

- **Evaluation and Experimentation**
  - Design, Deployment, and Validation
  - Scenario-based experimentation
- **Collaboration/Integration**
  - Research thrusts and projects
  - Tools and languages
- **Motivation**
  - **Red** team vs **Blue** team scenarios and challenges
- **Outreach**
  - Accessible tools and technologies on the web
- **Model libraries and repositories**

MODELING AND SIMULATION INTEGRATION

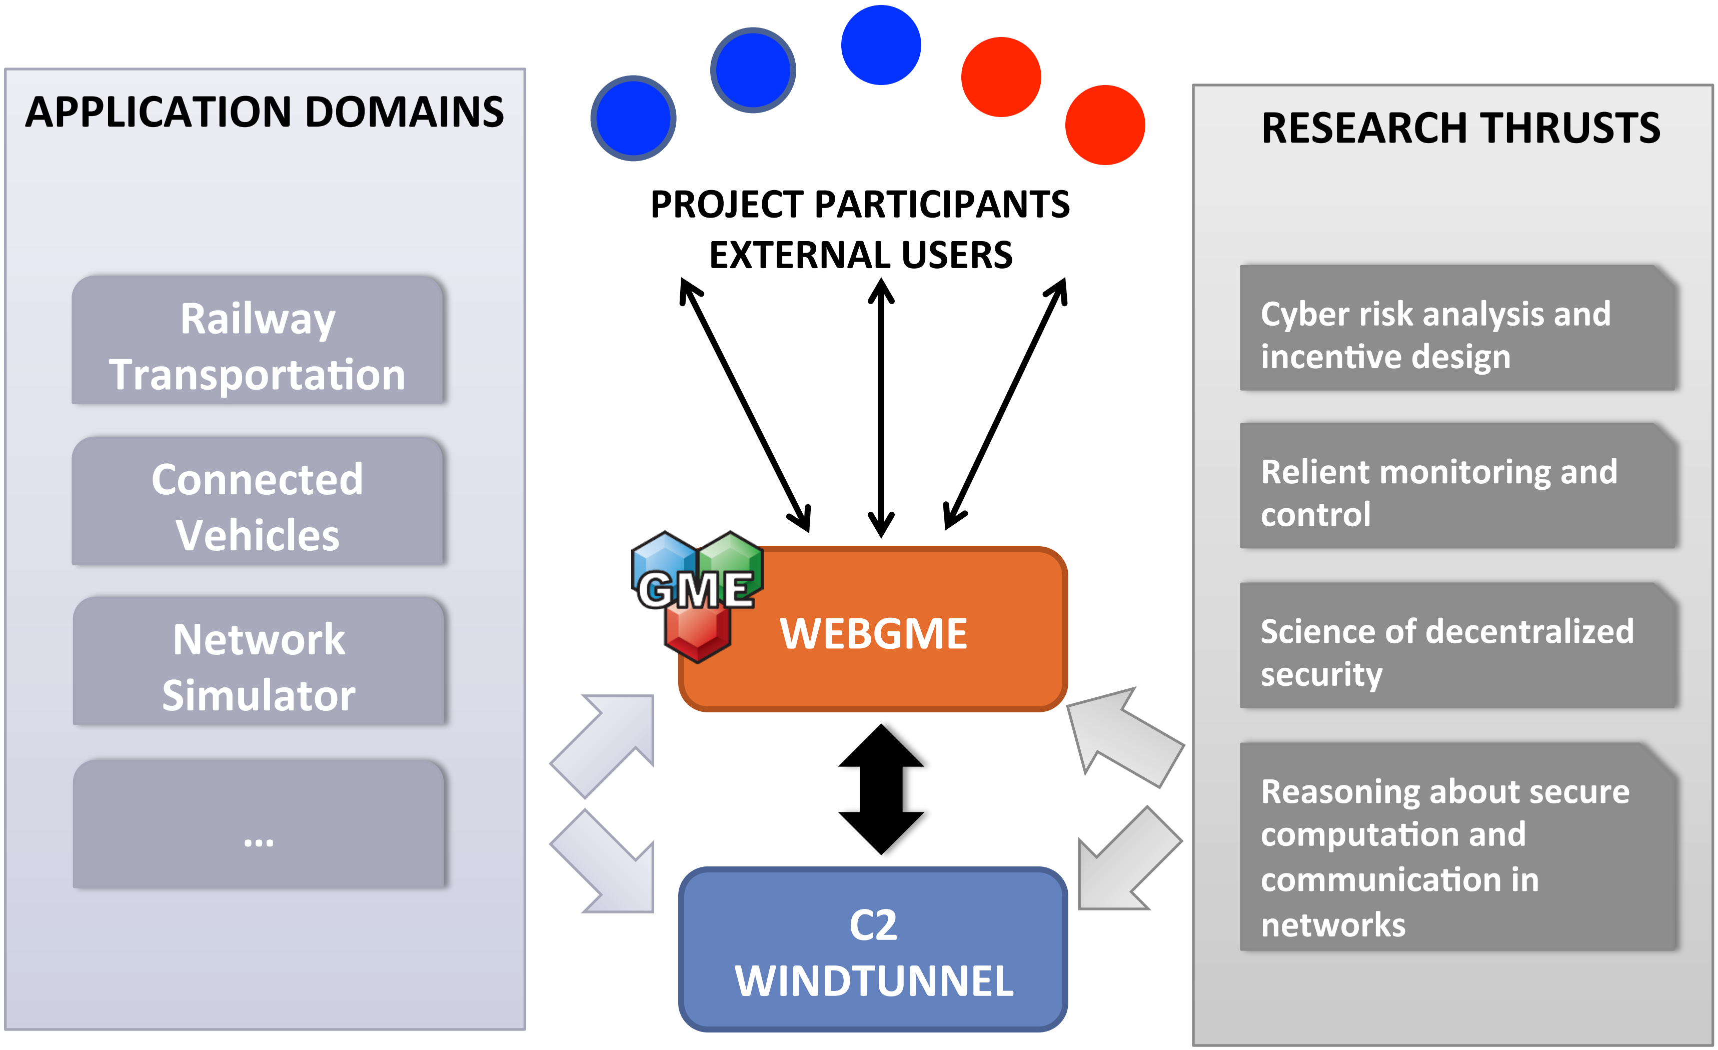

SCENARIO-BASED EXPERIMENTATION

Red Team vs Blue Team

- Pre-defined infrastructure model (transportation, IT … domains)
- Domain specific attack models, libraries, algorithms
- Red Team: design and deploy attacks
- Blue Team: design and deploy security and fail-over measures

(Red Team and Blue Team items: budget constraints)

- Cloud-based simulation tools configured and run w/o real-time user interactions
- Scoring, leaderboard

WEBGME

Meta-programmable collaborative on-line modeling environment

- **Scalable (number of contributors, size of models)**
- **Modern web-application framework**
- **Collaboration**
  - Immediate feedback
  - Branch and merge
- **Version management (git model)**
- **Clean and unified meta/DSML concepts**
- **Extensible, customizable GUI**
- **Cloud-based tool integration**
- **Live: http://webgme.org**

WEBGME

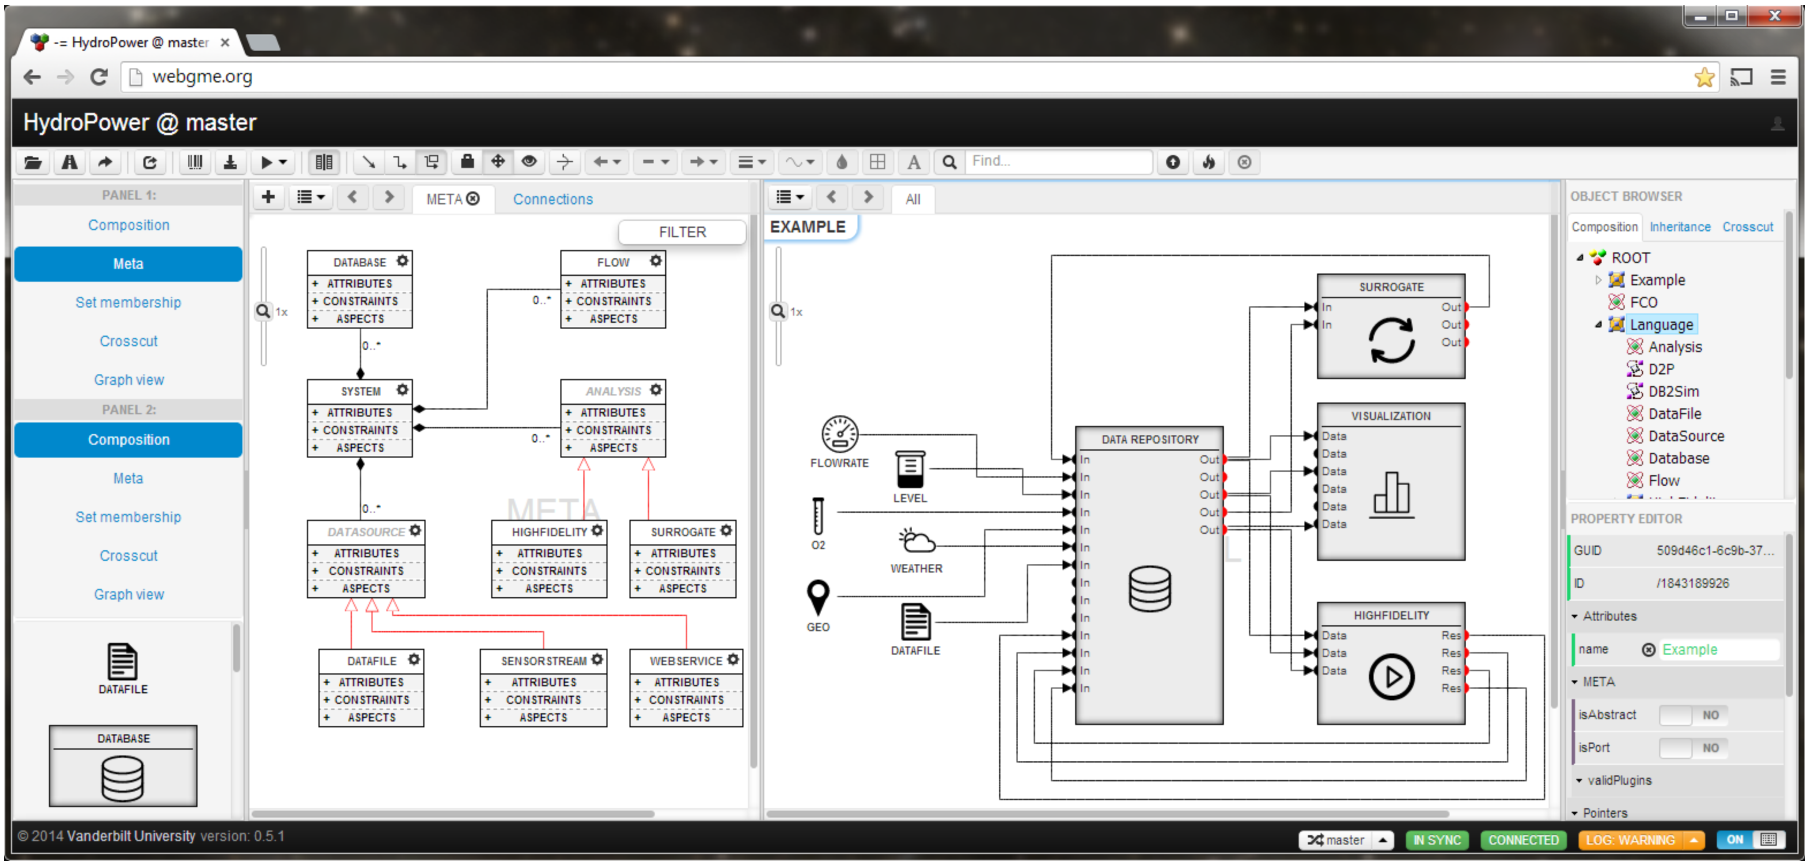

VANDERBILT UNIVERSITY

MULTI-MODEL INTEGRATION CHALLENGES

Integrating *models*

Heterogeneous models for different domains: human organizations, communication networks, C2 software systems, vehicle simulations, etc. These models need to talk to each-other somehow.

Needed: an overarching *integration model* that connects and relates these heterogeneous domain models in a logically coherent framework.

Integrating the *system*

Heterogeneous simulators and emulators for different domains: Colored Petri Nets, OMNET++, DEVS, Simulink/ Stateflow, Delta3D, etc.

Needed: an underlying *software infrastructure* that connects and relates the heterogeneous simulators in a logically and temporally coherent framework.

Key idea: Integration is about messages and shared data across system components. Why don't we model these messages and shared data elements and use these models to facilitate model and system integration?

C2 WINDTUNNEL

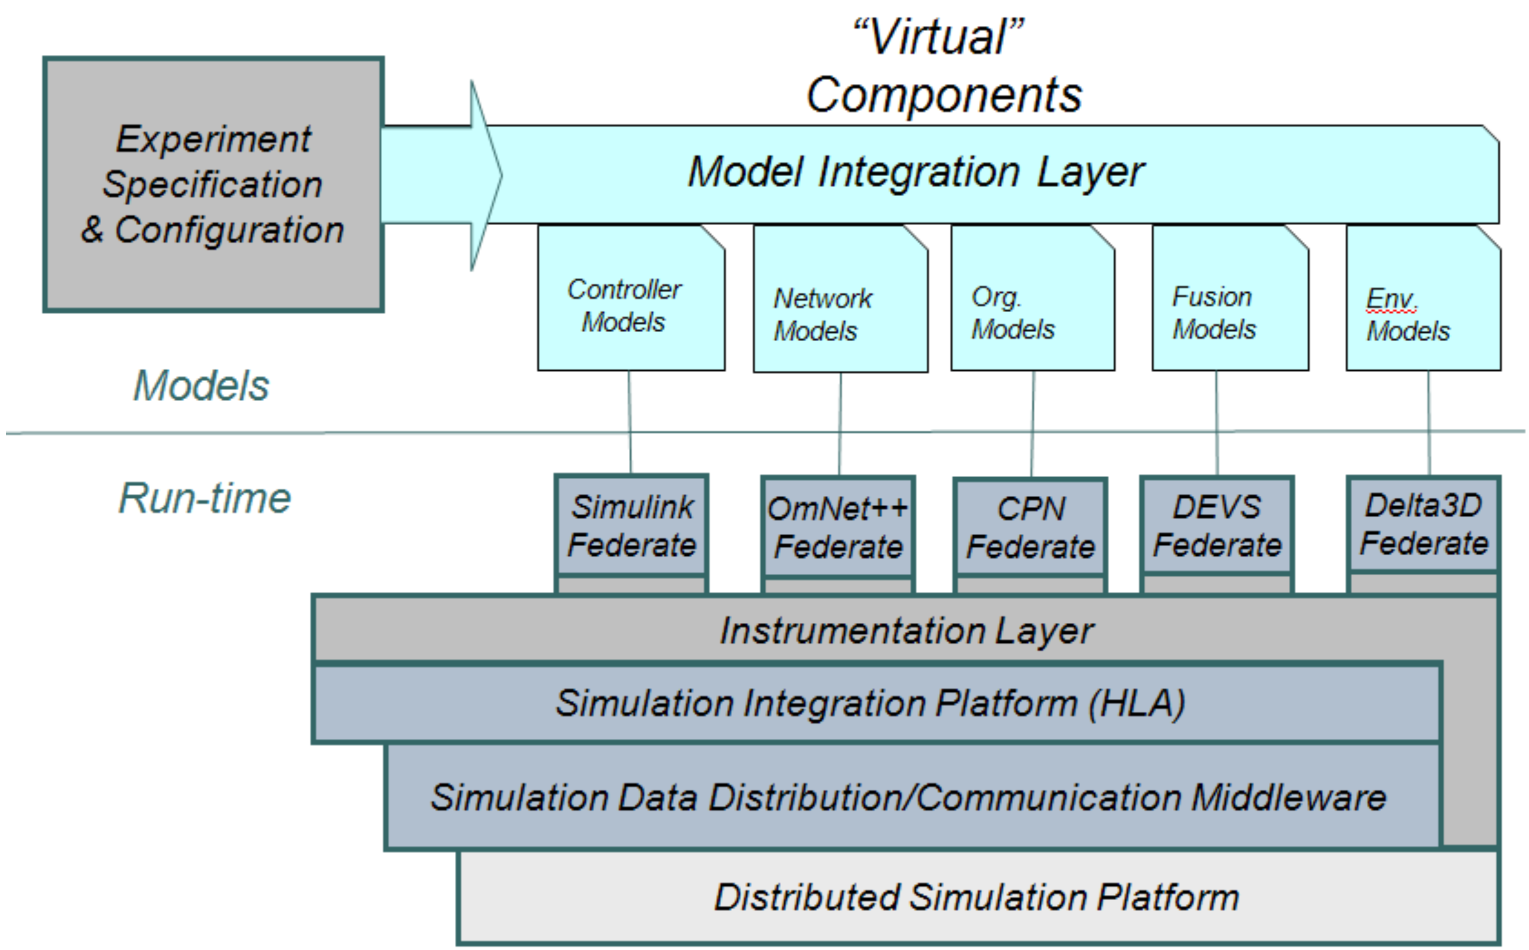

SMARTAMERICA CHALLENGE

Integrated CPS Testbed

Connected Corridors (CC)

+

High-fidelity simulation software (C2WT)

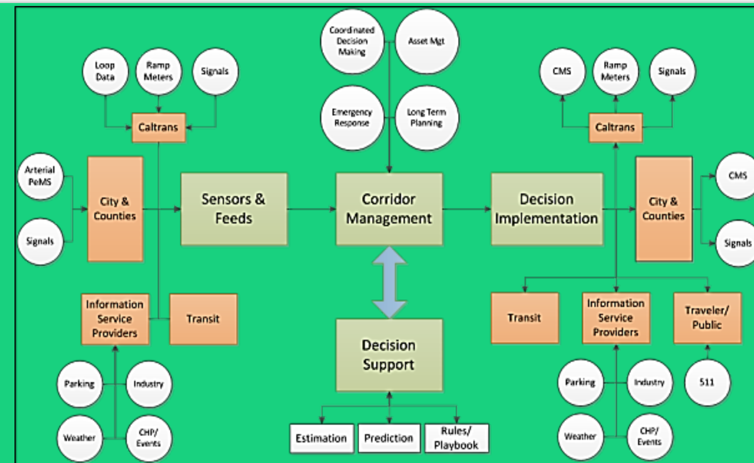

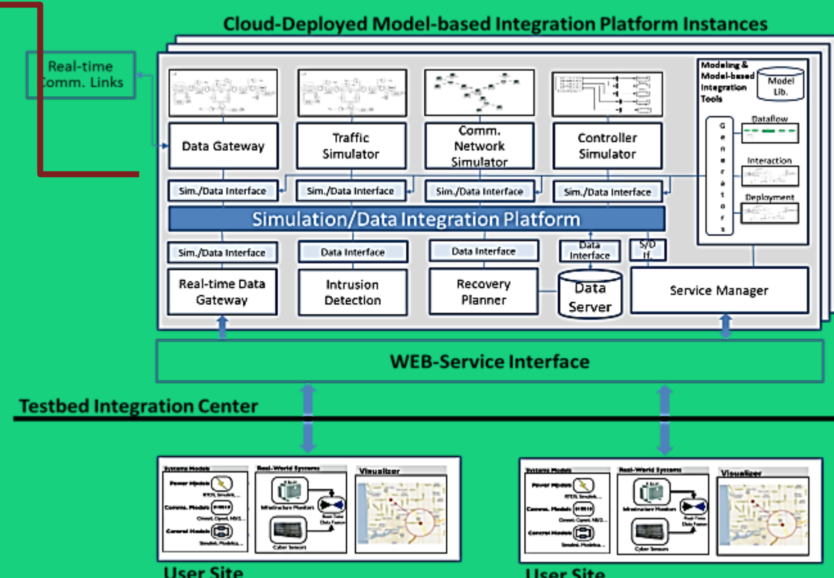

Well-managed and resilient traffic flows

SYSTEM SCIENCE OF SECURITY AND RESILIENCE OF CPS

SURE

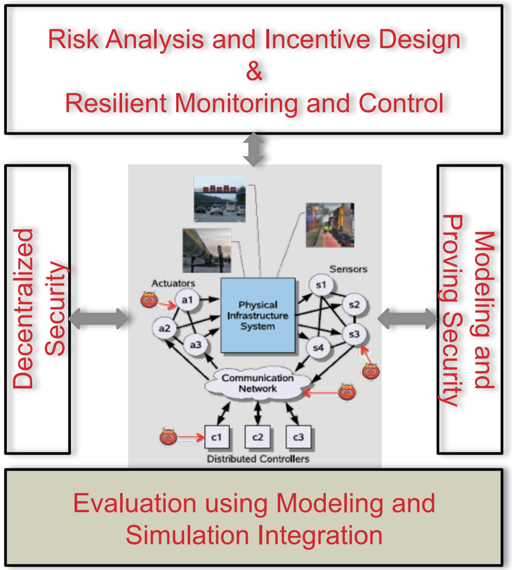

Key Ideas

1. Hierarchical Control and Coordination
   1. Risk analysis and incentive design that aim at developing regulations and strategies at the management level
   2. Resilient monitoring and control of the networked control system infrastructure
2. Science of decentralized security which aims to develop a framework that will enable reasoning about the security of all the integrated constituent CPS components
3. Reliable and practical reasoning about secure computation and communication in networks which aims to contribute a formal framework for reasoning about security in CPS
4. Evaluation and experimentation using modeling and simulation integration of cyber and physical platforms that directly interface with human decision makers.
5. Education and outreach

Impact

- Equip CPS designers and operators with foundations and theory-based comprehensive tools improve resilience against faults and intrusions
- Enable designers to take security decisions and allocate resources in a decentralized manner
- Enable experimentation, evaluation, and training using a modeling and simulation integration platform

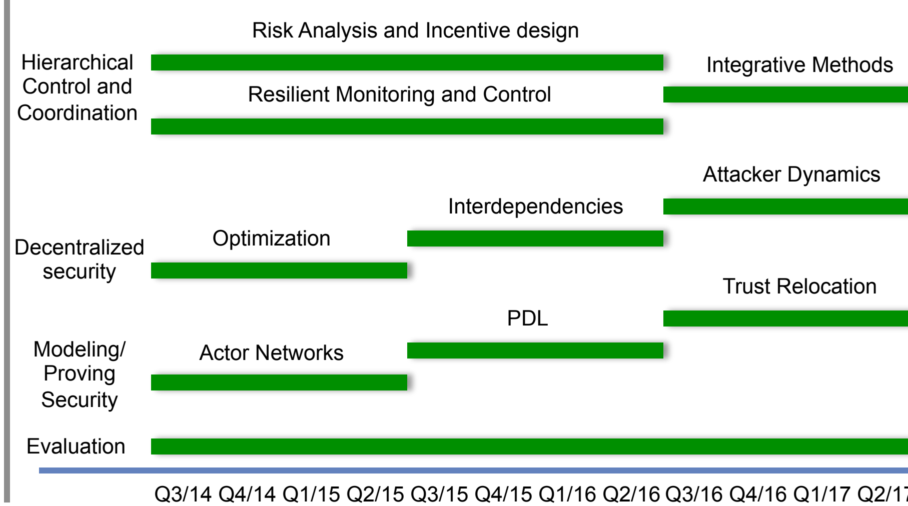

Berkeley UNIVERSITY OF CALIFORNIA

Massachusetts Institute of Technology

INFORMATION & COMPUTER SCIENCES UNIVERSITY of HAWAI'I at MĀNOA

VANDERBILT UNIVERSITY