



What You Always Wanted and Now Can: Hacking Chemical Processes

Marina Krotofil, Jason Larsen

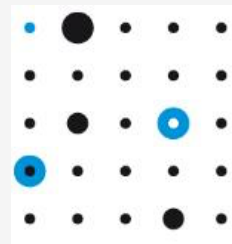
European Network for Cyber Security

IOActive

HITB, Amsterdam, Netherlands

29.05.15

IOActiveTM
Hardware | Software | Wetware
SECURITY SERVICES



ENCS

Who we are



(Ex)Academic



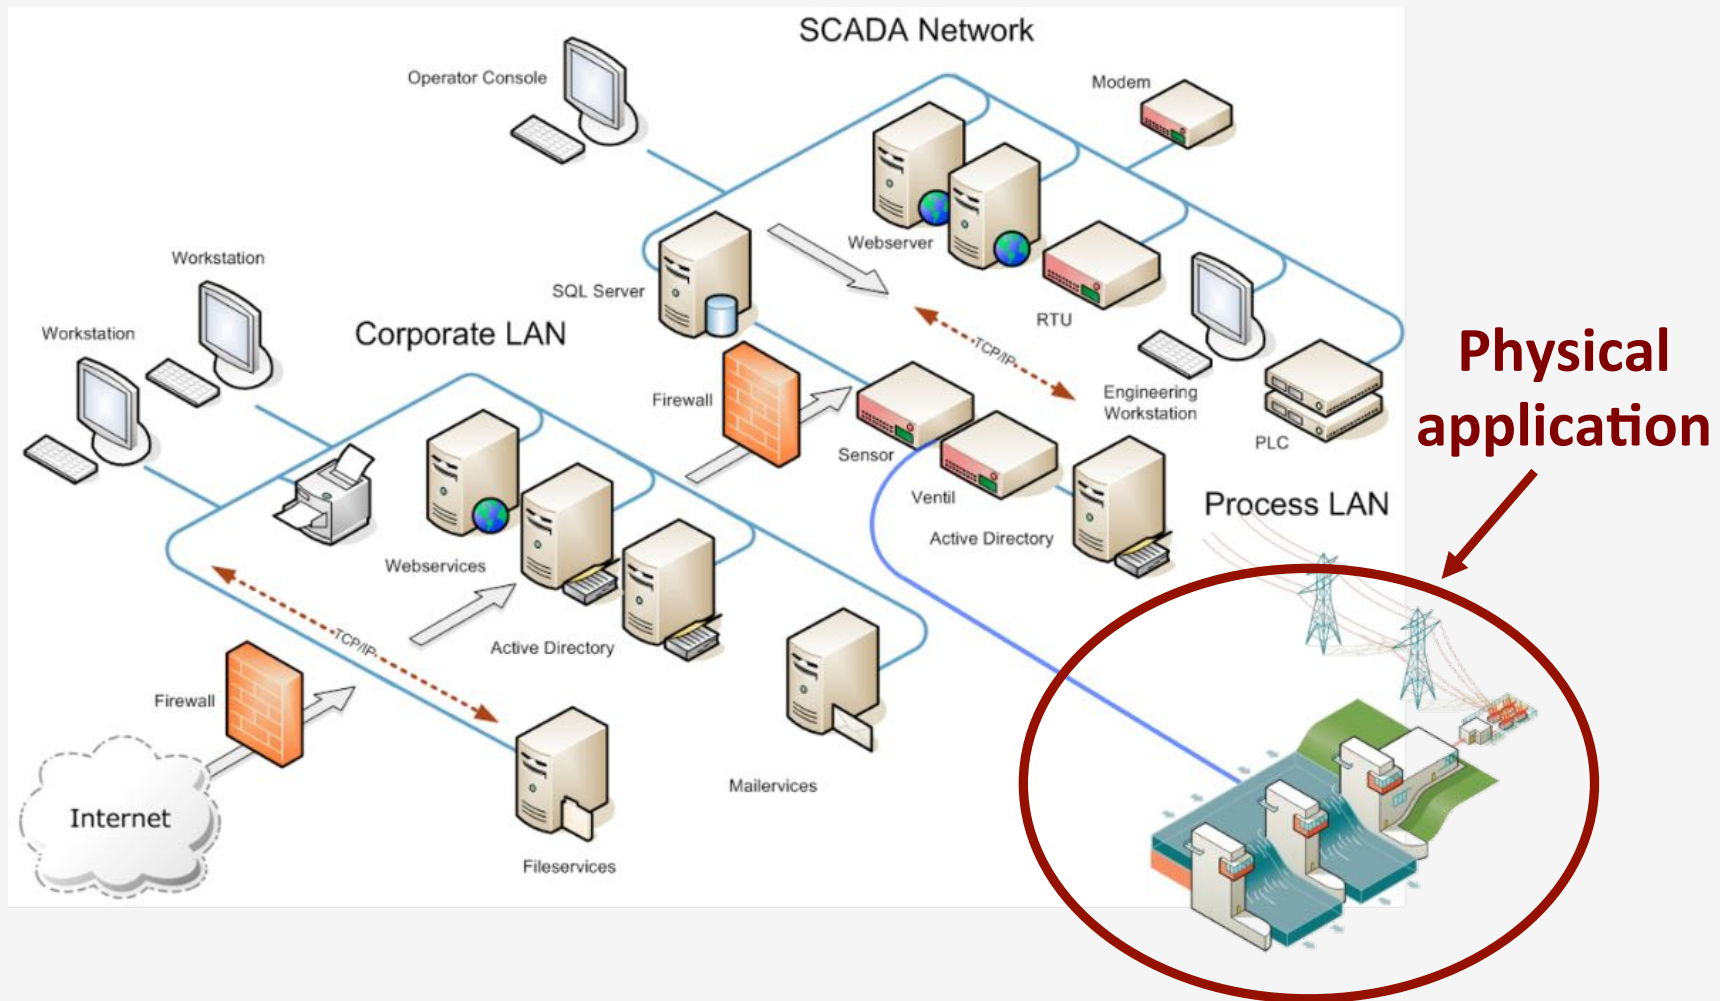
Hacker

- ☐ Countless Skypes and twice as that emails
- ☐ 5 joint publications
- ☐ Still finding each other awesome :-P



What is this talk about

Industrial Control Systems



Cyber-physical security

❑ Cyber-physical systems are IT systems “embedded” in an application in the physical world

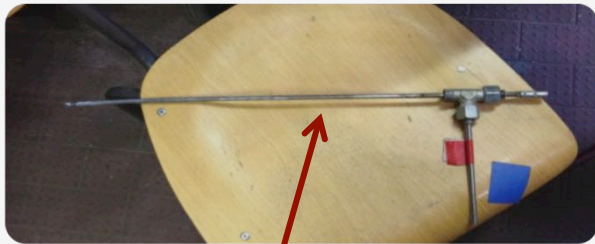
❑ Attacker's goals:

- Get the system in a state desired by the attacker
- Make the system perform actions desired by the attacker



Smart instrumentation

- ❑ Converts analog signal into digital form
- ❑ Pre-process the measurements
- ❑ IP-enabled (part of the “Internet-of-Things”)



**Old generation
temperature sensor**



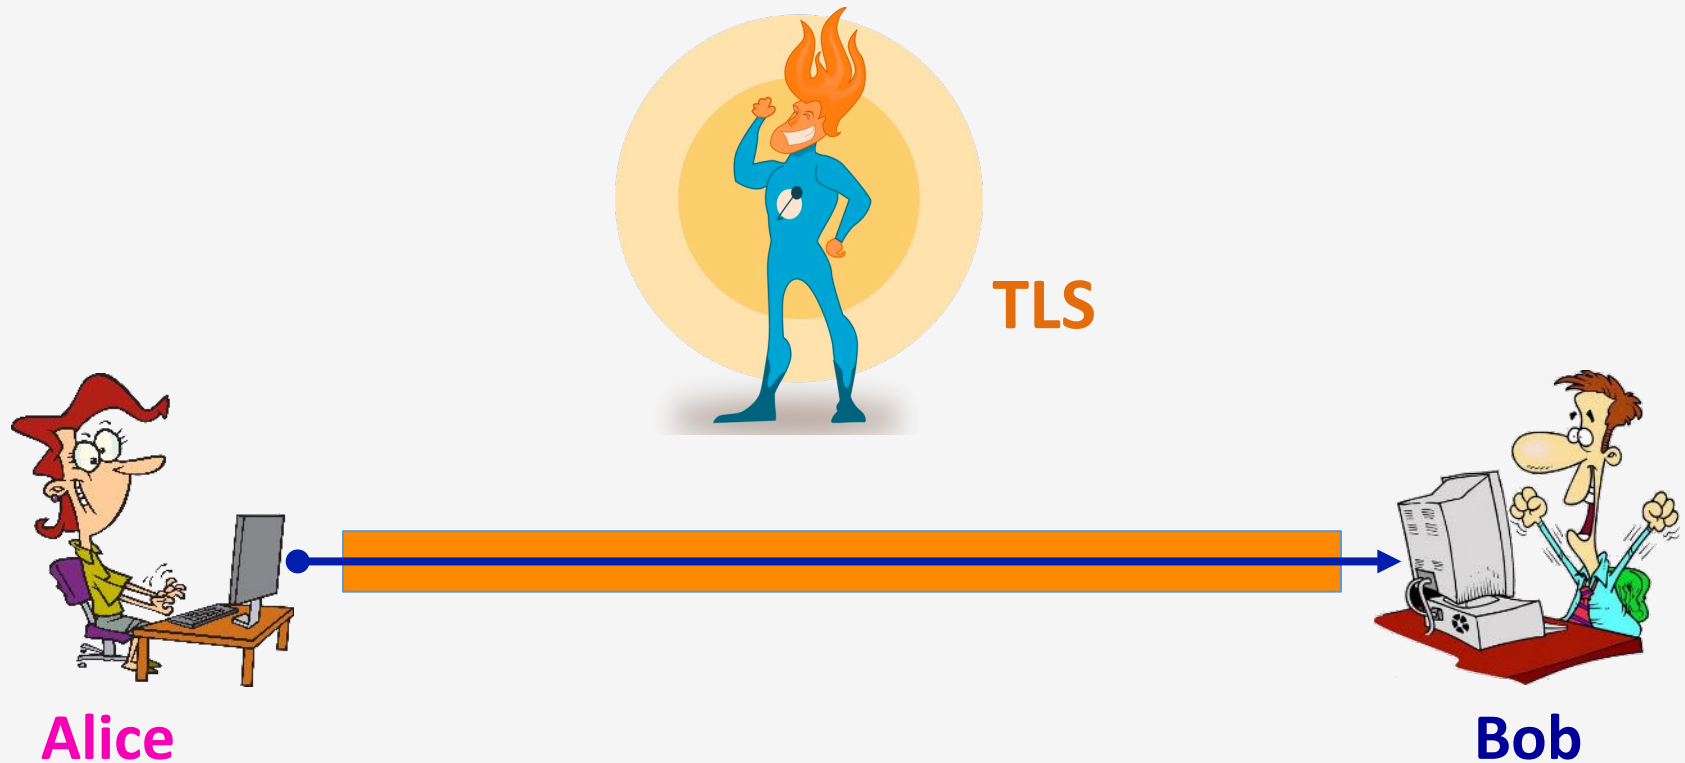
Sensor

**Computational
element**

Here's a plant. Go hack it.

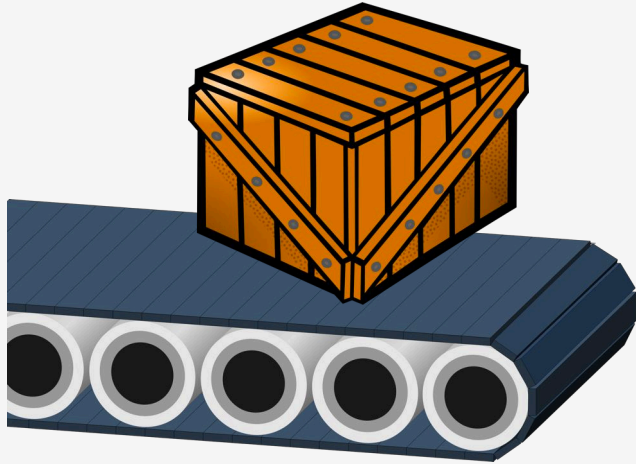


Traditional security



In traditional security TLS is the savior of all things

Encryption and safety



Help Me!



**Crypto Key Invalid:
Access Denied!!!!**

- ☐ Most of the time no one cares if you can read the data
- ☐ When the electronics stop, the physical process continues
- ☐ Rejecting a message is often not the “safe” thing to do

Physics don't stop and reboot



INERTIA

Your truck has breaks...
The massive hunk of stone doesn't

Key Take Aways



1

Industrial systems can be controlled without modifying the contents of the messages

- This can be effective even if the traffic is signed or even encrypted

2

Process data can be spoofed to make it look like everything is normal to mere humans

- This can be done despite all traditional communication security put in place

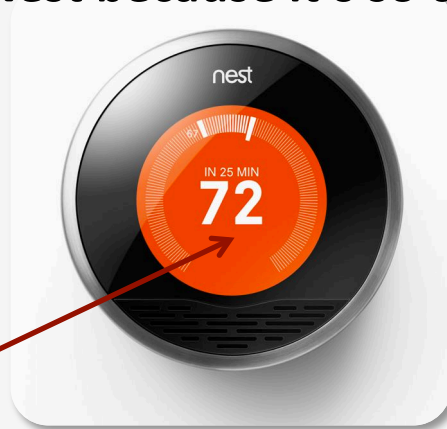


SCADA 101

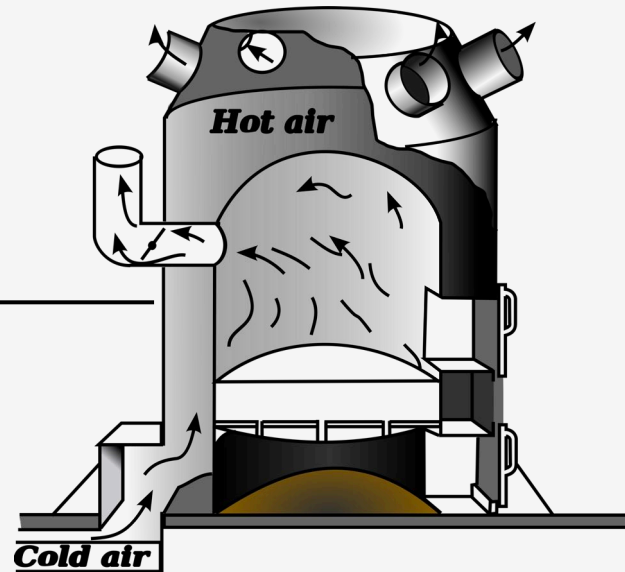
Automation

(Nest because it's so cute!)

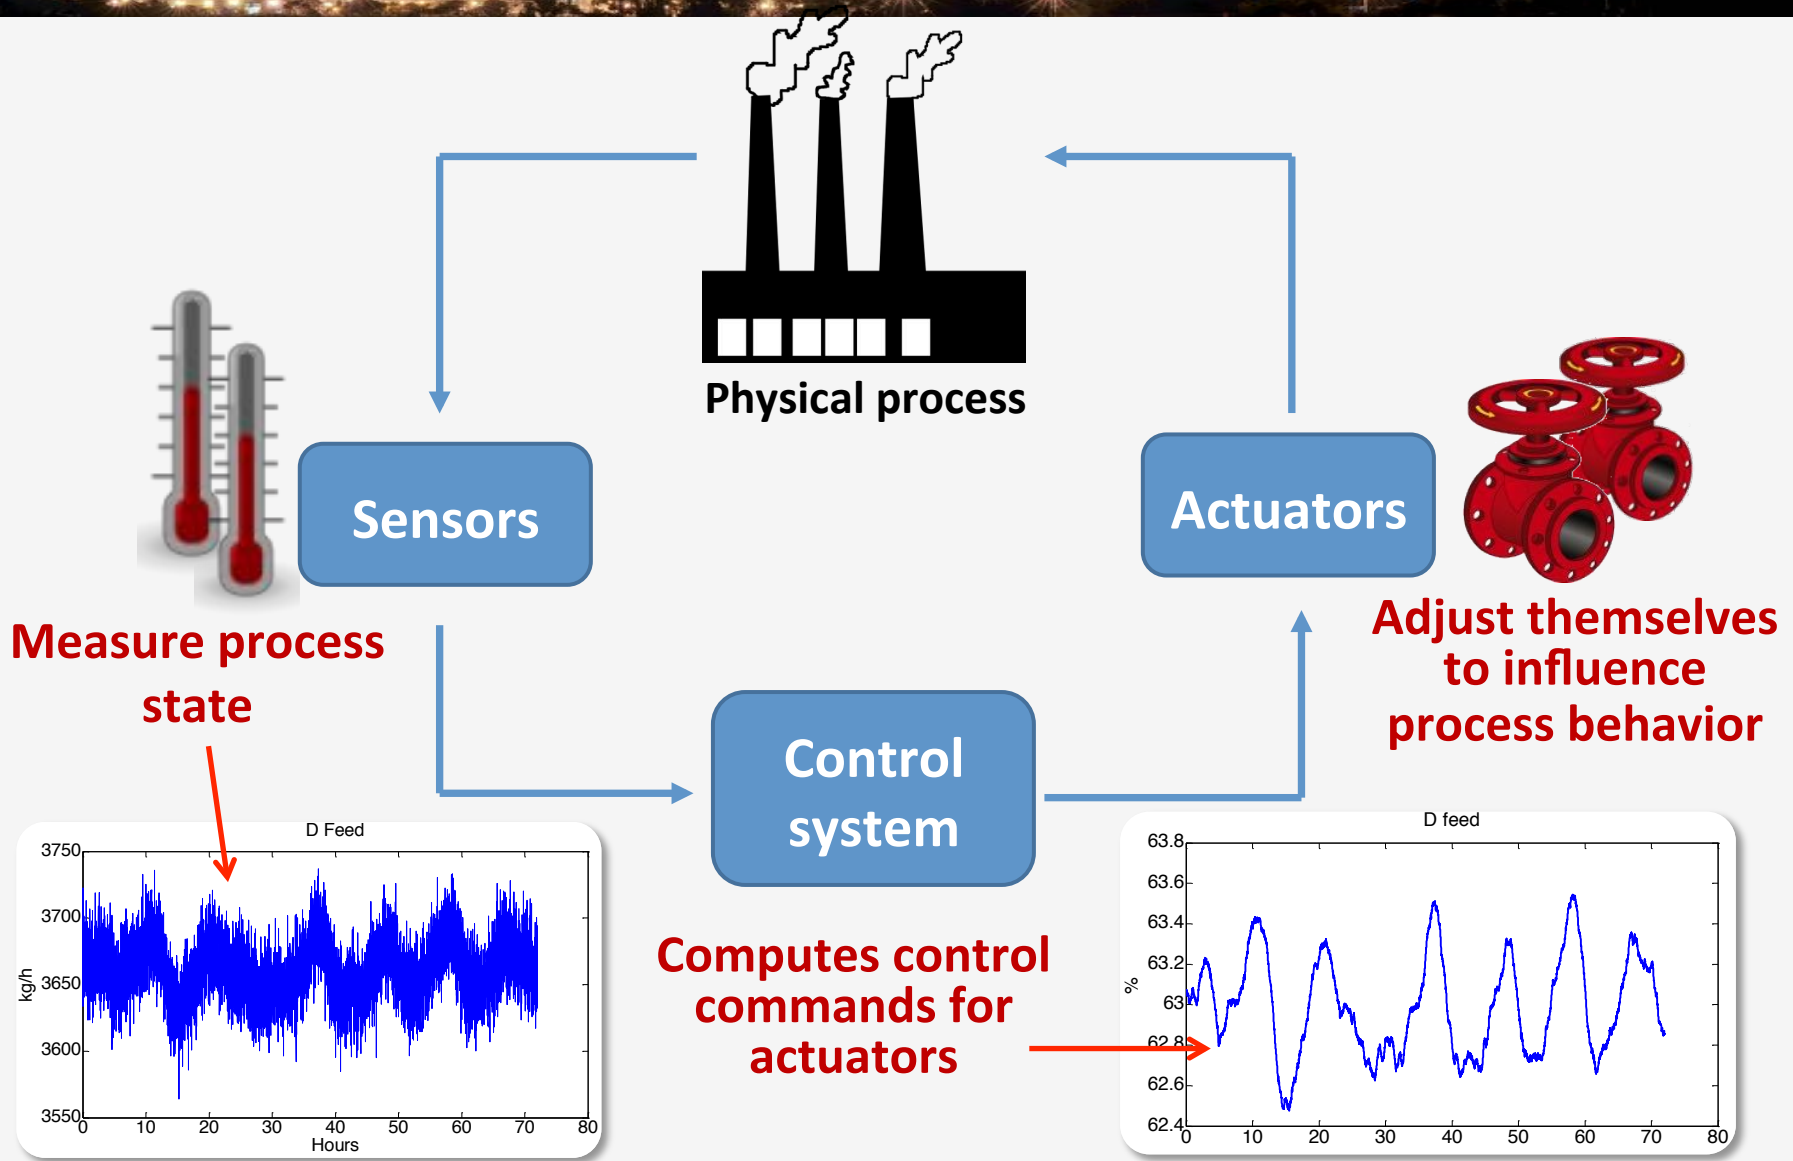
Set point



Running upstairs to turn on your furnace every time it gets cold gets tiring after a while so you automate it with a thermostat.



Control loop

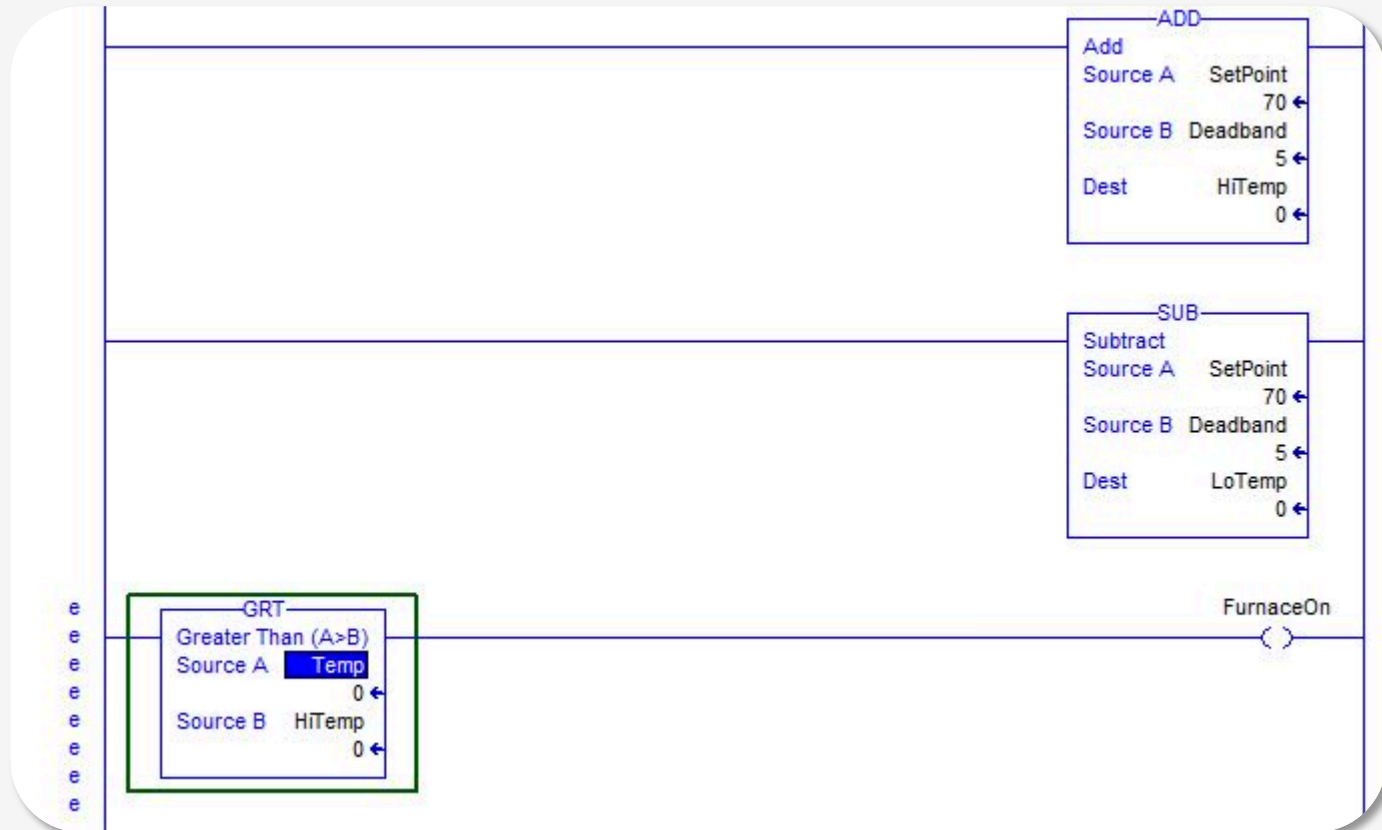


Control logic

- ❑ Obviously control logic gets more complex than your thermostat
- ❑ You'll need something bigger than a thermostat to handle it all
- ❑ Most of the time this is a programmable logic controller (PLC)
- ❑ It is programmed graphically most of the time



Control logic



Computer scientists: Noooooooooo!!!! Just give me a real language!

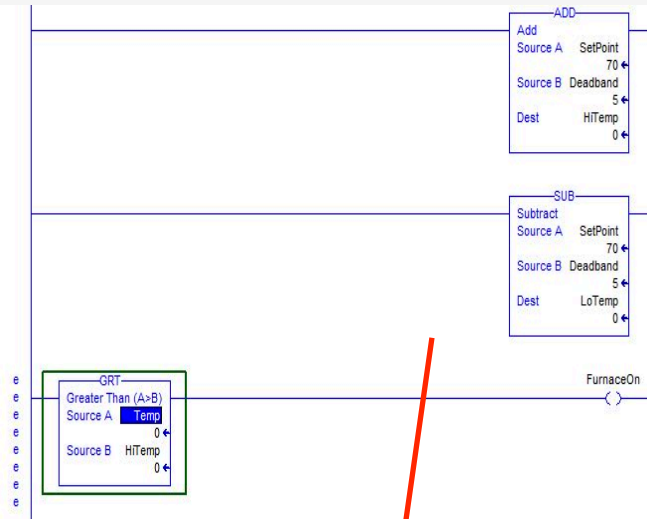
PLC internals

1. Copy data from inputs to temporary storage
2. Run logic
3. Copy from temporary storage to outputs

Sensors



Inputs



Outputs

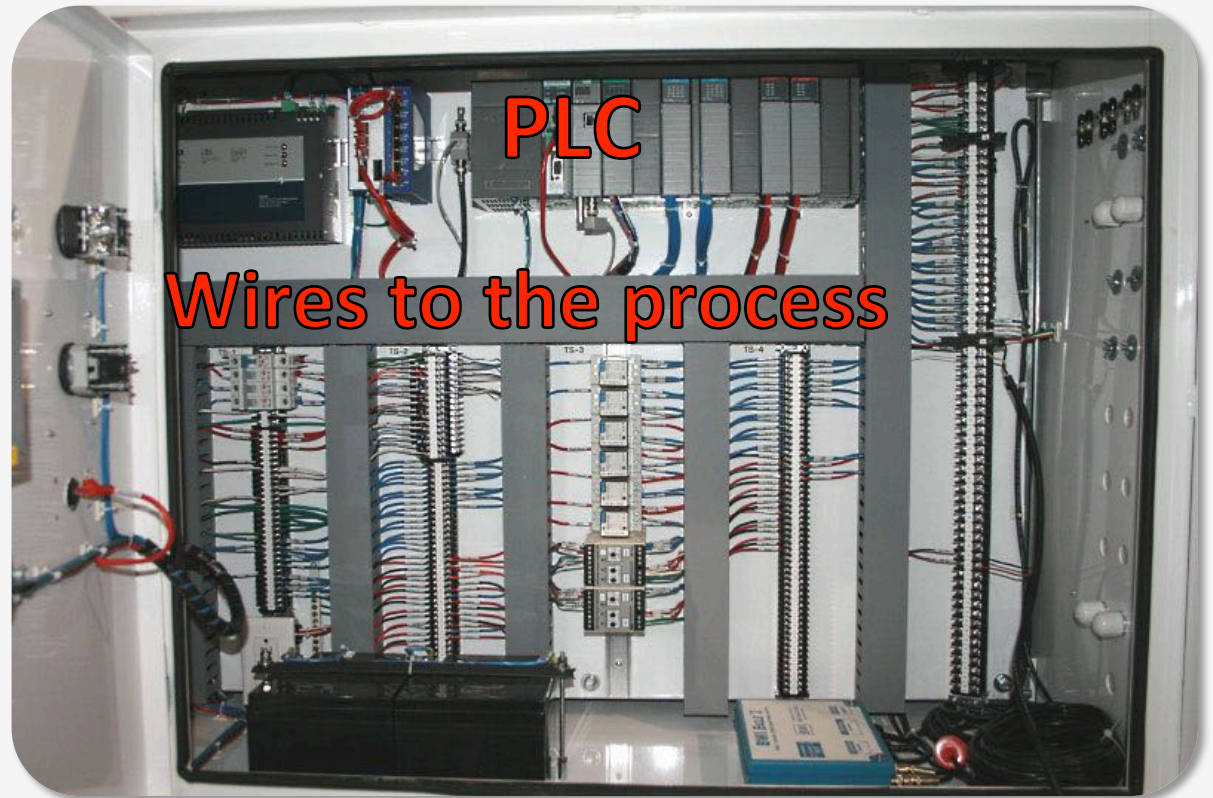
Actuators



Analog communication

- 4-20 mA
- 0-10 v
- Air pressure

Usually values
are scaled into
meaningful data
in the PLC



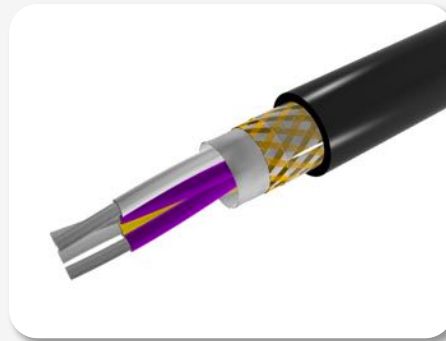
Wires are run from sensors and
actuators into wiring cabinets

Fieldbus based communications

Foundation Fieldbus



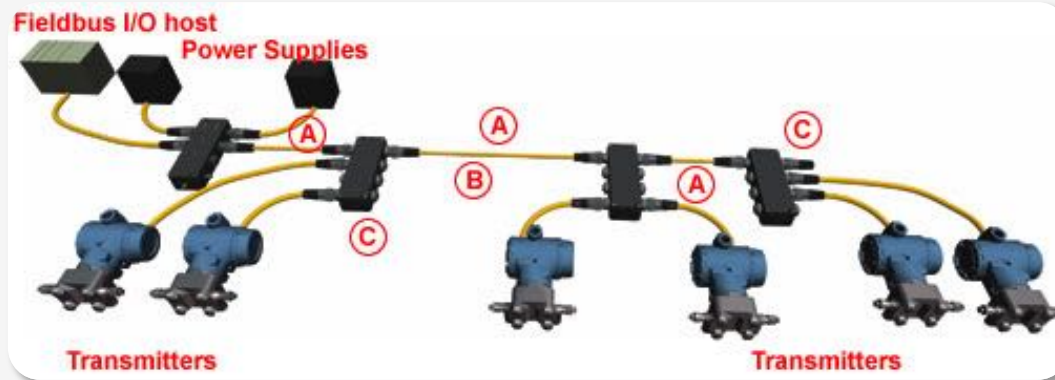
Hart



Profibus



You too can pay \$60 per meter for really bad cable



Custom network cables. Custom protocols. No TCP/IP here.

TCP/IP based communication



Industrial switch



Ethernet my old friend
(Hack meeeeeeee!!!)

Modbus, DNP, IEC850 are common protocols



Part 1

Process manipulation

Stale Data problem



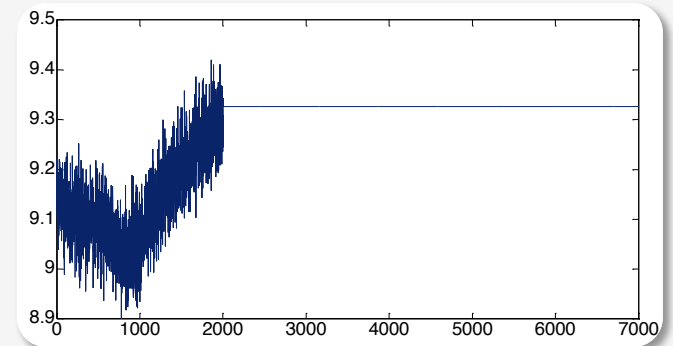
Stale Data attack

Reactor Pressure

— Without attack

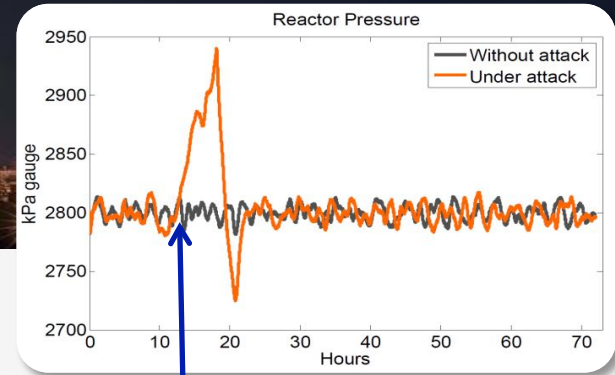
— Under attack

kPa gauge



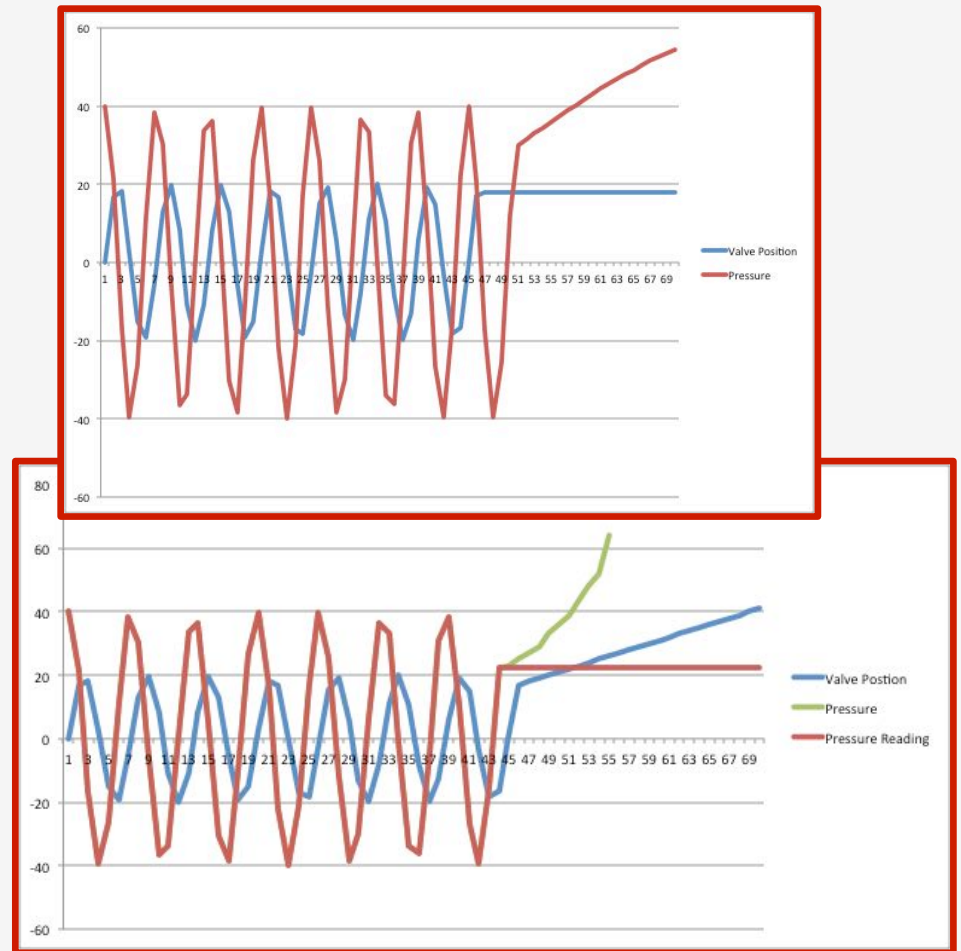
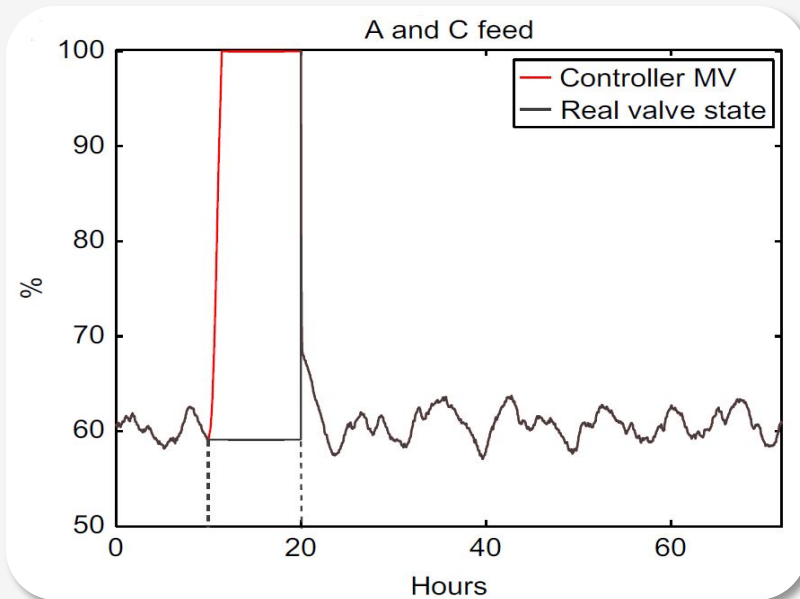
Stale Data attack

This figure illustrates a stale data attack on a reactor pressure sensor. The background shows a large industrial facility at night. The foreground features a line graph titled "Reactor Pressure" with the y-axis labeled "kPa gauge" ranging from 2800 to 2950. The graph compares two data series: "Without attack" (black line) and "Under attack" (orange line). The black line remains relatively stable around 2800 kPa. The orange line, representing the sensor under attack, shows a sharp spike to approximately 2940 kPa before returning to the baseline. A blue arrow points to the start of the attack, and a red arrow points to the end of the attack.



DoS attacks on cyber-physical systems

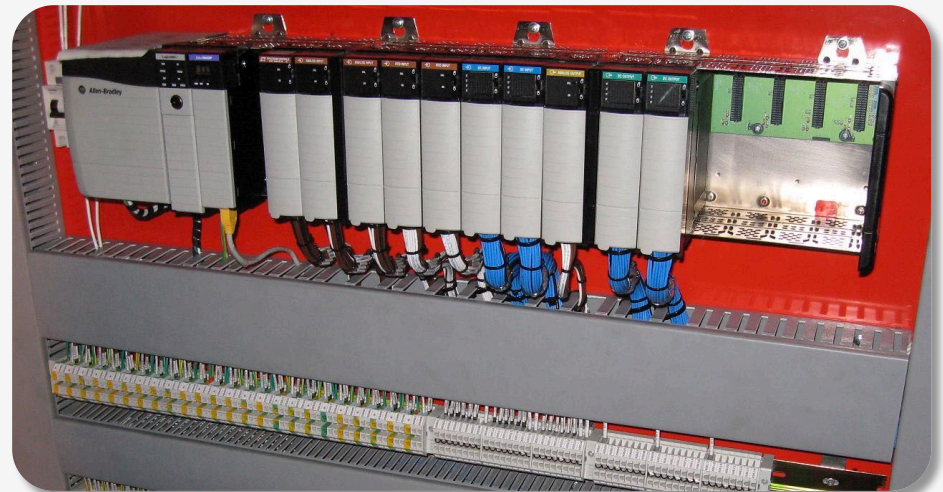
□ What to DoS: sensor or actuator?



Stale Data problem



- ❑ Process data doesn't show up every time around the logic
 - External racks may only report in every few cycles
 - **TCP/IP protocols are often report-by-exception**
- ❑ The input memory contains the last known good value



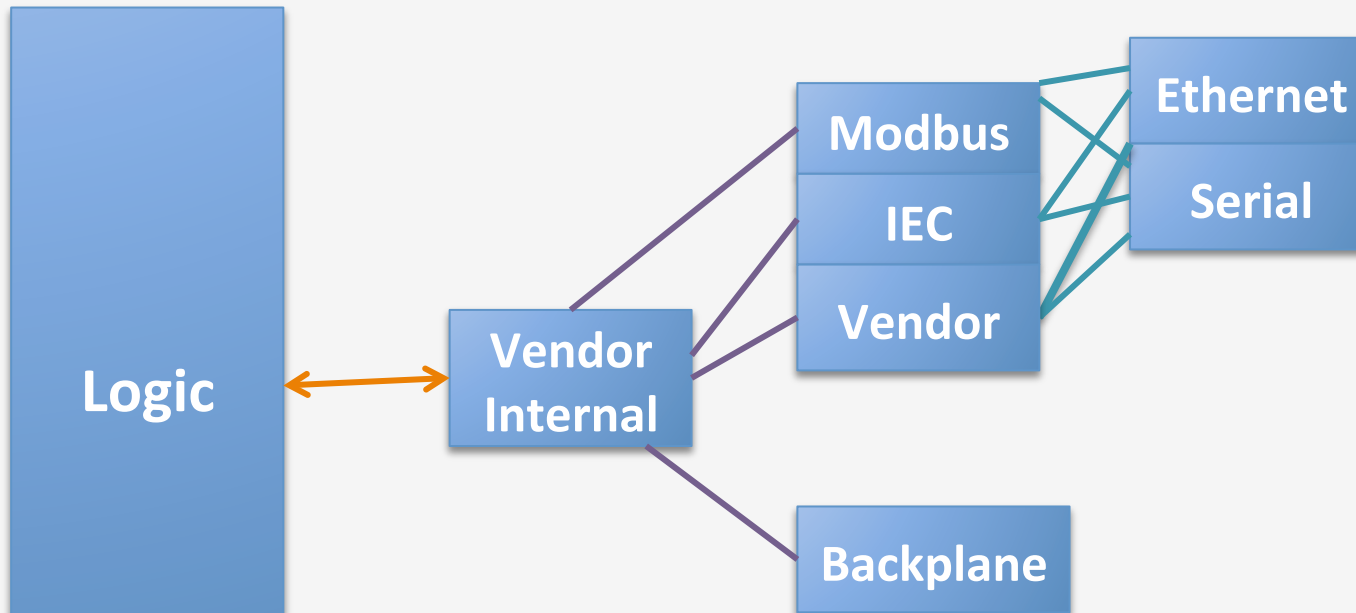
Case study



- ❑ Vendors please don't hate me again ☹
 - I kept your name secret
- ❑ This is actually a pretty typical example
- ❑ This vendor used the same style logic for all external data



Case study



Vendor Protocol Handshake - Session 4000

Vendor Protocol Handshake - Session 5000

Vendor Protocol Handshake - Session 6000

IEC Protocol Handshake

Vendor Protocol Handshake - Session 8000

Vendor Protocol Handshake - Session 9000

Case study



Length					Session ID				Sequence Number				
00	5A	5A	00	18	2F	A0	00	00	00	00	04	00	00
10	00	00	79	00	00	00	00	41	0E	08	0E	10	86
20	14	83	41	C3	0E	10	41	C6	0E	0C	01	42	30
										Number of Samples			
										Sample Value			

Case study



The result

- ☐ You can freeze all points for a particular session with a UDP packet by advancing the sequence number
- ☐ You can keep session alive and by sending a UDP packet every 30 seconds to any interface

DoS by Eireann Leverett & Matt Erasmus

- ❑ Eireann Leverett & Matt Erasmus showed bugs in industrial switches
 - With access to the switch only ACK messages could be passed
 - The link would show up as healthy
 - No data would be updated





All mighty DoS attacks

Plants for sale

From LinkedIn



+ Follow Tommy

Used VAM - Vinyl Acetate Monomer plant for sale & relocation! If any interest, please contact me!

Tommy Heino

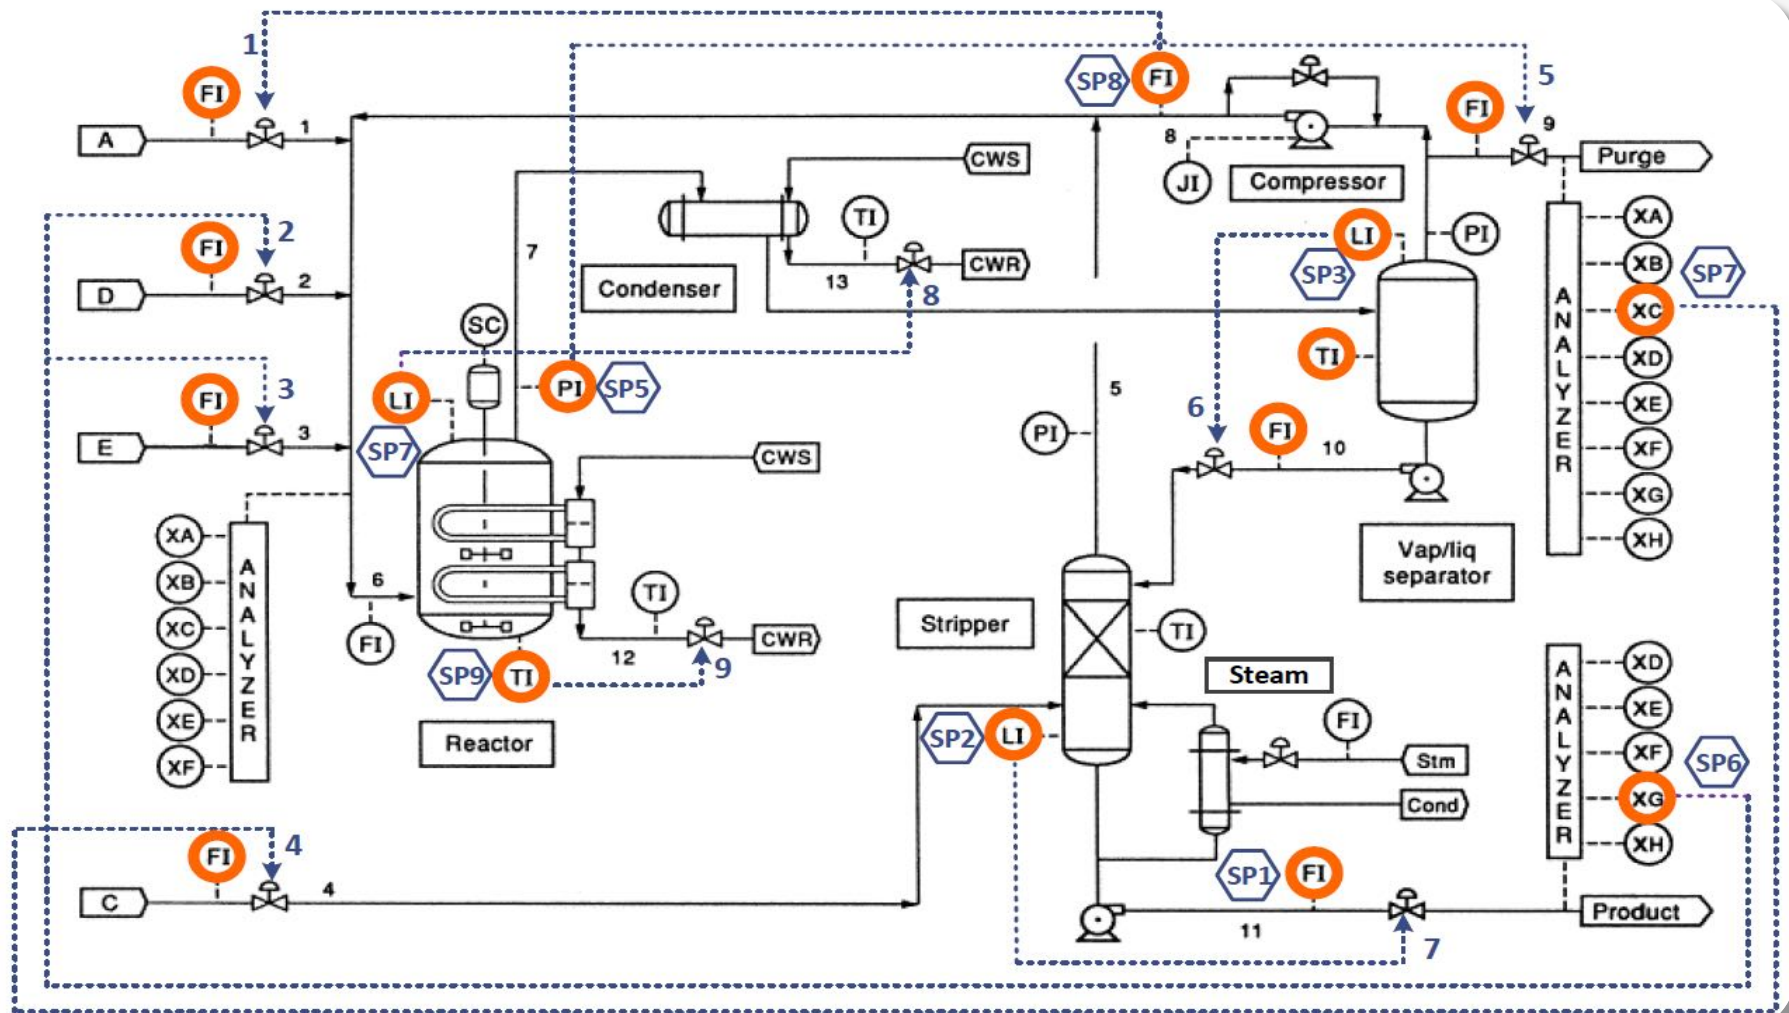
Industrialist & Entrepreneur, Owner, XHL Business Engineering

Top Contributor

Like • Comment (4) • Share • Follow • 3 m

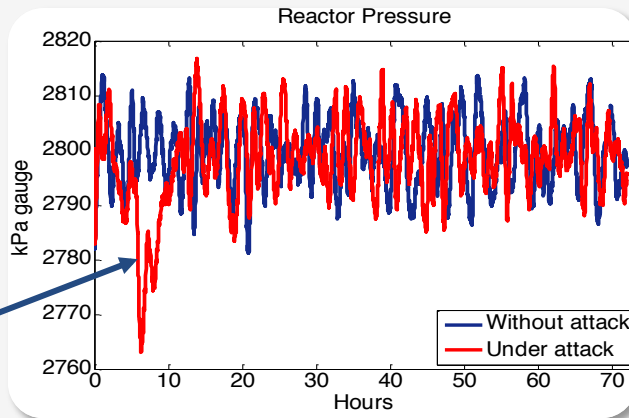


Tennessee Eastman (TE) chemical process

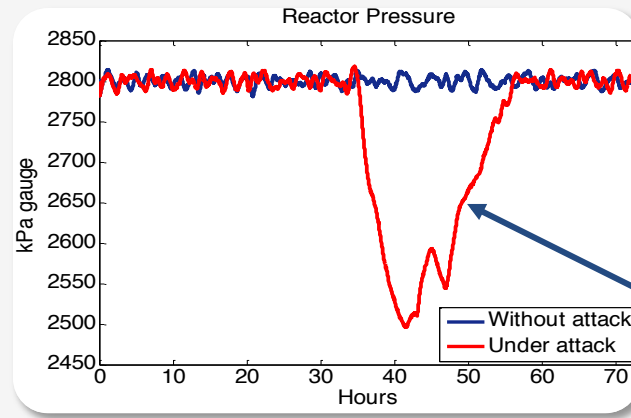


Timing of the DoS attack

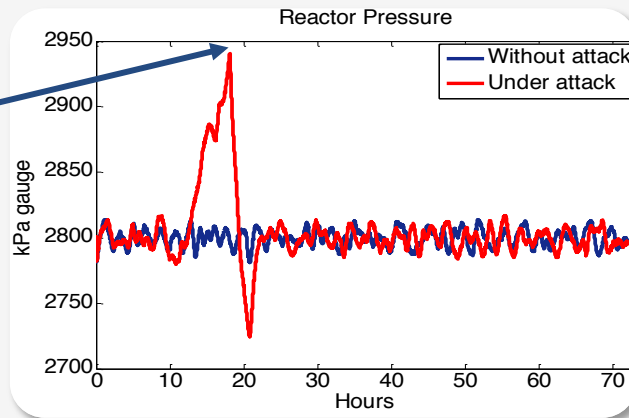
Ordinary
glitch



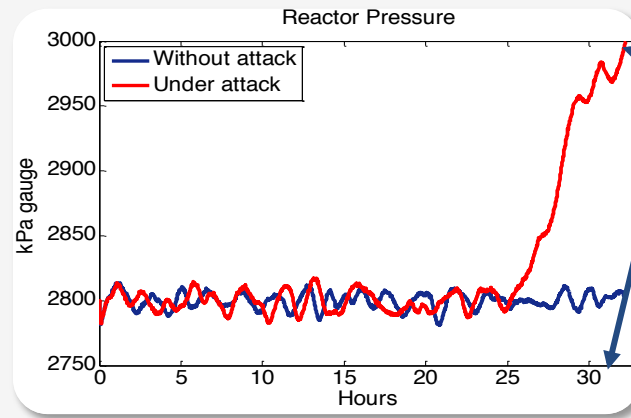
Economic
inefficiency



Near miss
(almost
safety
accident)



Safety
shutdown

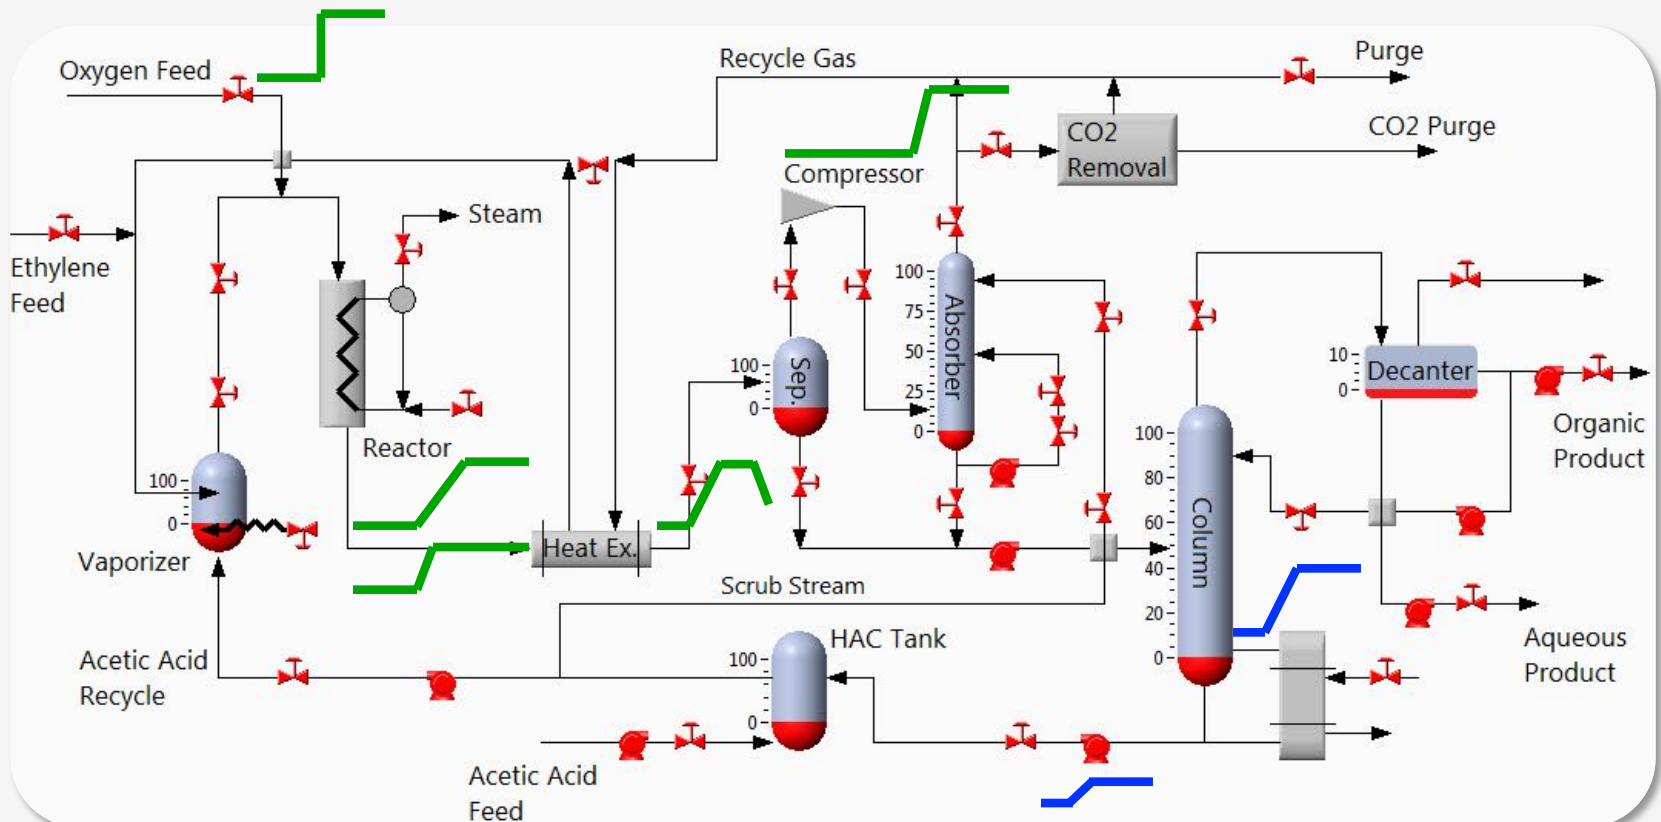


Impact of 8h long DoS attacks on reactor
pressure sensor at random time

Time to attack?

1 Derive a model a of the plant's dynamic behavior

- We have some ideas – ongoing research

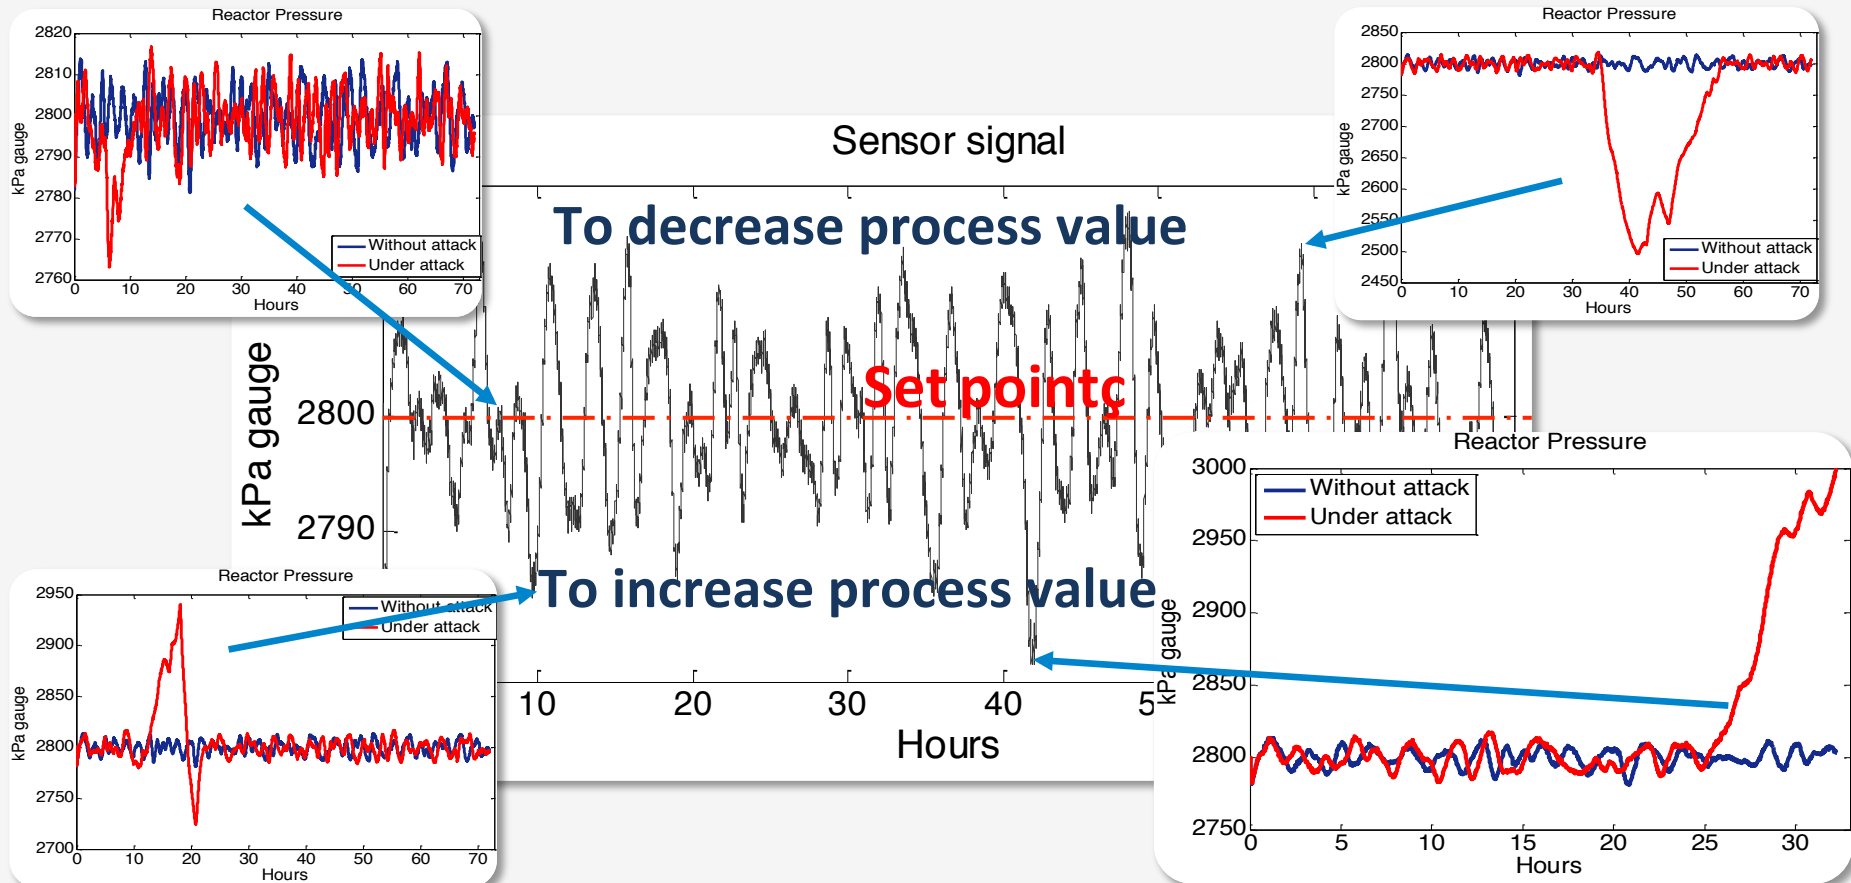


Time to attack?

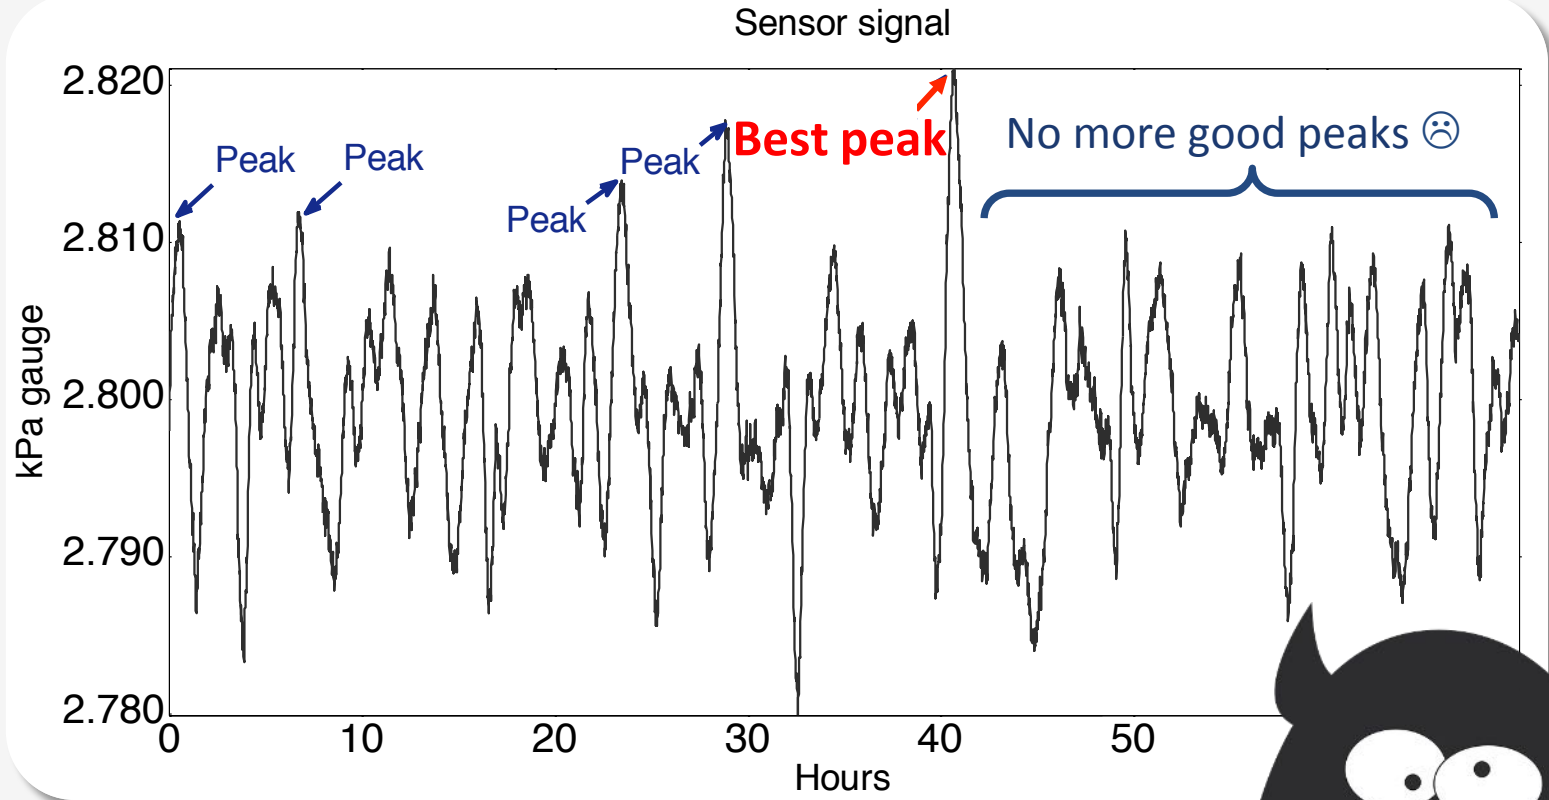
2

Educated guess

- Response of the process depends on the value of DoS value




Quest for the peak



- ☐ **REAL TIME** decision making problem
- ☐ Searching for the “**BEST**” peak
- ☐ Achieving results within some time horizon



Avocado problem



Not yet
Not yet
Not yet
Not yet
Not yet
EAT ME NOW
Too late.
- Avocados

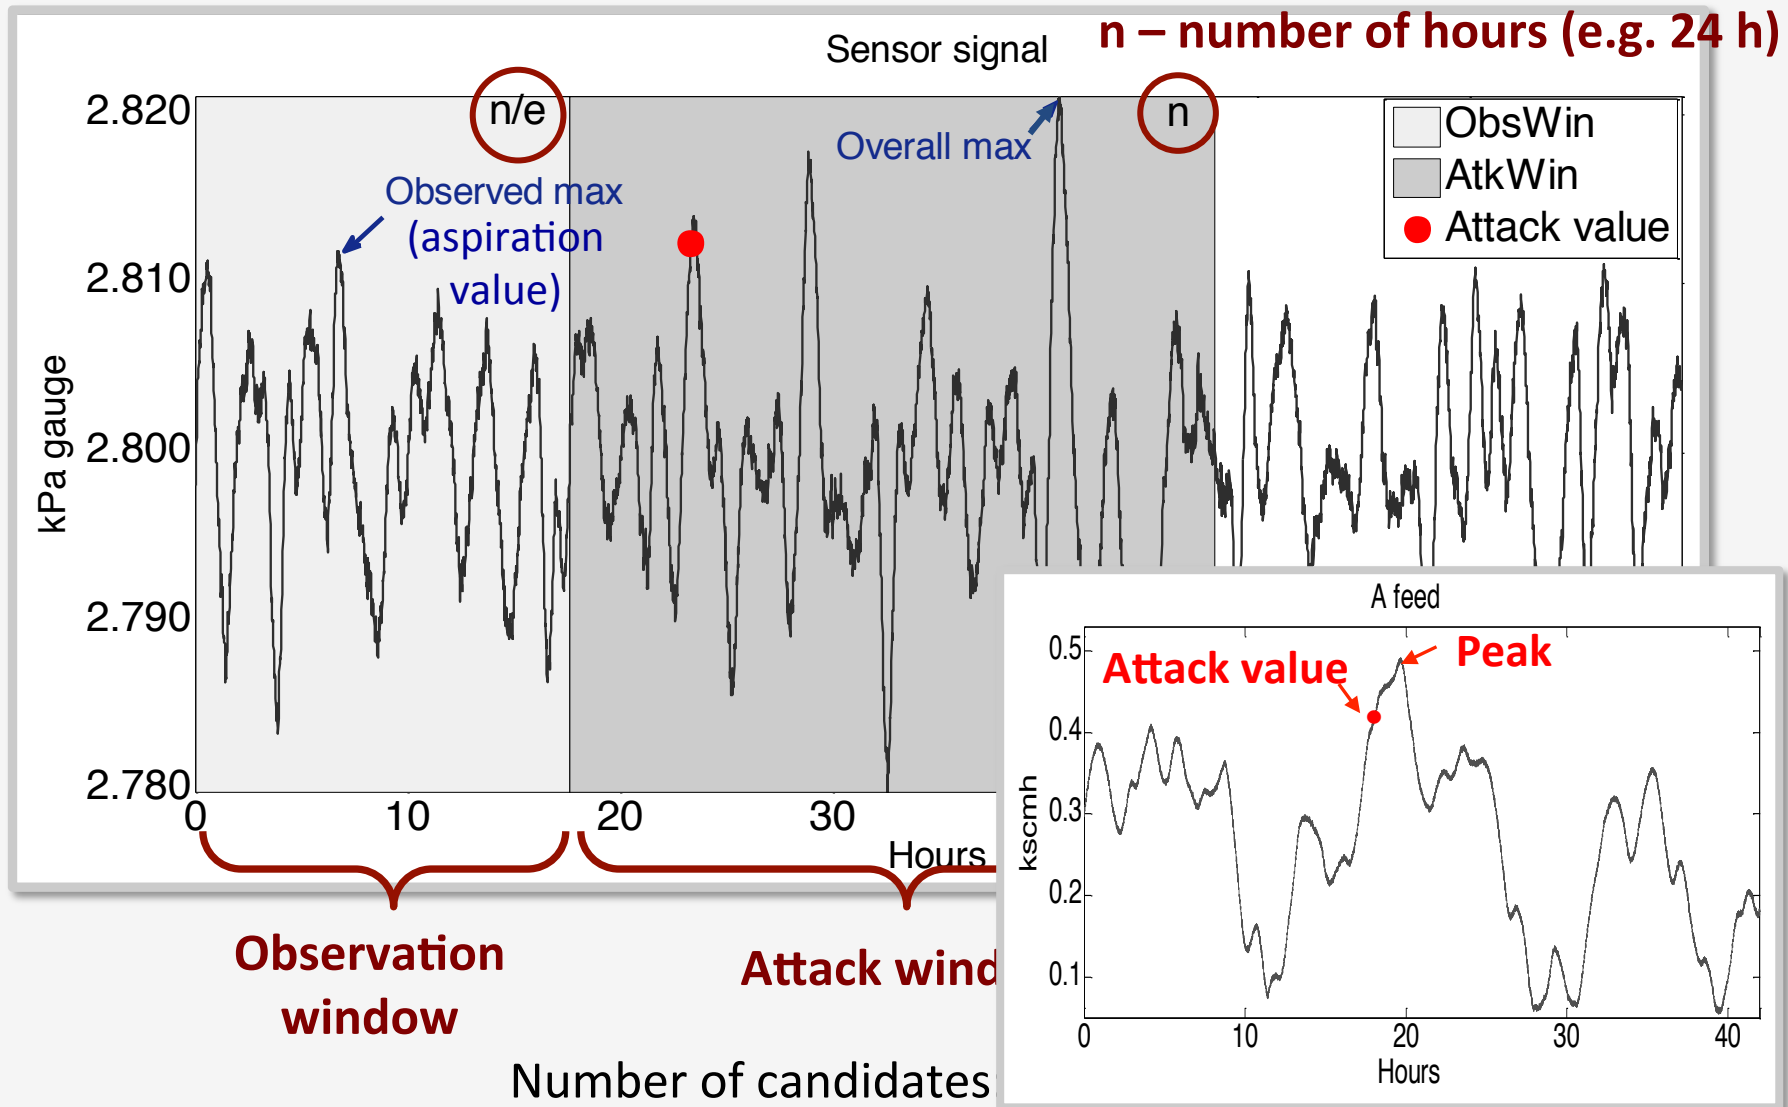
Best Choice Problems

- ❑ Problem of choosing the time to take a particular action
 - Based on sequentially observed random variables
 - In order to maximize an expected pay off
- ❑ Applied in a wide range of applications including financial
 - Best time to buy or sell stocks

Secretary Problem



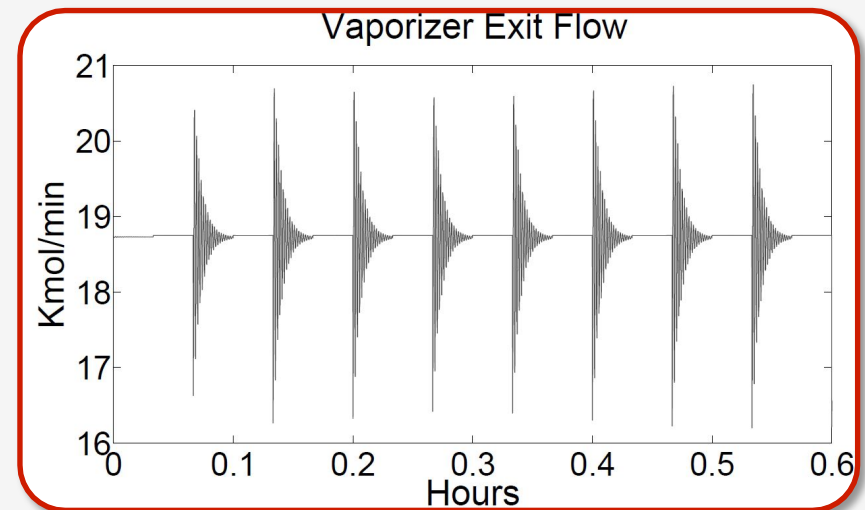
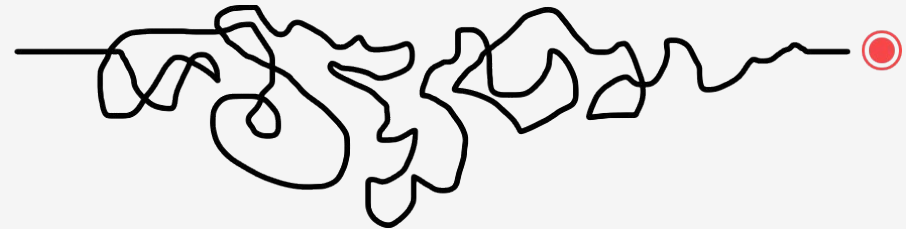
Secretary Problem: sensor signal



We are not successful yet

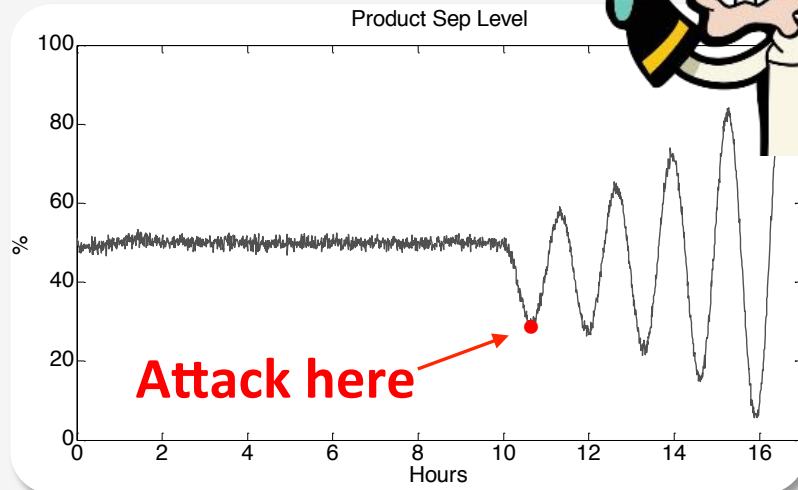
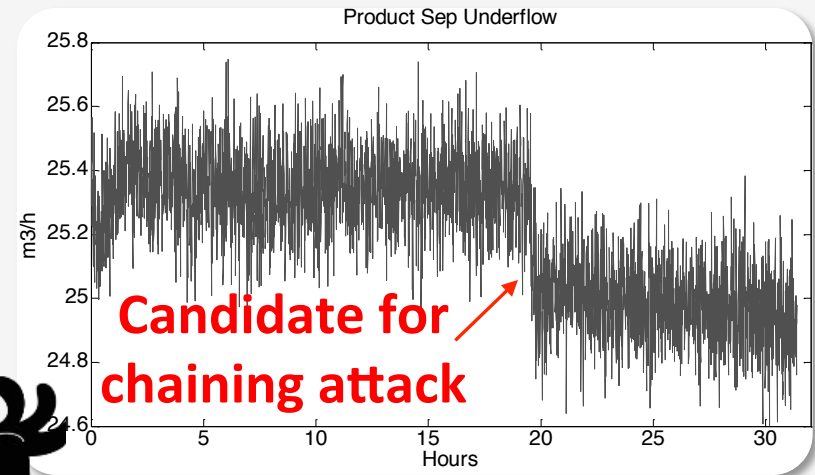
Sensor		Safety time, h
A-feed	min	22.22
	max	
E-feed	min	4.29
	max	2.83
Recycle flow	min	4.39
	max	9.17
Reactor pressure	min	8.56
	max	
Reactor level	min	2.37
	max	2.73
Reactor temperature	min	1.34
	max	0.65

❑ Process dynamic is highly non-linear (???)



Accelerate it: chaining attacks

- ❑ Chain DoS attacks: on sensors
- ❑ Use change detection algorithms (e.g. CUSUM) to detect state change



- ❑ Chain two DoS attacks: on sensor and actuator
- ❑ Safety time **3.43 h** vs. **12.03 h** in case of direct attack



Part 2

Attack concealment

Spooof scenarios

❑ „Record-and-play-back“

- Used in Stuxnet ;-)
- Storage requirements

❑ Derive process model

- Requires knowledge, CPU cycles and storage

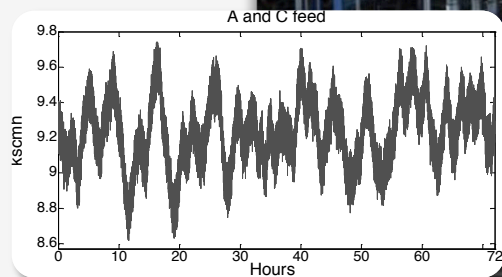
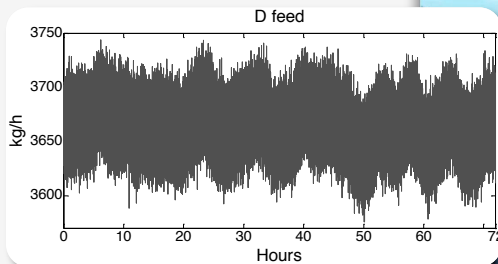
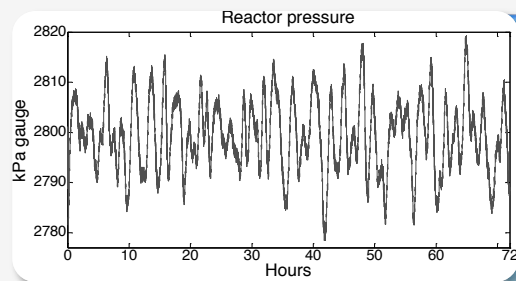
❑ Crafted sensor signals

- Reconstruction of sensor data features



Process data

Process data originates in the physical world

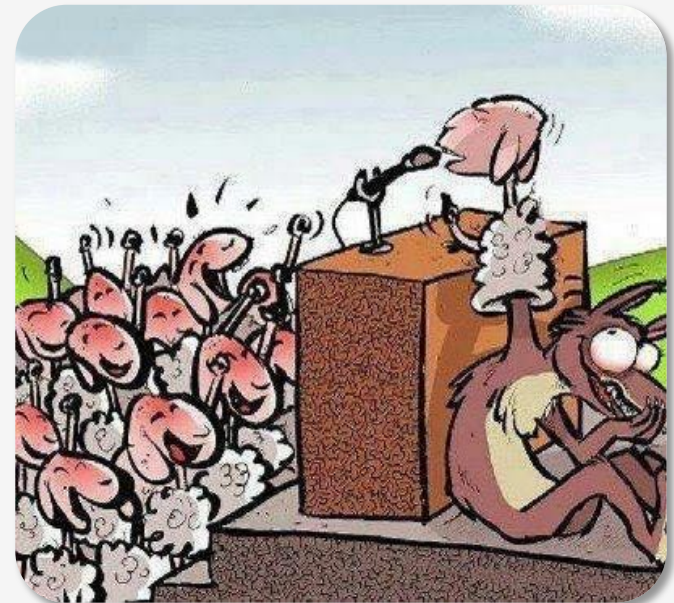


Data veracity

❑ So what if sensor readings are manipulated at source, **BEFORE** they are handed to the IT infrastructure?

- And wrong data securely transferred to the final destination (authenticated and integrity protected)...

Veracity: data security property that a statement about an aspect relevant in a given application truthfully reflects reality



As you always knew it: NEVER TRUST YOUR INPUTS!!

Data veracity violation

1

Manipulation of the physical process

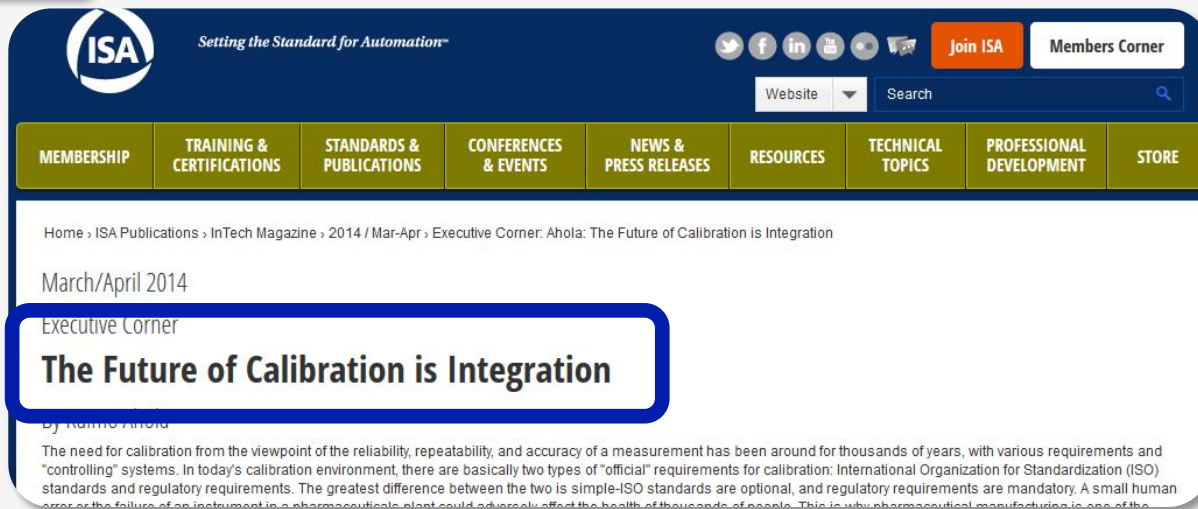
- ❑ Equipment connected to each to each other over digital communication and physics of the process
- ❑ Components can influence each other even if their control loops do not communicate electronically



Data veracity violation

2

Sensor miscalibration



International society of
Automation

InTech, ISA magazine
April 2014

HIMA presentation,
October 2014

Safety Instrumented Systems

- Due to a known bug at the engineering Software, all scaling of the SIS AI got altered to 0 to 100% automatically
- Altered values got loaded and activated automatically based on an unknown Bug at the same System

Data veracity violation

3

Data spoofing on microcontroller

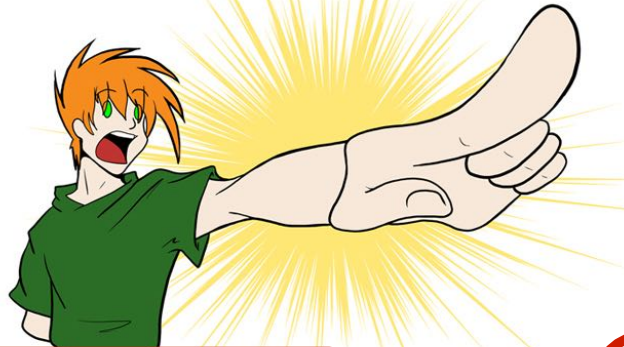
- ❑ Jason Larsen's presentation at Black Hat'14
- ❑ Hiding entire attack in a pressure meter
 - Kilobytes of memory (total)
 - Very little CPU power
 - Kilobytes of flash (total)



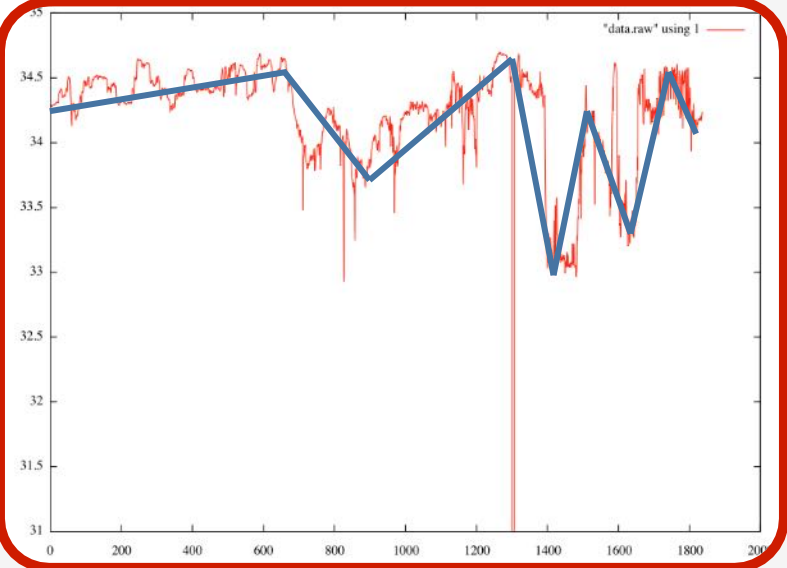
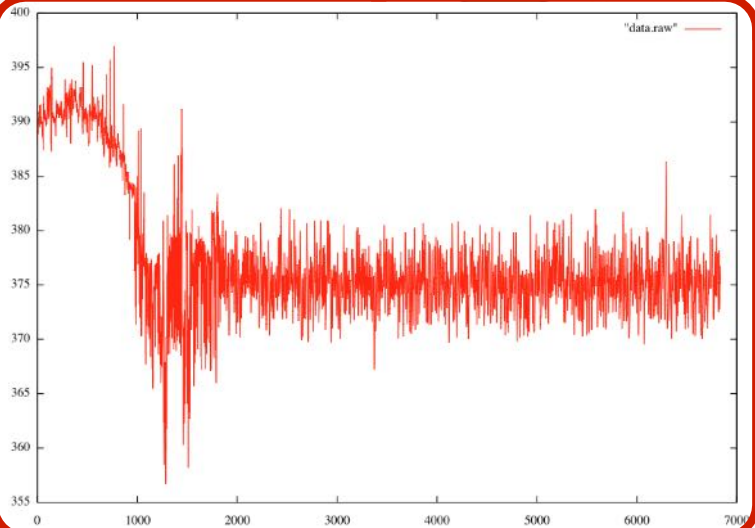
Black Hat

LOOK!

A Distraction!



Think of the process data as a set of triangles. Triangles are cheap and easy



Last year...

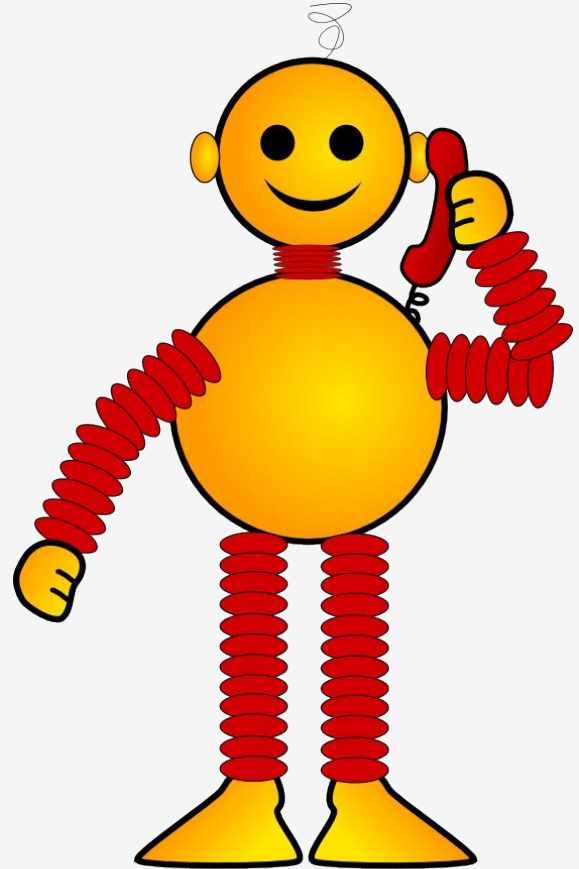


SCADA

A SERIES OF TRIANGLES



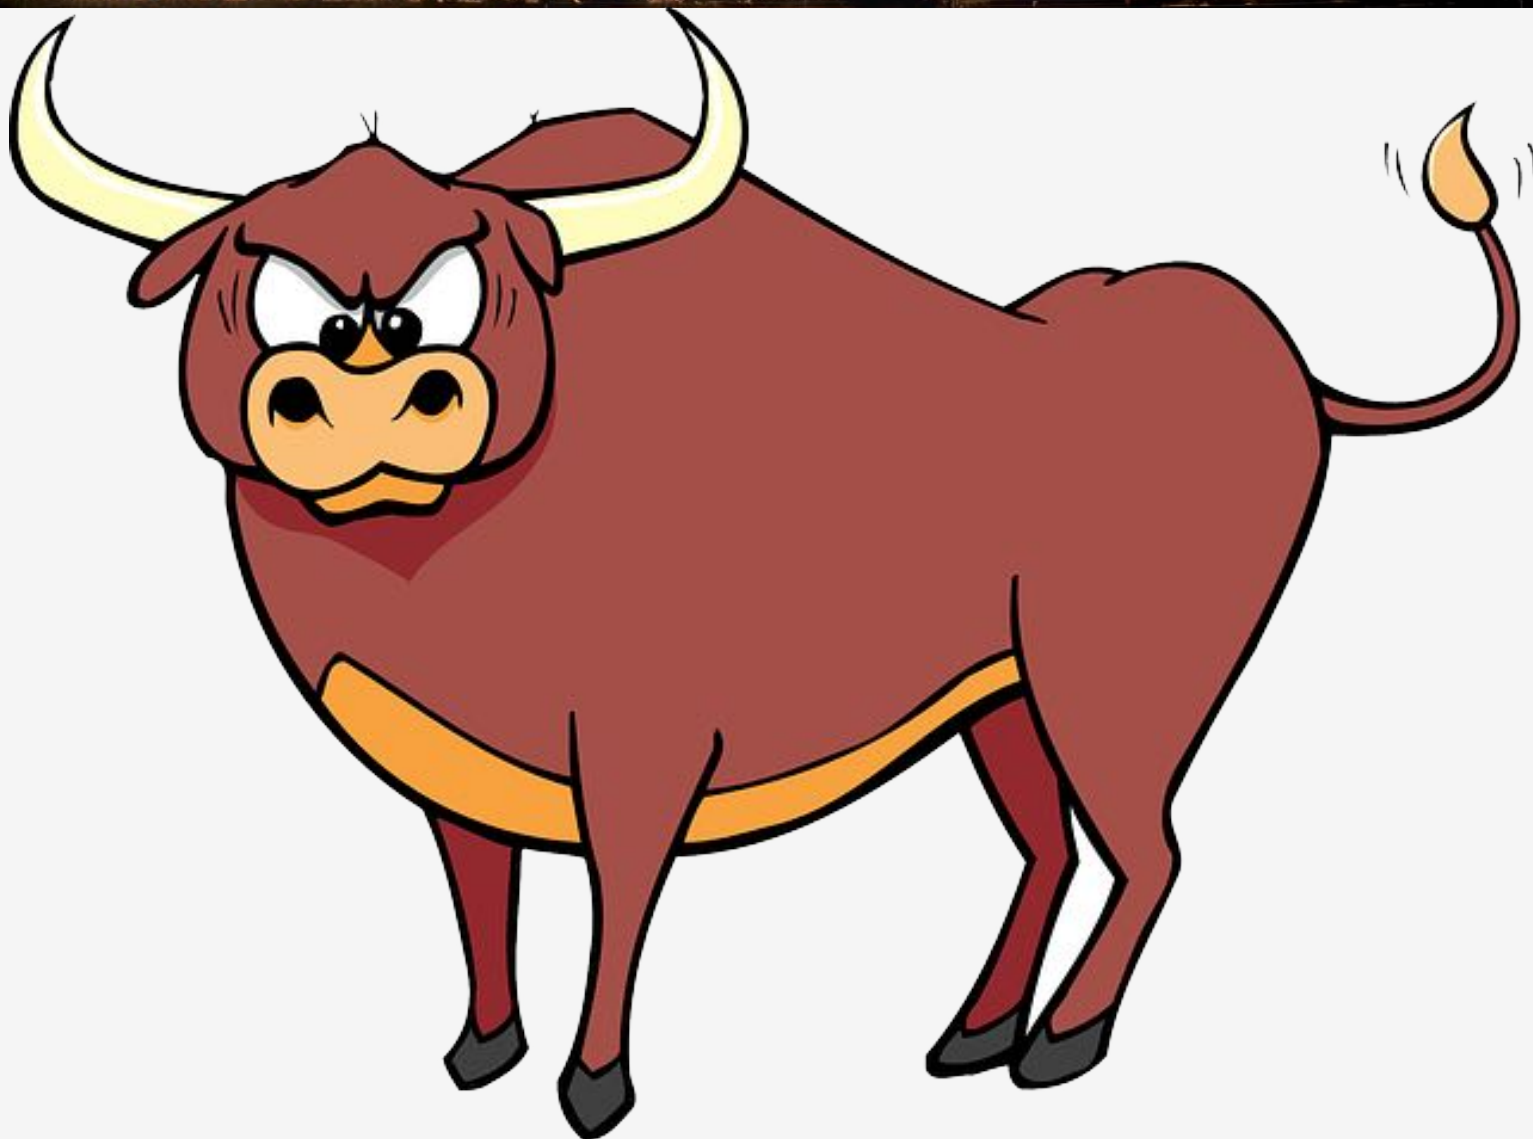
Some time after Black Hat



Some time after Black Hat



Some time after Black Hat



2 evil algorithms

Algorithm 1 Runs Analysis

```
1: procedure EXPLORE
2:   signal  $\leftarrow$  signal to analyse

3:   while not an end of signal do
4:     while moving up do
5:       runs ++
6:       value = sum(changes)
7:       if direction change then
8:         positivesruns(runs) ++
9:         positivesvalues(runs) = value

10:    while moving down do
11:      runs ++
12:      value = sum(changes)
13:      if direction change then
14:        negativesruns(runs) ++
15:        negativesvalues(runs) = value

16:    if no change then
17:      nils ++
18:  return runs, values, nils
```

▷ 1: analyse phase

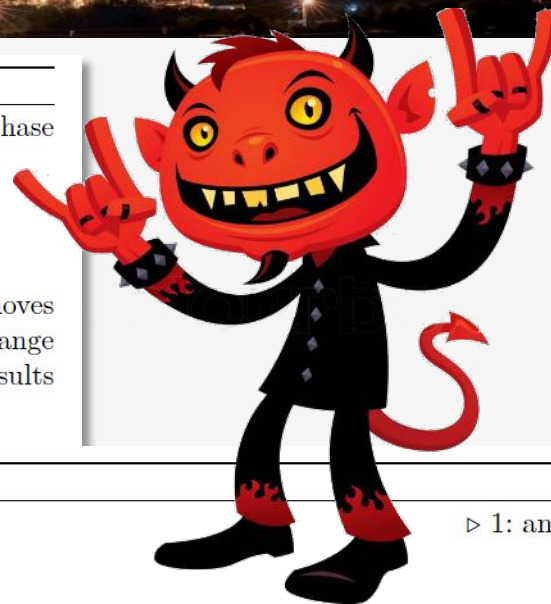
▷ count positives moves
▷ positive steps change
▷ save results

Algorithm 2 Triangles

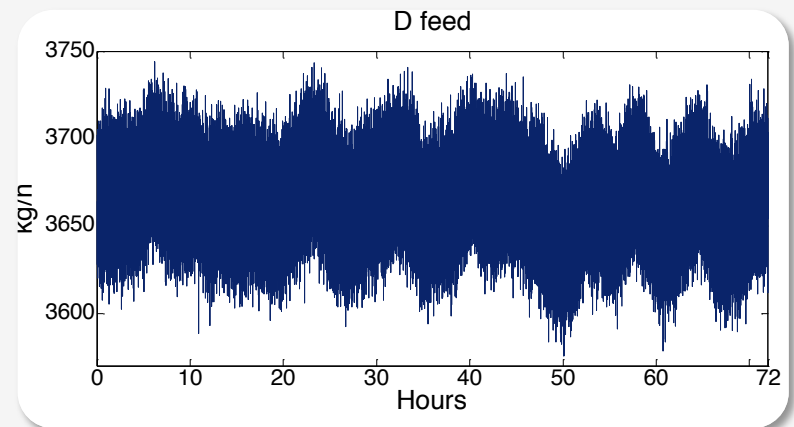
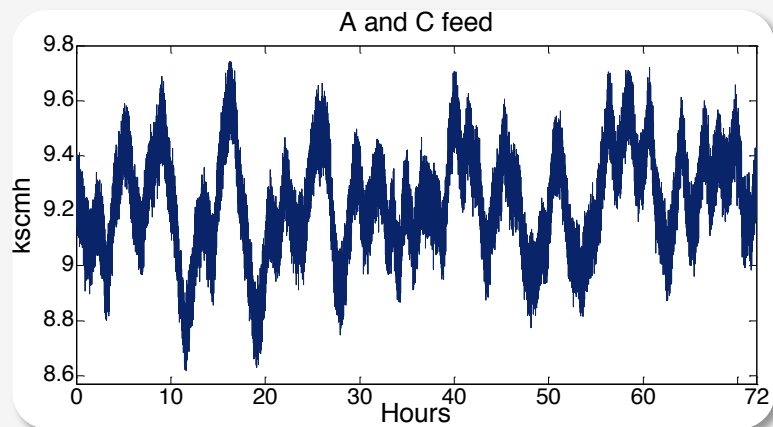
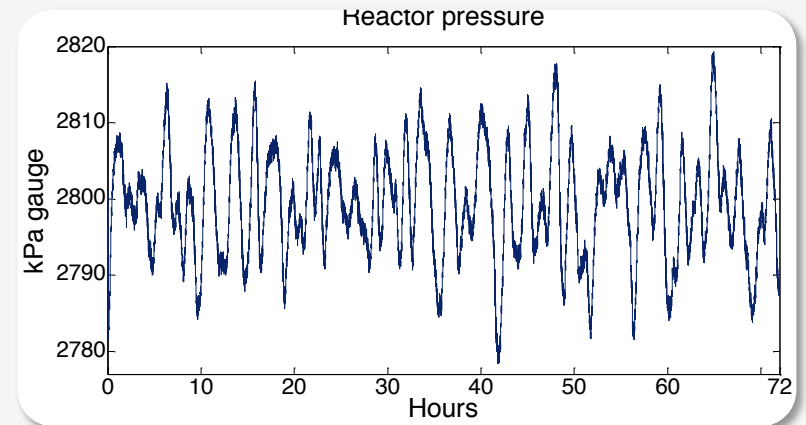
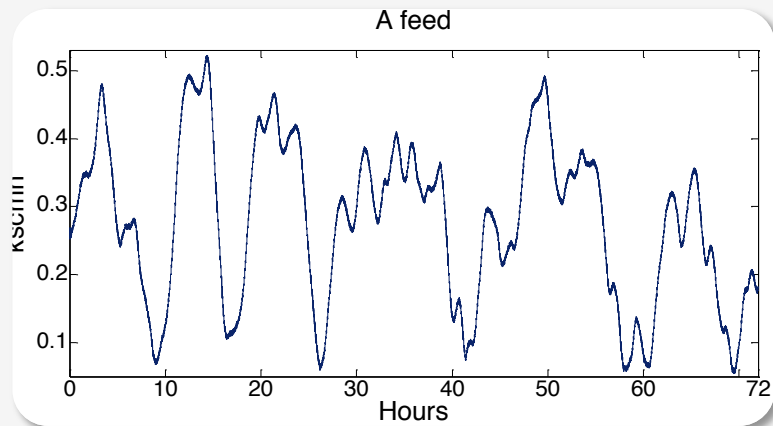
```
1: procedure EXPLORE
2:   signal  $\leftarrow$  signal to analyse
3:   window  $\leftarrow$  learning window
4:   noiselvl  $\leftarrow$  noise parameter

5:   step = window * 10
6:   topslope = -999.99
7:   bottomslope = 999.99
8:   while not an end of signal do
9:     if first elements then
10:      current = value
11:      index = 1
12:     while index < window do
13:       upperslope = (current - (last + noiselvl)) / index
14:       lowerslope = (current - (last - noiselvl)) / index
15:       if upperslope > topslope then
16:         topslope = upperslope
17:       if lowerslope < bottomslope then
```

▷ learning phase of *i* - *th* bucket

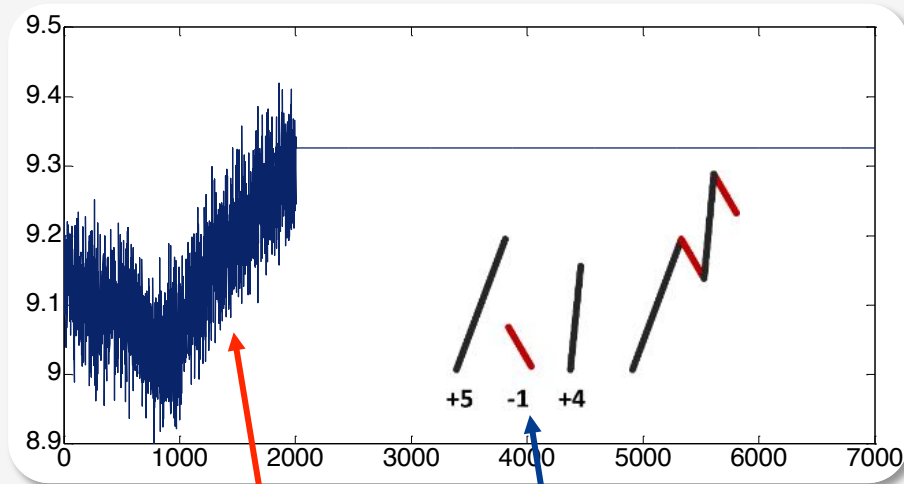
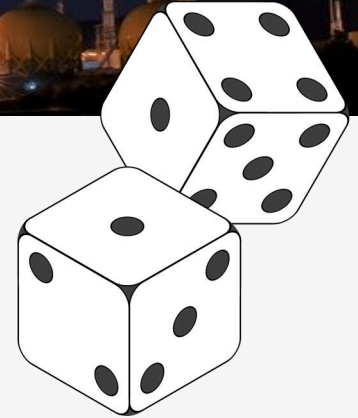


All sensor signals are not the same



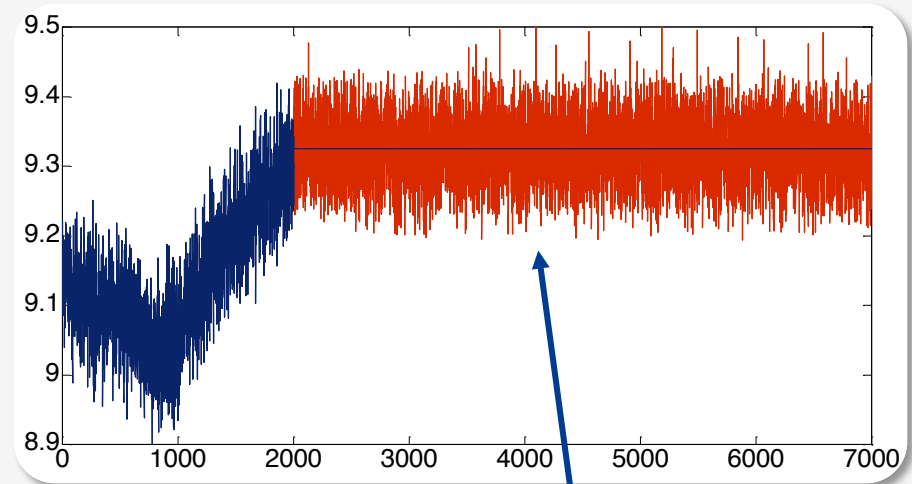
Sensor noise

- ❑ Runs analysis: treats noise as pseudo-random sequence



**Learning
phase**

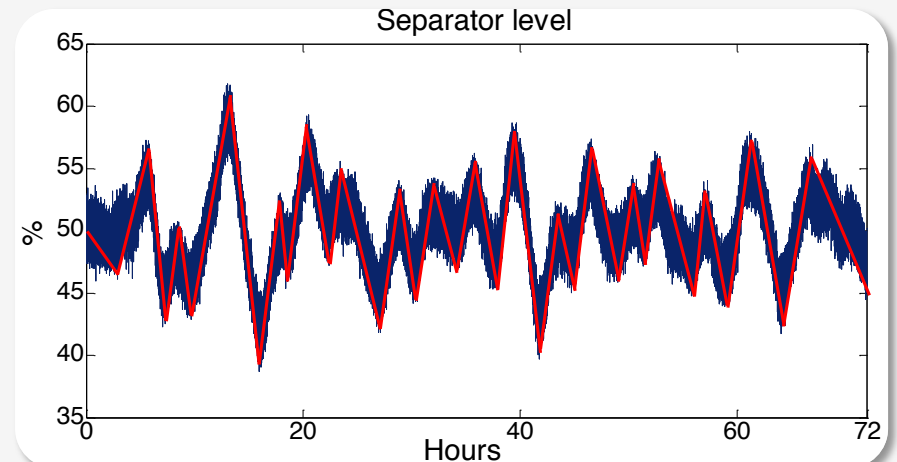
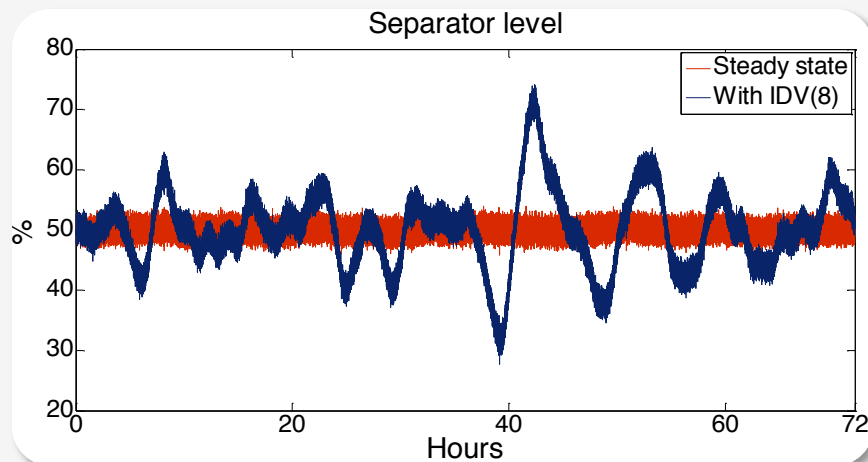
**Extracted
“runs”**



Believable noise

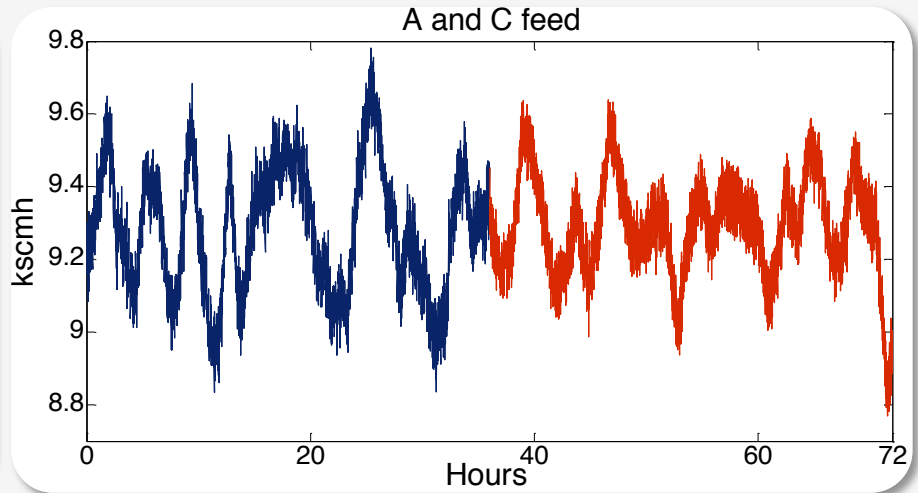
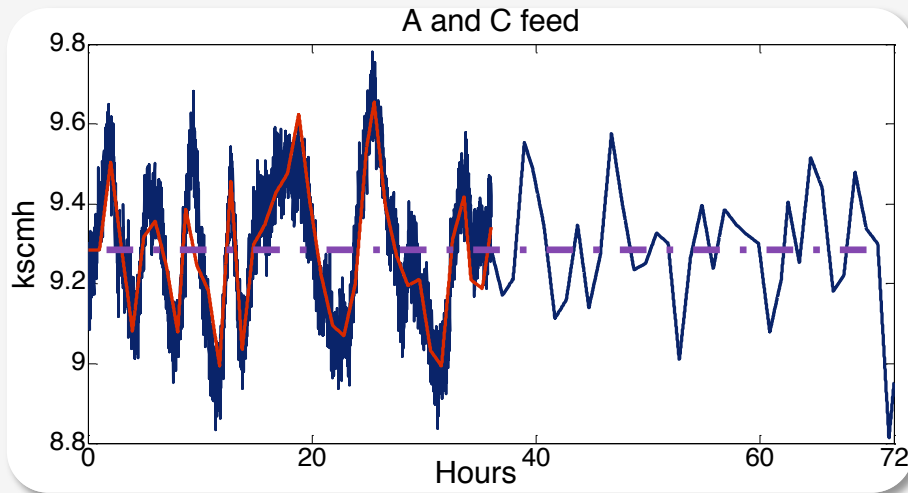
Sensor dynamic behavior

- Line segment (triangle) approximation for extracting process dynamic



Final result

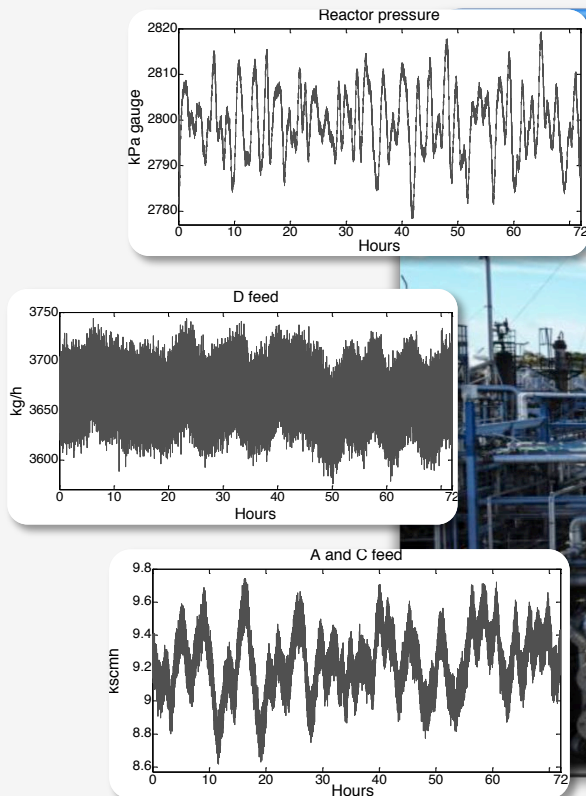
Resulted spoofed signals are extremely accurate



Find X differences

Semantic-free approach

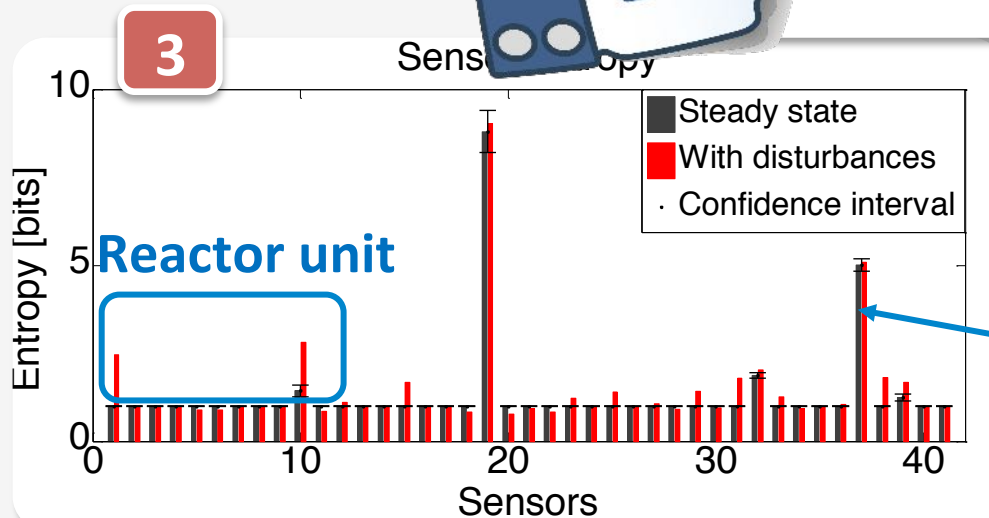
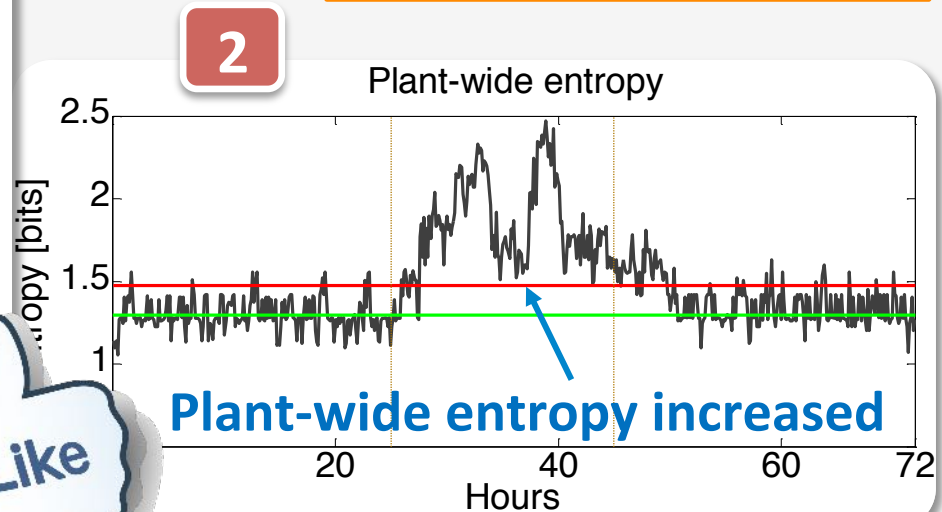
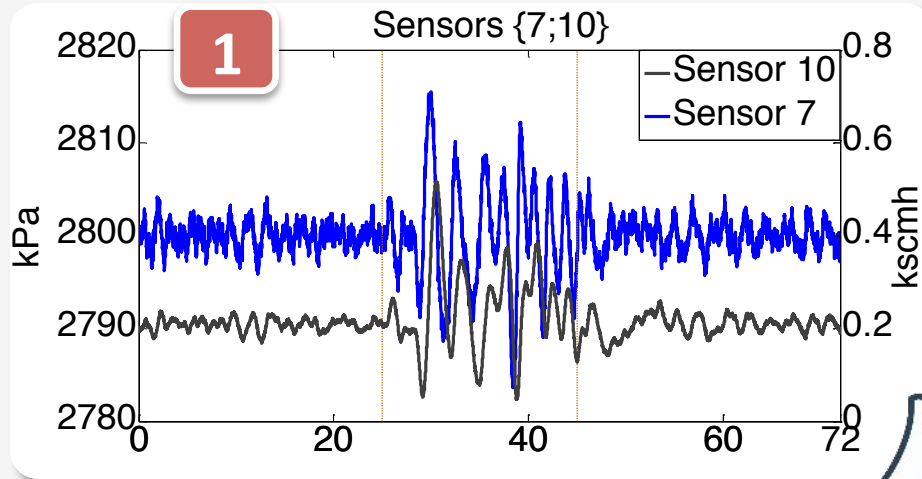
- ❑ Thousands of sensors signals in a facility
- ❑ All plants are unique



Anomaly detection

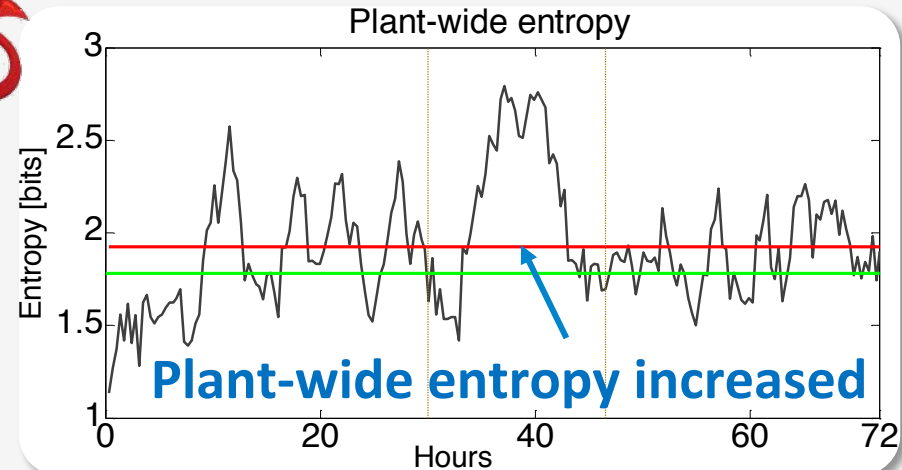
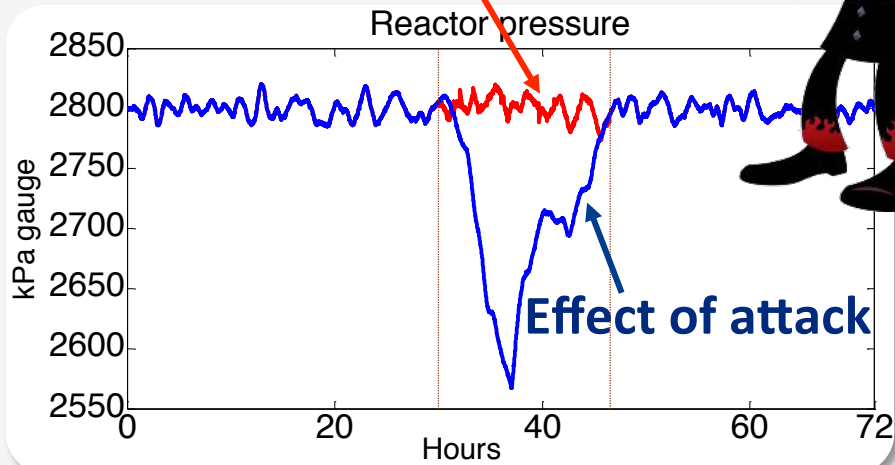
Entropy metric

$$H(X) = \sum_{x \in C_X} P(x) \log_a \frac{1}{P(x)}$$

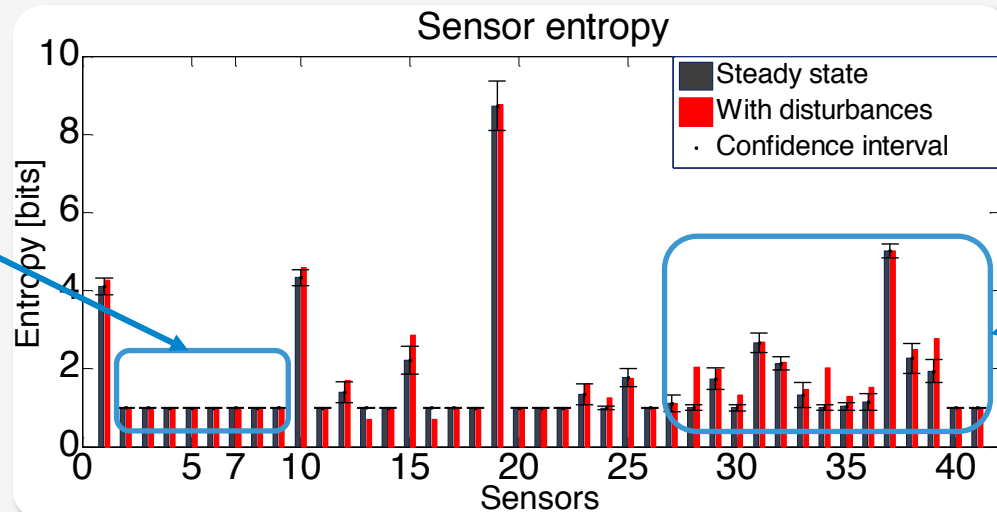


Detection

Spoofed signal

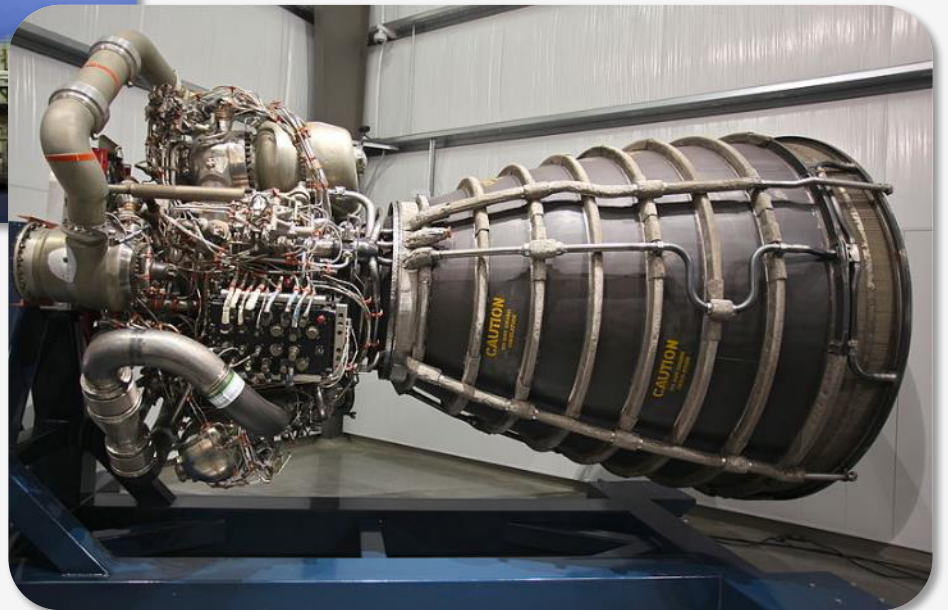


All OK in reactor unit

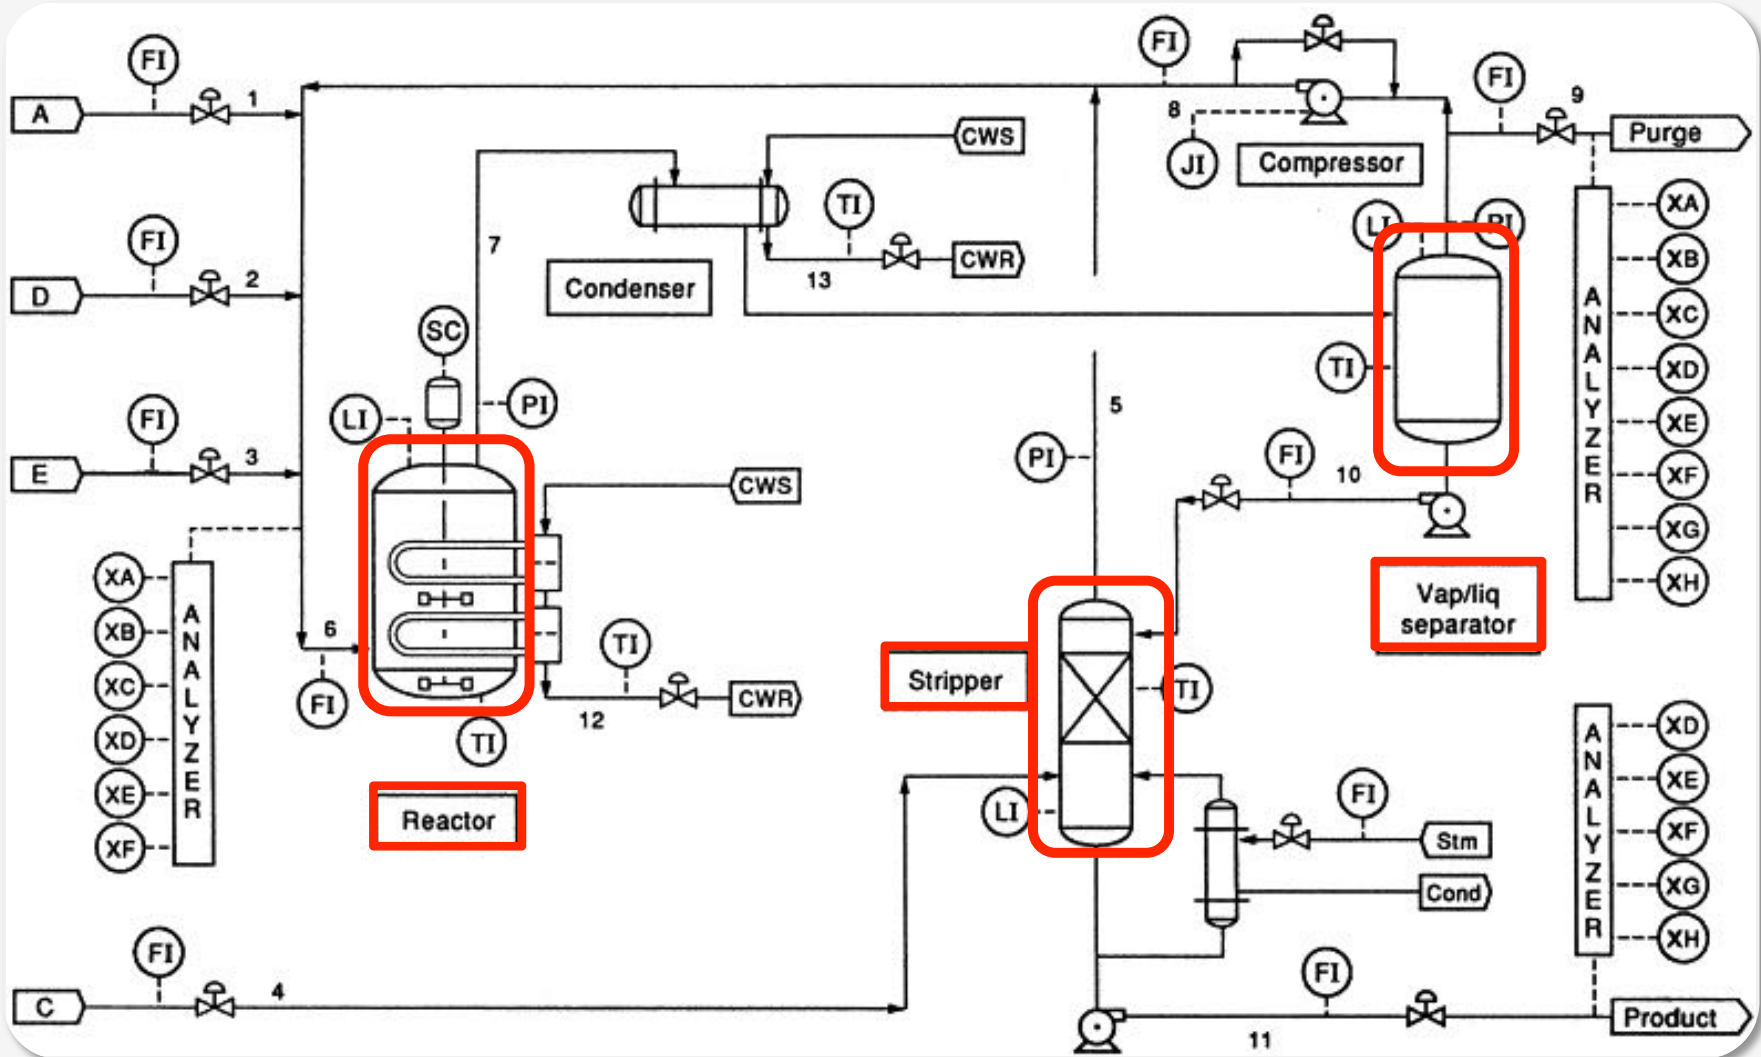


Product composition has changed

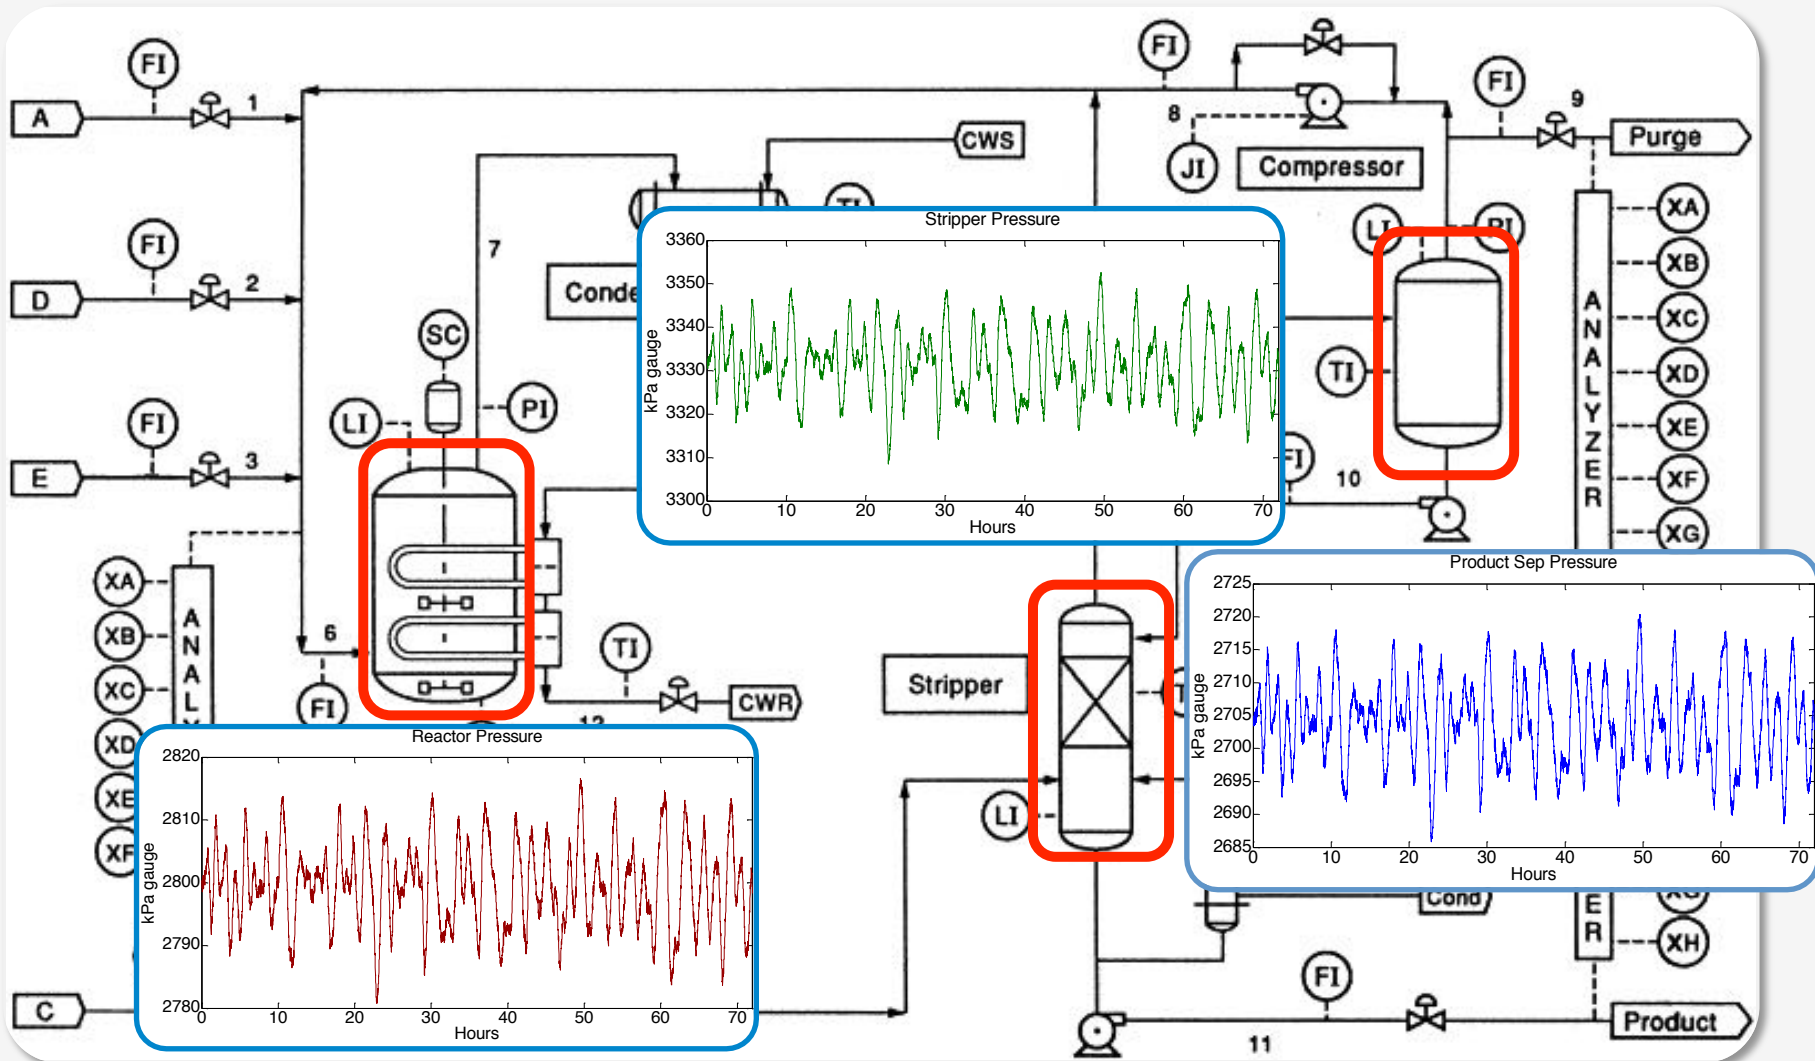
Approach



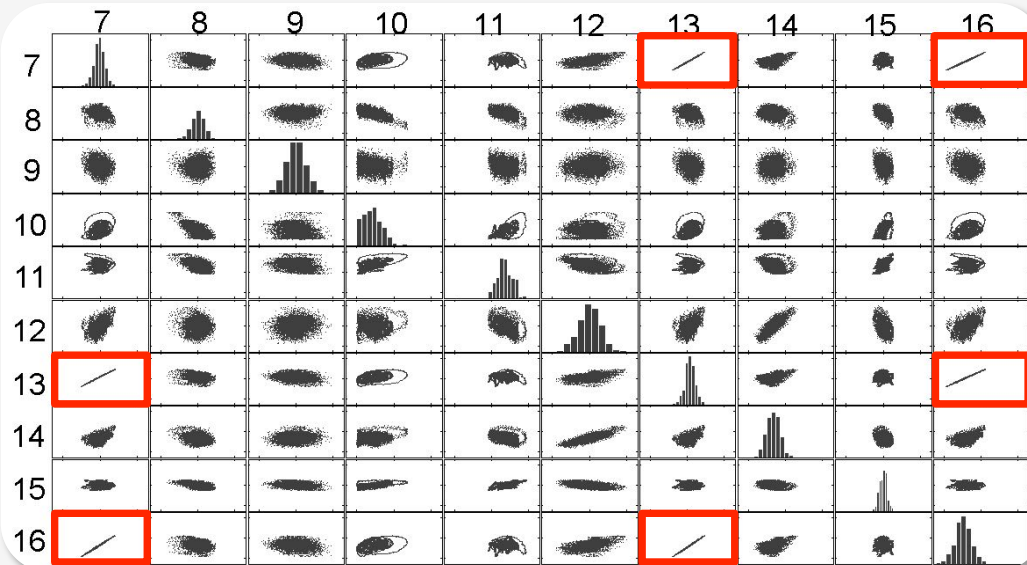
Correlated sensor signals



Correlated sensor signals

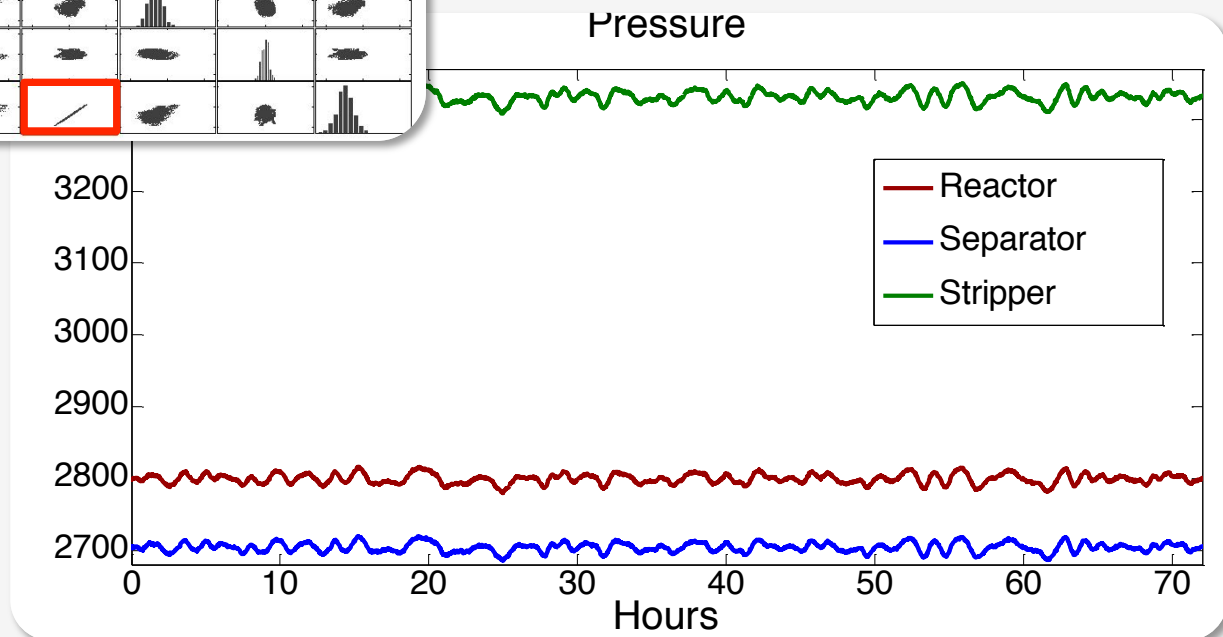


Sensor signal clustering

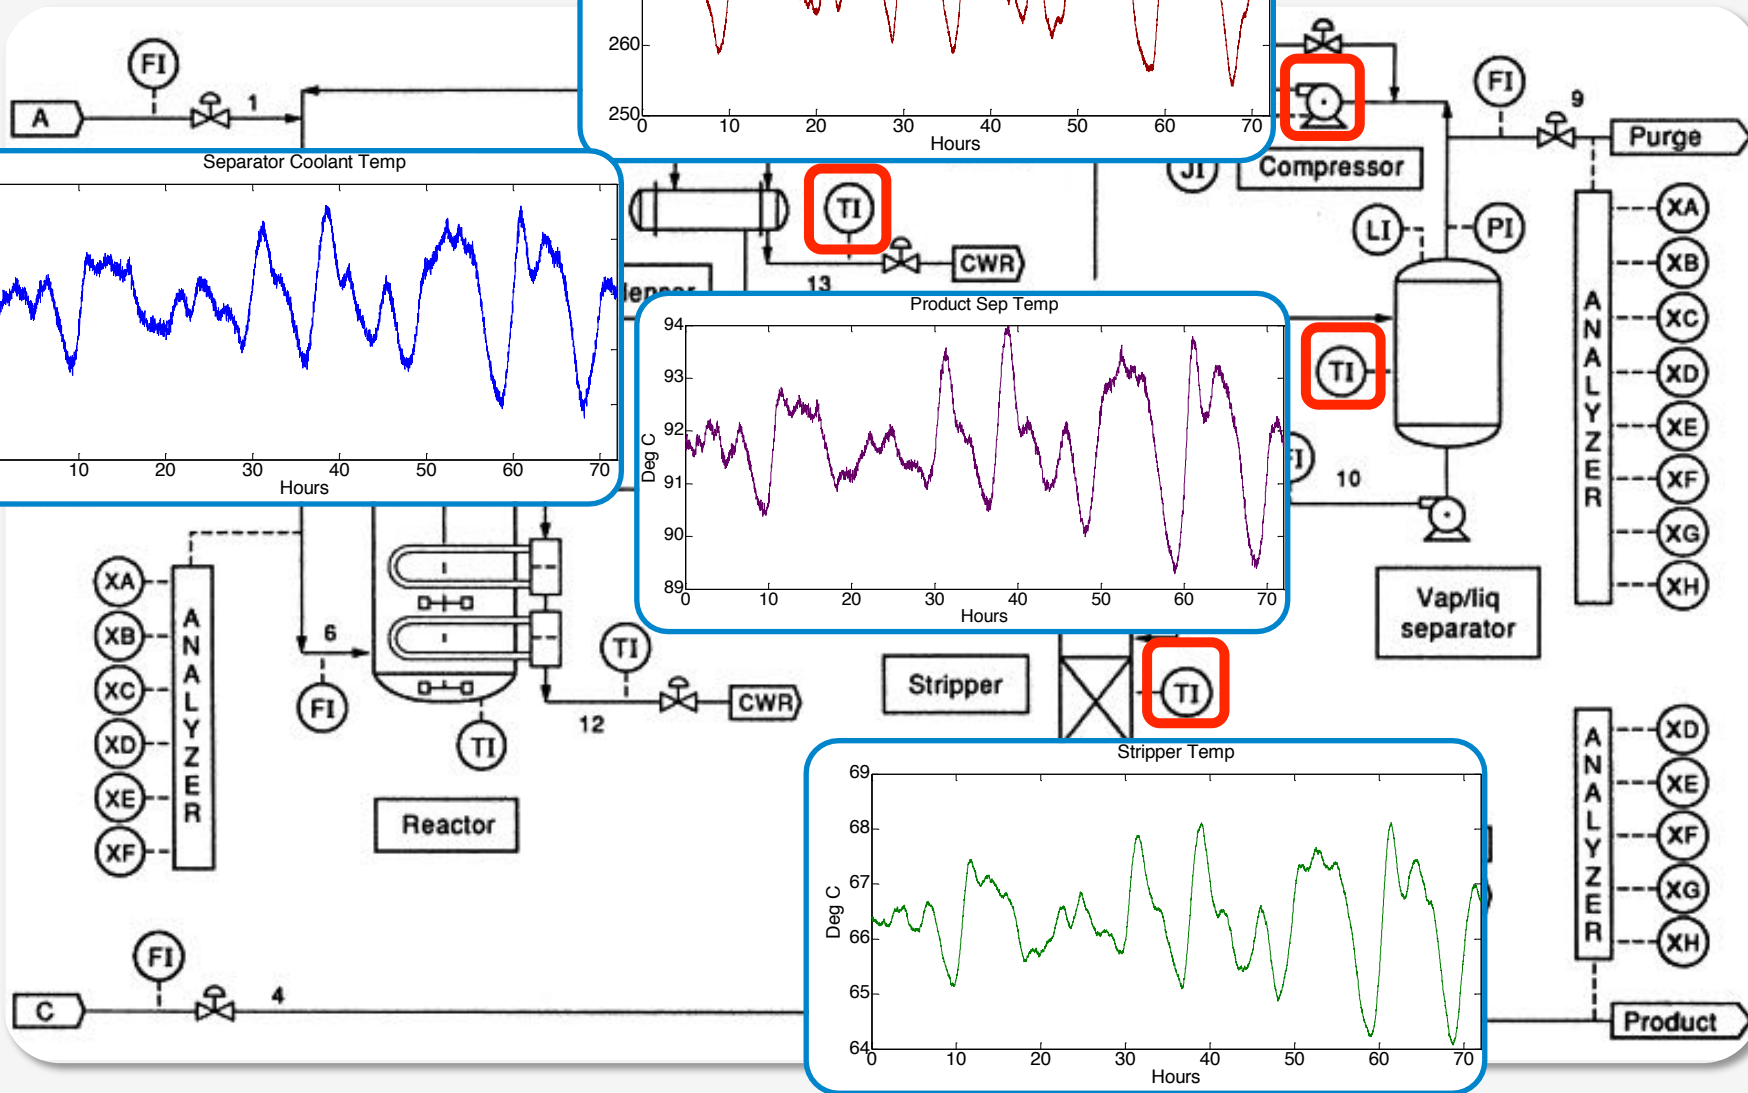
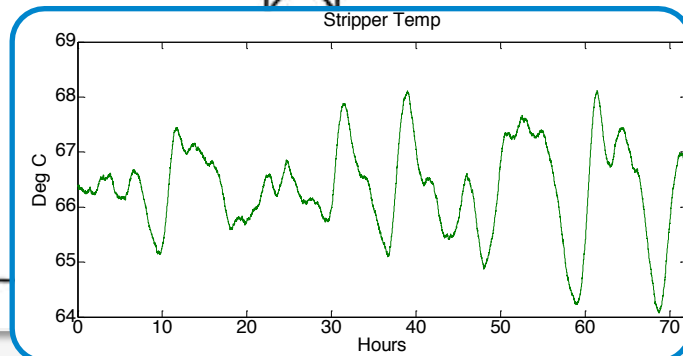
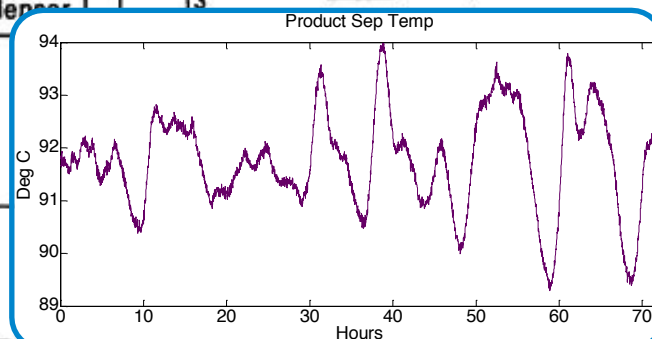
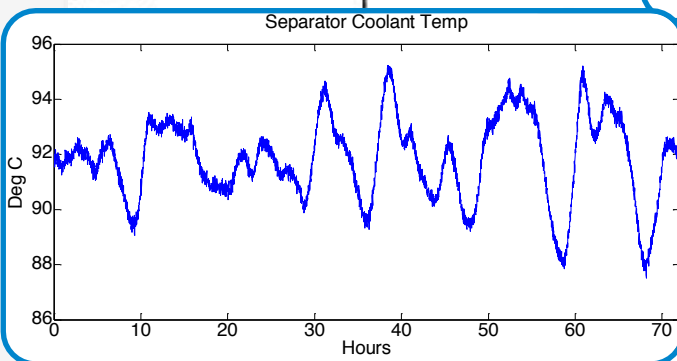
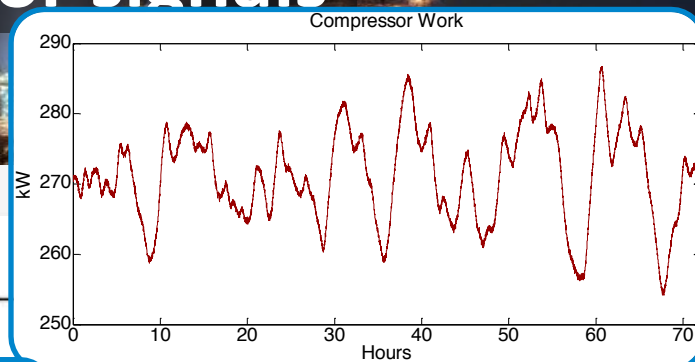


Scatter plot to visualize correlations between signals

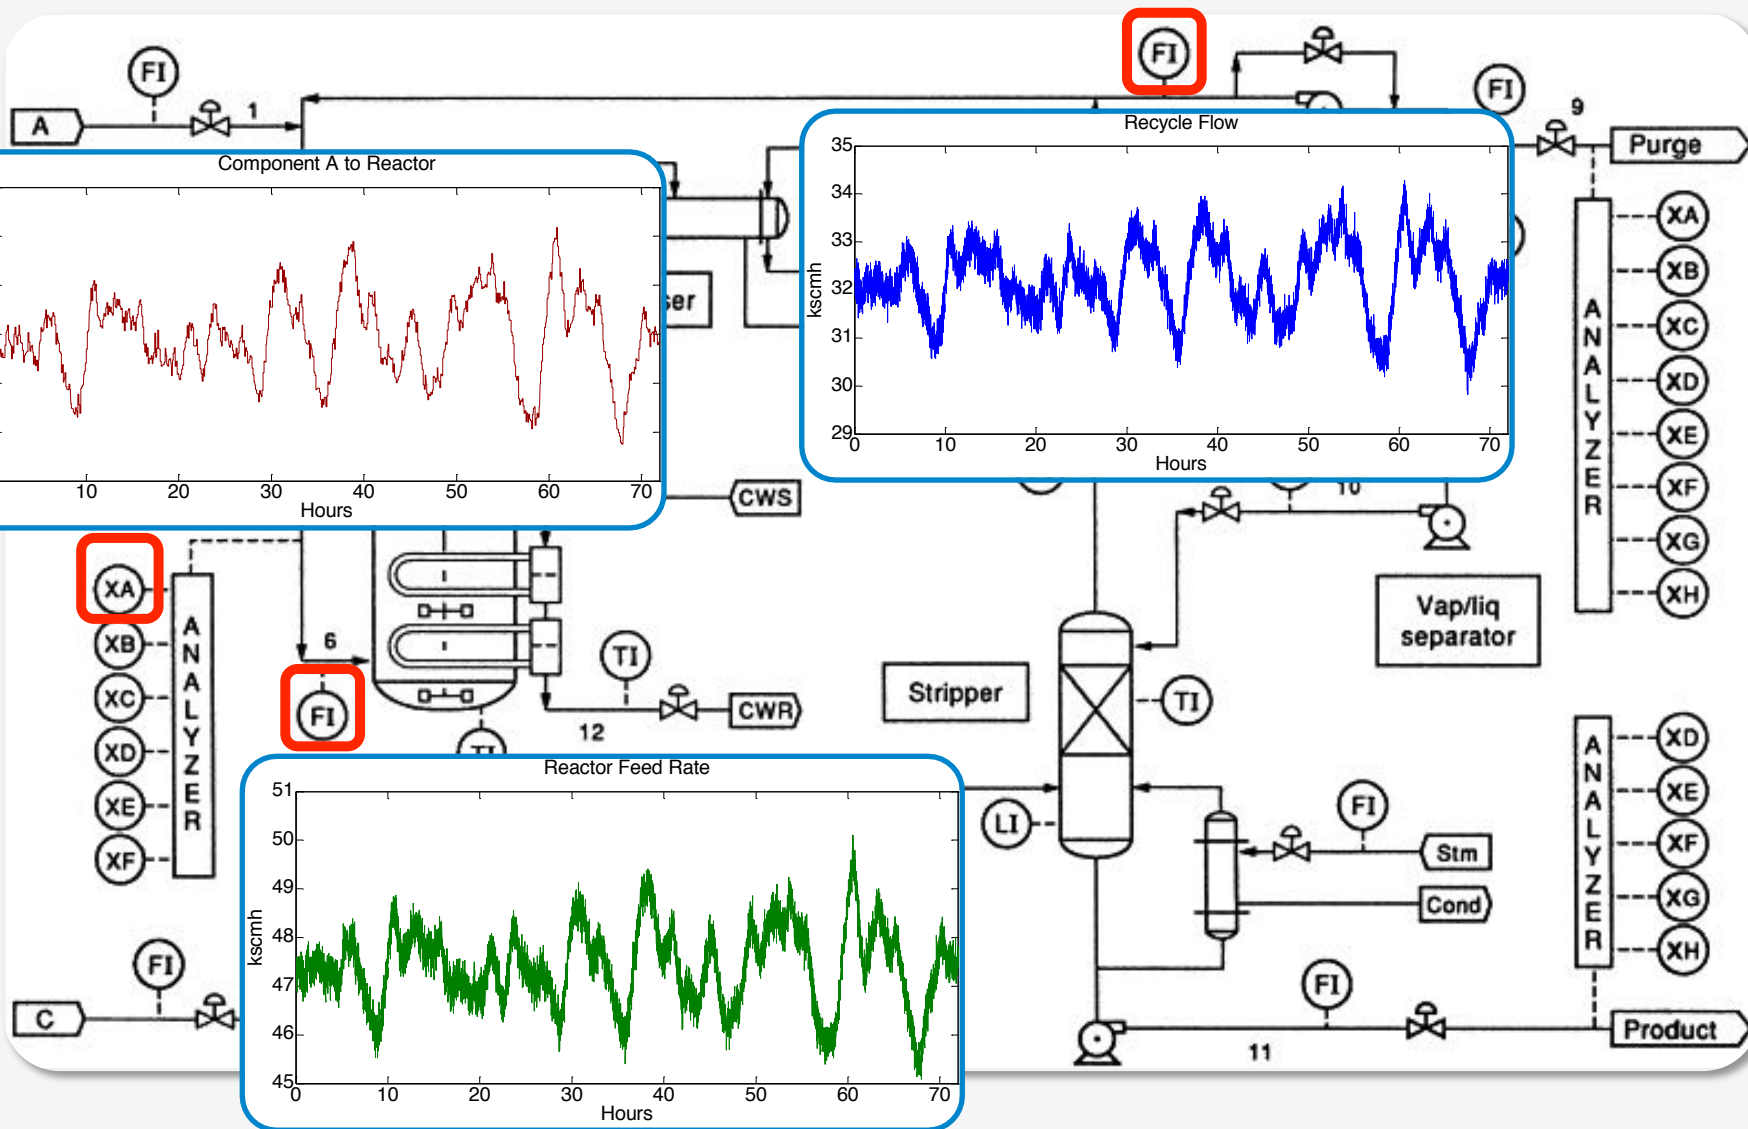
Metis tool kit: Graph partitioning for sensor clustering



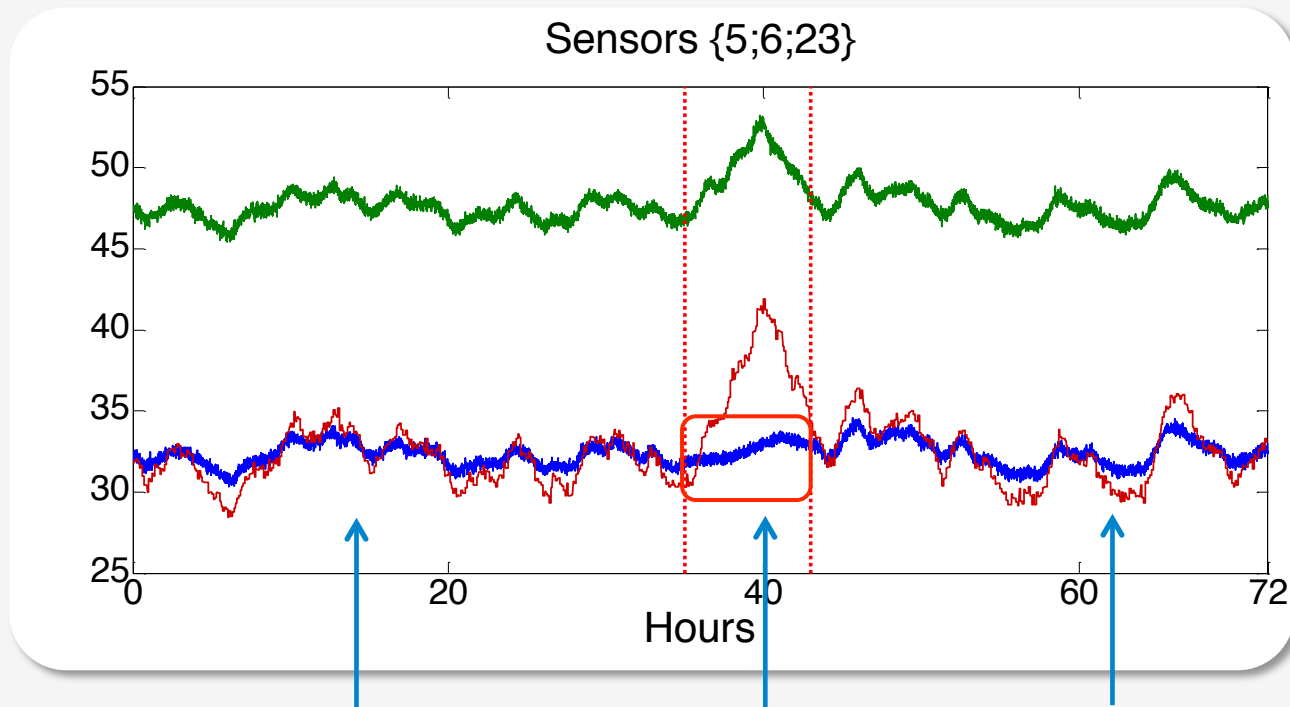
Correlated sensor signals



Correlated sensor signals



Cluster entropy



Signals correlation:

+

-

+

Correlation entropy:

LOW

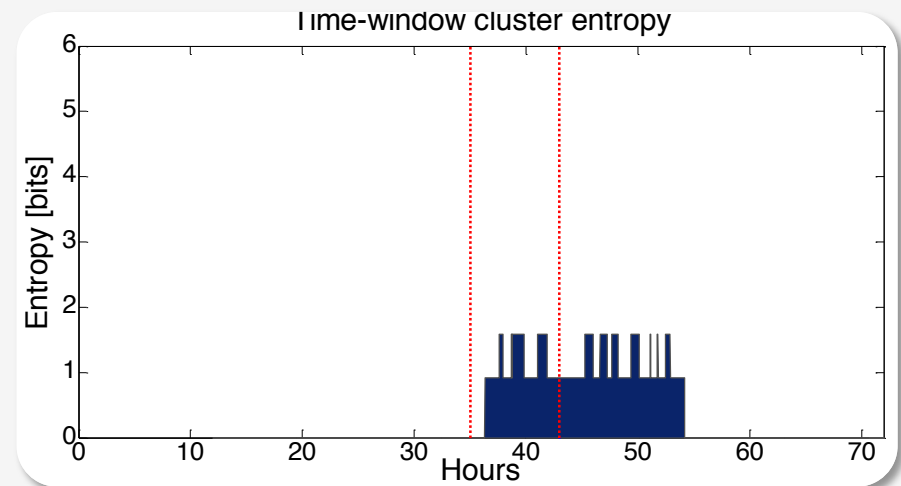
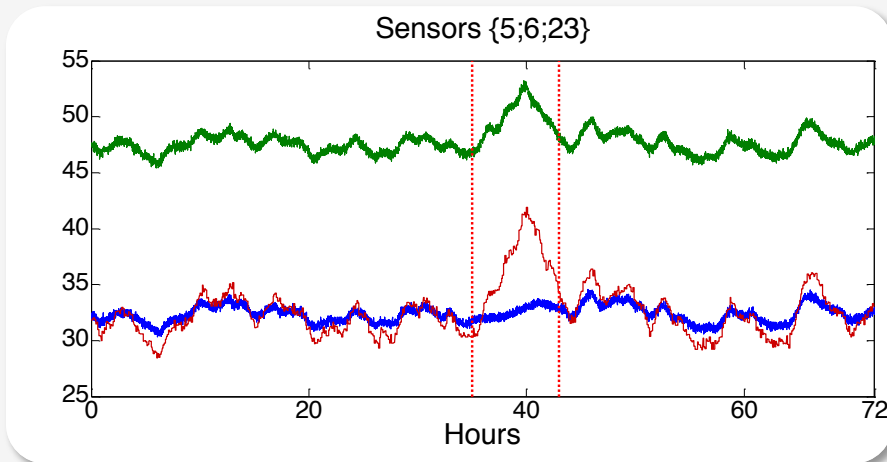
HIGH

LOW

Detection

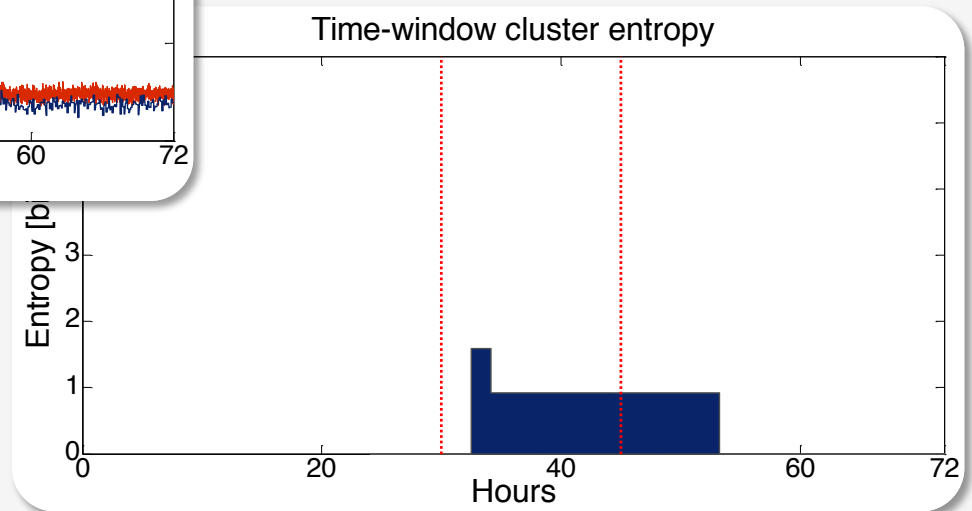
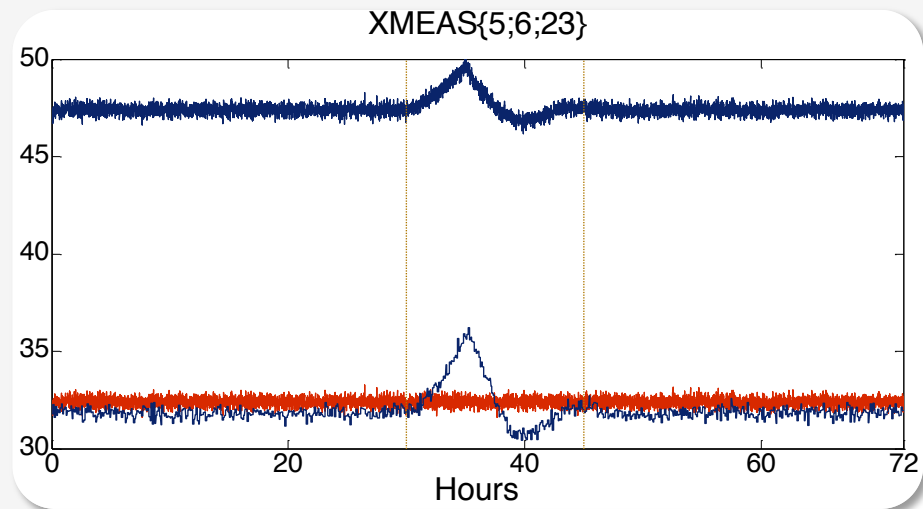


**Spoofed signals appears genuine at first glance.
But they are not be correlated with the rest of the signals in the
cluster of related sensors**



Detection in the presence of disturbance in reactor unit

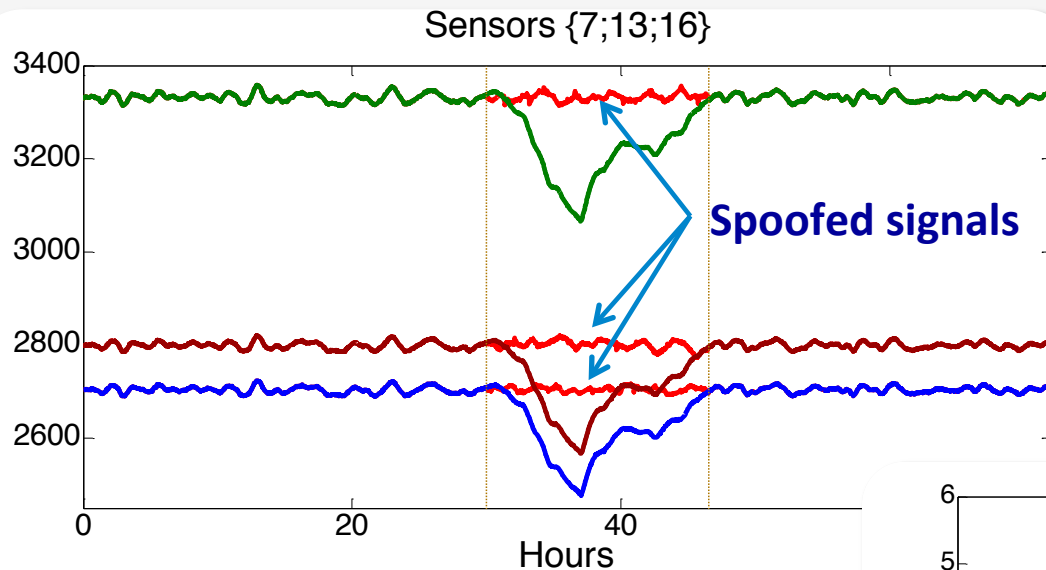
Detection



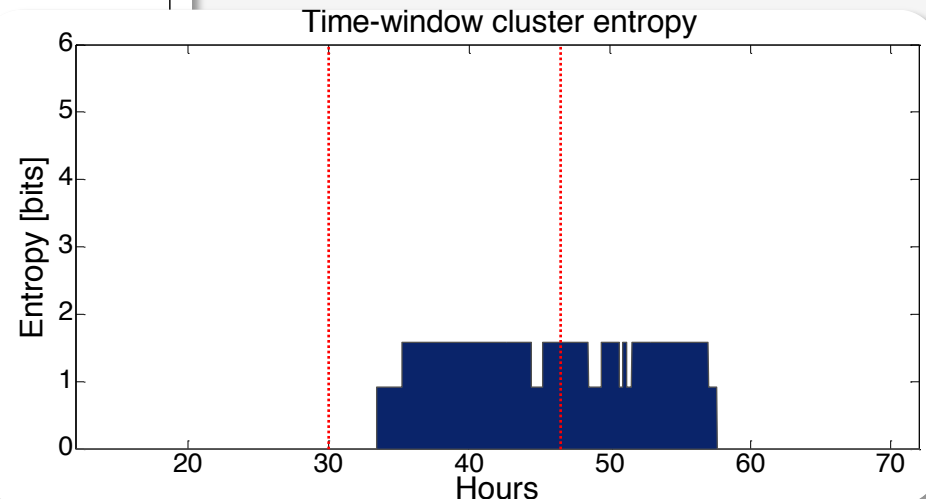
Detection in steady state (without disturbances)

Powerful attacker

He spoofed them all!!!

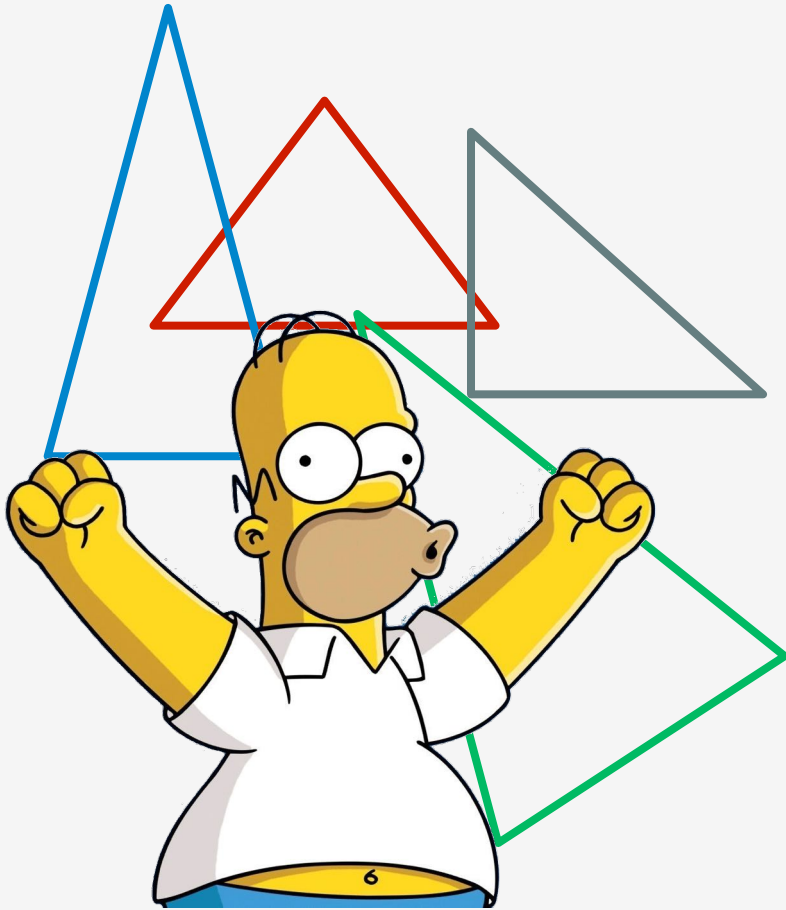


Bad luck!!!



**Spoofed signals will all look
genuine but won't be
correlated**

The END of SCADA triangles



Marina



Jason

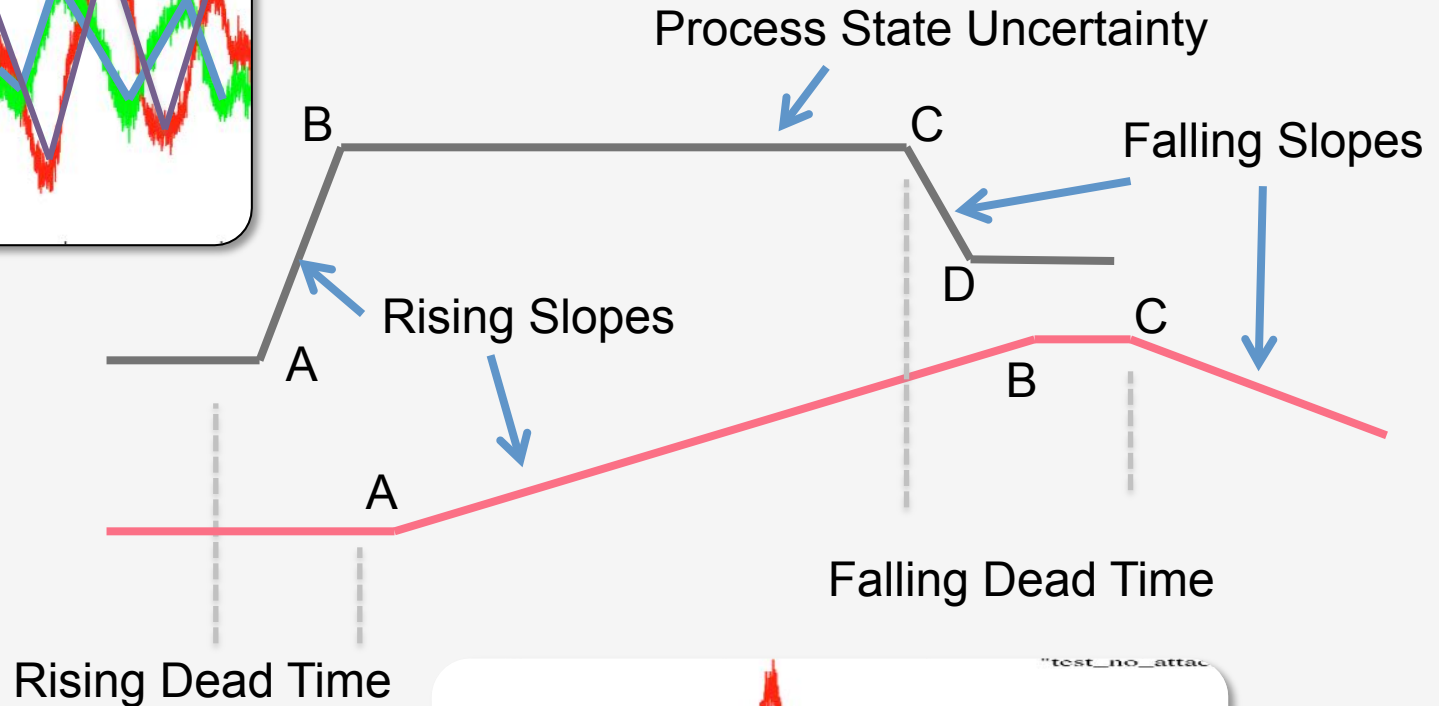
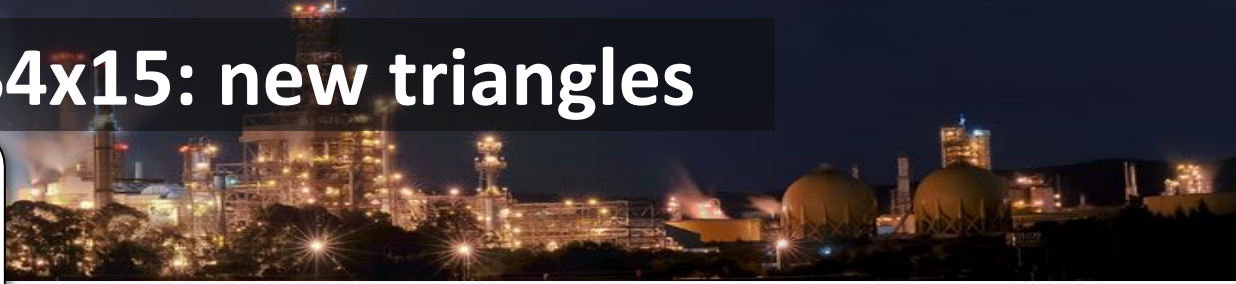
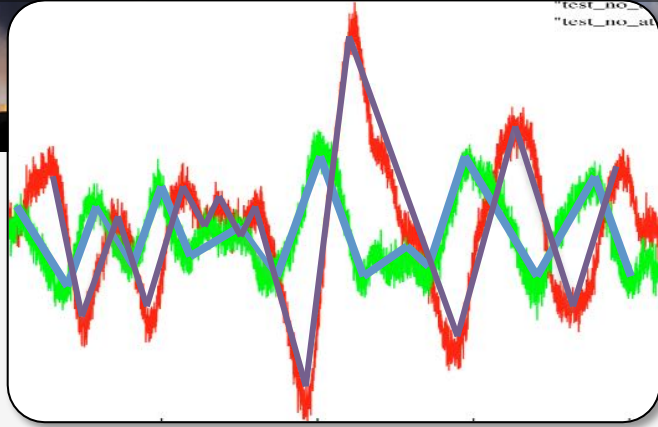


Arms race is on!!

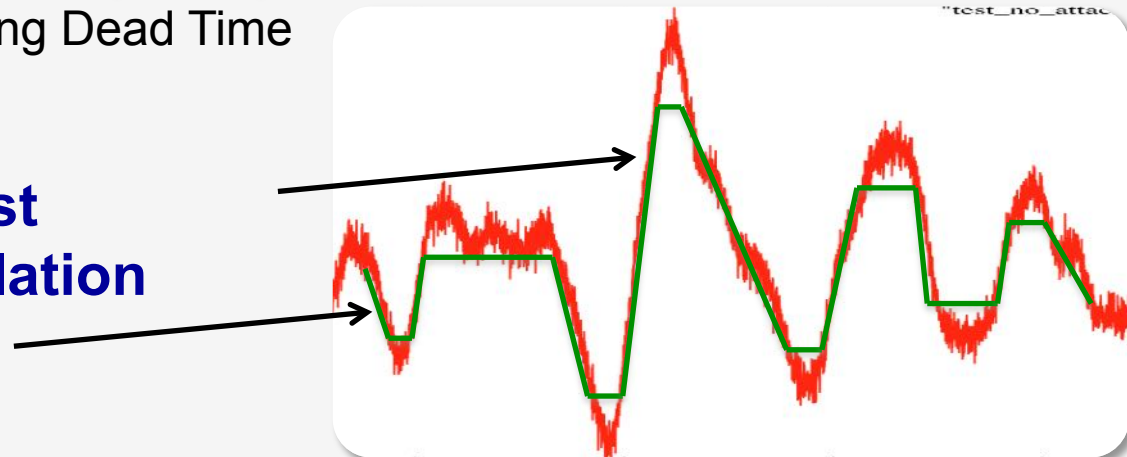
Jason Larsen was challenged



Jason Larsen at S4x15: new triangles



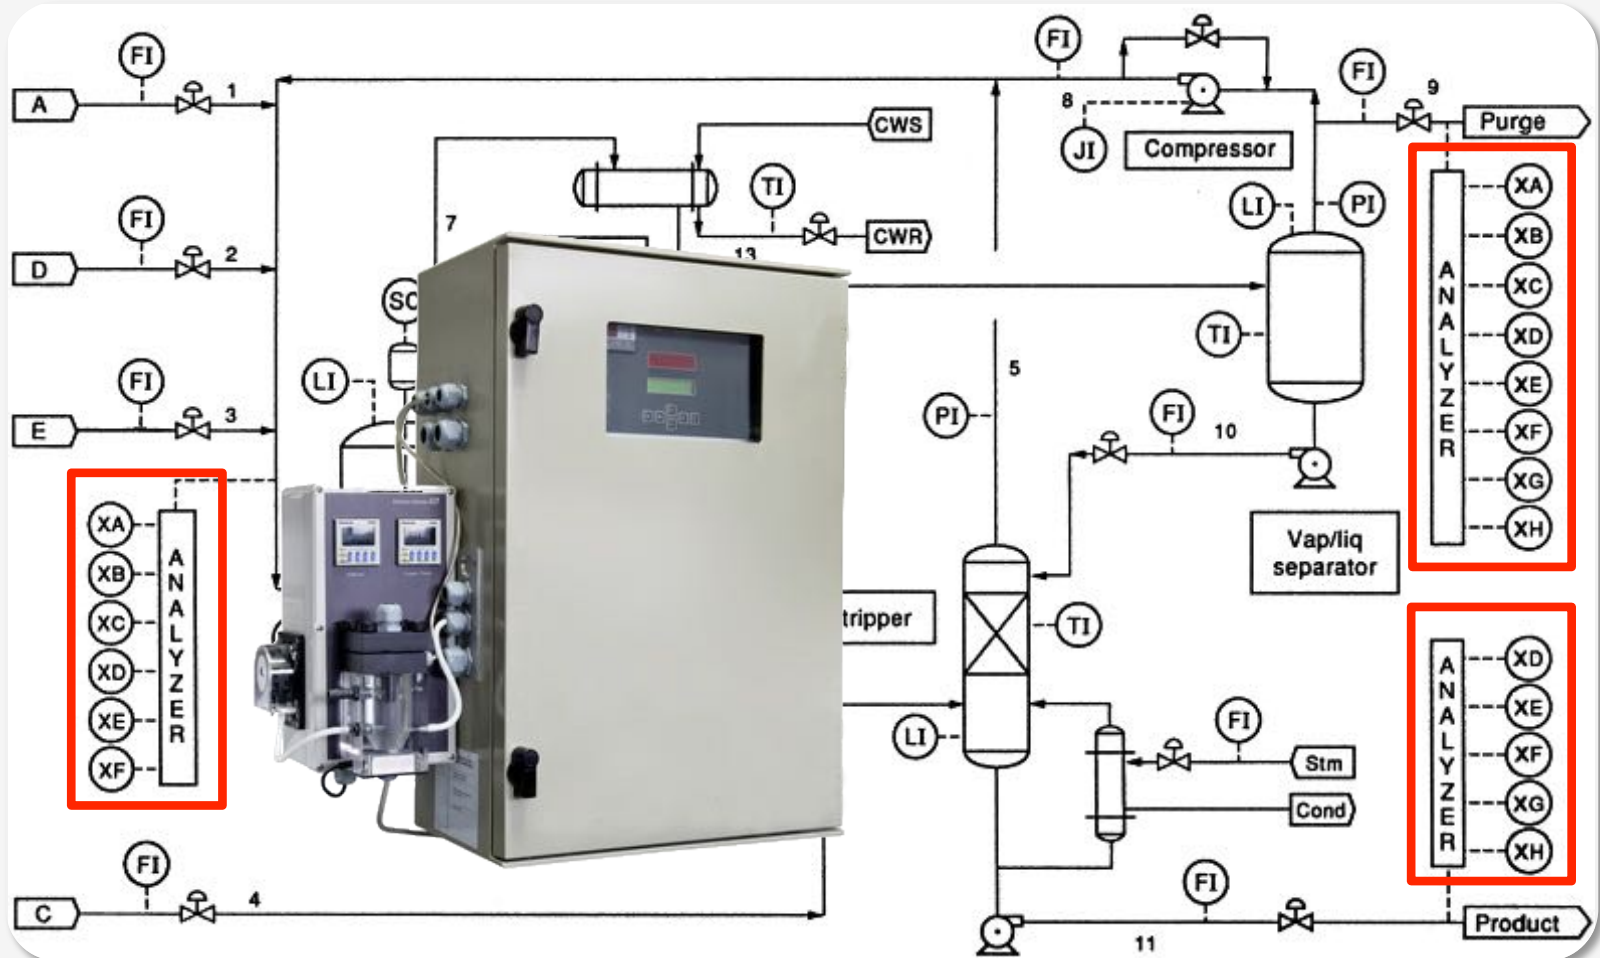
**Slopes: Most
confident correlation**



Marina Krotofil was concerned



Establishing trust: chemical analyzer



Dead time: 6-15 min

Establish trust

Include in clusters offline signals veracity of which can be guaranteed

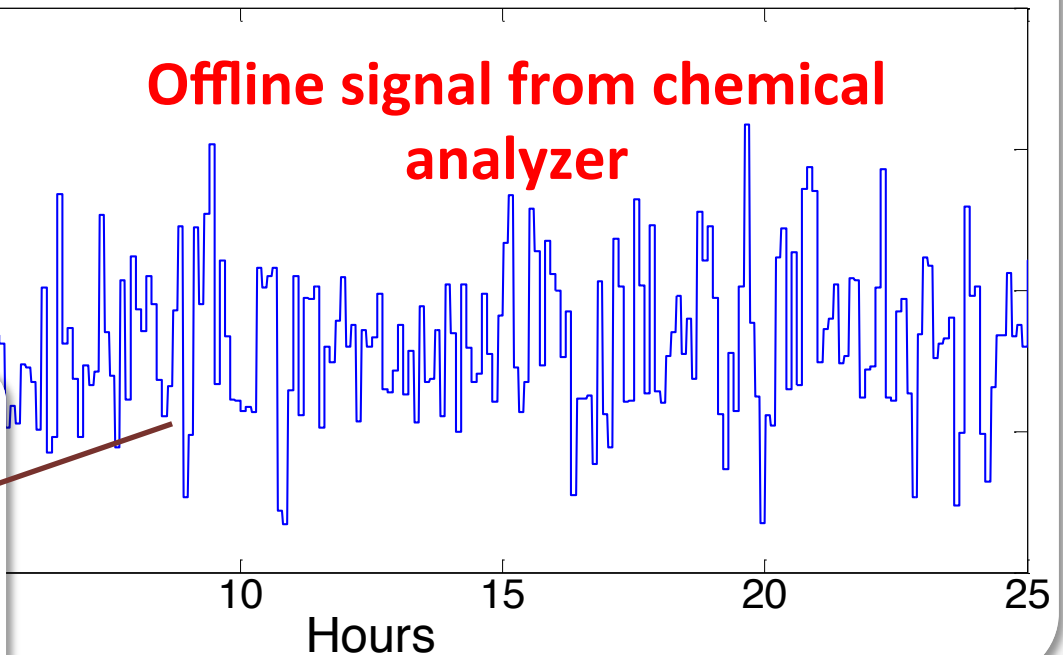
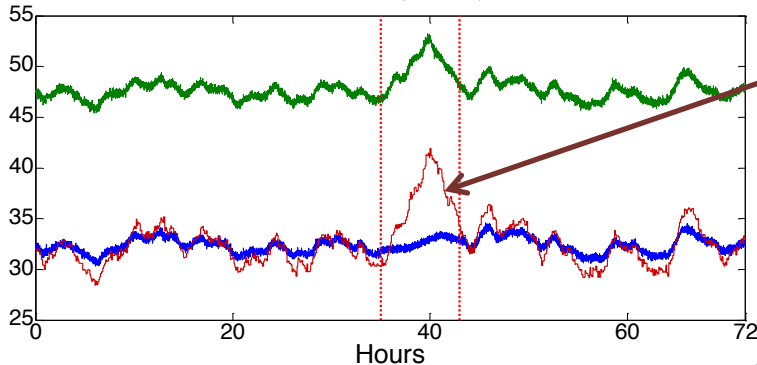


Component A in reactor feed

Offline signal from chemical analyzer

mol%

Sensors {5;6;23}



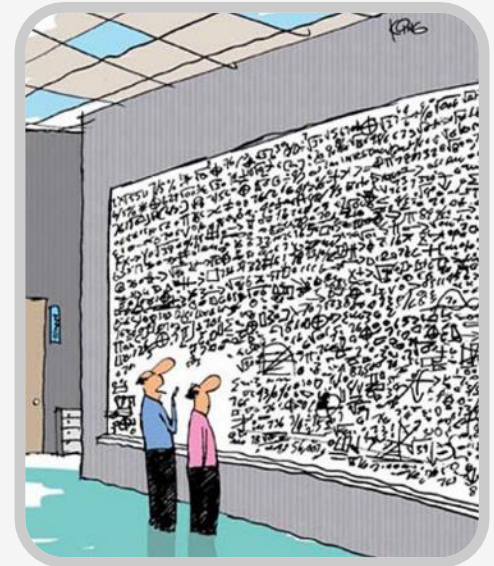


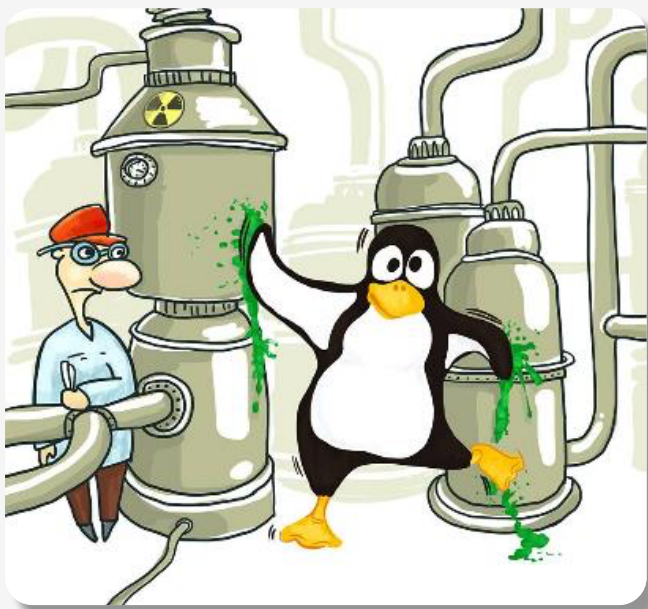
Conclusions

Good control vs. good crypto

Security controls must span all the way to the application

- ❑ Security specialists define required security protections
 - E.g. signatures for authentication and integrity protection
- ❑ Mathematicians do their magic and come up with strong cryptographic primitives and algorithms
- ❑ **It is no different with secure controls**
 - Specify the problem and a desired outcome
 - Let control guys do what they do best





**Damn Vulnerable
Chemical Process**

Thank you

marina.krotofil@encs.eu

jason.larsen@ioactive.com

IOActive  **ENCs**