Parameter-Invariant Monitor Design for Cyber-Physical Systems:
Part 3 – Implementation of Parameter-Invariant Monitors

James Weimer, Oleg Sokolsky, Insup Lee

Penn Engineering

PRECISE
PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING

Recall CPS Applications

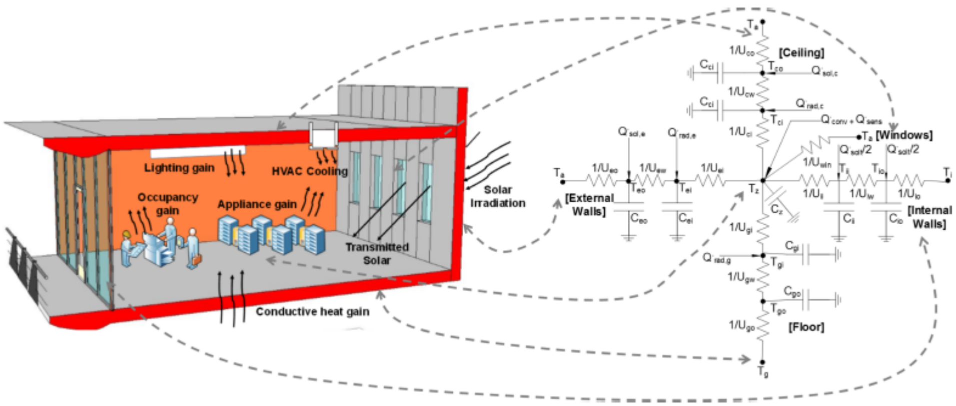

Recall the Monitor Design Problem

- Design a binary test between:
  - H_0 : null hypothesis
  - H_1 : event hypothesis

	H_0 is true	H_1 is true
test claims H_0	correct non-detection	missed detection
test claims H_1	false positive	true positive

- Performance constraints
  - bound false positive rate
  - maximize true positive rate

- Module 1 covered the fundamentals of parameter invariance:
  - LRT, GLRT, MI, and PAIN

- Module 2 covered the design of parameter invariant monitors:
  - general form: $\mathbf{y} = \mathbf{H}\theta + \sigma\mathbf{n}$

- This module presents the implementation of PAIN monitors
  - real-world applications

Outline

- Meal detection in type I diabetics
  - unknown linear time invariant systems
- Critical pulmonary shunt detection in infants
  - detection in structured linear systems with unknown parameters
- Building actuator fault detection
  - signal detection in unknown networked systems

Outline

- Meal detection in type I diabetics
  - unknown linear time invariant systems
- Critical pulmonary shunt detection in infants
  - detection in structured linear systems with unknown parameters
- Building actuator fault detection
  - signal detection in unknown networked systems

Meal Detection in Type I Diabetics

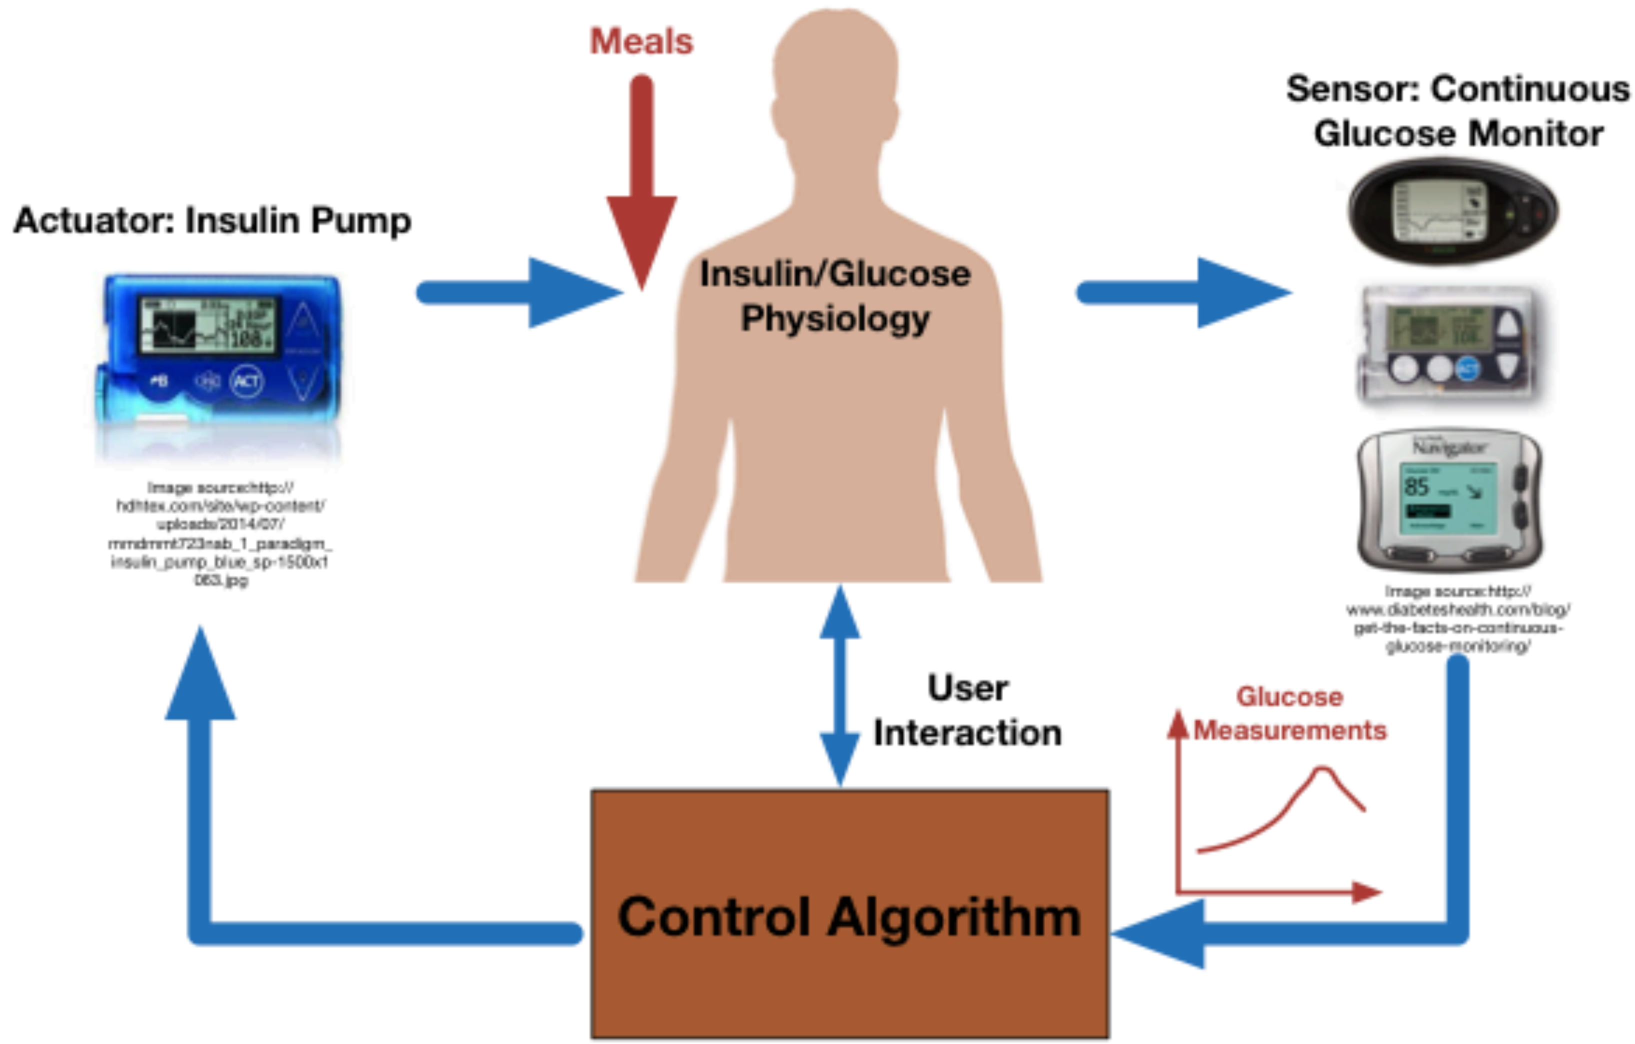

Meal Detection in Type I Diabetics

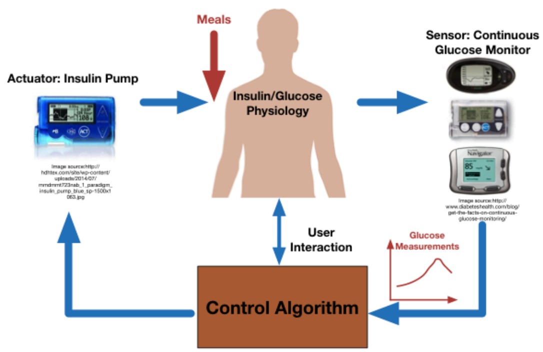

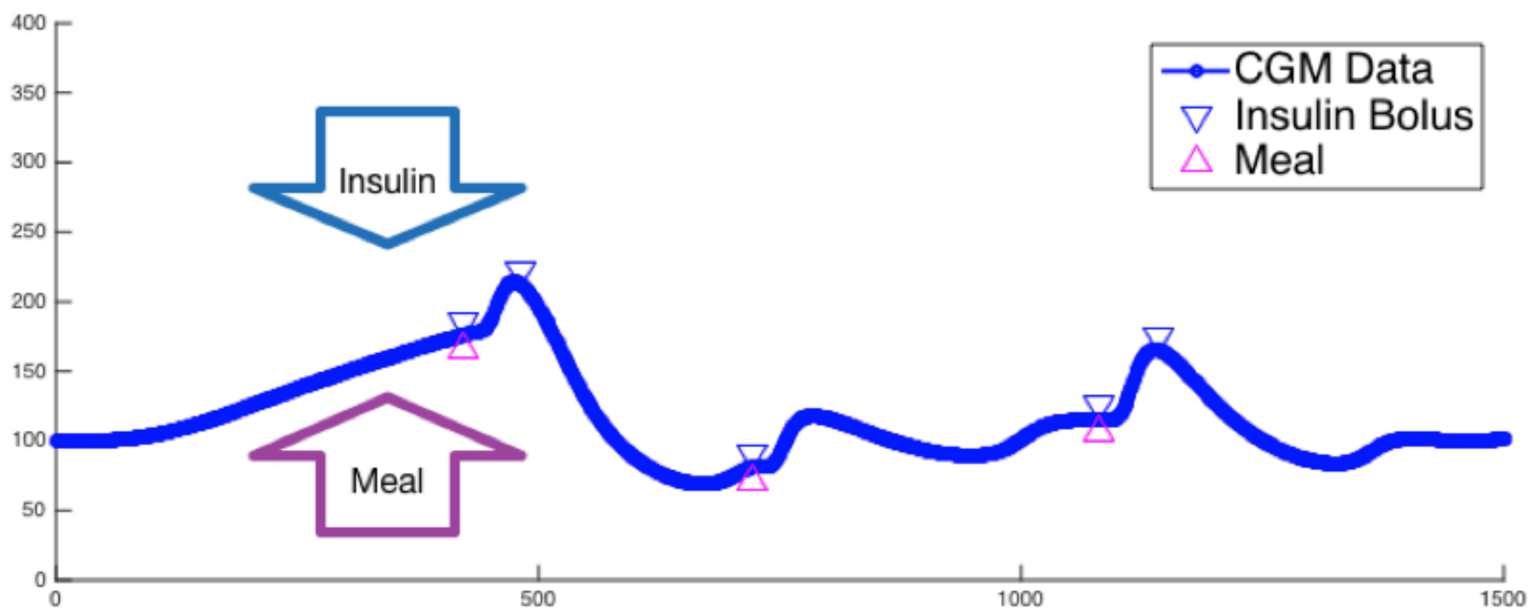

Meal Monitor Design Problem

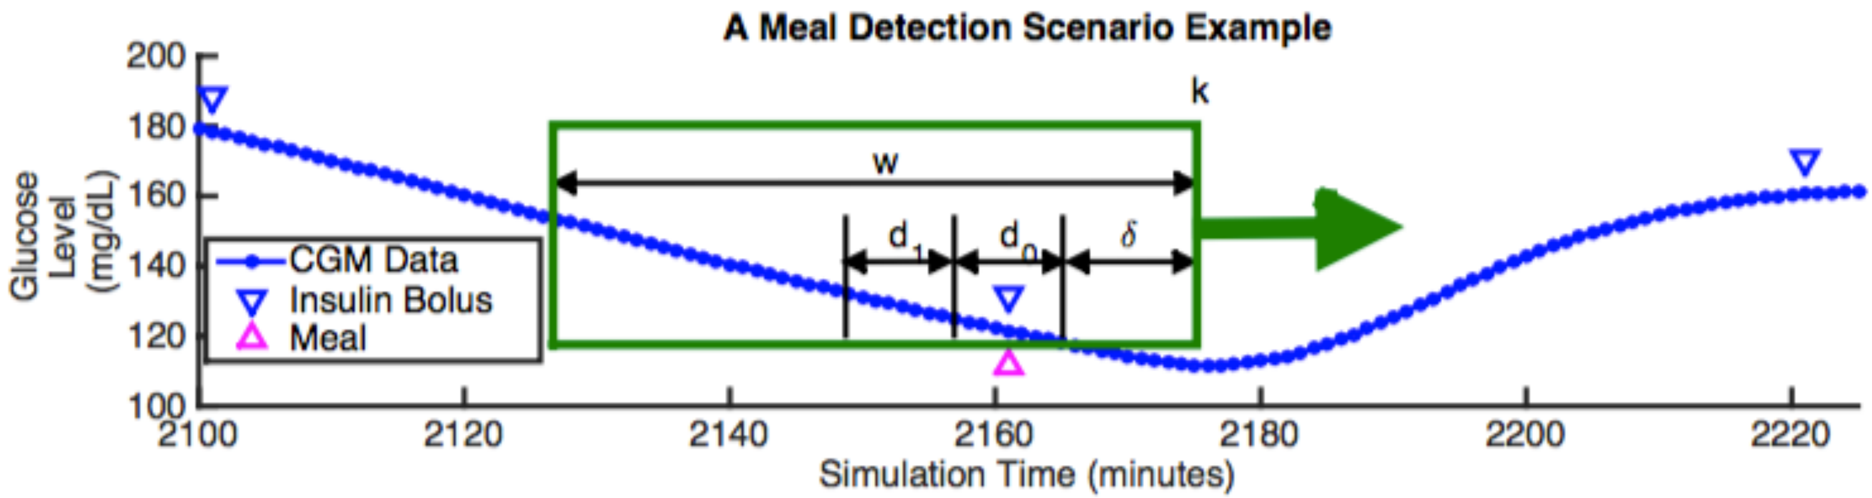

- hypothesis testing problem:
  - window of w measurements
  - test meal impulse happening in window d_1 or d_2
  - use the 2-sided PAIN approach
    - allows for the case where all hypotheses are incorrect

- What is the relationship between events and measurements?
  - i.e. What is the physical model?

Penn Engineering

PRECISE

Sequential Monitoring/Detection

- Sequential Monitoring of sequential events

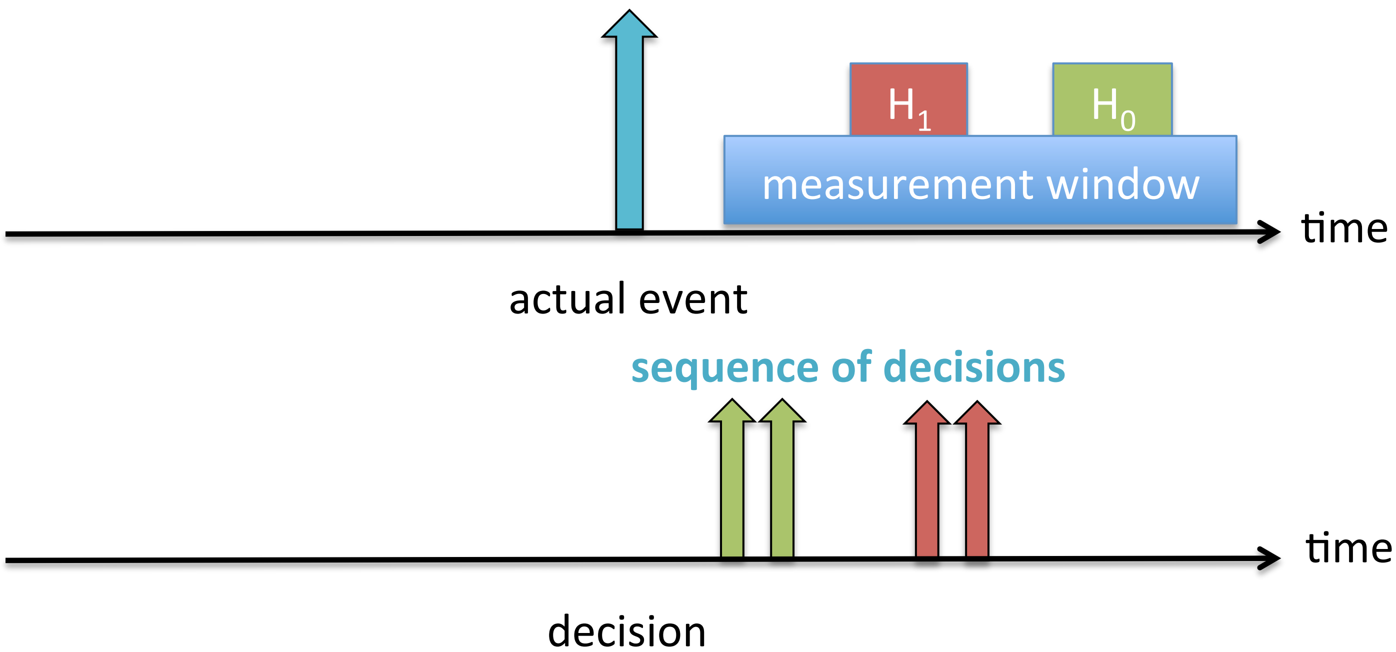

Physiological Modeling

- FDA accepted model
    - 12 states, 30 physiological parameters (unknown)
    - non-linear

- Bergman model – 5 states, linear – unknown physiological parameters
    - 5^{th} order model

$$
\begin{array}{r}
\text{Plasma Glucose} \rightarrow \\ \\ \\ \\ \text{Plasma Insulin} \rightarrow
\end{array}
\frac{d}{dt}\begin{bmatrix} G(t) \\ g(t) \\ m(t) \\ x(t) \\ I(t) \end{bmatrix} = \begin{bmatrix} p1 & 0 & 1 & 0 & p2 \\ 0 & \frac{-1}{t_G} & 0 & 0 & 0 \\ 0 & \frac{1}{t_G} & \frac{-1}{t_G} & 0 & 0 \\ 0 & 0 & 0 & -k_a & 0 \\ 0 & 0 & 0 & \frac{k_a}{V_d} & -k_e \end{bmatrix} \begin{bmatrix} G(t) \\ g(t) \\ m(t) \\ x(t) \\ I(t) \end{bmatrix} + \begin{bmatrix} p_3 \\ \frac{A_G}{t_G} D_G(t) \\ 0 \\ u(t) \\ 0 \end{bmatrix}
\begin{array}{l}
\\ \leftarrow \text{Meal Input} \\ \\ \leftarrow \text{Insulin Input} \\ \\
\end{array}
$$

- test signals – sequential ranges of hypothesized meal times
- disturbances:
    - reported meals = impulse at a time (amount unknown, effect unknown)
    - insulin = impulse at a time (amount known, effect unknown)
- measurements
    - plasma glucose

implement 2-sided PAIN monitor

PAIN monitor for Meal Detection

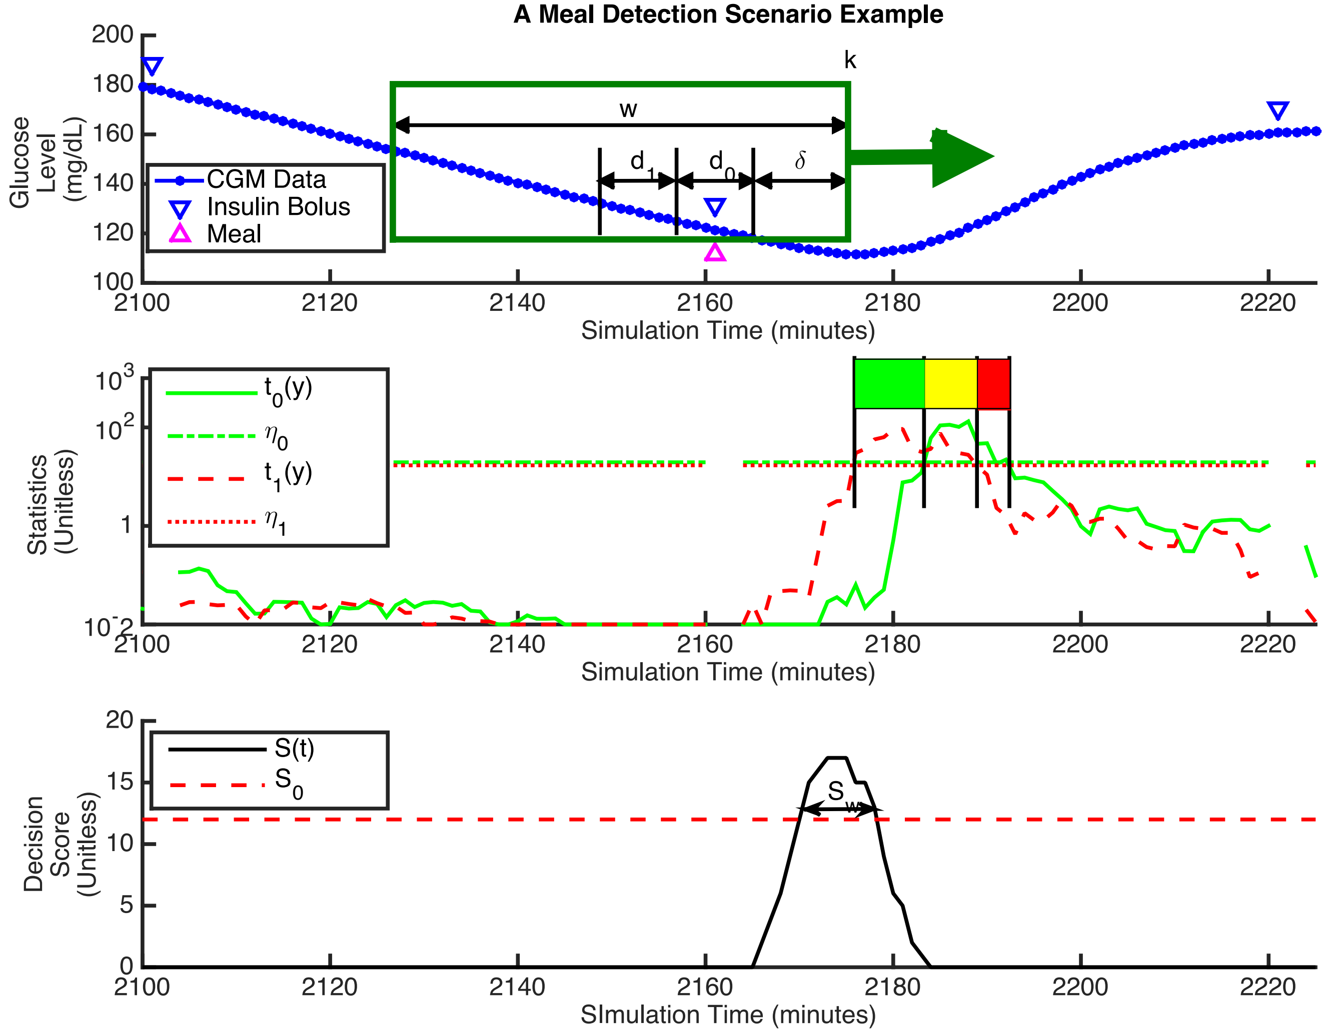

PAIN Meal Monitor Evaluation

- Generated 10,000 random virtual patients
  - parameters selected from a convex set of FDA-suggested physiological ranges

- Simulated each patient for 20 meals
  - using FDA-accepted T1DM simulator (maximal model, non-linear)

- Compared to prominent approaches in literature
  - Dassau et al. → Kalman, then rate-of-change (RoC) thresholding
  - Lee et al. → a priori specified FIR filter, then RoC thresholding
  - Harvey et al. → multi-stage filter, then RoC thresholding

- Evaluate on the criteria:
  - false positive rate vs. true positive rate
  - time-to-detection (when correct)
  - number of false positives per patient

PAIN Meal Monitor Performance

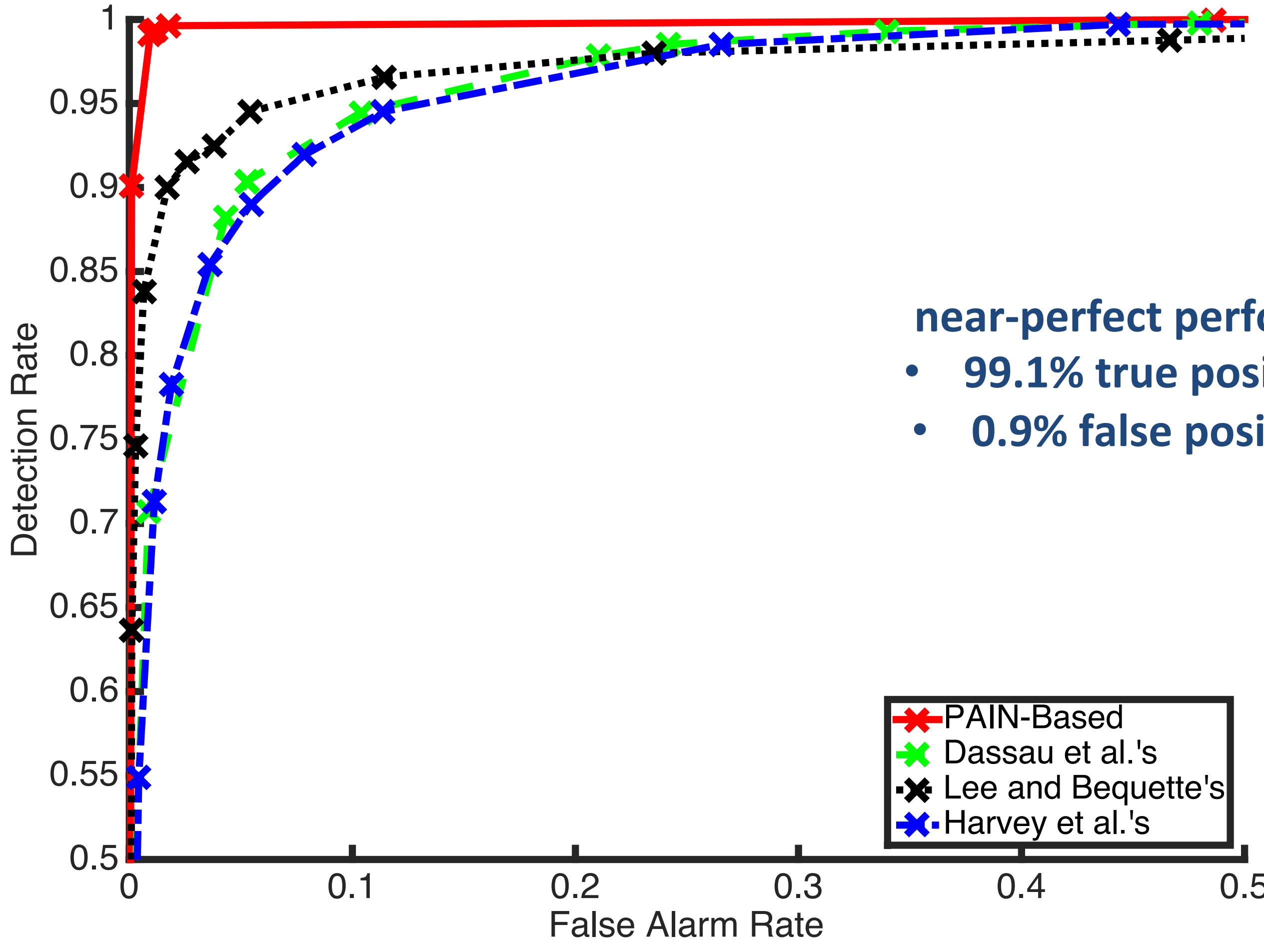

near-perfect performance

- **99.1% true positive rate**
- **0.9% false positive rate**

PAIN Meal Monitor Performance

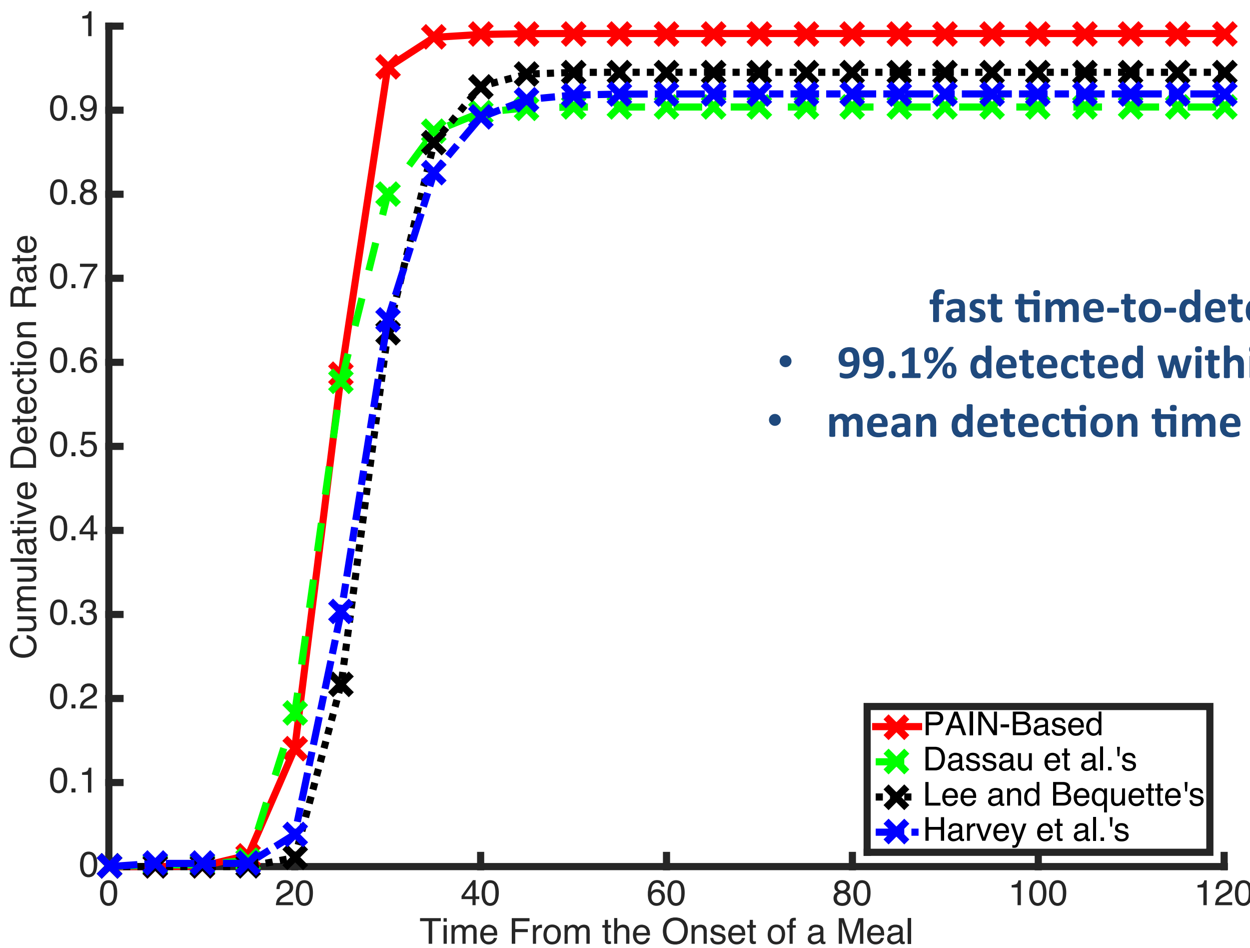

fast time-to-detection

- **99.1% detected within 40 minutes**
- **mean detection time of 24 minutes**

PAIN Meal Monitor Performance

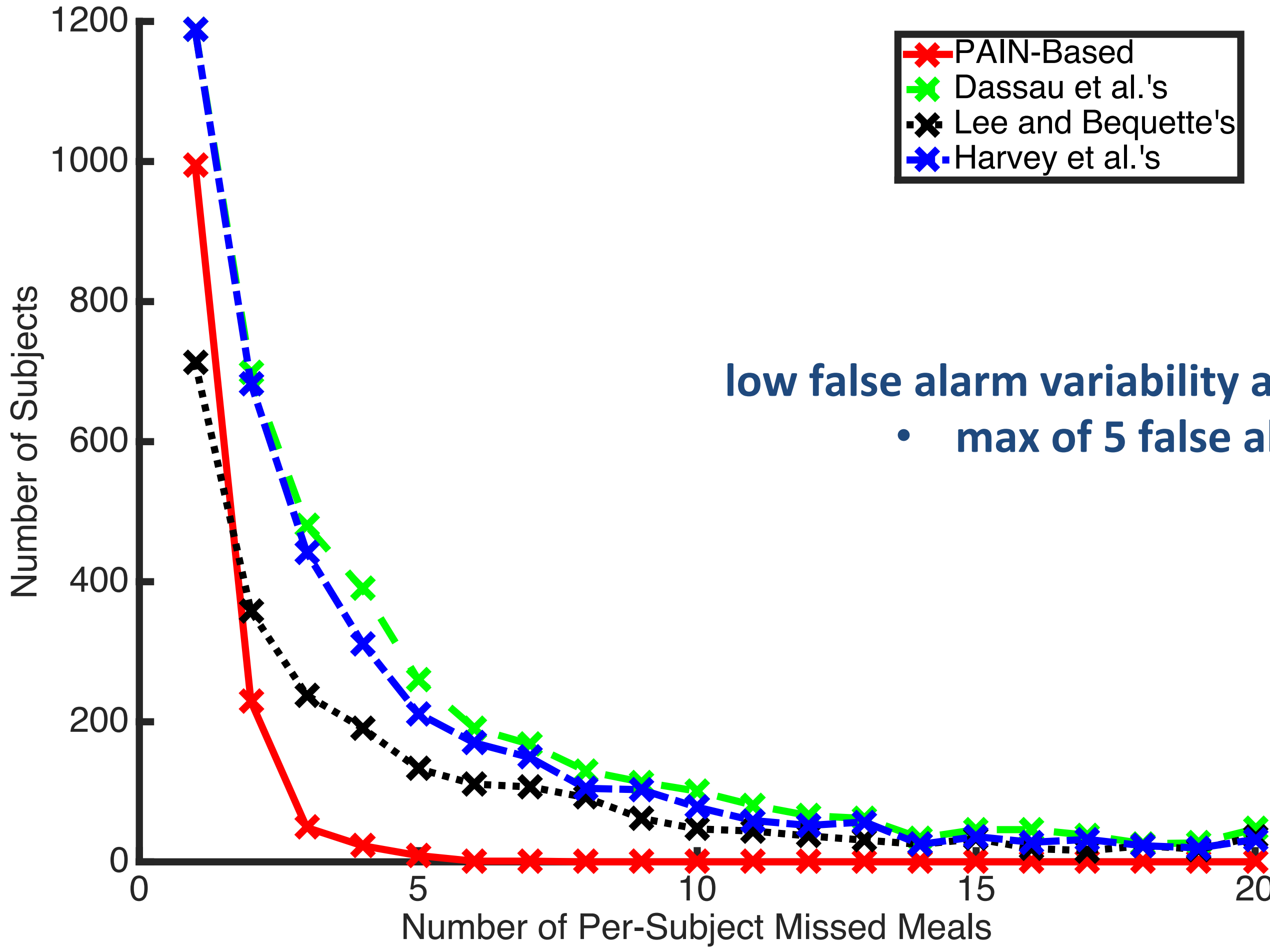

low false alarm variability across patients

- **max of 5 false alarms**

Summary : Detection with Unknown LTI models

- Sequential detection with sequential inputs is powerful
  - works very well for meal-detection
  - dominates rate-of-change approaches in literature

- Diabetic meal detection is not a new problem (over 15 years old)
  - No classical "machine learning" solution in literature
  - why? ... possibly because of physiological variability between patients

- What if the system has some structure which can be exploited?

Outline

- Meal detection in type I diabetics
  - unknown linear time invariant systems
- Critical pulmonary shunt detection in infants
  - detection in structured linear systems with unknown parameters
- Building actuator fault detection
  - signal detection in unknown networked systems

Detecting Critical Pulmonary Shunts in Infants

Both lungs participating in pulmonary exchange

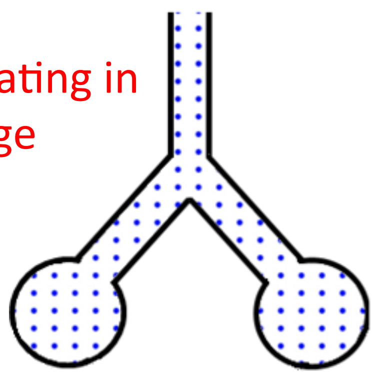

Shunt

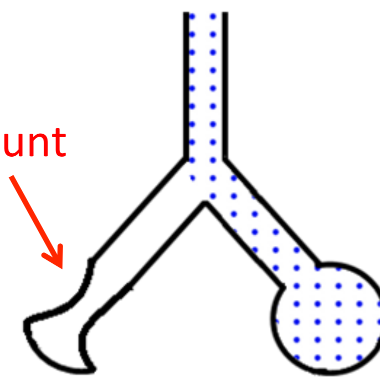

One lung participating in pulmonary exchange

Critical Pulmonary Shunt Detection Problem

- Option A: hypothesize the shunt as an input
  - use the unknown LTI system monitor (as before)

- Option B: build a "structured" model of the dynamics when:
  - a shunt is present
  - a shunt is not present

- Both options require some model information
  - where does this come from?

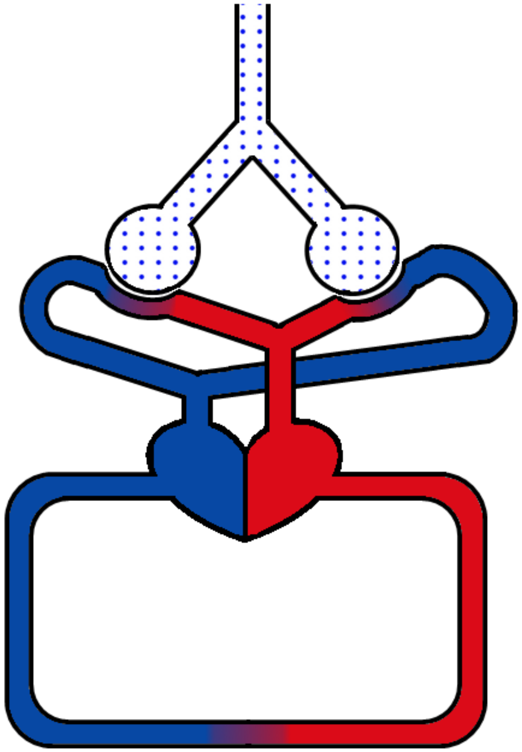

Compartmental Modeling

- Option A: hypothesize the shunt as an input (use LTI approach)
  - requires little domain expertise
- Qualitative heuristic for option A: add dimension(s) to the LTI model when:
  - physical separation (+1 per degree separation)
  - time-delay (+1 per unit delay)
  - test signal is not “really” an impulse (+ model_order_needed)
  - critical shunt detection: model order = 4
    - diffusion → +1, circulation delay → +2, sustained event → + 1
- Concept extends beyond physiology
  - networks (degree of separation)
  - any dynamically coupled linkage
    - e.g. fluid transfer in automotive transmission

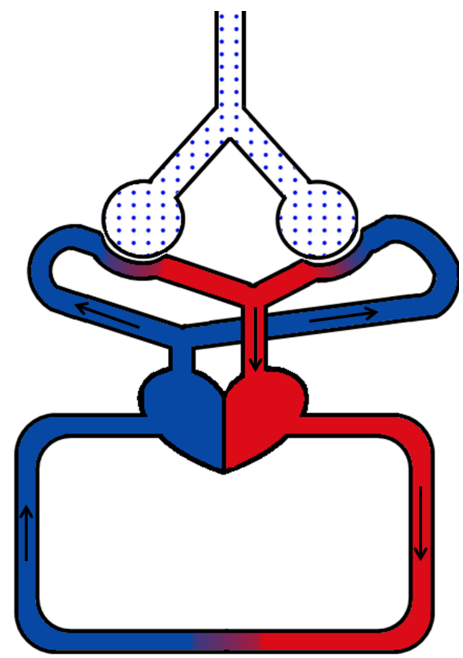

Apply 2-sided PAIN monitor as before

Penn Engineering

PRECISE

Compartmental Modeling

- Option B: build a structured model of the dynamics
  - requires significant domain expertise

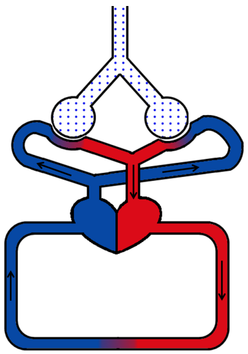

no shunt dynamics

$$\begin{bmatrix} x^L(k)\\ x^R(k)\end{bmatrix} = \begin{bmatrix}\frac{\alpha}{V(k)} & \frac{\alpha}{V(k)}\\ \frac{\alpha}{V(k)} & \frac{\alpha}{V(k)}\end{bmatrix}\begin{bmatrix}x^L(k-\kappa)\\ x^R(k-\kappa)\end{bmatrix} + \begin{bmatrix}\frac{2\alpha}{V(k)} & n^L(k)\\ \frac{2\alpha}{V(k)} & n^R(k)\end{bmatrix}\begin{bmatrix}\mu\\ \sigma\end{bmatrix}$$

$$y(k) = \begin{bmatrix}\frac{1}{2} & \frac{1}{2}\end{bmatrix}\begin{bmatrix}x^L(k)\\ x^R(k)\end{bmatrix}$$

diffusion coefficient

$$\mathcal{H}_j :\ \mathbf{y} = \mathbf{H}_j\theta + \sigma_j\mathbf{n} \qquad \theta = \begin{bmatrix}\alpha\\ \alpha\mu\end{bmatrix}$$

metabolism times diffusion coefficient

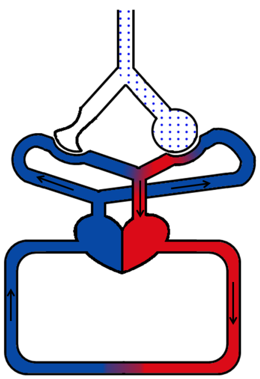

$$\begin{bmatrix} x^{NS}(k)\\ x^S(k)\end{bmatrix} = \begin{bmatrix}\frac{\alpha}{2V(k)} & \frac{\alpha}{2V(k)}\\ \frac{1}{2} & \frac{1}{2}\end{bmatrix}\begin{bmatrix}x^{NS}(k-\kappa)\\ x^S(k-\kappa)\end{bmatrix} + \begin{bmatrix}\frac{\alpha}{V(k)} & n^{NS}(k)\\ 1 & 0\end{bmatrix}\begin{bmatrix}\mu\\ \sigma\end{bmatrix}$$

$$y(k) = \begin{bmatrix}1 & 0\end{bmatrix}\begin{bmatrix}x^{NS}(k)\\ x^S(k)\end{bmatrix}$$

shunt dynamics

- **pros: potential gains in performance**
- **cons: difficult to design**

PAIN Critical Shunt Monitor Evaluation

- 209 human patients considered (all children)
  - 61 patients experiencing with potential critical shunts
    - annotations are unreliable
  - 148 patients without a shunt
- Compare the following approaches
  - dd-PAIN_{TS} → option A with trained thresholds
  - $\text{PAIN}_{\text{PHYS}}$ → option B without trained thresholds
  - dd-$\text{PAIN}_{\text{PHYS}}$ → option B with trained thresholds
  - $\text{GLRT}_{\text{PHYS}}$ → physiology based GLRT with trained thresholds
- Evaluate on the criteria:
  - false positive rate variability between patients (false positive rate vs. patient)
    - using patients without a shunt
  - predictive capability of the detector (true positive rate vs. time)
    - using patients with a shunt

Penn Engineering

PRECISE
PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING

Critical Shunt Monitor Performance

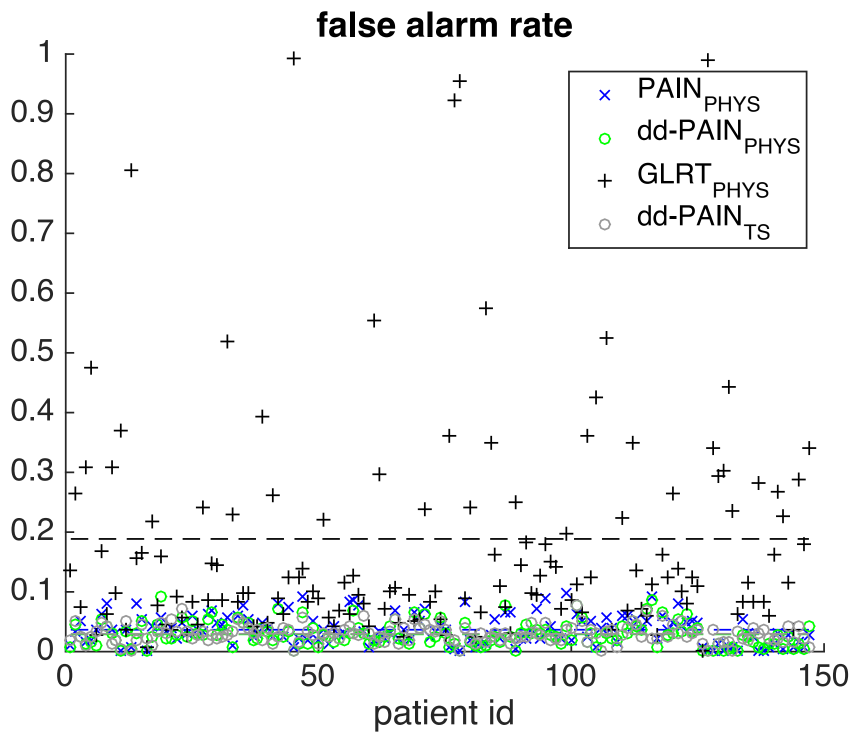

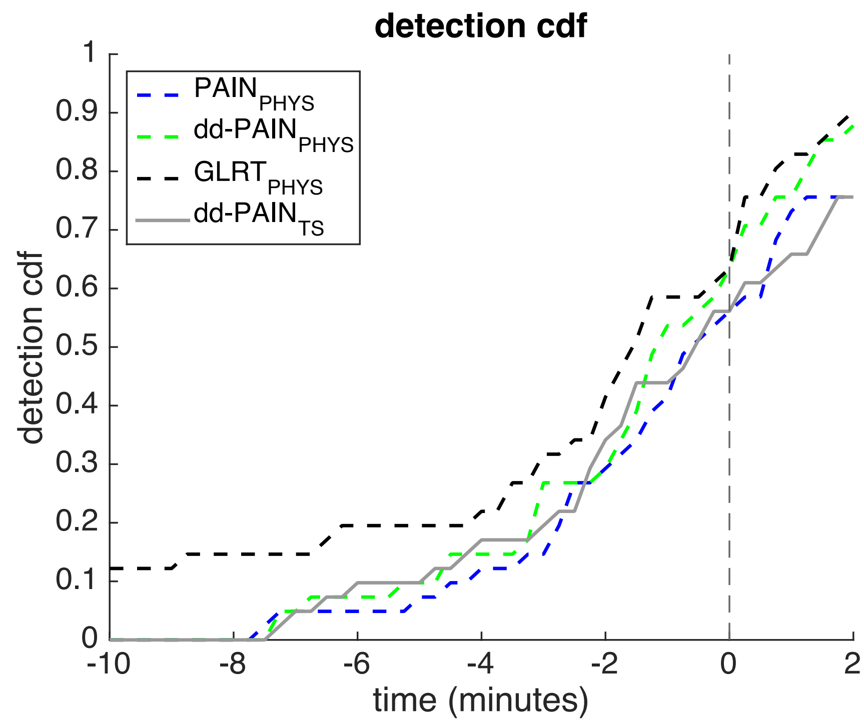

- **trained option B is the "best"**
- **trained option A is still good**
- **GLRT has wide variance in false positive rate across patients**

Summary : Detection in Structured Linear Systems

- Improved performance achievable by including physical model knowledge
  - sequential detection of sequential events approach still can be useful
- GLRT can not bound the false positive rate in all applications
  - e.g. critical shunt detection
  - statement generalizes to other classical data-driven approaches
    - e.g. detection/classification via ARMAX features
- Are there any physical model invariances that are easy to exploit?
  - Doesn't require domain knowledge to build a model.

Outline

- Meal detection in type I diabetics
  - unknown linear time invariant systems
- Critical pulmonary shunt detection in infants
  - detection in structured linear systems with unknown parameters
- Building actuator fault detection
  - signal detection in unknown networked systems

Detecting Building Actuator Faults

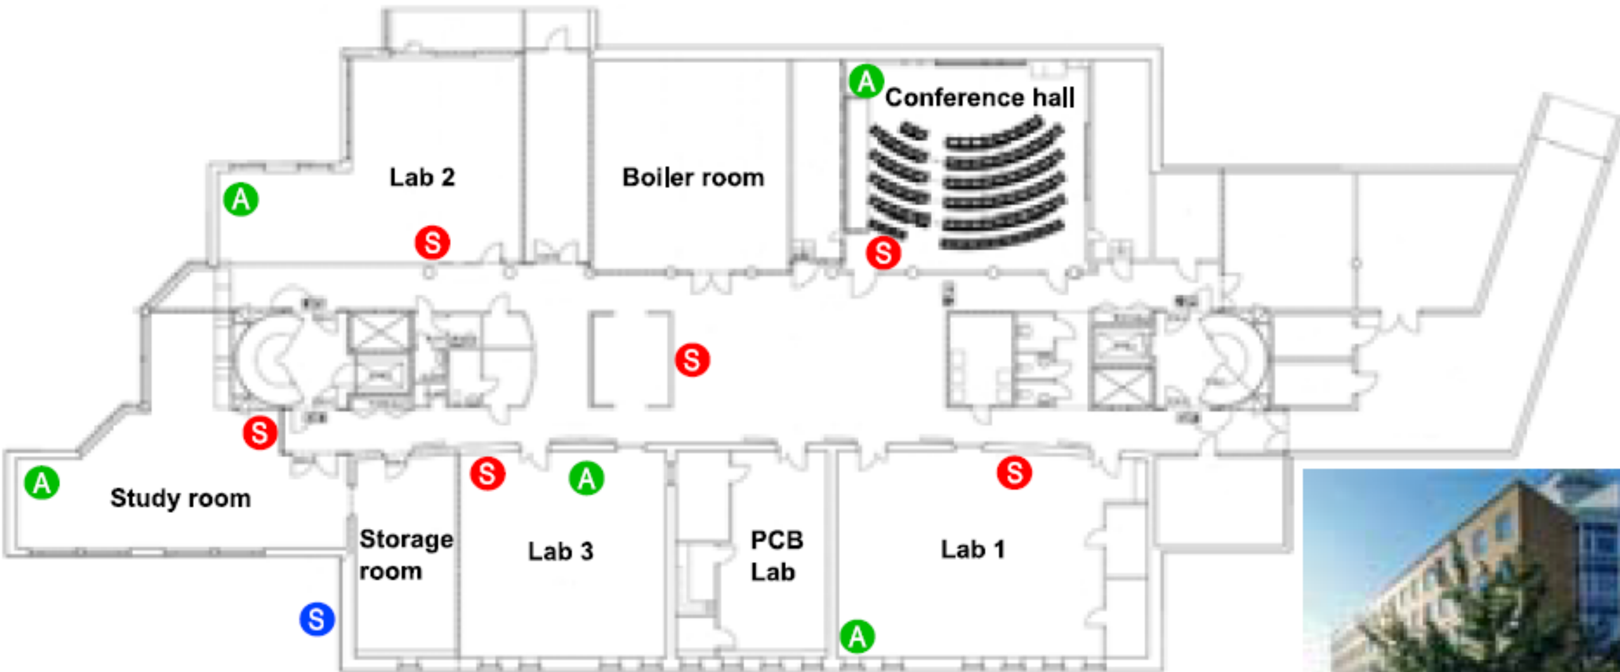

Building Actuator Fault Detection Problem

- Test signals:
  - H_0 : applied actuator voltage
  - H_1 : a constant voltage
    - captures “zero” applied voltage (electrical failure)
    - captures stuck in a position (mechanical failure)

- Dynamics are well approximated by a network system
  - dynamics has a unit eigenvalue corresponding to sum of values in network
    - a natural invariant to dynamics
    - still has unknown model error

punch line:
reduces to the standard form

$$\mathcal{H}_j : \ \mathbf{y} = \mathbf{H}_j \theta + \sigma_j \mathbf{n}$$

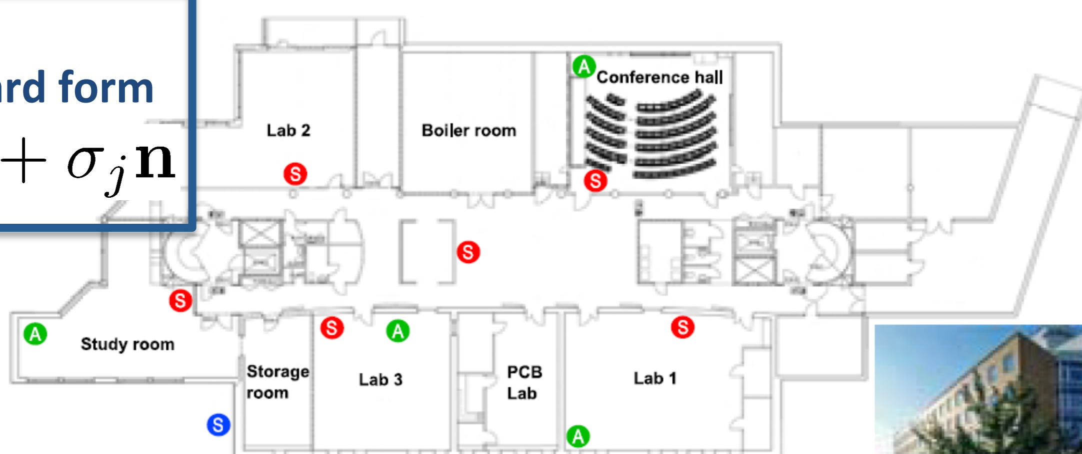

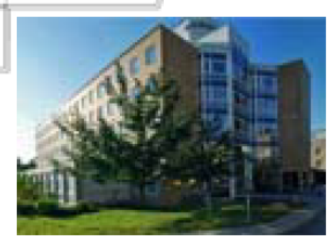

Penn Engineering

PRECISE

Building Fault Detector Performance

- near constant false alarm rate, detection rate improves with time

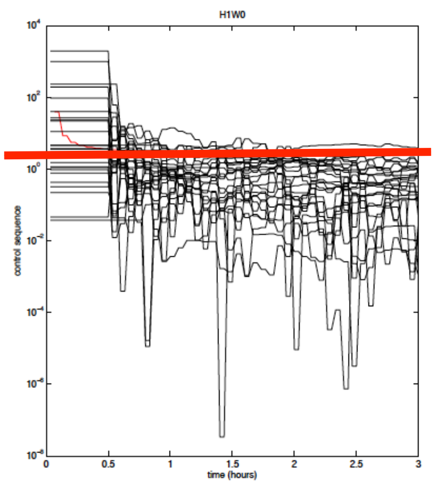

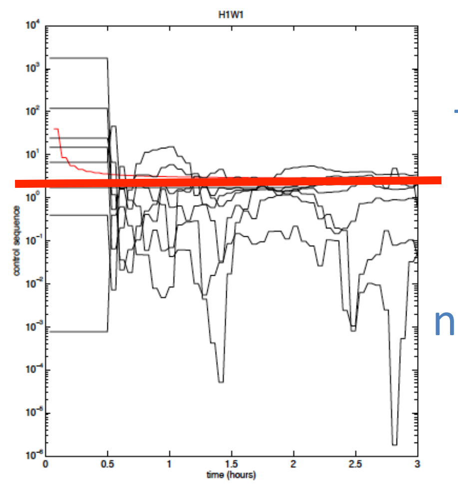

fault

no fault

window closed (3 hours) window open

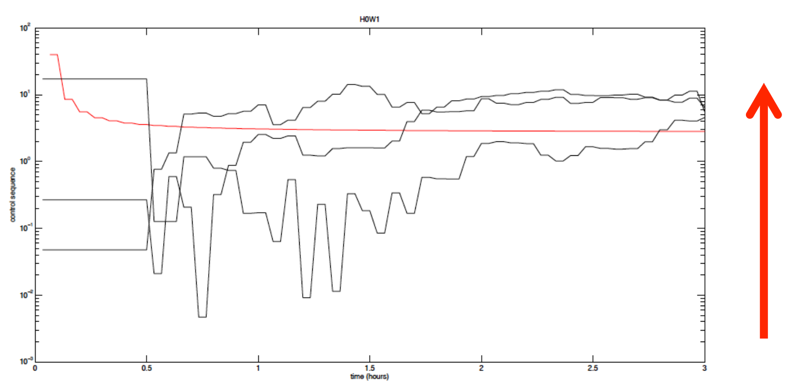

increasing true positive rate

Penn Engineering

PRECISE

Summary: Detection in Unknown Networked Dynamics

- Exploiting natural invariances can be useful
  - reduction in model error

- Many other “systems” are well approximated by networked dynamics
  - power transmission dynamics
  - epidemics
  - social dynamics

Closing Remarks and Insight

- parameter-invariant monitoring is a structured approach to monitor design that addresses variability in CPS applications.
  - can address some difficulties with classical monitor design
- The general form presented herein is not the only statistic:
  - statistics to deal with missing measurements
  - cases when parts of the model are known
    - e.g. model error is known
- Machine learning + Parameter Invariant statistics
  - use parameter invariant techniques to generate feature
    - invariant to variability
  - learn the best classifier over the parameter invariant features
    - can boost performance
- See all our work at: https://rtg.cis.upenn.edu/parameter-invariant.html

Penn Engineering

PRECISE