The More the Merrier: Adding Hidden Measurements for Anomaly Detection and Mitigation in Industrial Control Systems

HOTSOS 2020

Jairo Giraldo
THE UNIVERSITY OF UTAH

David Urbina
UT DALLAS The University of Texas at Dallas

CheeYee Tang
NIST National Institute of Standards and Technology U.S. Department of Commerce

Álvaro A. Cárdenas
UNIVERSITY OF CALIFORNIA SANTA CRUZ

Acknowledgments: This work was performed under the financial assistance award 70NANB17H282N from U.S. Department of Commerce, National Institute of Standards and Technology (NIST).

The Growing Threat to Industrial Control Systems

KIM ZETTER SECURITY 03.03.16 7:00 AM

INSIDE THE CUNNING, UNPRECEDENTED HACK OF UKRAINE'S POWER GRID

MIT Technology Review

Sign in Subscribe

Computing / Cybersecurity

Triton is the world's most murderous malware, and it's spreading

The rogue code can disable safety systems designed to prevent catastrophic industrial accidents. It was discovered in the Middle East, but the hackers behind it are now targeting companies in North America and other parts of the world, too.

by **Martin Giles** March 5, 2019

KIM ZETTER 11.03.14 06:30 AM

An Unprecedented Look at Stuxnet, the World's First Digital Weapon

False Data Injection Attacks (FDI)
Actuator or Sensor attacks

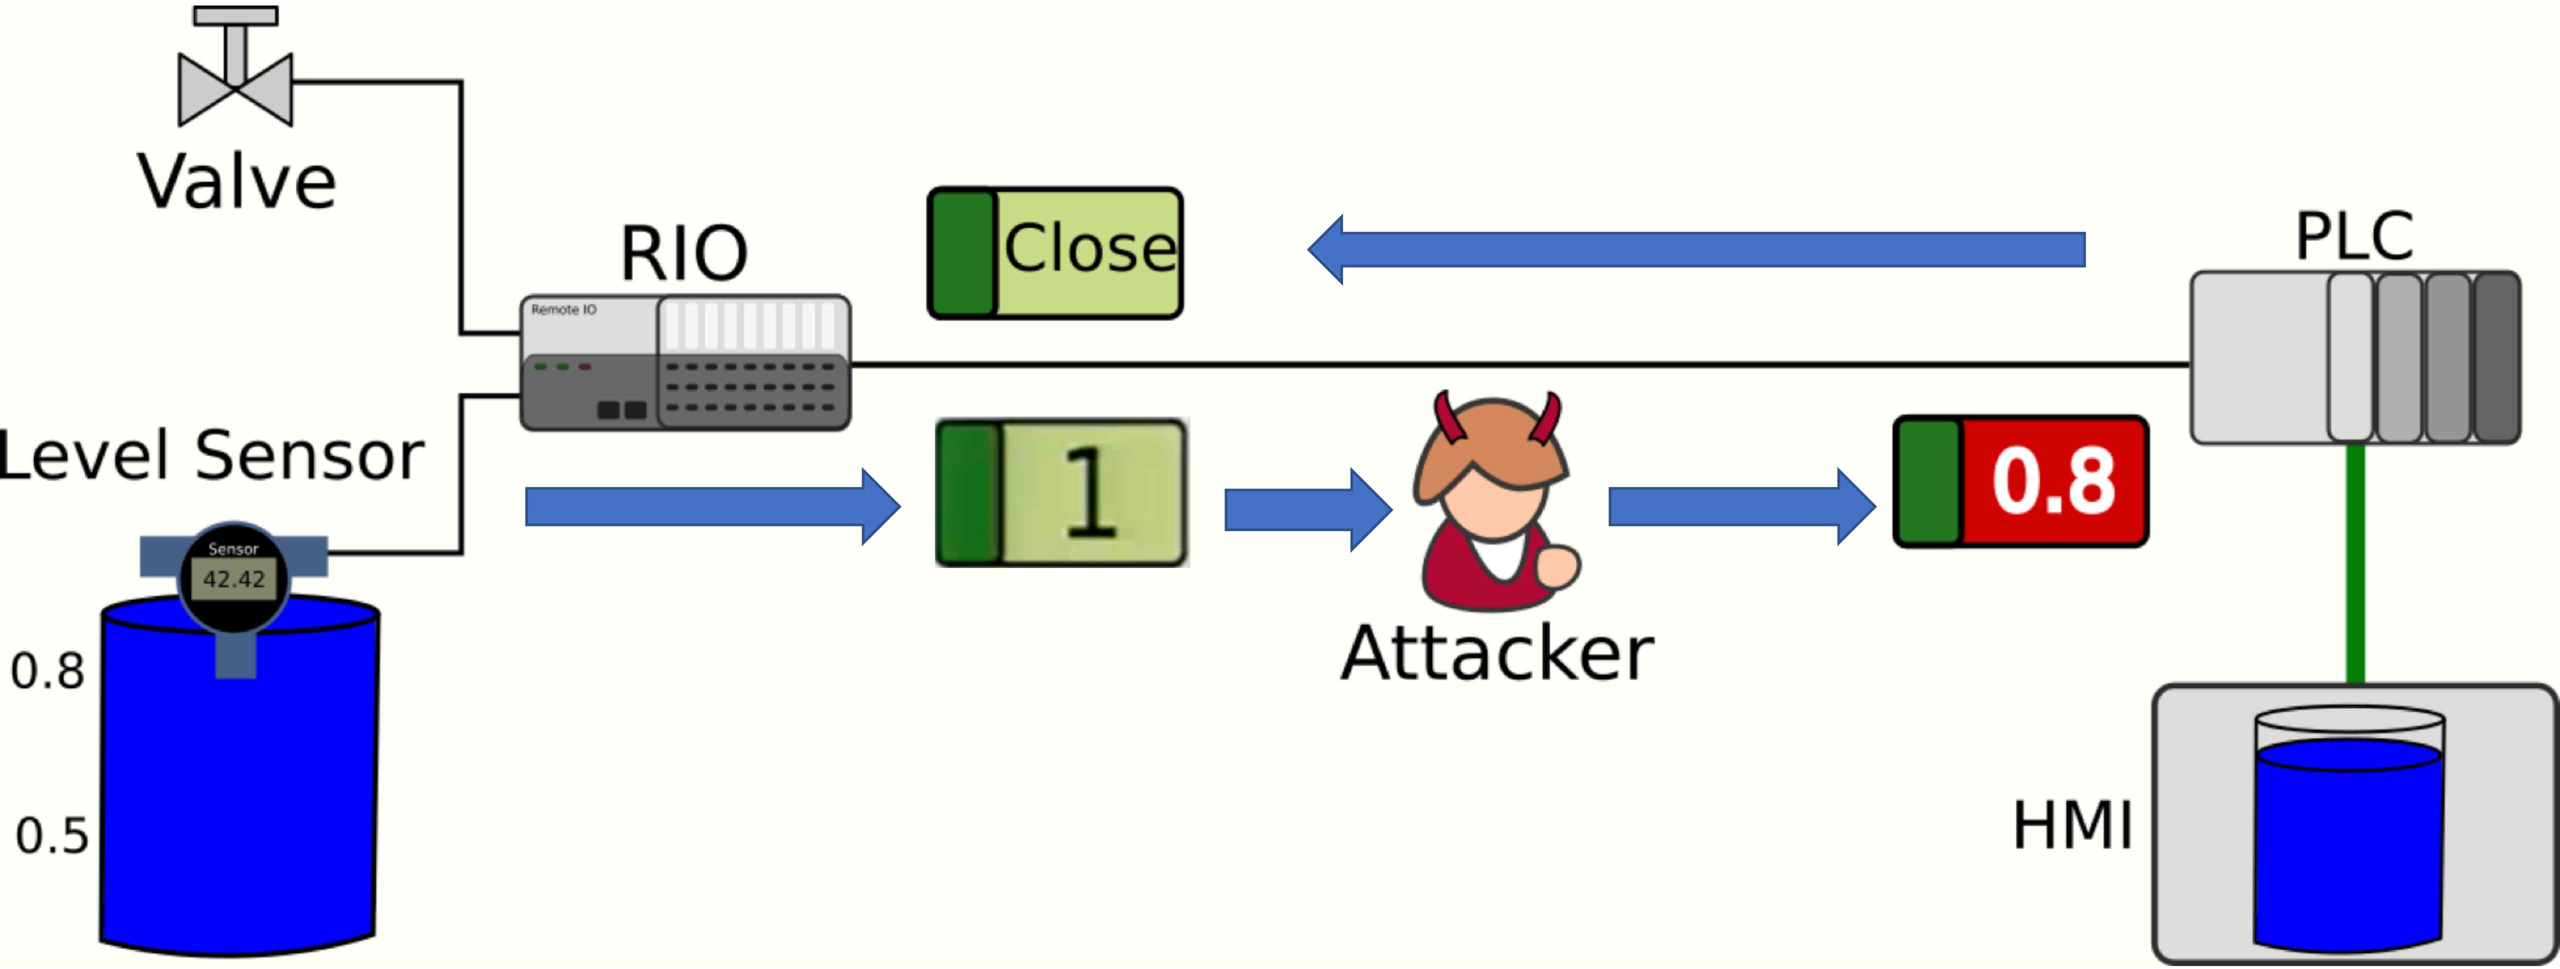

There is a lot of Work on Physics-Based Anomaly Detection (PBAD) to prevent FDI

Venue: Control	[8] Bobba *et al.*	[53] Sandberg *et al.*	[58] Teixeira *et al.*	[6] Bai, Gupta	[45] Mo *et al.*	[43] Mo, Sinopoli	[7] Bai *et al.*	[42] Miao *et al.*	[22] Hou *et al.*	[16] Eyisi *et al.*	[44] Mo *et al.*	[49] Pasqualetti *et al.*	[60] Teixeira *et al.*	[30] Kwon *et al.*	[59] Teixeira *et al.*	[15] Do *et al.*	[1, 2] Amin *et al.*	[55] Smith	[24] Kerns *et al.*
Detection Statistic																			
stateless	●	●	●	-	-	-	●	●	●	◐	●	●	●	●	●	-	●	◐	●
stateful	-	-	-	◐	⊛	⊛	-	-	-	-	-	-	-	-	-	●	-	-	-
Physical Model																			
AR	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
SLS	●	●	◐	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
LDS	-	-	-	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
other	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Metrics*																			
impact	-	●	-	●	-	-	●	-	●	●	-	●	●	●	●	-	-	-	●
statistic	-	-	●	-	●	-	-	●	●	-	-	●	●	●	●	-	●	-	-
TPR	-	-	-	-	●	●	-	-	-	-	●	-	-	-	-	●	-	-	-
FPR	-	-	-	-	●	-	-	-	-	-	●	-	-	-	-	●	-	-	-
Validation																			
simulation	-	●	●	●	●	●	●	●	●	●	●	●	●	●	-	●	-	●	●
real data	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	●	-	-
testbed	-	-	-	-	-	-	-	-	-	-	-	-	-	-	●	-	-	-	●

Venue: Smart/Power Grid	[32] Liang *et al.*	[19] Giani *et al.*	[13] Dan, Sandberg	[27] Kosut *et al.*	[25] Kim, Poor	[14] Davis *et al.*	[56] Sridhar, Govindarasu	[28] Koutsandria *et al.*
Detection Statistic								
stateless	-	●	●	●	●	●	-	●
stateful	⊛	-	-	-	-	-	⊛	-
Physical Model								
AR	-	-	-	-	-	-	-	-
SLS	●	●	●	●	●	◐	-	-
LDS	-	-	-	-	-	-	-	-
other	-	-	-	-	-	-	●	●
Metrics*								
impact	●	-	●	-	●	-	-	●
statistic	●	-	-	●	-	-	-	●
TPR	-	●	-	●	-	-	●	-
FPR	-	-	-	●	-	-	●	-
Validation								
simulation	●	-	●	●	●	●	●	◐
real data	-	-	-	-	-	◐	●	-
testbed	-	-	-	-	-	-	-	-

Venue: Security	[39] Mashima *et al.*	[34, 35] Liu *et al.*	[48] Parvania *et al.*	[33] Lin *et al.*	[54] Shoukry *et al.*	[20] Hadziosmanovic *et al.*	[10] Cardenas *et al.*	[66] Wang *et al.*	[41] McLaughlin	[52] Sajjad *et al.*	[64] Valente, Cardenas	[29] Krotofil *et al.*
Detection Statistic												
stateless	-	●	◐	●	-	●	-	◐	◐	-	-	●
stateful	●	-	-	-	⊛	-	●	-	-	●	●	●
Physical Model												
AR	●	-	-	-	-	●	-	-	-	-	-	-
SLS	-	●	-	-	-	-	-	-	-	-	-	-
LDS	-	-	-	-	●	-	●	-	-	●	-	-
other	-	-	◐	◐	-	-	-	●	●	-	●	●
Metrics*												
impact	●	●	-	●	-	-	-	-	-	●	-	-
statistic	-	●	-	-	-	-	-	-	-	-	●	●
TPR	-	-	-	-	●	●	-	●	-	-	-	-
FPR	●	-	-	-	●	-	-	●	-	-	-	-
Validation												
simulation	-	●	-	●	-	-	●	●	-	●	-	●
real data	●	-	-	-	-	●	-	-	-	-	●	-
testbed	-	-	◐	-	●	●	-	-	●	-	-	-

Venue: Misc.	[65] Vukovic, Dan	[46] Morrow *et al.*	[12] Cui *et al.*	[9] Carcano *et al.*	[21] Hei *et al.*	[26] Kiss *et al.*
Detection Statistic						
stateless	●	●	●	●	-	●
stateful	-	-	●	-	⊛	-
Physical Model						
AR	-	-	-	-	-	-
SLS	◐	◐	●	-	-	-
LDS	-	-	-	-	-	-
other	-	-	-	●	◐	●
Metrics*						
impact	●	-	●	-	●	-
statistic	●	-	●	-	-	●
TPR	-	-	-	●	-	-
FPR	-	-	-	●	-	-
Validation						
simulation	●	●	●	-	-	●
real data	-	-	-	-	●	-
testbed	-	-	-	●	-	-

CCS 2016 & ACM Computing Surveys 2018

PBAD uses Physical Models with Statistics

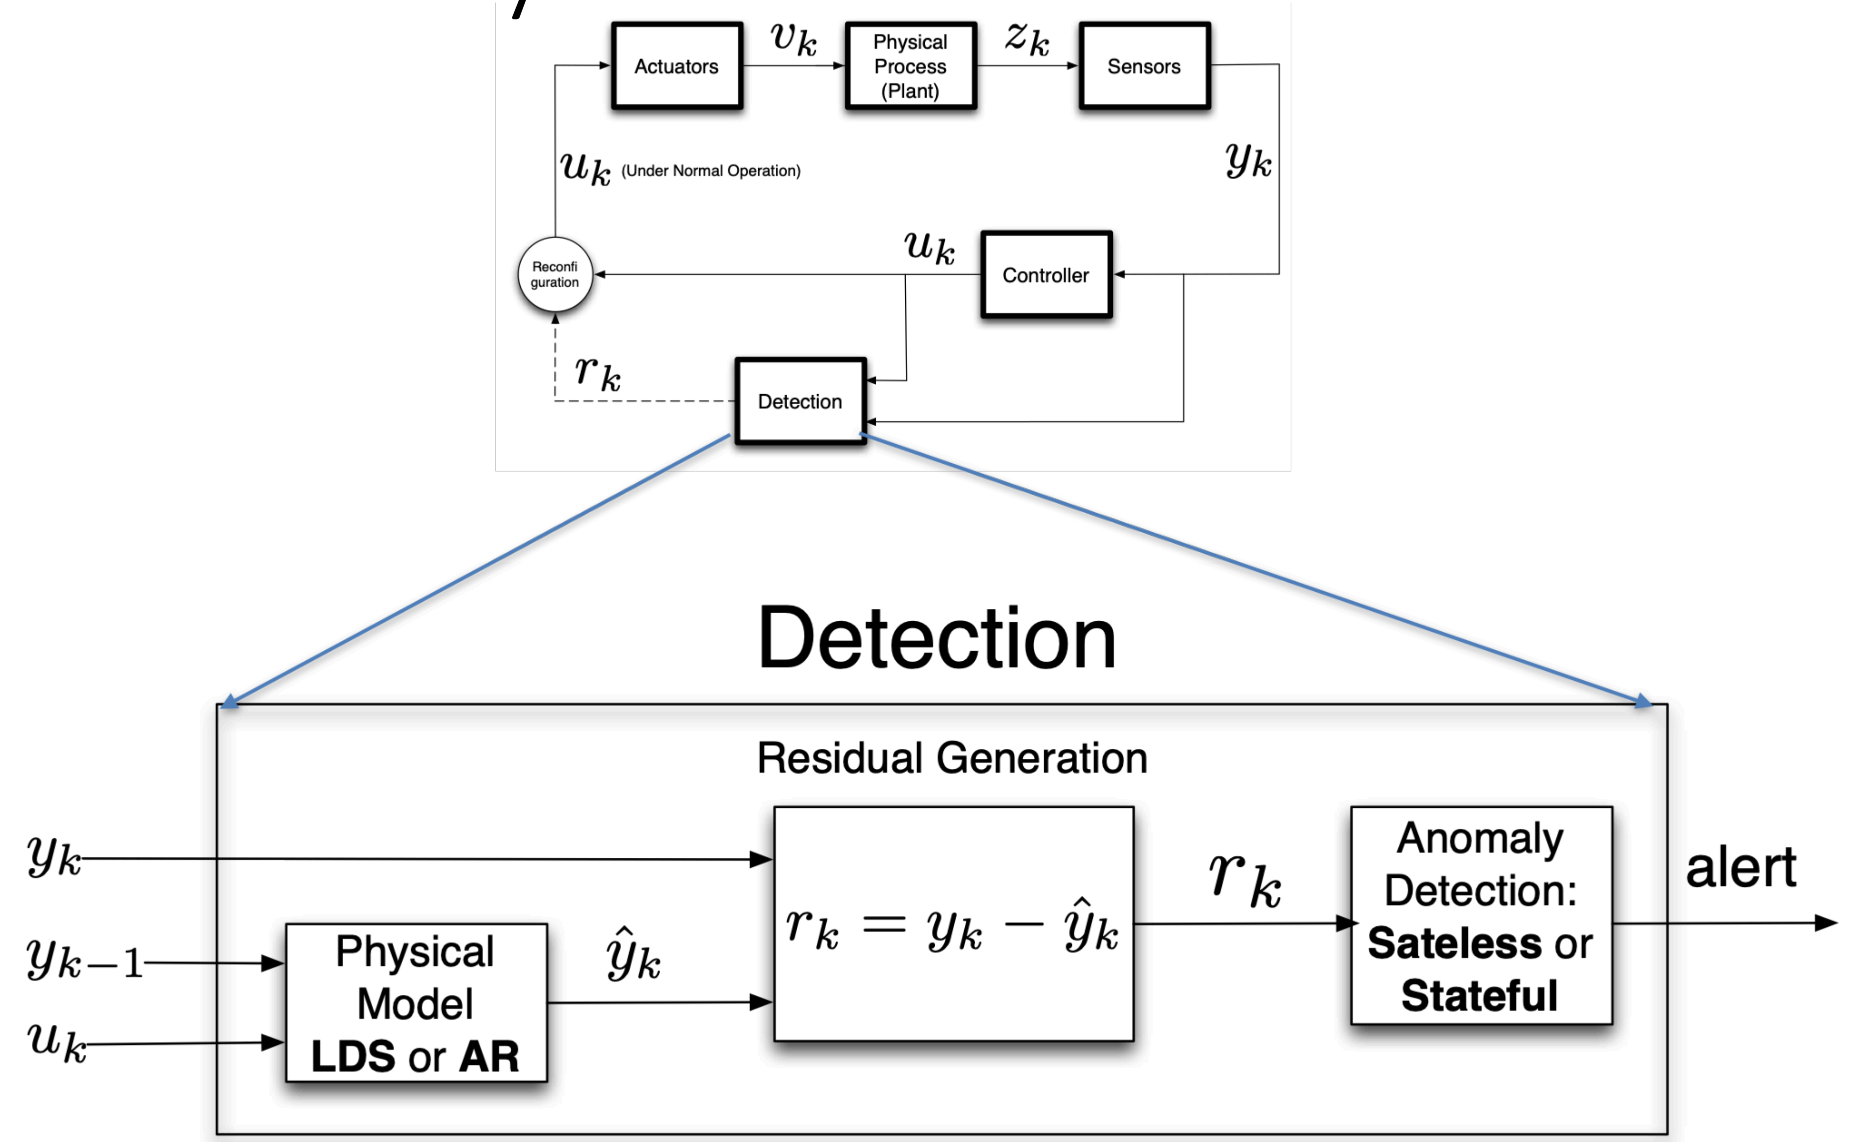

Key Insights of our Paper

- Previous work leverages only the sensors that are already in place for operations of the system.
- These sensors only measure a limited number of physical quantities
- Our idea: hidden sensor measurements
  - Add new sensors to measure new physical quantities
  - These sensors are not used for operations, only for security
  - They can even be used for attack-response

Hidden Sensor Measurements

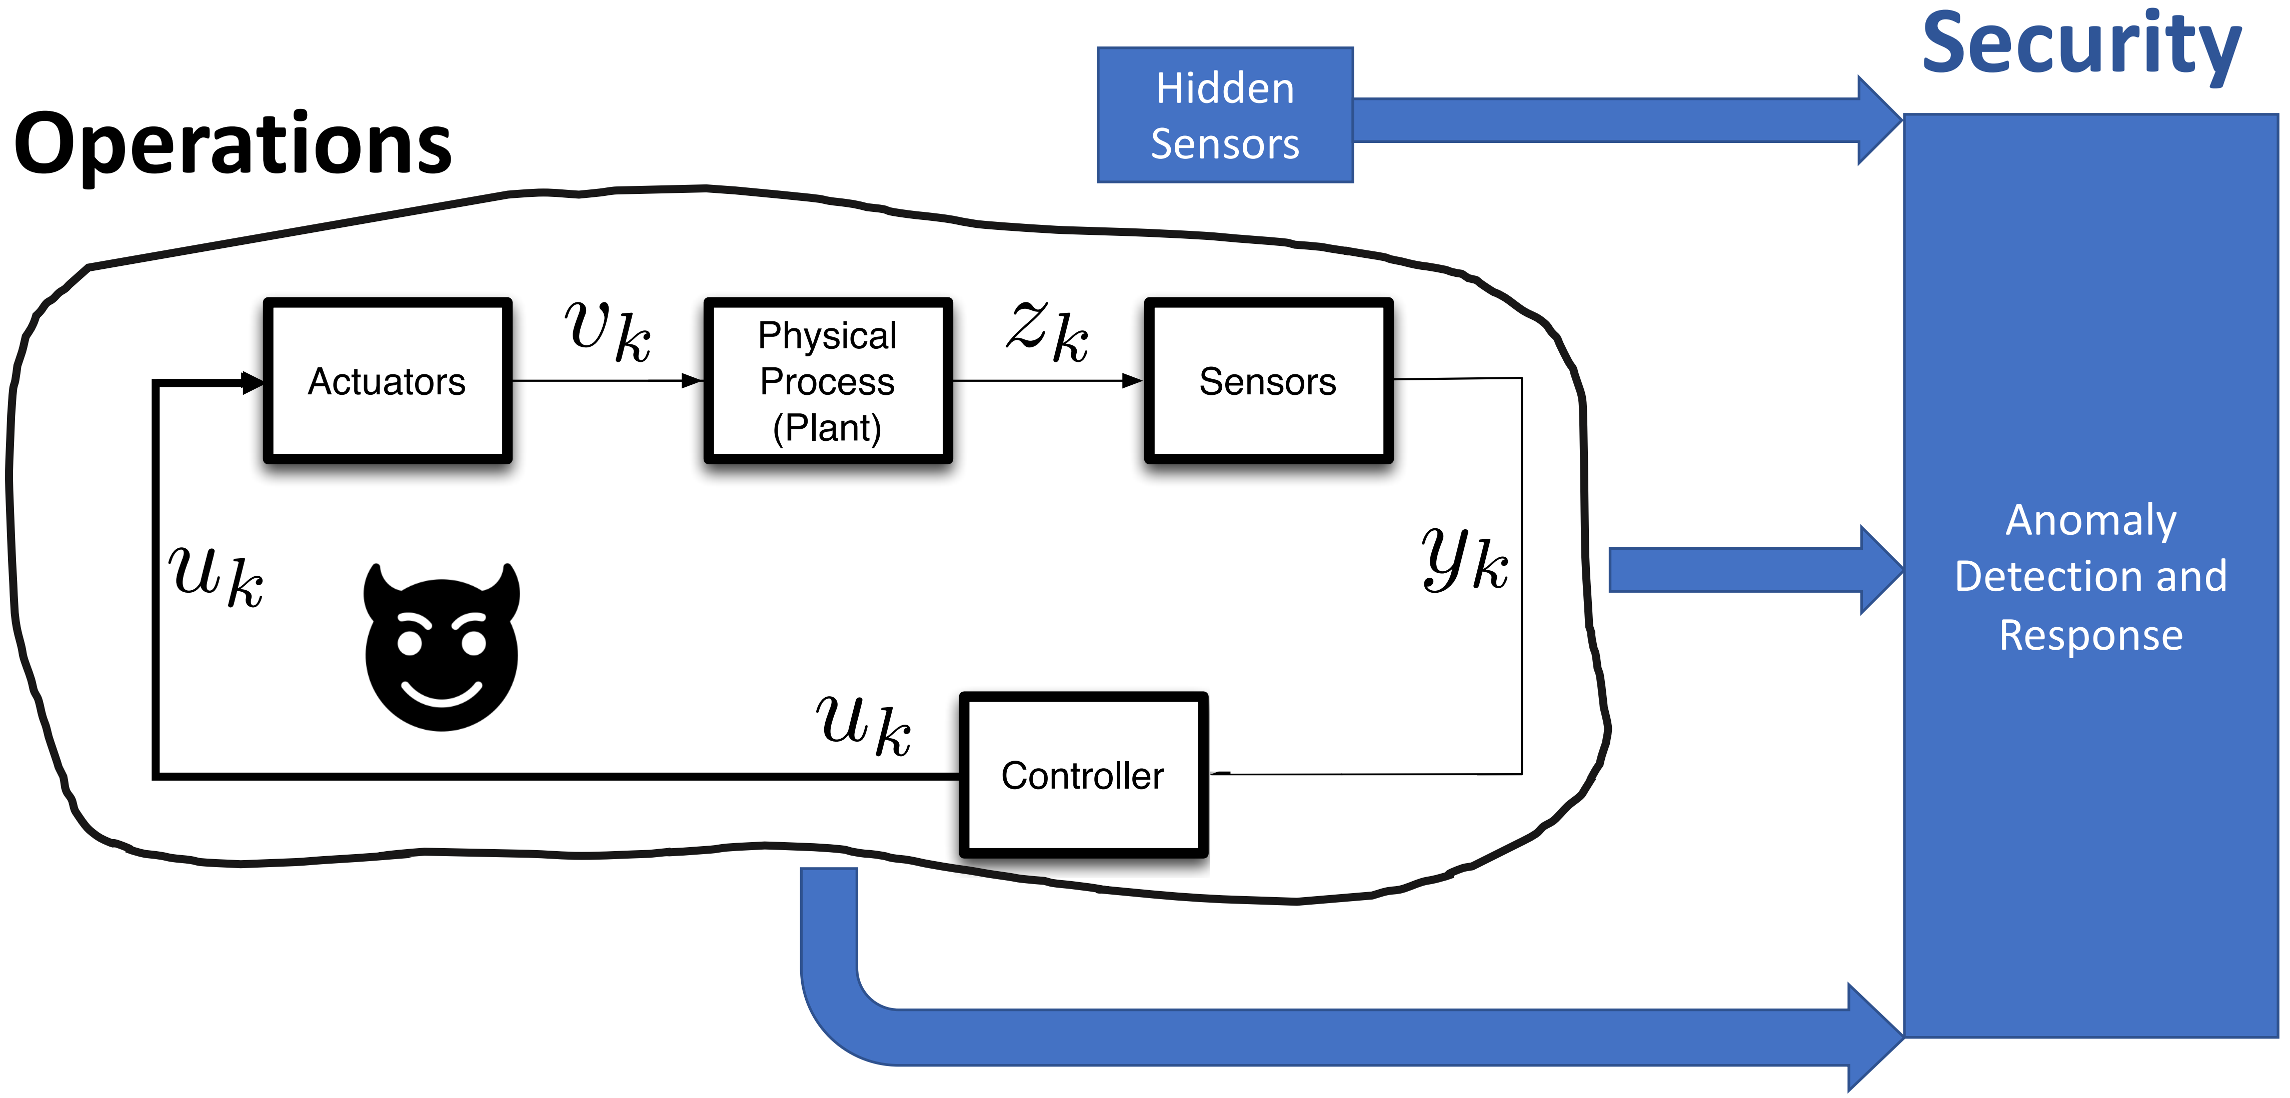

Industrial Control Systems

- ICS are typically composed by multiple control loops.
- The dynamics of a physical process can be summarized as

$$\dot{\boldsymbol{x}}(t) = F(\boldsymbol{x}(t), \boldsymbol{u}(t)),$$
$$\boldsymbol{y}(t) = H(\boldsymbol{x}(t), \boldsymbol{u}(t))$$

- $x(t)$: System states (e.g., pressure, temperature)
- $y(t)$: measurable variables.
- $u(t)$ control command

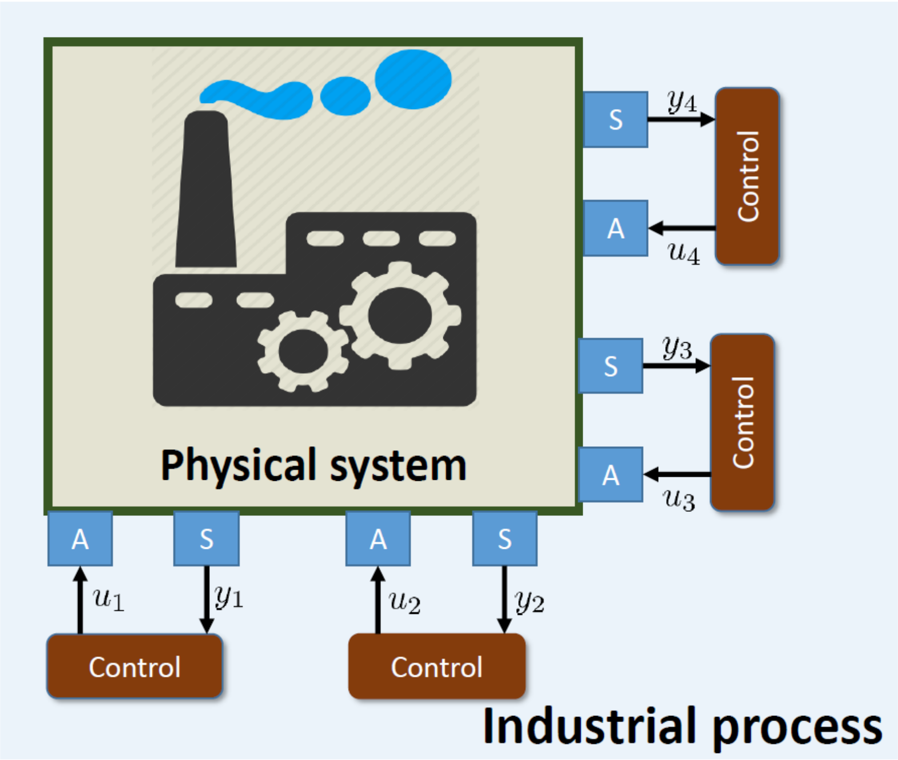

Hidden Sensors and Correlation

- Measurable variables can be divided into
  - I^{CL} Set of operational sensors used to compute control commands u_i
  - I^{rd} Set of physical quantities that can be potentially measured

- Attacks in y_i, for $i \in I^{CL}$ will cause variations in u_i, which will lead the system to deviate from its nominal operation.
- An attack in an operational sensor can be reflected in a hidden one.

- We can compute the correlation coefficient to find, *for each visible measurement* y_i, a group of correlated hidden sensors y_j for $j \in I^{rd}$

Multi-Variable attack-Detection

- For each operational sensor, there is an MVD.
- MVD computes a prediction based only on y_i and a group of correlated hidden sensors (e.g., y_j^{rd}, y_l^{rd})
- In order to generate stealthy attacks, it would be necessary to compromise y_i as well as y_j^{rd}, y_l^{rd}.

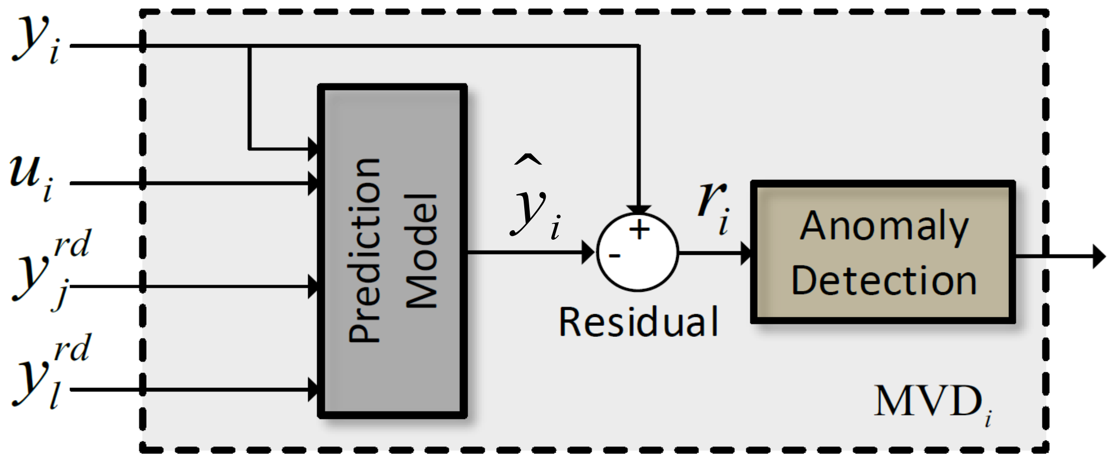

System Reconfiguration for Attack Response

- From a correlated hidden sensor y_j^i, we can estimate the operational sensor $\tilde{y}_i = G(y_j{}^i)$
- When an attack is detected, y_i is replaced by $\tilde{y}_i$.
- This guarantees the system continues operating while avoiding the impact of the attack

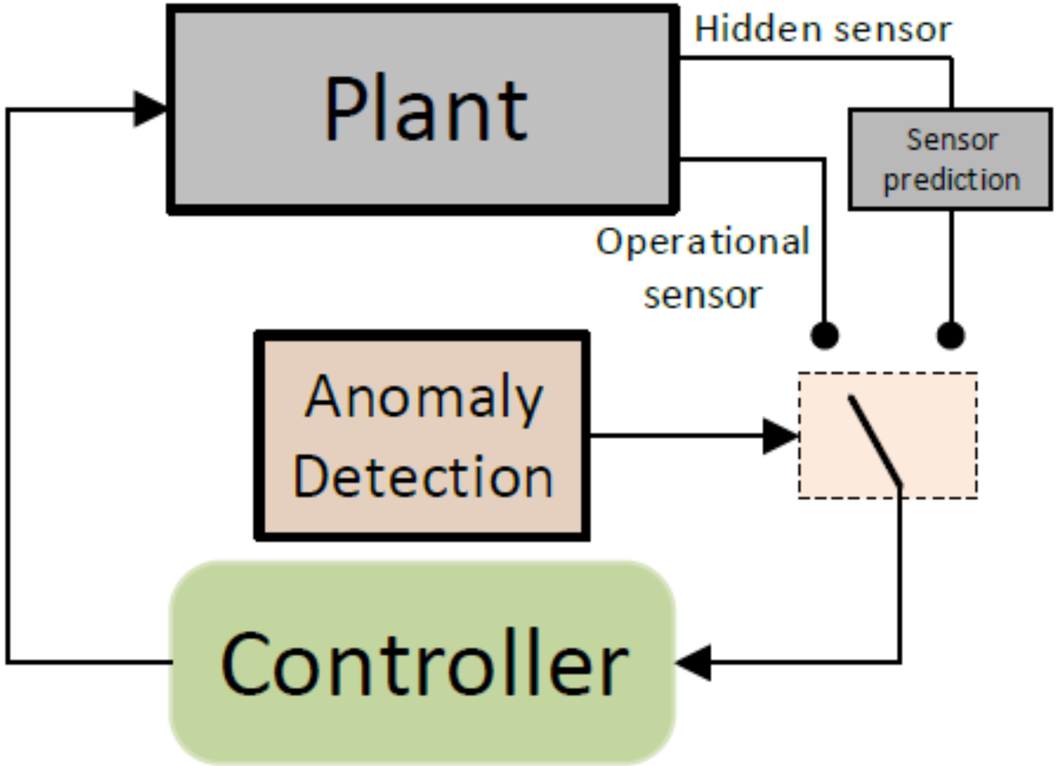

Case Study: NIST Cybersecurity for Manufacturing Systems Testbed

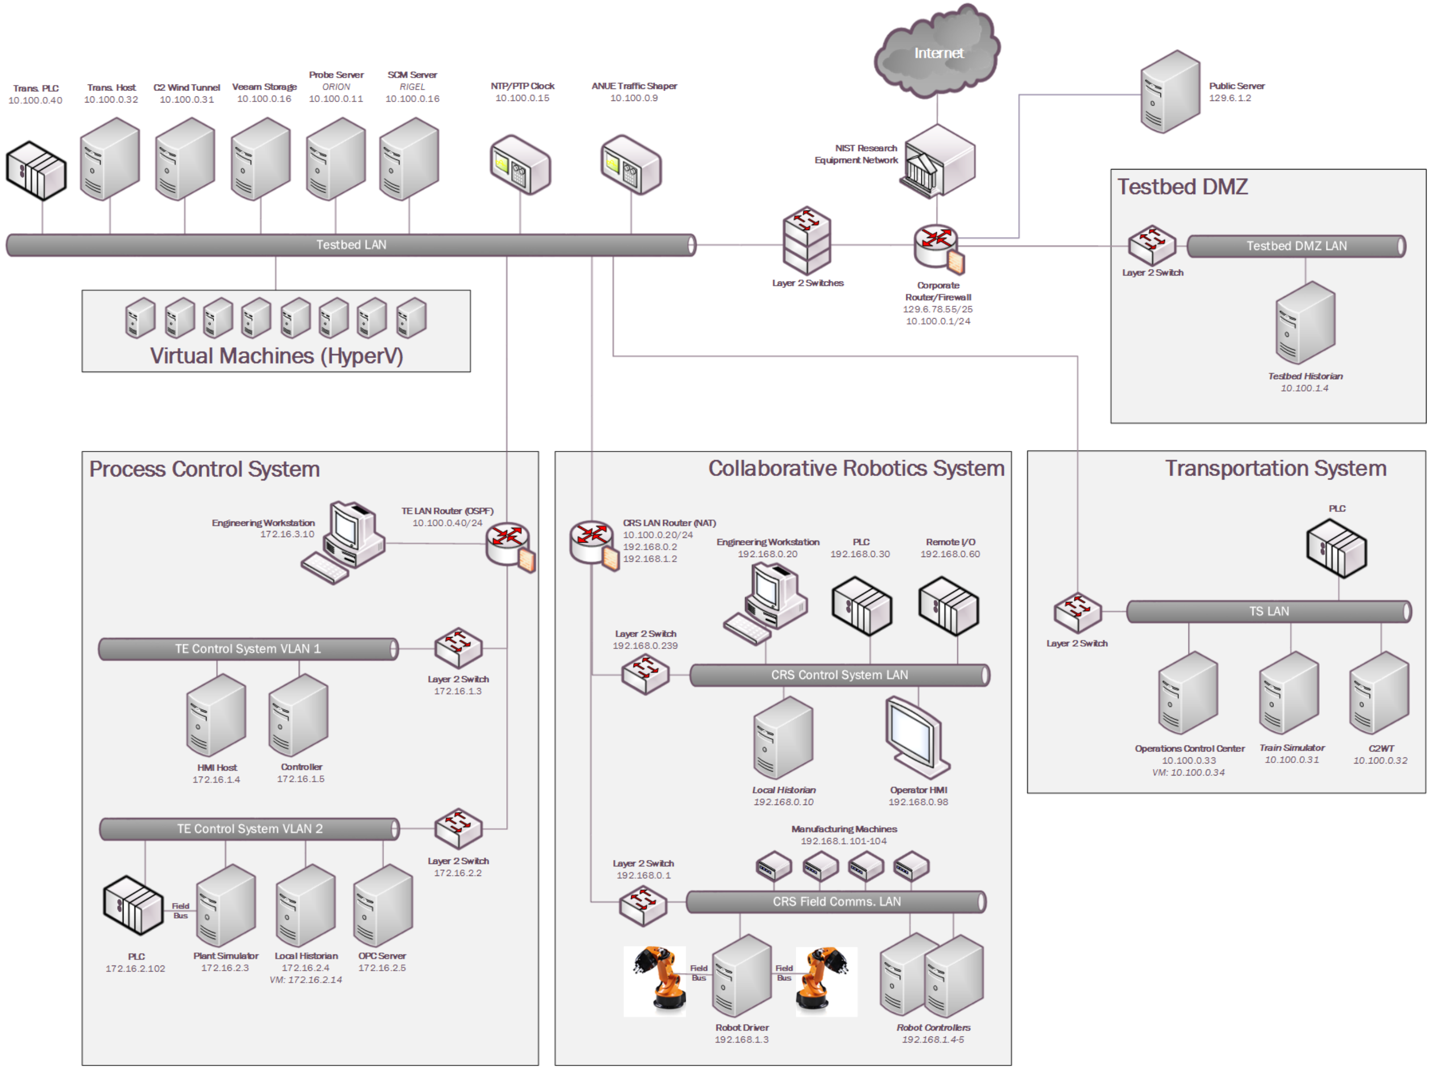

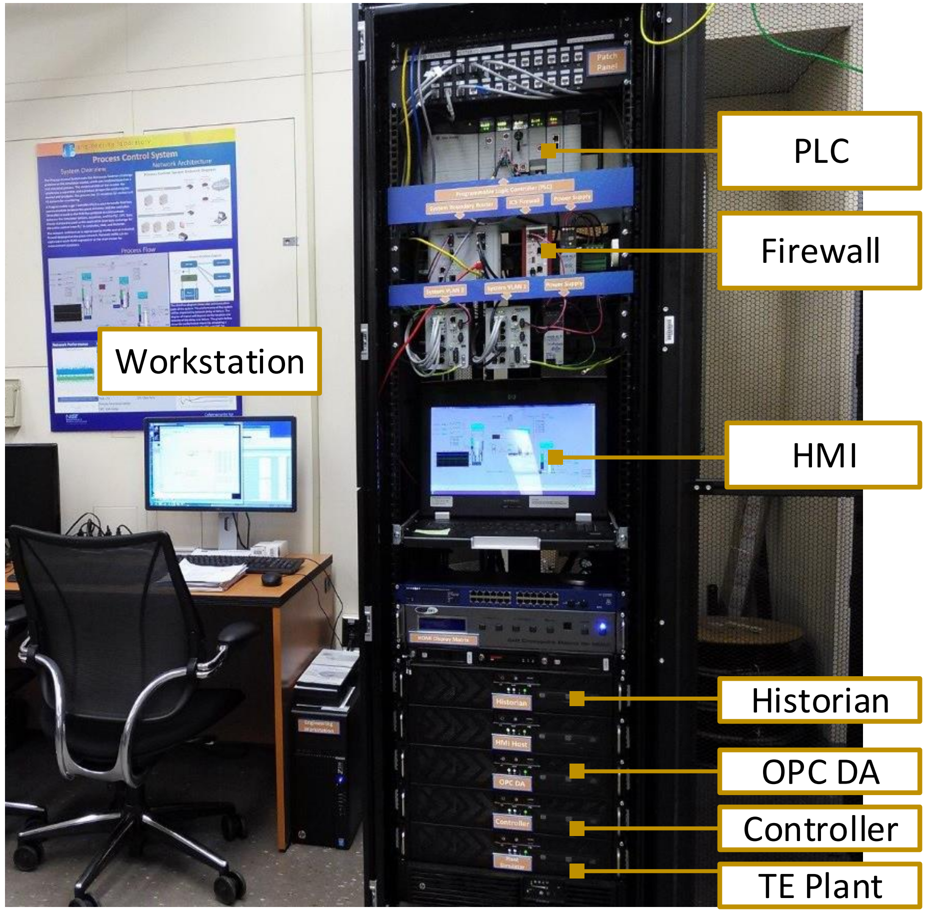

Tennessee Eastman Process

- Classical Control Exemplar based on a real-world chemical reaction system.
  - Use in security in our earlier work*
- Ricker’s control algorithm.
  - Uses only 11 measurements for operations
  - But there are 42 possible measurable physical quantities
- Using correlation analysis, we can identify which of these unused physical quantities are correlated to our critical values

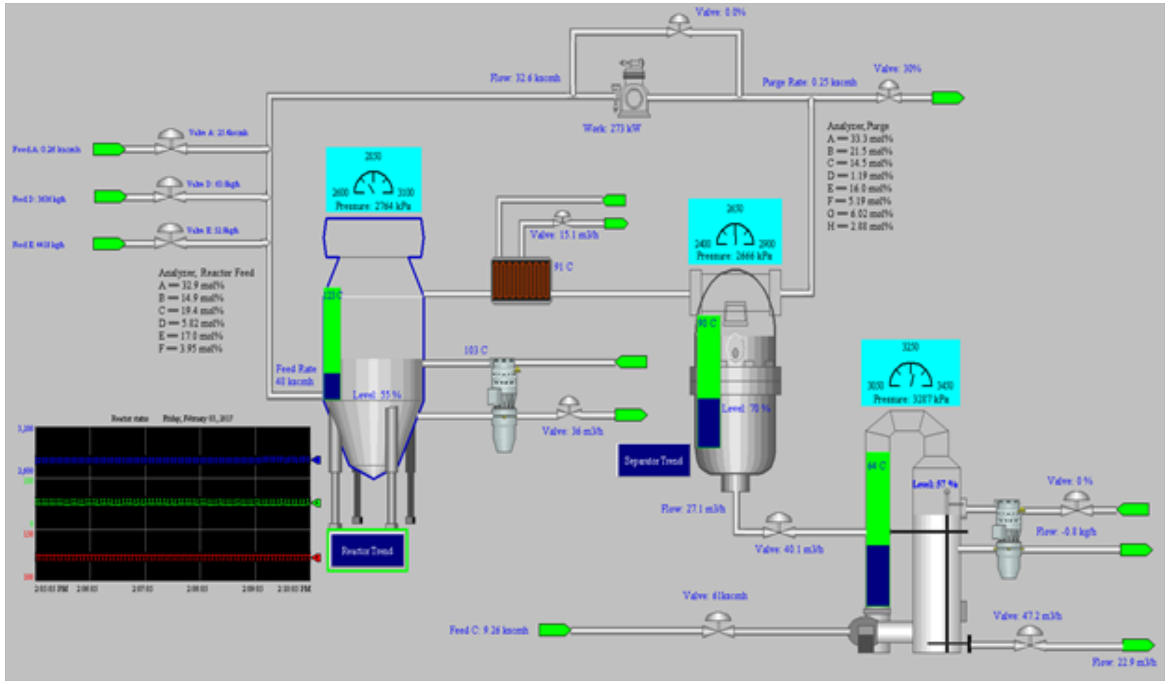

*Huang, Yu-Lun, et al. *International Journal of Critical Infrastructure Protection.* 2009.
*Cárdenas, Alvaro A., et al. ASIACCS 2011.

Hidden Measurements Found in Different Control Loops

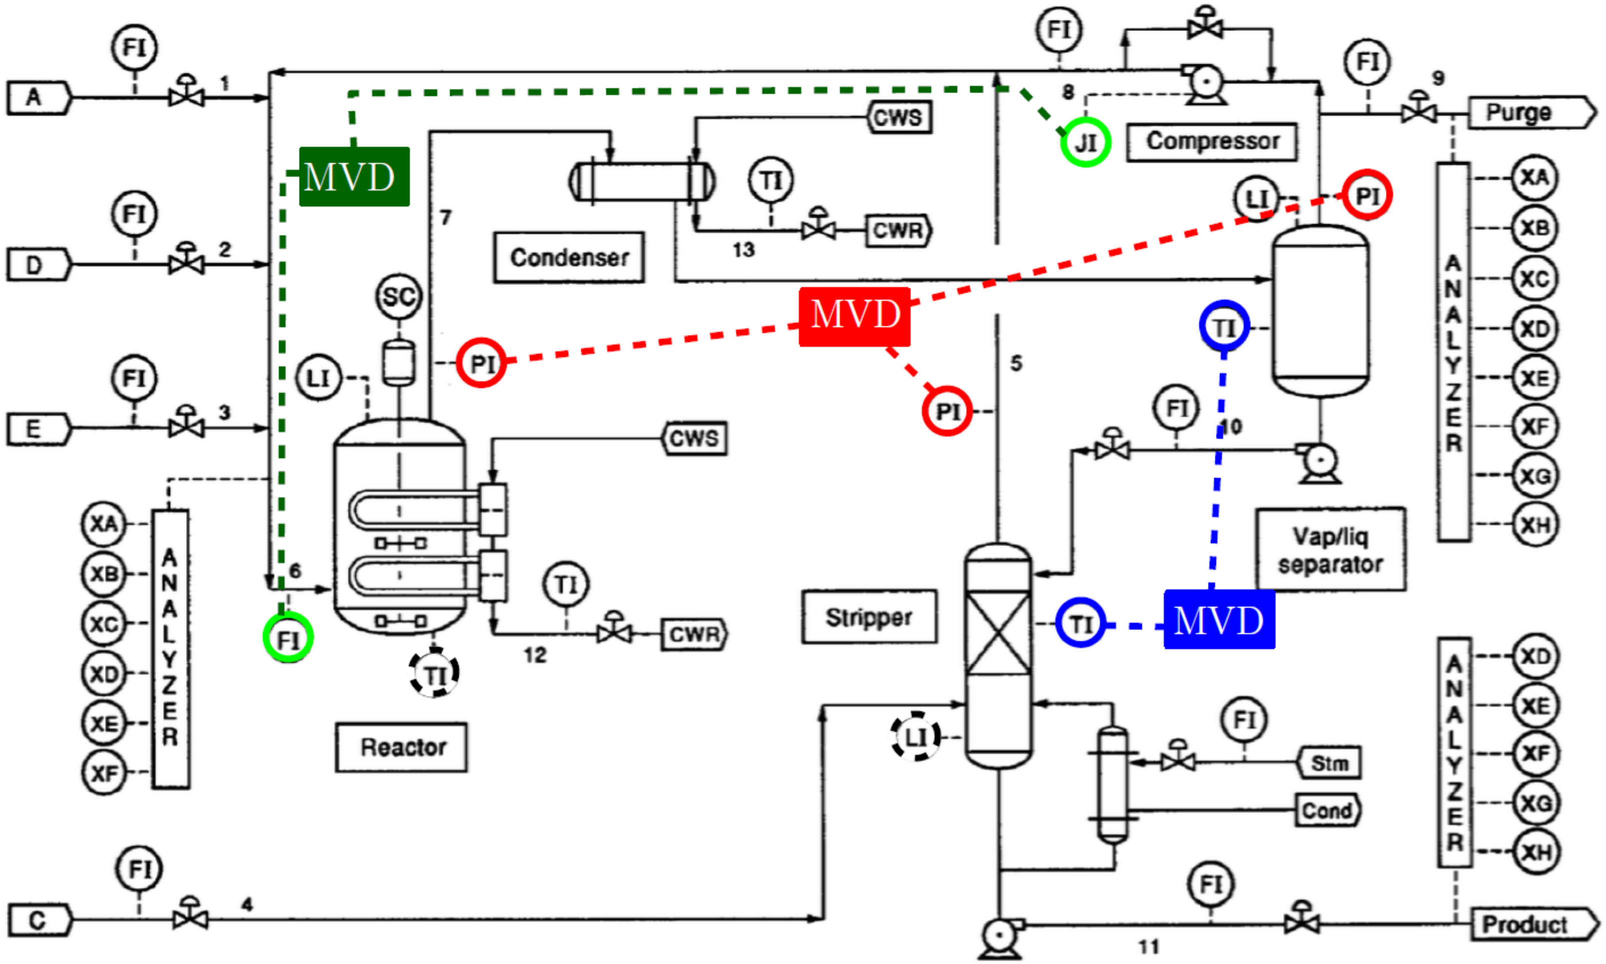

- Using system identification we obtain Hammerstein-Weiner models and use them to compute correlation coefficients to find hidden measurements; e.g.,
  - Reactor pressure is correlated to unmeasured product separator pressure.

State-of-the-art Works well for non-stealthy attacks

- **Control loop**: Reactor pressure (y_7) and purge rate (u_6).
- A predictor is constructed only with y_7 and u_6.
- Bias attack in reactor pressure (y_7) is easily detected.

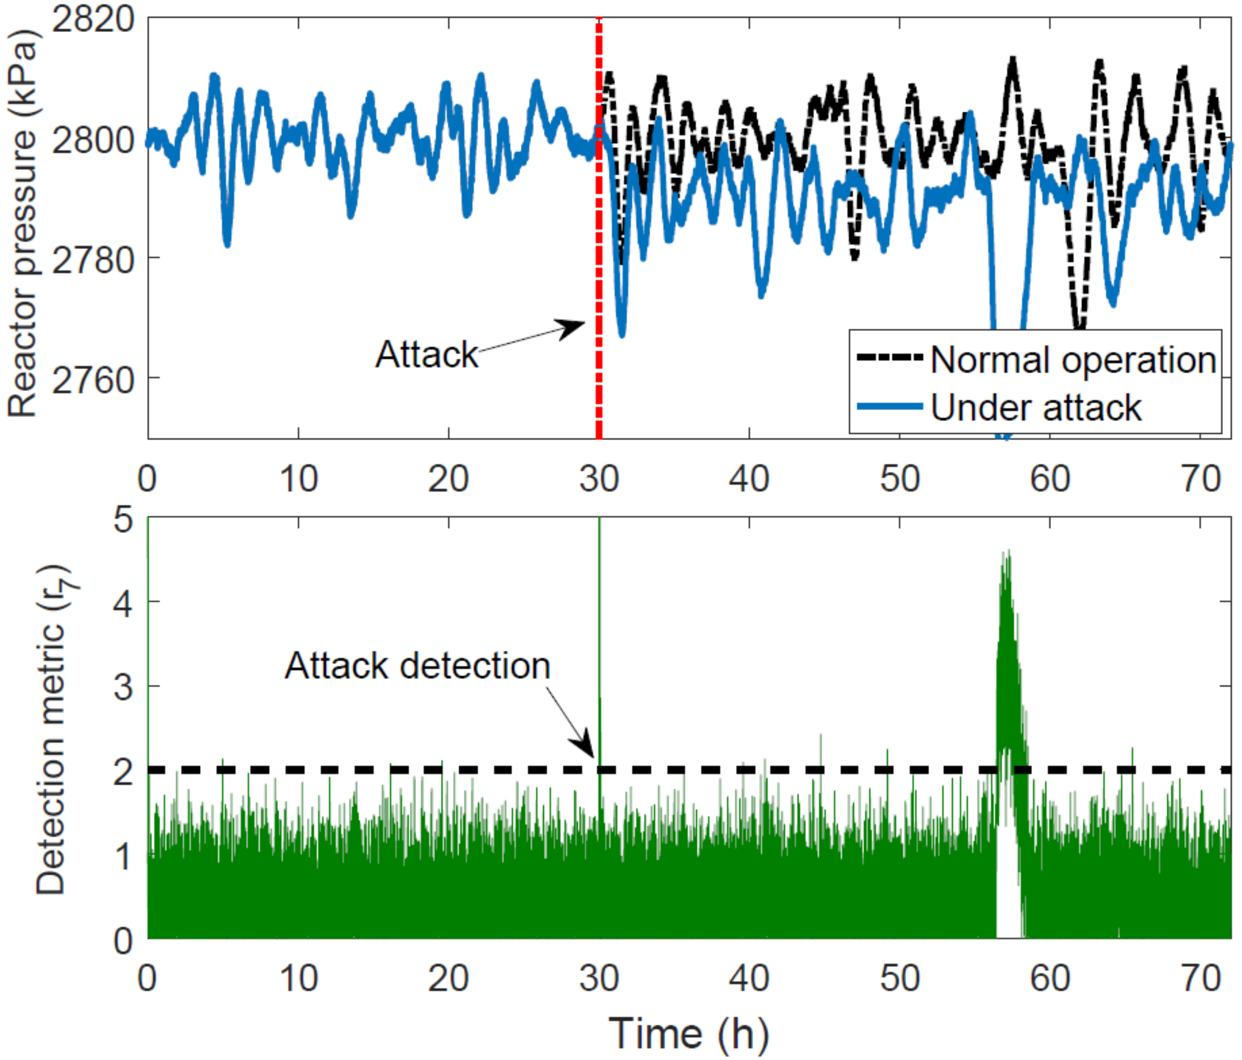

Hidden Measurements Can Detect "stealthy" attacks to state-of-the-art proposals

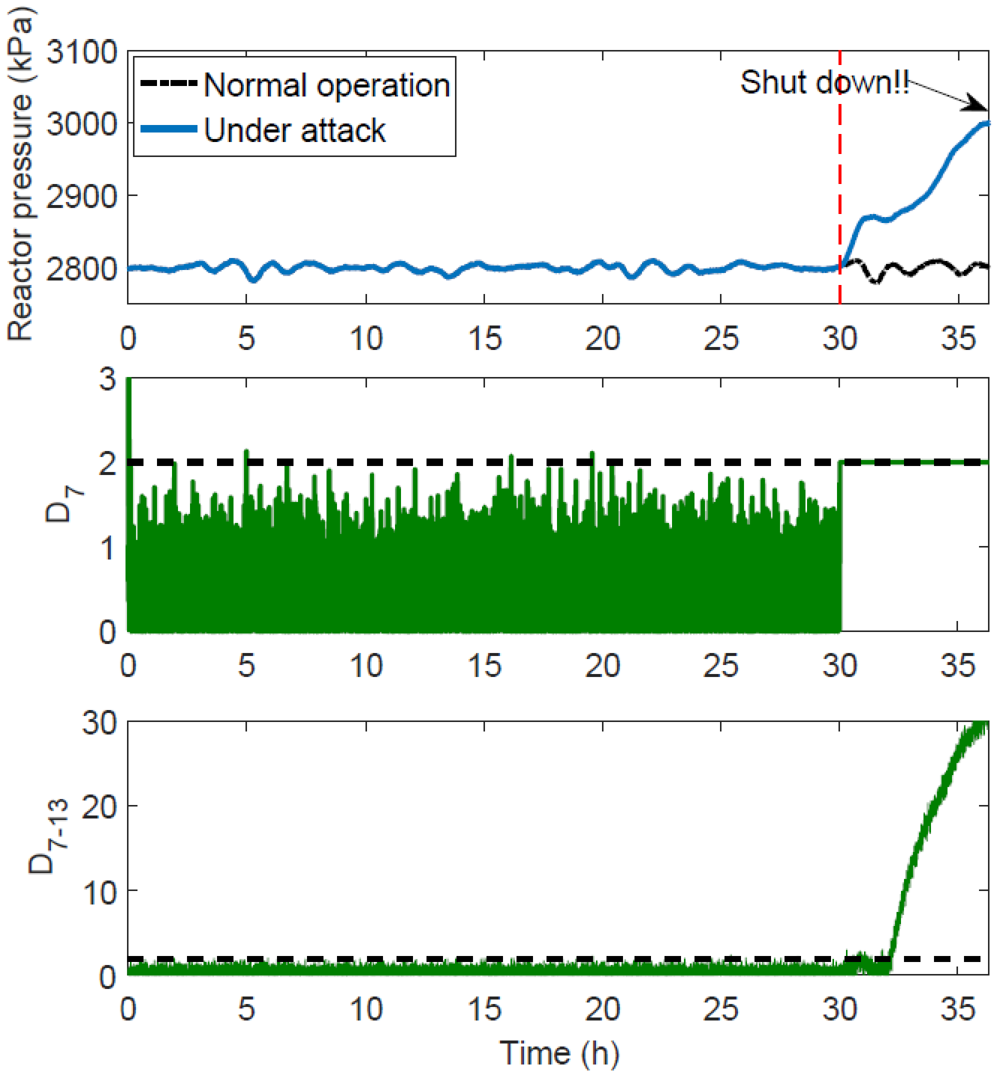

- For a detection mechanism that uses only reactor pressure, a stealthy attack causes a shut down.
- However, an MVD that also includes the product separator pressure (y_{13}) can detect this attack.
- The correlation index between y_7 and y_{13} is 0.96

Hidden Measurements Can Limit Strong Stealthy Attacks

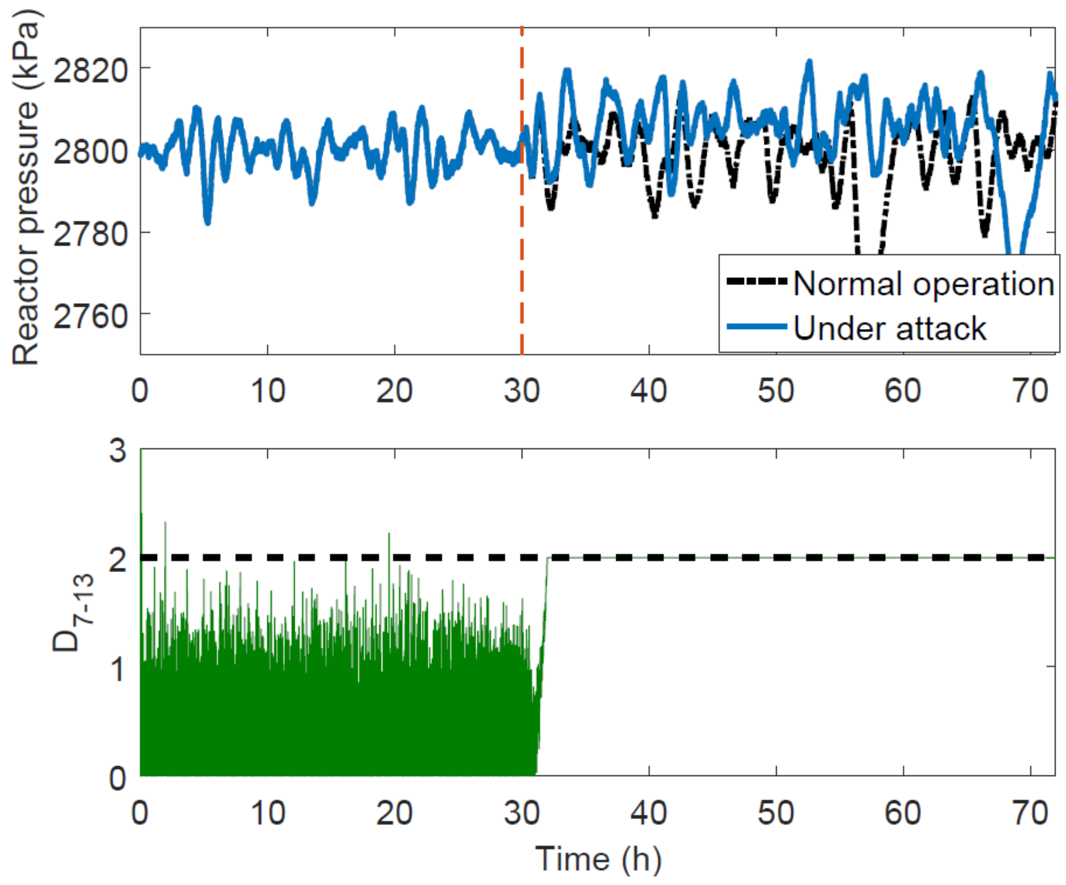

- In the previous slide we assume the attacker does not know our hidden measurement
  - But if an attacker knows our hidden measurement and wants to remain undetected, it can launch a strong stealthy attack
- However, a strong-stealthy attack for an MVD has a small impact on the system (the pressure never reaches unsafe levels, as in the slide before)

Hidden Measurements can be used for Attack Response

- From a correlated hidden sensor y_j^i, we can estimate the operational sensor
- When an attack is detected, y_i is replaced by $\tilde{y}_i$.
- The mitigation strategy offers a second layer of security.

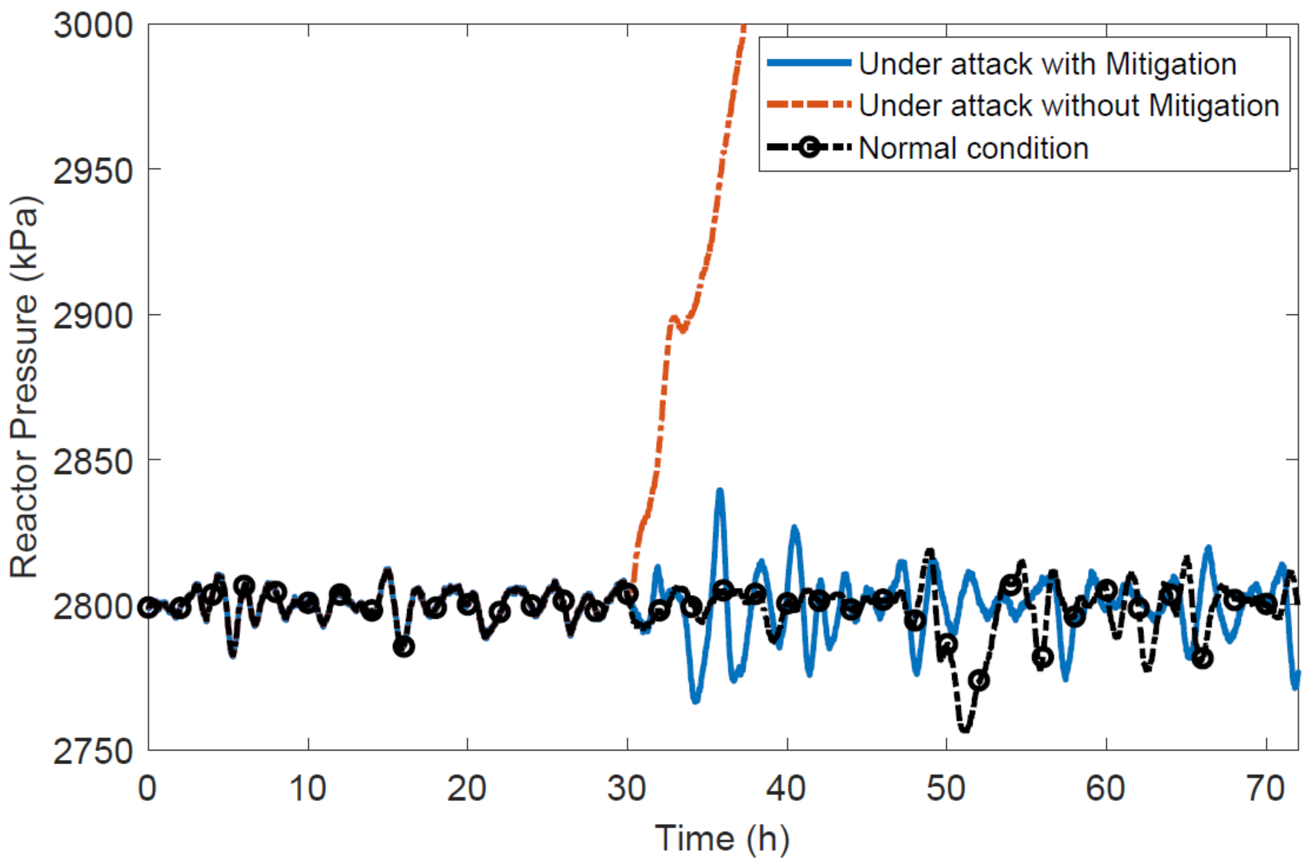

How Many Hidden Sensors are Good Enough?

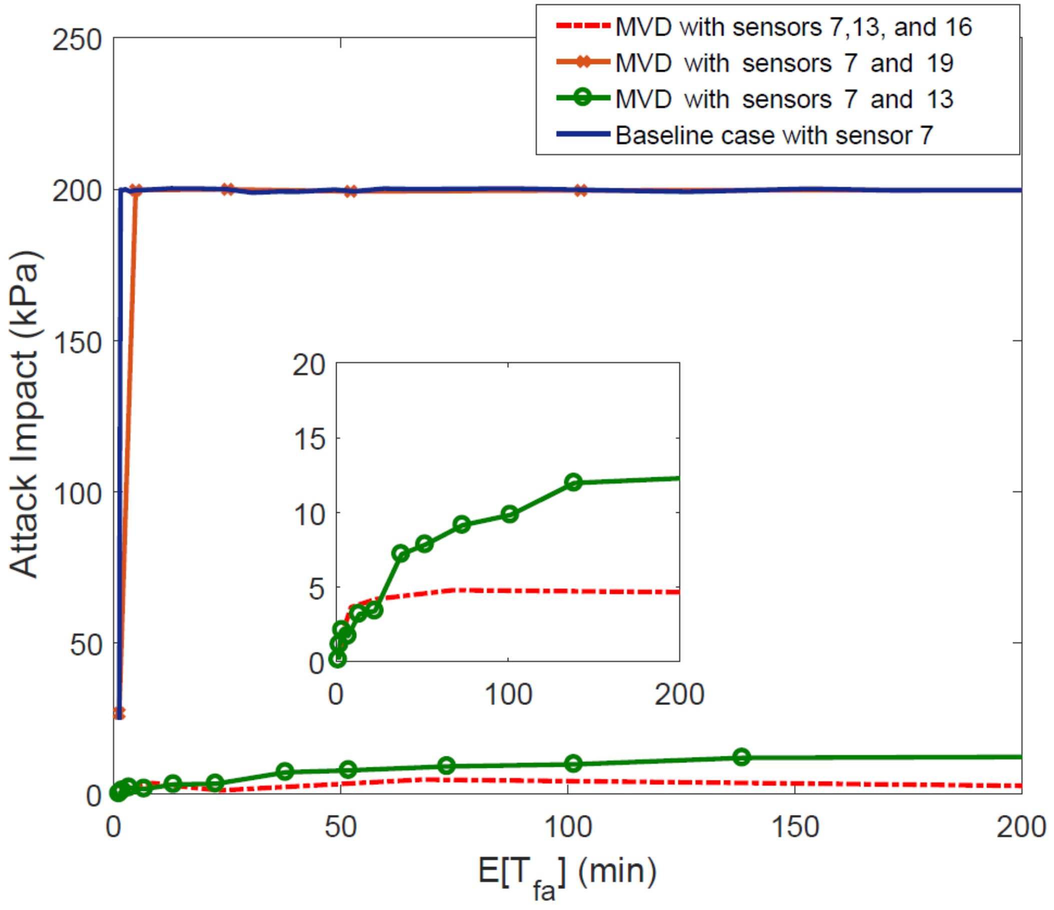

- Should we just create a full model with all possible physical quantities?
  - Not necessary
- We can estimate the impact of "strong" stealthy attacks
  - In this case an MVD with sensors y_7, y_{13}, appear to be enough
  - y_{16} might not be needed

Conclusions

- Hidden Measurements Can
  - Help us detect attacks that cannot be detected using state-of-the-art proposals
  - Limit the impact of stealthy attacks from attackers that know of our “hidden” measurements
  - Provide new measurements to respond to attacks
- Limitations:
  - Trust model, we assume these hidden sensors are not compromised