



Resilient Monitoring and Control of Distributed Cyber-Physical Systems

Xenofon Koutsoukos

Waseem Abbas, Sajal Bhatia, Anirban Bhattacharjee, Arul Moondra, Aron Laszka, Goncalo Martins

Gabor Karsai, Janos Sztipanovits, Yevgeniy Vorobeychic

Vanderbilt University/ISIS



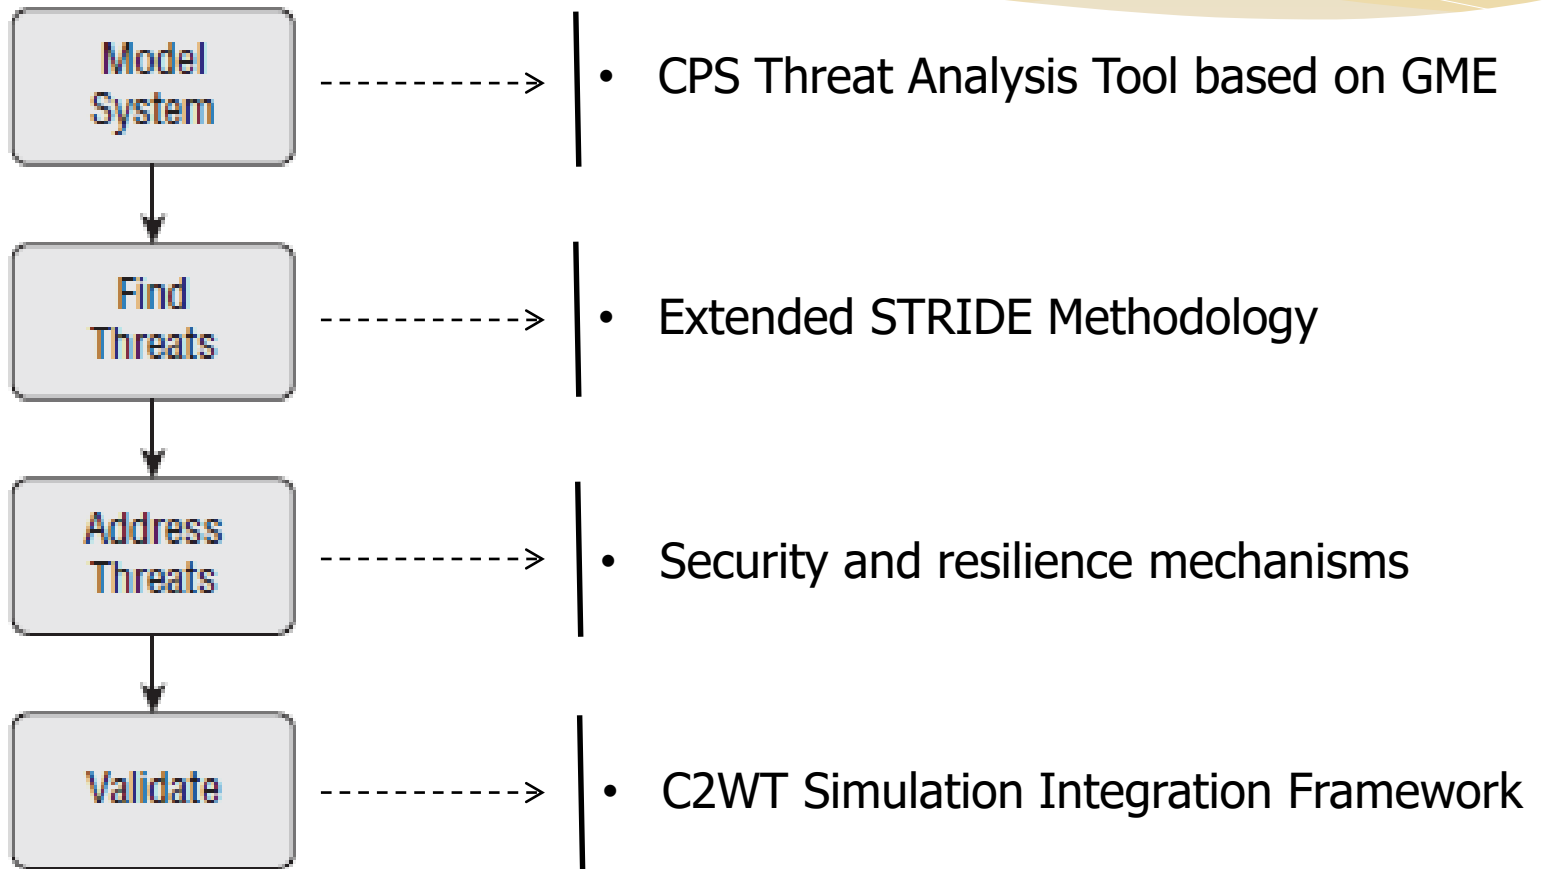
Overview

- * Threat Modeling for CPS Security
- * Performance Impact of Authentication in Time-Triggered Networked Control Systems
- * Resilient Consensus Protocols with Trusted Nodes
- * Resilient Observation Selection
- * Conclusions

Overview

- * Threat Modeling for CPS Security
- * Performance Impact of Authentication in Time-Triggered Networked Control Systems
- * Resilient Consensus Protocols with Trusted Nodes
- * Resilient Observation Selection
- * Conclusions

Threat Modeling for CPS Security



STRIDE

- Threat
- Mitigation
- Technology Summary

Spoofing

Authentication

To authenticate principals:

- Cookie authentication
- Kerberos authentication
- PKI systems such as SSL/TLS and certificates

To authenticate code or data:

- Digital signatures

Tampering

Integrity

- Windows Vista Mandatory Integrity Controls
- ACLs
- Digital signatures

Repudiation

Non Repudiation

- Secure logging and auditing
- Digital Signatures

Information Disclosure

Confidentiality

- Encryption
- ACLS

Denial of Service

Availability

- ACLs
- Filtering
- Quotas

Elevation of Privilege

Authorization

- ACLs
- Group or role membership
- Privilege ownership
- Input validation

Threat Modeling for CPS Security

■ Address Threats

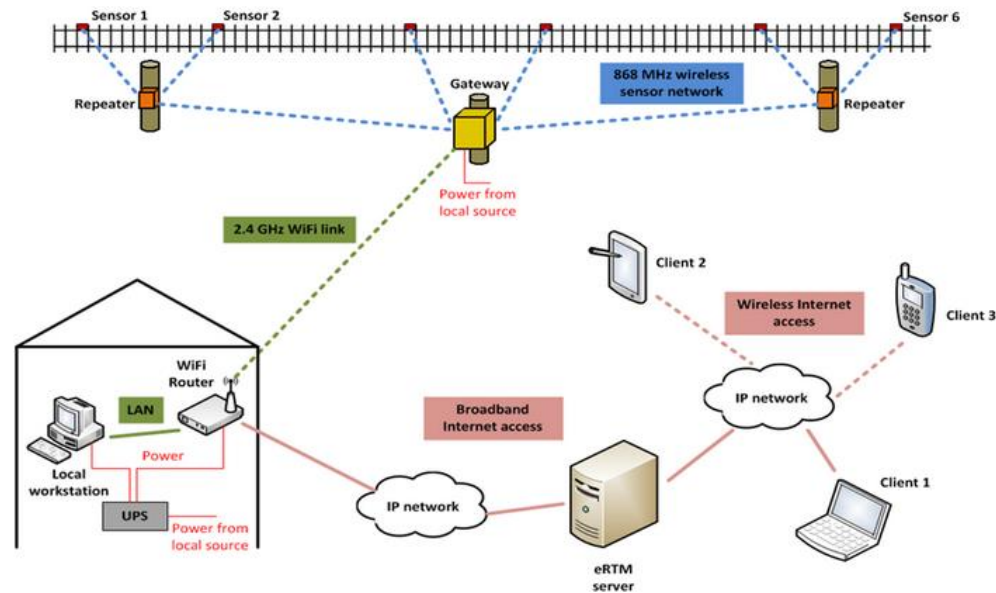


Mitigation Strategies:

- **Do nothing**
(for example, hoping for the best or not applicable)
- **Inform about the risk**
(for example, warning user population about the risk)
- **Mitigate the risk**
(for example, by putting countermeasures in place)
- **Accept the risk**
(for example, after evaluating the impact of the exploitation)
- **Transfer the risk**
(for example, through contractual agreements and insurance)
- **Terminate the risk**
(for example, shutdown, turn-off, unplug or decommission the asset)

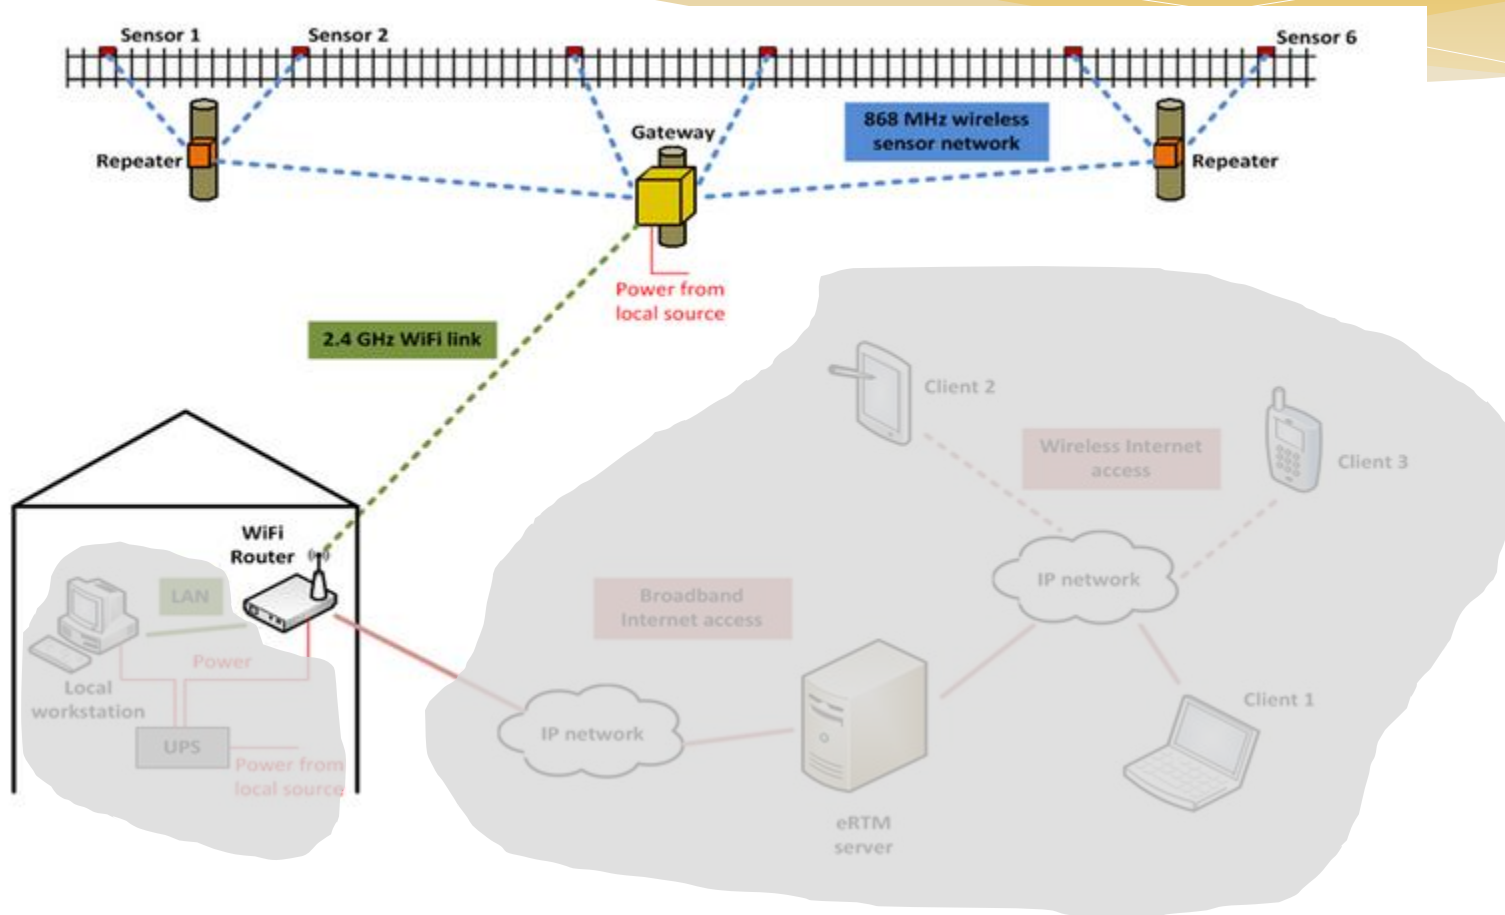
Case Study: eRTM – Railway Monitoring

- * Battery powered temperature measuring modules are connected via wireless
- * The communication is organized via repeater and gateway units
- * The gateway units collect the data of the network and transmit to the central processing server
- * Monitoring data records are accessed by browser and smartphone applications
- * Alarm messages are sent to specified clients according to temperature limit settings



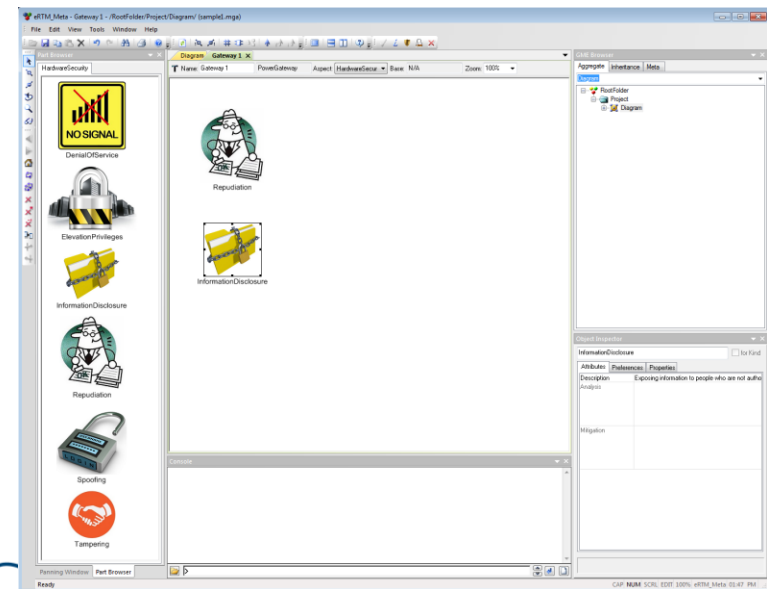
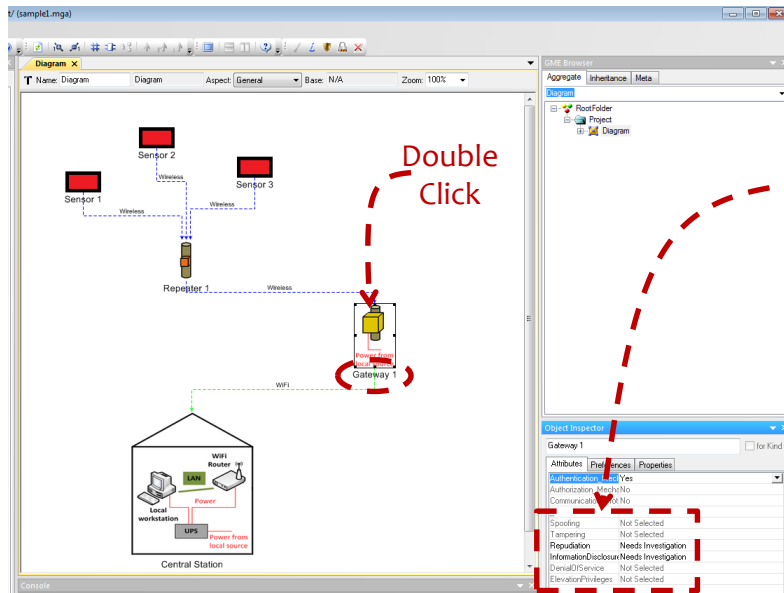
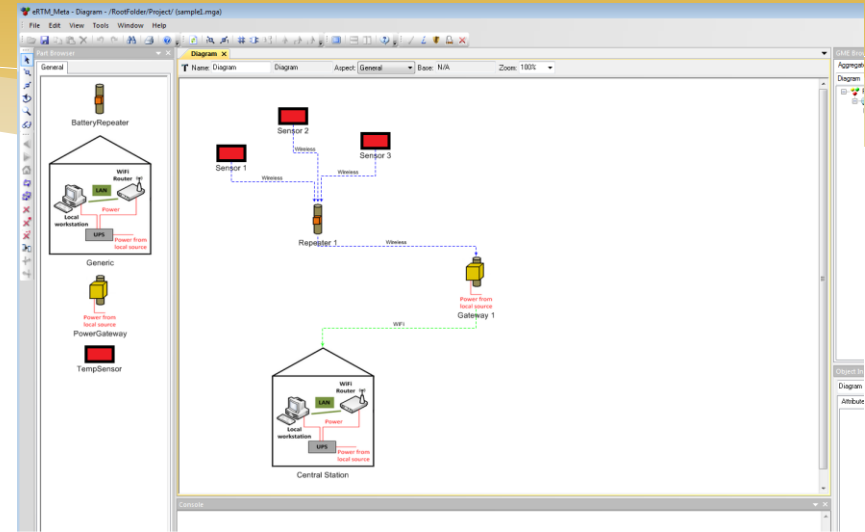
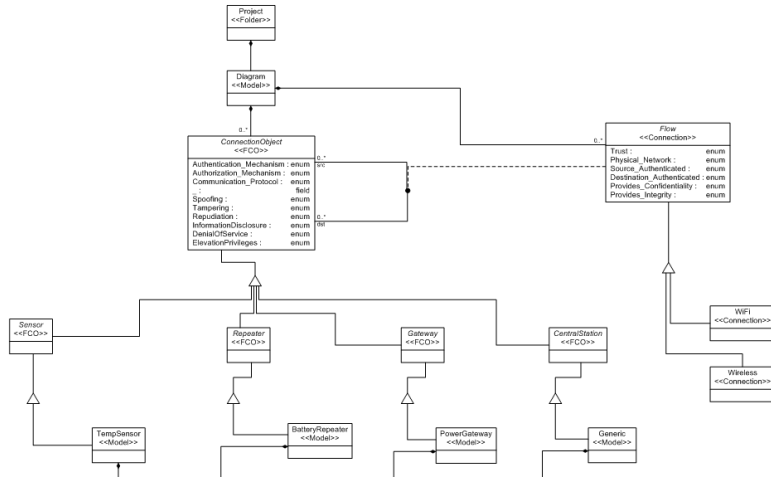
- * Monitoring rail temperature distribution
 - * Prediction of buckling
 - * Measurement-based control of speed limits
 - * Measurement-based control of switch heating

eRTM – System Architecture



System under consideration – CPS section of eRTM (excluding the IP network part)

Meta-Model and eRTM Model



Threat Modeling and Reporting

eRTM Threat Modeling Report
(Software Security)

Connection: WiFi
Physical Network: 2.4 GHz
Source: Gateway 1
Destination: Central Station

1.1 Spoofing [State: Not Started] [Priority: High]

Threat: Spoofing the Central Station
Description: Central Station may be spoofed by an attacker and this may lead to information disclosure by Gateway 1. Consider using a standard authentication mechanism to identify the destination process.
Mitigation: <no mitigation provided>

1.2 Spoofing [State: Not Started] [Priority: High]

Threat: Spoofing the Gateway 1
Description: Gateway 1 may be spoofed by an attacker and this may lead to information disclosure by Central Station. Consider using a standard authentication mechanism to identify the destination process.
Mitigation: <no mitigation provided>

1.3 Tampering [State: Not Started] [Priority: High]

Threat: Potential Lack of Input Validation for Central Station
Description: Data flowing across WiFi may be tampered with by an attacker. This may lead to a denial of service attack or an elevation of privilege attack against Central Station or an information disclosure by Central Station. Failure to verify that input is as expected a root cause of a very larger number of exploitable issues. Consider all paths and the way they handle data. Verify that all inputs are verified for correctness using an approved list input validation approach.
Mitigation: <no mitigation provided>

1.4 Repudiation [State: Not Started] [Priority: High]

Threat: Potential Data Repudiation Central Station
Description: Central Station claims that it did not receive data from a source outside the trust boundary. Consider using logging or auditing to record the source, time, and summary of the received data.
Mitigation: <no mitigation provided>

1.5 Information Disclosure [State: Not Started] [Priority: High]

Threat: Data Flow Sniffing
Description: Data flowing across WiFi may be sniffed by an attacker. Depending on what type of data an attacker can read, it may be used to attack other parts of the system or simply be a disclosure of information leading to compliance violations. Consider encrypting the data flow

Summary and Future Work

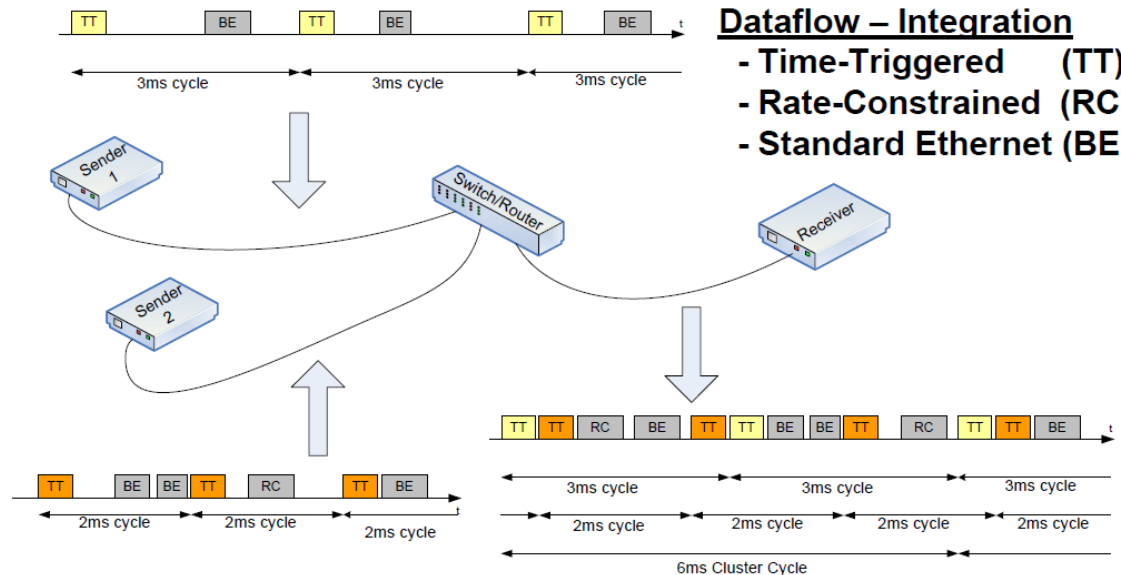
- * Threat modeling for CPS security
 - * Modeling language based on STRIDE
 - * Threats analysis & report generation (under development)
 - * Track threats visually in the model (under development)
- * Future work
 - * Modeling of security and resilience mechanisms for addressing threats
 - * Secure information flow policies
 - * Integration with simulation tools for evaluation and validation

Overview

- * Threat Modeling for CPS Security
- * Performance Impact of Authentication in Time-Triggered Networked Control Systems
- * Resilient Consensus Protocols with Trusted Nodes
- * Resilient Observation Selection
- * Conclusions

Performance Evaluation of HMAC in Time-Triggered Ethernet

Integrated Dataflow Example



- * 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

	Theoretical Values			Hardware Values			
	60 Bytes	80 Bytes	1514 Bytes	60 Bytes	80 Bytes	1514 Bytes	
NF_{Max}	48	48	31	NF_{Max}	23	20	11
$Frame_{Time}$ (ms)	0.0048	0.0064	0.12	$Frame_{Time}(Tx_{Max})$ (ms)	0.115	0.150	0.347
Max_{TT} (ms)	0.2096	0.2128	0.44	Max_{TT} (ms)	0.43	0.5	0.894


Scalability Analysis

- What to do if the application requires more nodes ?

- Solutions:

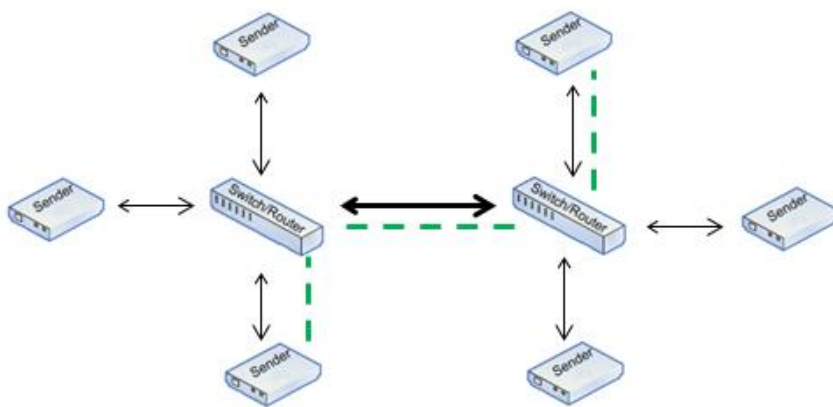
- Get a switch that has enough ports for the required number of nodes;

Impractical



- Connect additional switches in a cascade topology;

Analyze the implications



Scalability Analysis

- Connect additional switches in a cascade topology
- Guard period added, to the TDMA frame, per each additional switch

$$NF_{max} = \frac{BP}{(Frame_{time}) + (GP * N_{Switch})}$$

- HMAC can be implemented with no impact on NF_{max} and on the number of end nodes that can transmit in the same TDMA frame;
- However, after adding the 3rd switch the number of nodes that can be connected to the switch is bigger than the actual number of messages that can be transmitted per BP;

Table 2: Platform A - Scalability Theoretical Results [@BP = 10 ms]

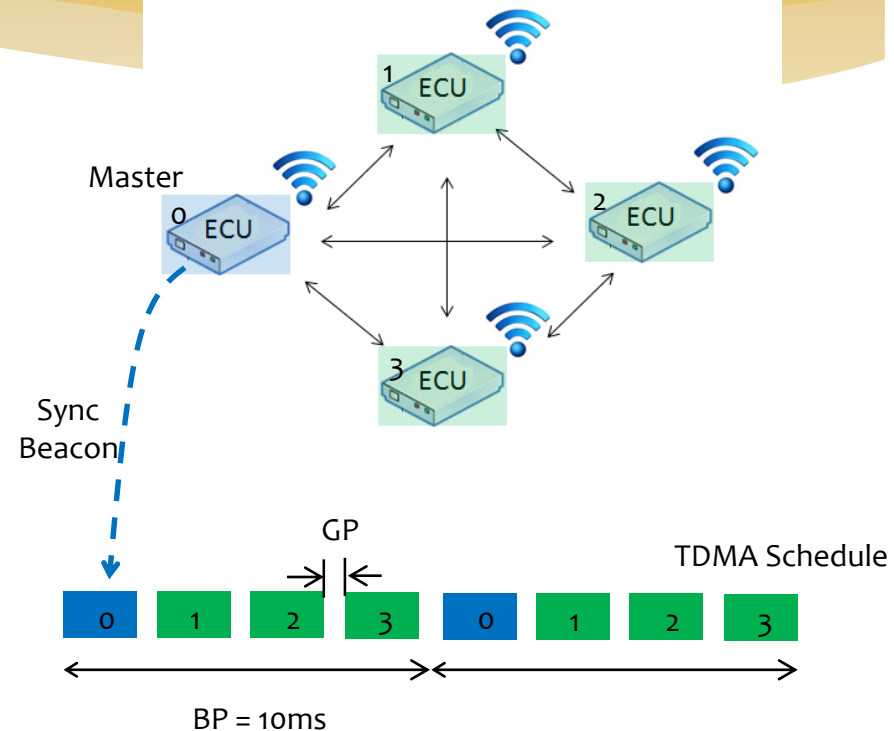
GP = 0.2 ms					
60 bytes			80 bytes		
N_{switch}	NF_{max}	# Nodes	N_{switch}	NF_{max}	# Nodes
1	48	8	1	48	8
2	24	14	2	24	14
3	16	20	3	16	20
4	12	26	4	12	26
5	9	32	5	9	32

Wireless Time Triggered Networks

- * Mesh Network
- * Nodes connected in Ad Hoc mode using WiFi 802.11g at 54 Mbps

TTA Wireless Implementation

- Each node has a unique identifier (ID)
- The node with ID = 0 is the Master
- All packets are broadcast
- Master node sends the sync beacon every 10 ms (BP)
- The TDMA schedule is defined off line
- Remaining nodes adjust their clocks upon receiving the sync beacon and send the respective packet at their designated slot




$$Frame_{Time} = \frac{Packet_{Size}}{Transmission_{Rate}}$$

$$NF_{Max} = \frac{BP}{(Frame_{Time} + GP)}$$

Overhead of HMAC Tag


Theoretical Values		
	60 Bytes (without tag)	80 Bytes (with tag)
NF_{Max}	1125	843
Frame _{Time} (ms)	0.0089	0.0119

No Guard Period

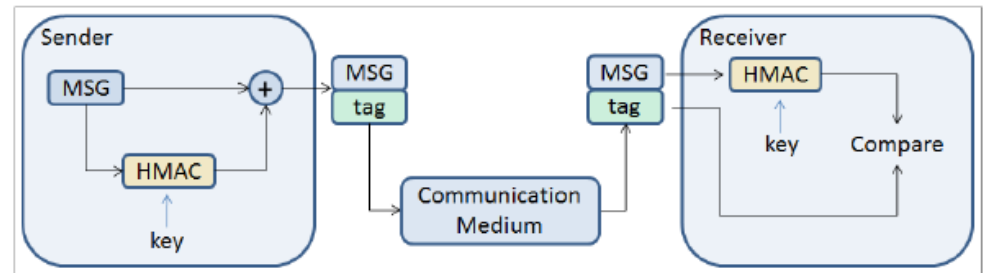


Theoretical Values		
	60 Bytes (without tag)	80 Bytes (with tag)
NF_{Max}	47	47
Frame _{Time} + GP (ms)	0.2089	0.2119

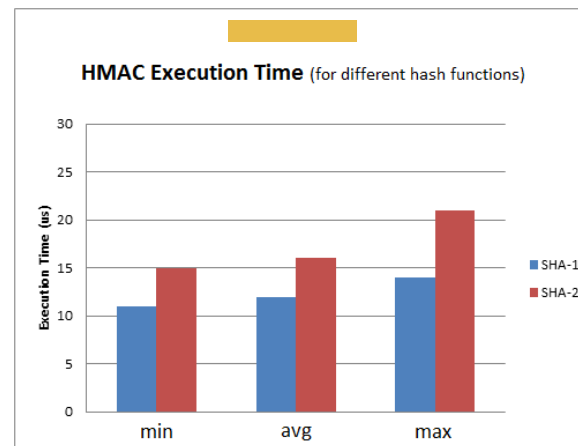
Guard Period = 0.2 ms



Hardware Platform: Trimslice




- * Dual Core ARM Cortex-A9
- * Linux kernel 3.1.10-14.r16.02
- * Crypto library



Average Execution time:
(SHA-1) 12 us
(SHA-2) 16 us

Propagation Time Measurements

$$Frame_{time} = Packet\ Transmission\ Time + Send\ Function\ Time$$



Send Function - Network Device Driver Time

Packet Size (bytes)	Min (us)	Max (us)
60	208	240
80	208	240
1280	240	320
1300	240	320

Packet Transmission Time

Packet Size (bytes)	Min (us)	Max (us)
60	432	584
80	432	592
1280	696	888
1300	700	896

Negligible
Time (1us)

Impact on System Performance

Performance Impact Results

	Theoretical		Measured	
	60 bytes	80 bytes	60 bytes	80 bytes
NF_{max}	47	47	9	9
Max_{TT} (ms)	0.2089	0.2119	1.024	1.032

- * Decrease of the number of frames: NF_{max}
 - * Theoretical values do not include the send function time;
 - * External Factors: Electrical interference, other wireless networks;
 - * In reality the WiFi throughput is less than 54 Mbps; typically mid-20 Mbps

Conclusions

- * The overhead time introduced by the kernel module implementing HMAC reduces the effective number of frames per base-period (BP)
- * Wireless Time Triggered Networks
 - * There is a significant impact on the maximum number of frames per BP between the theoretical and practical values;
 - * WiFi throughput and external factors are nondeterministic
 - * There is no impact by increasing the packet size from 60 to 80 bytes (tag)



Overview

- * Threat Modeling for CPS Security
- * Performance Impact of Authentication in Time-Triggered Networked Control Systems
- * **Resilient Consensus Protocols with Trusted Nodes**
- * Resilient Observation Selection
- * Conclusions

Resilient Distributed Consensus

Consensus
(All nodes behave normally)

No Consensus
(Malicious nodes)



Resilient Consensus

Agents **agree** upon at a common value even in the presence of some **malicious** nodes or **adversaries**.

Resilient Consensus with Trusted Nodes

Resilient Consensus
(Various approaches to resilience)

Connectivity
of underlying graph

Node's capabilities
(security levels)

Key idea

Add few **trusted nodes** in the network

Questions

- Can we **relax** connectivity based conditions for resilient consensus?
- By how much can we **improve** network resilience by adding few trusted nodes?

Some nodes might be more secured against adversarial attacks.

Trusted nodes

- Highly secure
- Very unlikely to be compromised

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:

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



Connected Dominating Set:

Nodes in the dominating set induce a **connected** subgraph



Graph Domination and Resilience

Connected Domination Number

γ_c = Connected domination number

- Widely studied in graph theory, sensor networks (backbone)
- If the number of trusted nodes is **at least** γ_c , the network can be made resilient against **any** number of attacks.

Question

By how much can we improve the resilience of networks if


No. of trusted nodes $< \gamma_c$

Observation

Interestingly, sometimes adding as many as $(\gamma_c - 1)$ trusted nodes does not improve the resilience.

Graph Domination and Resilience

Example



(2,2)-robust



Resilient against a **single** attack
(with no trusted nodes)


$\gamma_c = 4$



Resilient against **any**
number of attacks
(with 4 trusted nodes)

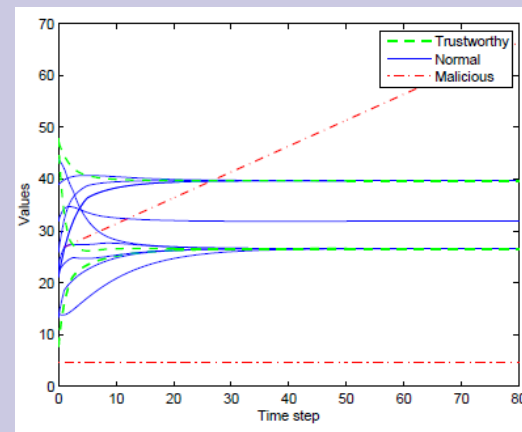
Observation

With any **three trusted** nodes, the graph is **not** resilient against **two adversaries**



Attacked nodes {3,8}

Trusted nodes {5,12,14}



Erdos-Renyi (ER) Graphs


$G_{n,p}$: n is the total no. of nodes

p is the probability of the existence of an edge between a pair of nodes.

Let F_l be the maximum no. of malicious nodes that can be present in the neighborhood of any node.

No trusted node: $F_l \leq \frac{np - \ln n}{\ln \ln n} \longrightarrow$ Consensus is achieved

Trusted nodes: Conn. dominating set \longrightarrow Consensus is achieved for **any** F_l



Example:

- For an ER graph with $n=200$, and $p=0.2$, consensus is achieved if $F_l \leq 11$ and **no trusted node**.
- However, consensus is achieved for **any** F_l whenever **5%** of the nodes are **trusted**.


Preferential Attachment (PA) Graphs

A new node is connected to m existing nodes with probability proportional to their degrees.

Let F_l be the maximum no. of malicious nodes that can be present in the neighborhood of any node.

No trusted node: $F_l \leq \frac{m-1}{2}$ \longrightarrow Consensus is achieved

Trusted nodes: Conn. dominating set \longrightarrow Consensus is achieved for **any** F_l




Example:

- For a PA graph with $m=5$, consensus is achieved if $F_l \leq 2$ and **no trusted node**.
- However, consensus is achieved for **any** F_l whenever **10%** of the nodes are **trusted**.

Resilient Fault Diagnosis in Flow Networks

Distributed Hypothesis Testing Using Belief Consensus



$$p(h_a | Z) = ap(h_a) \prod_{i=1}^n p(z(i) | h_a)$$

$$x(t+1) = (I - eL)x(t), \quad x_i(0) = \log(p(z(i) | h_a))$$

- Sensors may be compromised
- The system can be made more resilient using additional sensors
- Distributed hypothesis testing can be performed using resilient consensus algorithms
- What is the optimal sensor network design?
- What are the dependencies on the system model?

Overview

- * Threat Modeling for CPS Security
- * Performance Impact of Authentication in Time-Triggered Networked Control Systems
- * Resilient Consensus Protocols with Trusted Nodes
- * **Resilient Observation Selection**
- * Conclusions

Short Term Forecasting using Gaussian processes



- Large area to be monitored
- Only a limited number of sensors can be placed
- Where to place the sensors?

- **Observation Selection**

Y predictor variable

$\mathcal{V} = \{X_1, \dots, X_M\}$ set of possible sensor locations

$$S \hat{=} \min_{\mathcal{V}: |S|=N} S_{Y|S}^2; \quad S_{Y|S}^2 = S_Y^2 - S_{YS} S_S^{-1} S_{SY}$$

- **An attacker may try to disable some of the sensors**

- Sensor placement has to be resilient to such attacks

$$S \hat{=} \min_{\mathcal{V}: |S|=N} \max_{\mathcal{A}: |\mathcal{A}|=K} S_{Y|(S \setminus \mathcal{A})}^2$$

Summary

- * Threat Modeling for CPS Security
- * Performance Impact of Authentication in Time-Triggered Networked Control Systems
- * Resilient Consensus Protocols with Trusted Nodes
- * Resilient Observation Selection