FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Resilient Monitoring and Control of Distributed Cyber-Physical Systems

Xenofon Koutsoukos

Waseem Abbas, Sajal Bhatia, Anirban Bhattacharjee, Aron Laszka, Goncalo Martins, Abhishek Dubey

Gabor Karsai, Janos Sztipanovits, Yevgeniy Vorobeychic

Vanderbilt University/ISIS

Berkeley UNIVERSITY OF CALIFORNIA
Massachusetts Institute of Technology
NSF
MICHIGAN
VANDERBILT UNIVERSITY

Distributed Parameter Estimation

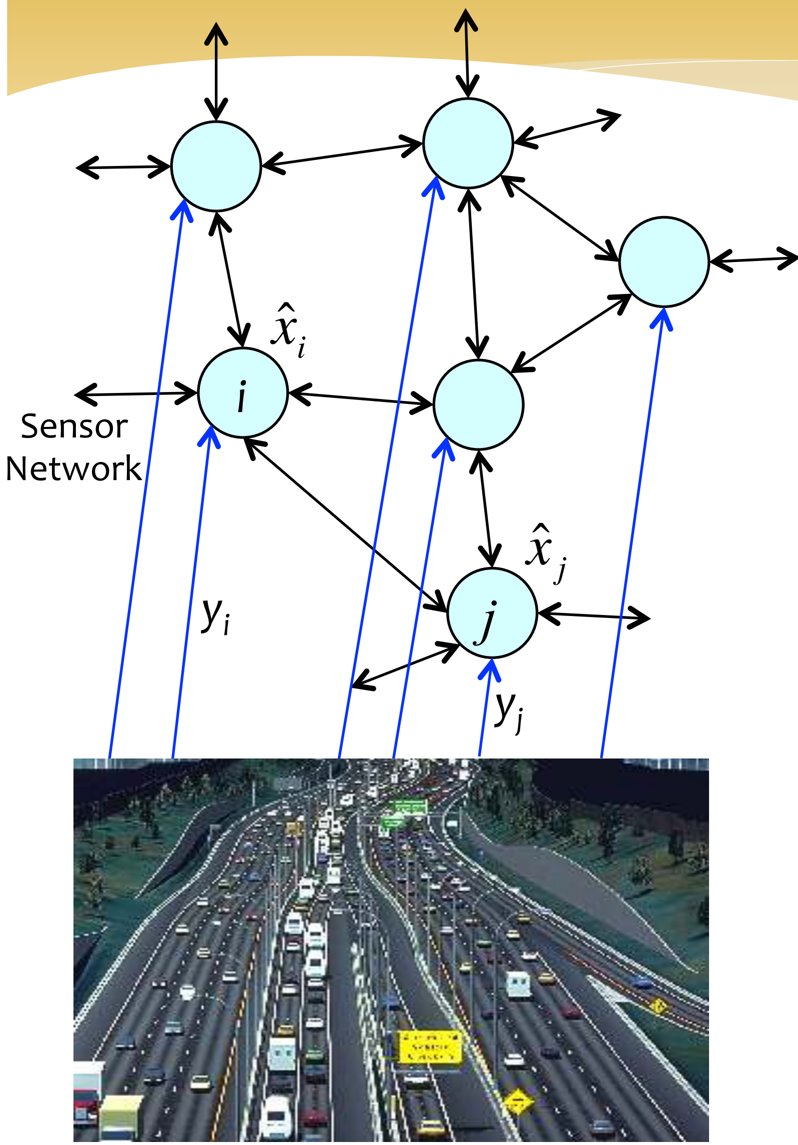

- All sensors measure independently some physical phenomenon with some error due to noise
  $y_i = \theta + v_i, v_i \sim N(0, \sigma_i^2), i = 1, 2, \ldots, n$
- The sensors improve their estimate by averaging the measurements
- Minimum variance estimate

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

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

Distributed Control of Multi-Agent Systems

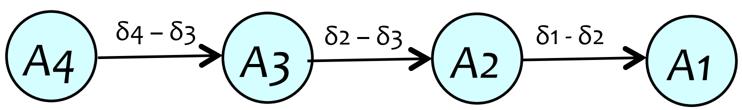

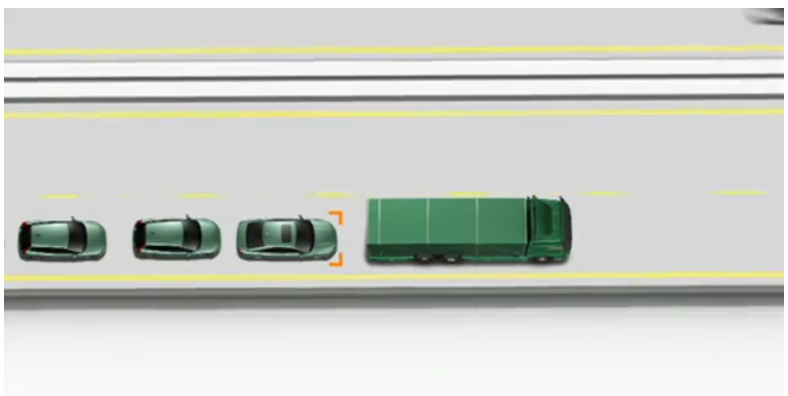

Consensus-based Formation Control

* Distributed consensus
  * Each vehicle updates its state based on the states of its local neighbors
  * The final state of each vehicle converges to a common value
* Distributed Consensus Applications in CPS
  * Vehicle rendezvous
  * Formation control
  * Parameter estimation
  * Least squares data regression
  * Sensor calibration
  * Time synchronization
  * Kalman filtering

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Resilient Consensus in the Presence of Adversaries

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

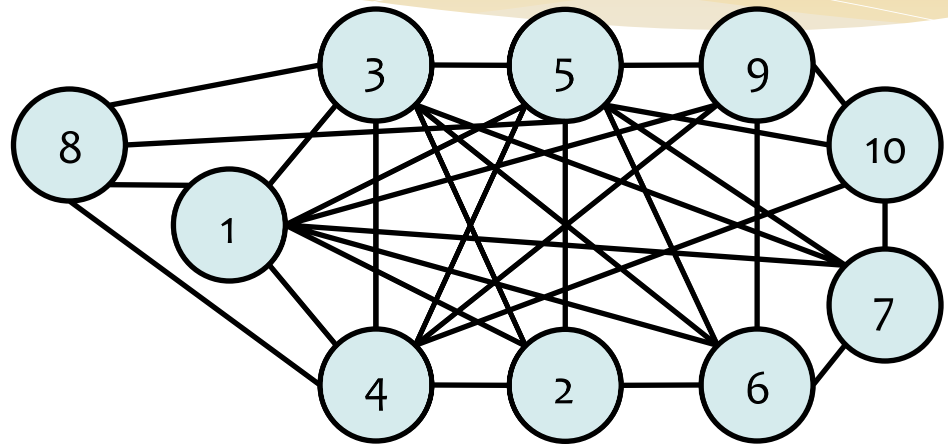

* Adversarial Consensus Protocol
* Adversary models
  * Threat
  * Scope
* Robust network topologies
  * Local redundancy
* Resilience requires high degree of redundancy
* Can we relax the redundancy requirements?

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Overview

* Performance Impact of Authentication in Time-Triggered Networked Control Systems
  * Theoretical analysis of performance impact
  * Experimental validation
* Resilient Consensus Protocols with Trusted Nodes
  * Connected Dominating Set
  * Trusted Nodes and Network Robustness
* Stochastic Message Authentication
  * Game Theoretic Model
  * Trade-off Between Computation and Security
* Conclusions

Overview

* Performance Impact of Authentication in Time-Triggered Networked Control Systems
  * Theoretical analysis of performance impact
  * Experimental validation
* Resilient Consensus Protocols with Trusted Nodes
  * Connected Dominating Set
  * Trusted Nodes and Network Robustness
* Stochastic Message Authentication
  * Game Theoretic Model
  * Trade-off Between Computation and Security
* Conclusions

Time-Triggered Ethernet Overview

Integrated Dataflow Example

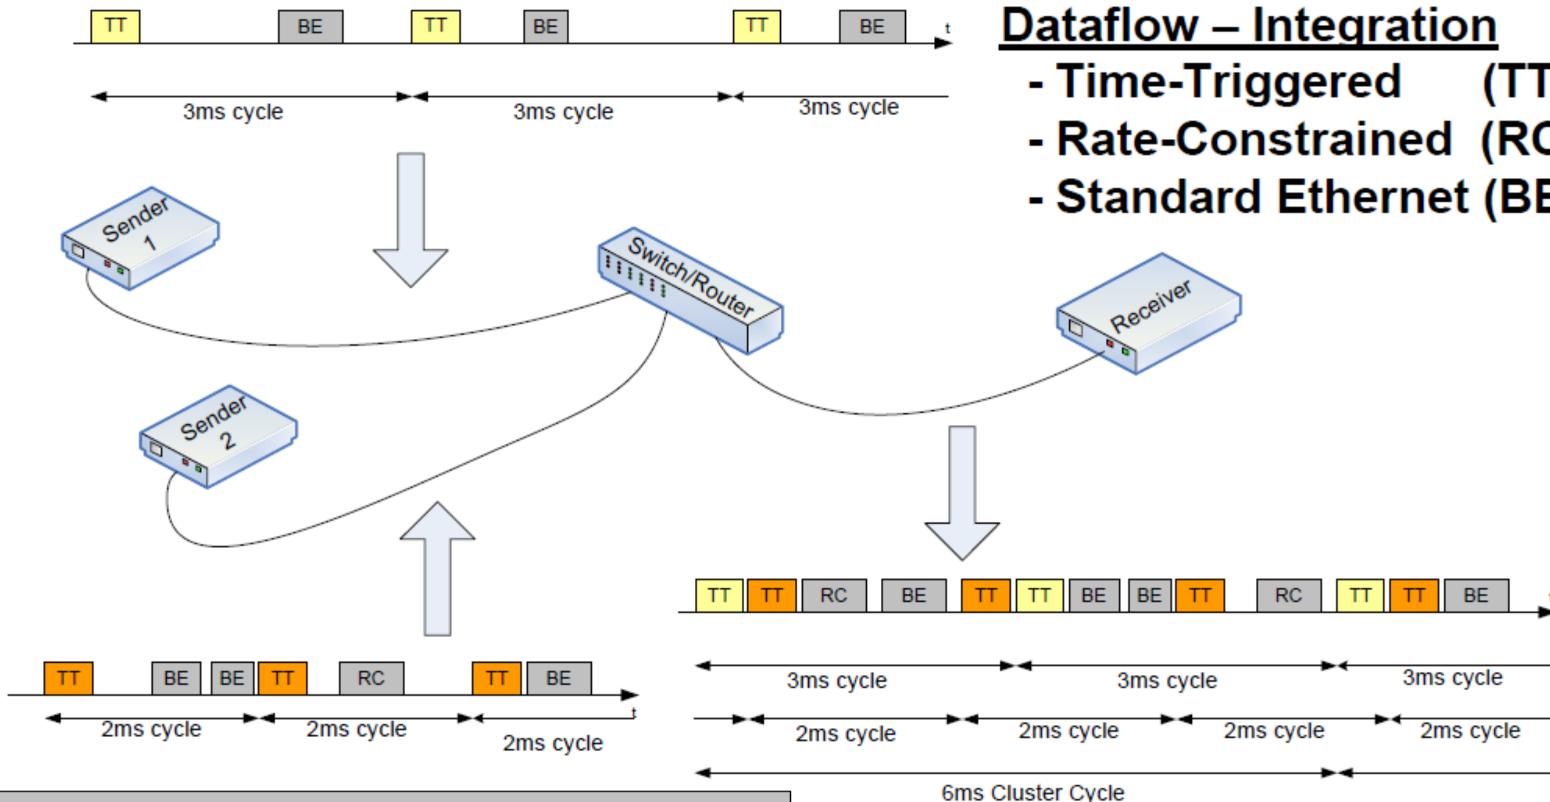

Dataflow – Integration

- **Time-Triggered (TT)**
- **Rate-Constrained (RC)**
- **Standard Ethernet (BE)**

In parallel, two types of Ethernet communications:

- **Synchronous (TDMA-style) Communication: TT**
- **Asynchronous (event-triggered style): RC + BE**

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Protected Message Transmission

Hash-based message authentication code

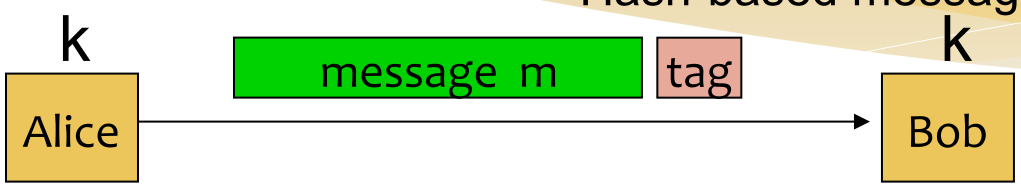

Generate tag:

tag ← S(k, m)

Verify tag:

V(k, m, tag) = `yes'

MAC I = (S,V) defined over (K,M,T) is a pair of algorithms:

- S(k,m) outputs t in T
- V(k,m,t) outputs `yes' or `no'

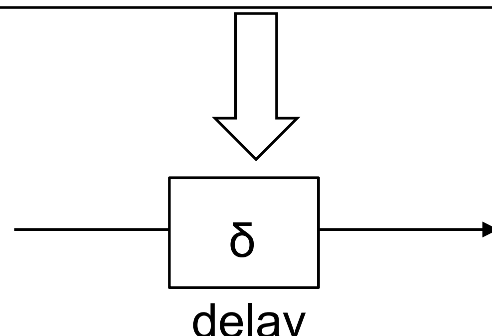

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Analysis of Performance Impact

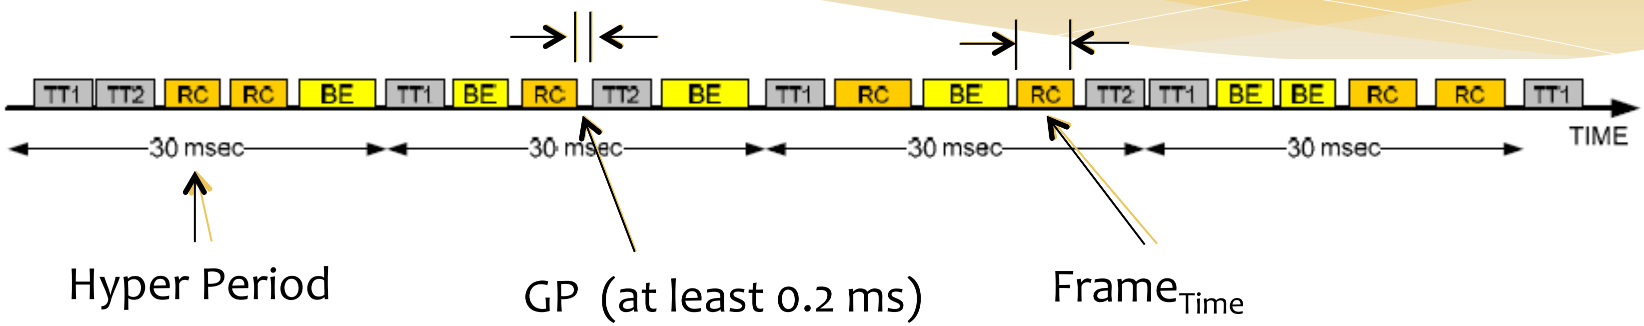

- Max Number of Frames per Hyper Period (NF_Max)

$$NF_{Max} = \frac{HP}{\left(Frame_{Time} + GP\right)} \qquad Frame_{Time} = \frac{Packet_{Size}}{Transmission_{Rate}}$$

- Example:

$$NF_{Max} = \frac{10\,(ms)}{\left(Frame_{Time} + 0.2\,(ms)\right)} \cong 48 \qquad Frame_{Time} = \frac{60\,(bytes)}{\left[\left(100\,(Mbits)/8\right)\right]\left(bytes/s\right)}$$

Effect of Packet Size

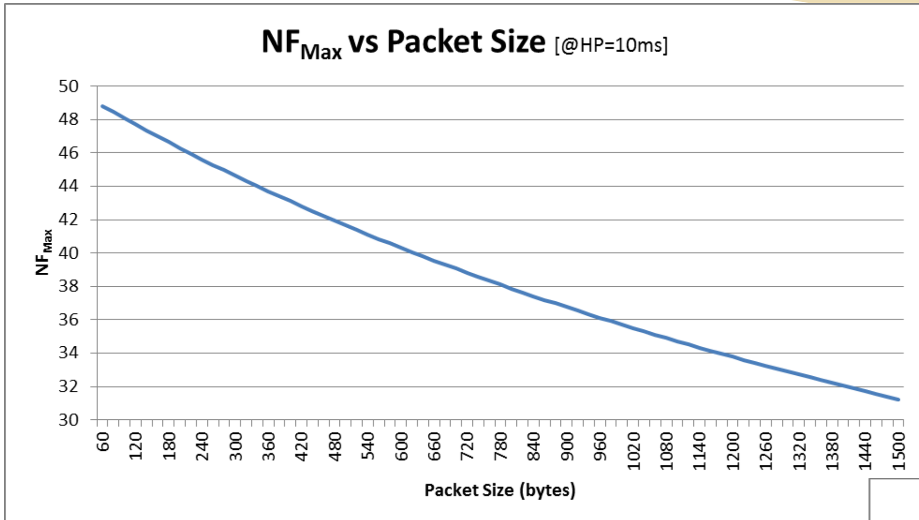

Guard Period = 0.2 ms

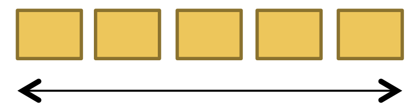

HP = 10 ms

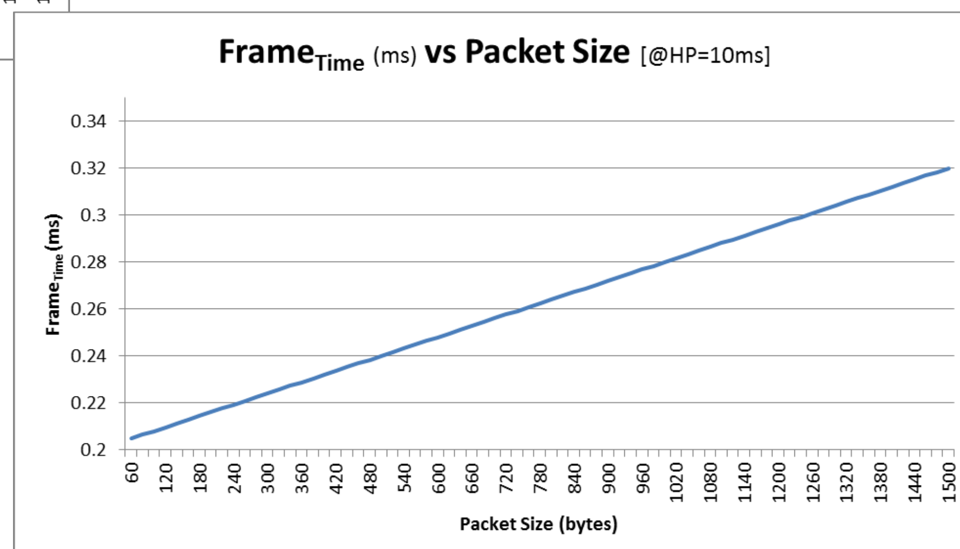

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Hardware Platform: IBX-530W

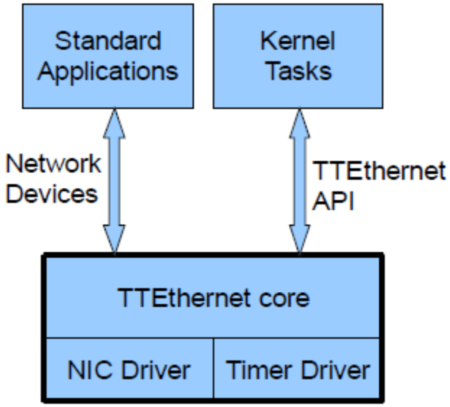

Figure 1: Structure of a TTEthernet end system

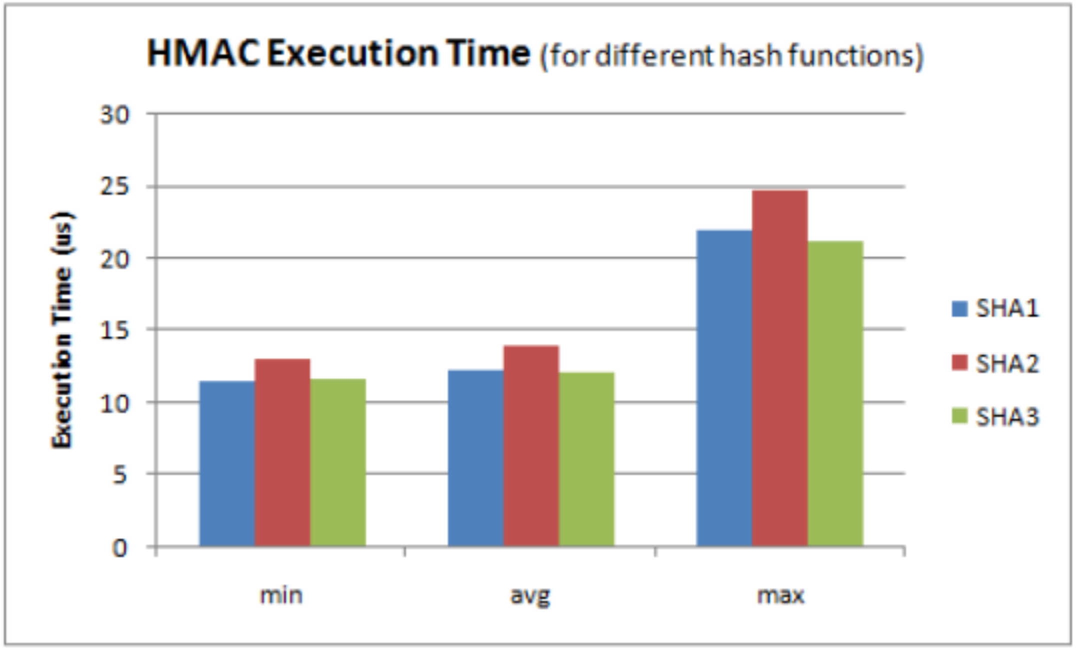

* Intel Atom Processor, 1.6GHz
* Linux 2.6.24-24-rt kernel
* Crypto library

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Impact on System Performance

- Original Physical Setup

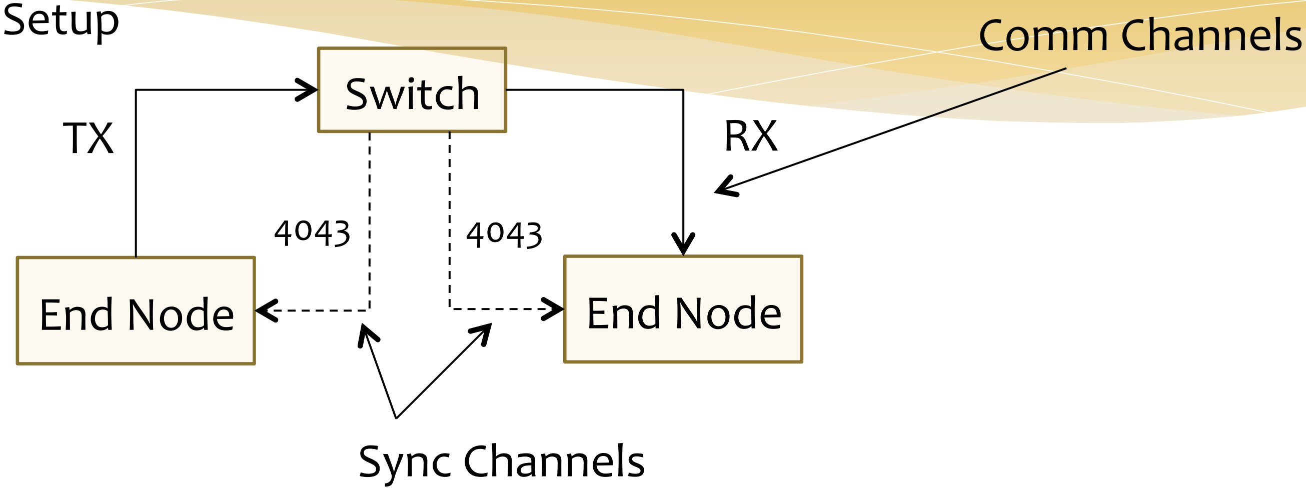

- WireShark (WS) Physical Setup

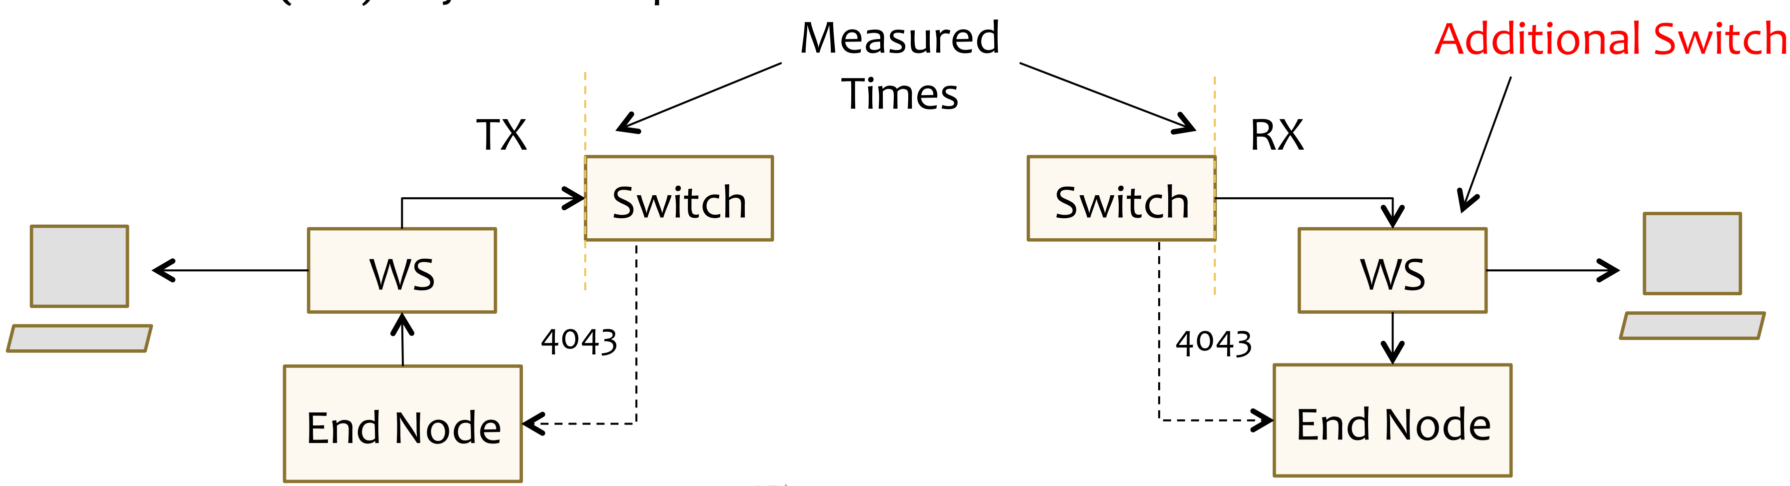

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Experimental Results (WireShark)

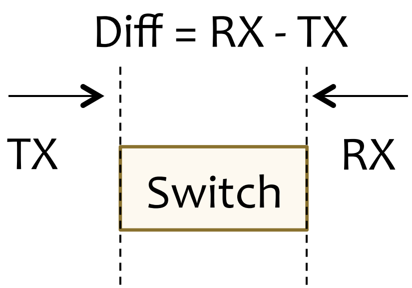

	60 bytes			1514 bytes		
	Min	Avg	Max	Min	Avg	Max
Tx (ms)	0	0.008	0.115	0.110	0.222	0.347
Rx (ms)	0.274	0.386	0.673	0.494	0.594	0.826
Diff	0.274	**0.378**	0.558	0.384	**0.372**	0.479

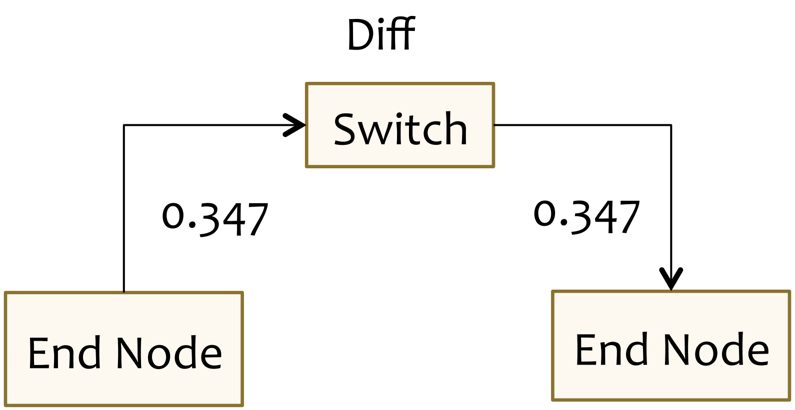

- Max Total Transmission Time (Max_{TTT})

$$Max_{TTT} = (2 * Frame_{Time}) + Diff$$

(0.347 * 2) + 0.375 = 1.069 ms

With WirelessShark Switch

(0.347 * 2) + 0.2 = 0.894 ms

With assumed GP

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Comparison

- TTTech – TDMA Theoretical Results

Guard Period = 0.2 ms

Theoretical Values			
	60 Bytes	80 Bytes	1514 Bytes
NF_{Max}	48	48	31
$Frame_{Time}$ (ms)	0.0048	0.0064	0.12
Max_{TTT} (ms)	0.2096	0.2128	0.44

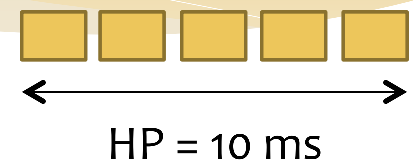

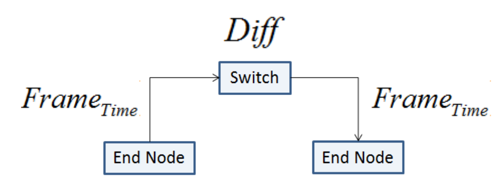

- TTTech – TDMA Hardware Results

$$Max_{TTT} = (2 * Frame_{Time}) + Diff$$

GP

Hardware Values			
	60 Bytes	80 Bytes	1514 Bytes
NF_{Max}	23	20	11
$Frame_{Time}$ (Tx_{Max}) (ms)	0.115	0.150	0.347
Max_{TTT} (ms)	0.43	0.5	0.894

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Experimental Analysis: Conclusions

* The overhead time introduced by the kernel module implementing HMAC reduces the effective number of frames per hyper-period (HP)
* There is a small impact on the maximum number of frames per HP by increasing the packet size from 60 to 80 bytes (tag)
* Experimental results are consistent with the theoretical analysis
  * Overhead time spent by the kernel module to transmit data to the physical medium is not considered by the theoretical analysis

[Martins et al., ISRCS 2014, Submitted]

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Overview

* Performance Impact of Authentication in Time-Triggered Networked Control Systems
  * Theoretical analysis of performance impact
  * Experimental validation
* Resilient Consensus Protocols with Trusted Nodes
  * Connected Dominating Set
  * Trusted Nodes and Network Robustness
* Stochastic Message Authentication
  * Game Theoretic Model
  * Trade-off Between Computation and Security
* Conclusions

Resilient Consensus Protocol 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:

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

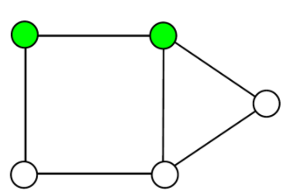

Connected Dominating Set:

Nodes in the dominating set induce a connected subgraph

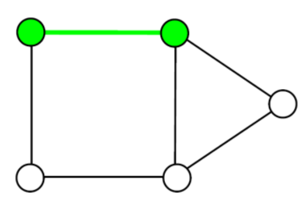

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Resilient Consensus Protocol 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)

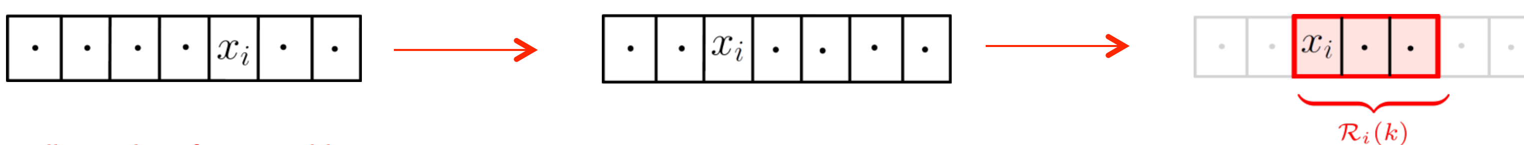

Collect values from neighbors

Sort

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

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

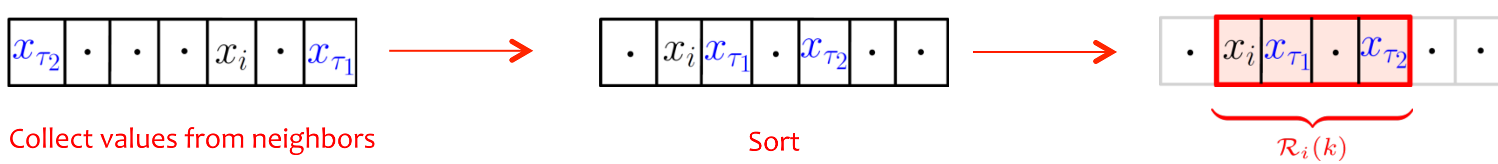

Collect values from neighbors
(x_{τ_1}, x_{τ_2} are trustworthy nodes' values)

Sort

Examples

Example 1: (Tree – Sparse network)

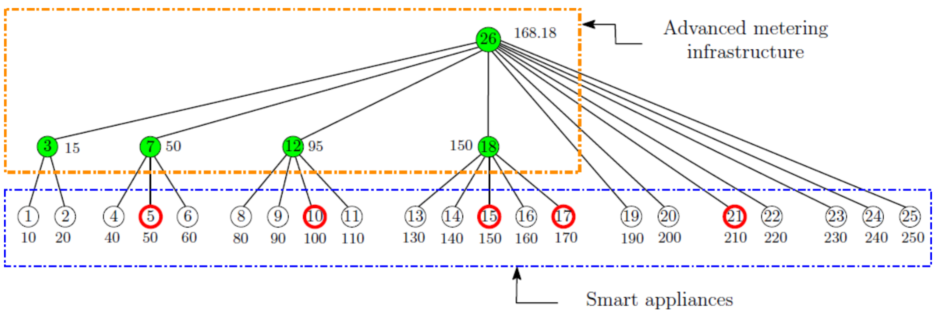

Example 2: (2,2) Robust graph

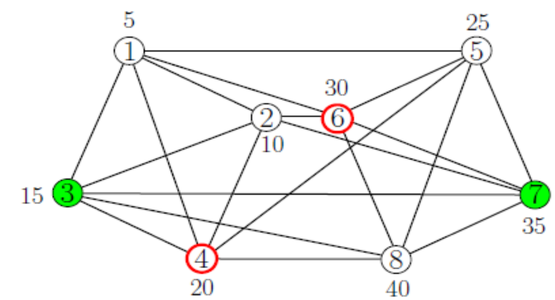

RCP-T

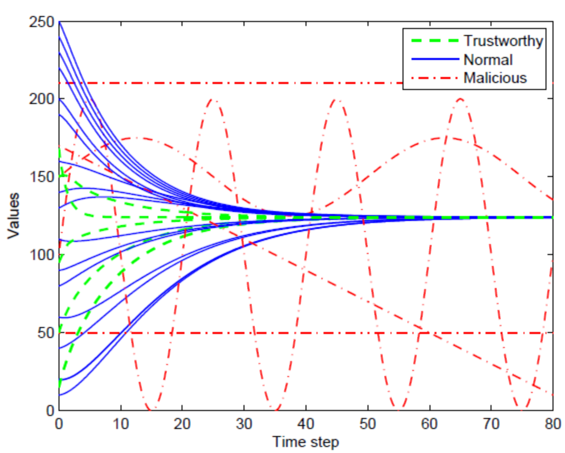

RCP-T

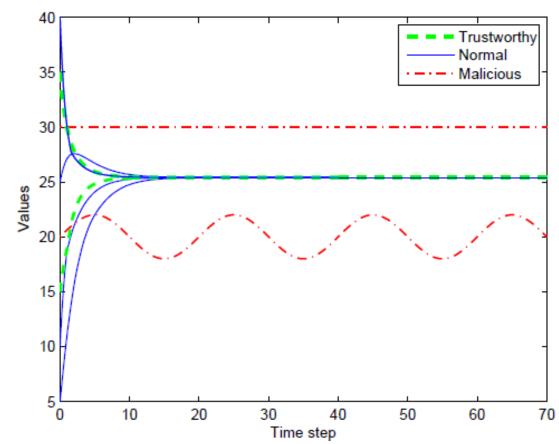

[Abbas et al., ISRCS 2014, Submitted]

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Overview

* Performance Impact of Authentication in Time-Triggered Networked Control Systems
  * Theoretical analysis of performance impact
  * Experimental validation
* Resilient Consensus Protocols with Trusted Nodes
  * Connected Dominating Set
  * Trusted Nodes and Network Robustness
* Stochastic Message Authentication
  * Game Theoretic Model
  * Trade-off Between Computation and Security
* Conclusions

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Motivation

* Computational demand of cryptographic primitives can be too high for **resource-bounded** devices
  * legacy devices in supervisory control systems
  * embedded or battery-powered devices (RFID tags, sensors)
* "Lightweight" cryptographic primitives
  * Decision to secure a system is still **binary**: either security is employed, incurring some fixed overhead, or it is not
* Our approach: General-purpose framework for trading off security and computational resources using an existing MAC scheme

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Stochastic Message Authentication

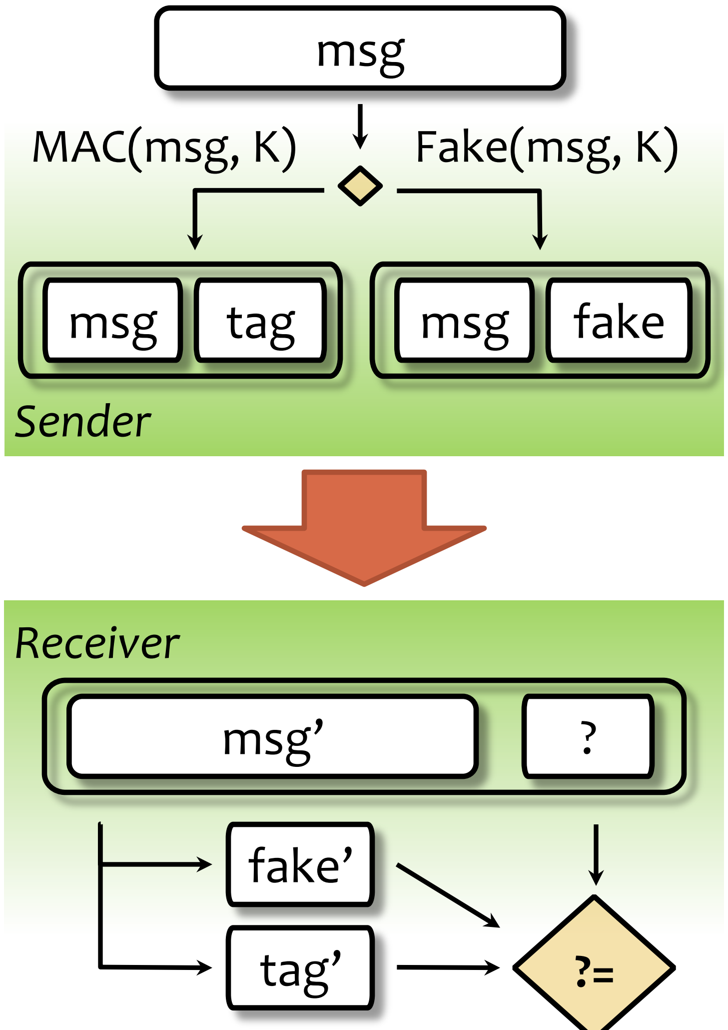

* For some messages, the sender computes a "fake tag", which is computationally less demanding, but does not protect integrity
* Adversary cannot distinguish fake tags from correct tags
* Receiver can verify if a message has a fake or a correct tag efficiently
  → detect attacks with high probability

Game-Theoretic Model

* Stackelberg security game
  * Divide messages into C classes based on their potential to cause damage

		Defender	Attacker
Strategy choice		*for each class c, the probability of authentication p_c*	*for each class c, the number of modified / inserted messages a_c*
Detection probability		$1 - \prod(1 - p_c)^{a_c}$	
Payoff	attack undetected	*- (amount of total damage $\sum a_c L_c$)*	*amount of total damage $\sum a_c L_c$*
	attack detected	*zero*	*"punishment" $-F$*

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Results

Trade-off between computation and security

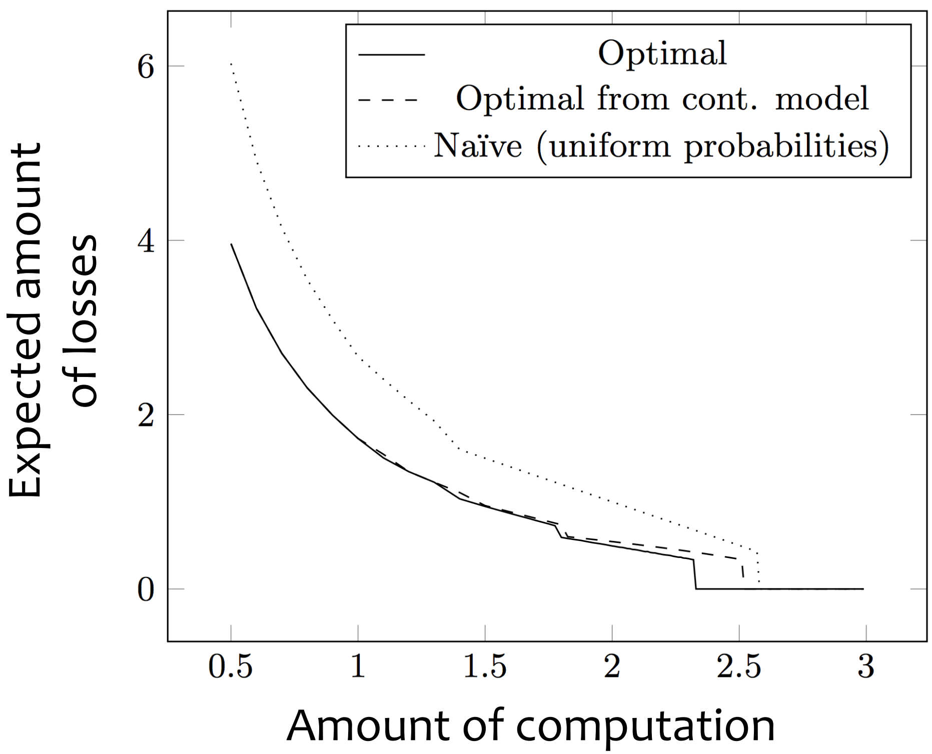

Proof-of-concept implementation using SHA-1 HMAC on an ATmega328P microcontroller

[Laszka et al., CCS 2014, Submitted]

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Conclusions

* Resilient Distributed Consensus Protocols in the Presence of Adversaries
  * Exploit local information redundancy
* Performance Impact of Authentication Mechanisms
  * Theoretical analysis and experimental validation
* Resilient Distributed Consensus Protocols with Trusted Nodes
  * Trusted nodes form a connected dominating set
* Stochastic Message Authentication
  * Trade-off between computation and security