Design and Optimisation Problems in Wireless Sensor Networks

Anurag Kumar

Department of Electrical Communication Engineering
Indian Institute of Science (IISc), Bangalore 560012, India

Outline of Talk

1. Wireless Sensor Networks (WSNs): Quick Overview
2. Optimum Event Detection
3. Other Algorithm Design Work

Outline of Talk

1. Wireless Sensor Networks (WSNs): Quick Overview
2. Optimum Event Detection
3. Other Algorithm Design Work

Wireless Networks of Multifunction Smart Sensors (WSNs)

- A smart sensor node is popularly called a *mote*
- Sensing: temperature, chemicals, light, infrared, biosensors, strain, sound, vibration (often using MEMS technology)
- Processing: e.g., 16-bit, 8 Mhz, 48KB flash, 10 KB RAM
  - with a simple OS
- Digital radio: e.g., ISM band; a few Kbps
- Battery: e.g., 100mAh ("button" batteries) to 2000mAh (2 AA batteries)

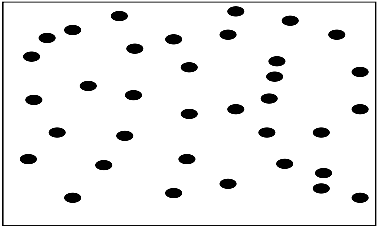

The Structure and Operation of a WSN

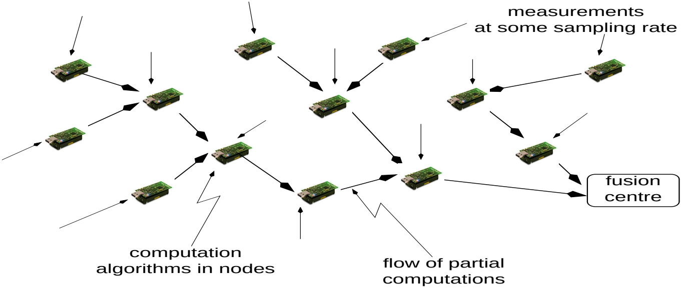

- Motes are deployed and self-organise into a mesh network
- Nodes contain pieces of a distributed signal processing algorithm
- Computations are performed on the measurements
- Partial results flow over the links
- Downstream nodes combine their inputs and forward the results
- The MAC schedules the transmissions in the network
- Embedded distributed smart instrumentation

The Structure and Operation of a WSN

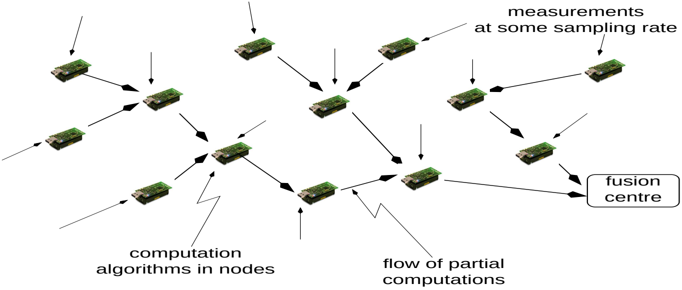

- Motes are deployed and self-organise into a mesh network
- Nodes contain pieces of a distributed signal processing algorithm
- Computations are performed on the measurements
- Partial results flow over the links
- Downstream nodes combine their inputs and forward the results
- The MAC schedules the transmissions in the network
- Embedded distributed smart instrumentation

Potential Applications of Wireless Sensor Networks

- Environment monitoring
- Self-monitoring structures
- Industrial sensor networks
  - Chemical factories, refineries, power distribution yards
  - Energy and environment management networks in large buildings
  - Emerging ISA 100 standard (IEEE 802.15.4 PHY, TDMA MAC)
- Ecological monitoring and wild life management
- Monitoring mobile patients in hospitals and homes
- Locating people in large buildings or public transportation systems
- Detecting intruders in border areas (smart fields, electronic "trip wires")

Design and Optimisation Problems in WSNs

- Node placement
  - Sensing coverage, and radio connectivity
- Topology design
  - Discover a topology that has good throughput and delay performance
  - Related to node placement
- Node localisation: GPS might not work in certain domains
- Packet forwarding:
  - Nodes sleep-wake cycle; table driven routing not practical
  - Discover routes on the fly
- Time synchronisation across the network
- Distributed signal processing
  - Event detection and classification
  - Object tracking

Outline of Talk

1. Wireless Sensor Networks (WSNs): Quick Overview
2. Optimum Event Detection
3. Other Algorithm Design Work

The Classical Change Detection Problem (CDP)

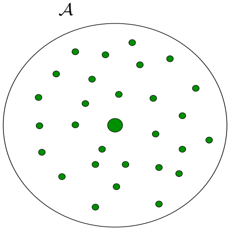

Before Change $X_k^{(i)} \sim f_0$

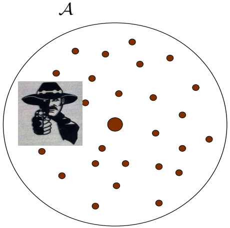

After Change $X_k^{(i)} \sim f_1$

- **The Setting:** Discrete time, $k \geq 0$. The change occurs at a random time $T \geq 0$.
  - Need to detect the change
  - We have a prior on T, and the observations $X_k^{(i)}$, $1 \leq i \leq N, k \geq 1$
    - Conditionally i.i.d. in time and space
- **The Problem:** Determine a stopping time τ, such that

$$\begin{aligned} \min \quad & \mathsf{E}\Big[(\tau - T)^+\Big] \\ \text{s.t.} \quad & \mathsf{P}(\tau < T) \leqslant \alpha \end{aligned}$$

- **The Solution:** The optimal τ is given by [Shiryaev '63]

$$\begin{aligned} \tau &= \inf\left\{k : \Pi_k \geqslant \Gamma_\alpha\right\} \\ \text{where } \Pi_k &= \mathsf{P}\left(T \leqslant k \mid \mathbf{X}_{[1:k]}\right) \end{aligned}$$

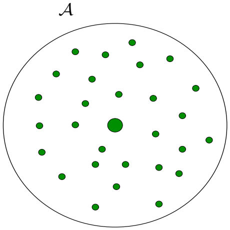

Before Change $X_k^{(i)} \sim f_0$

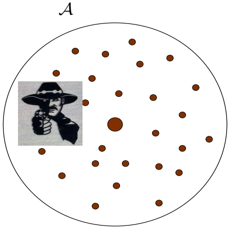

After Change $X_k^{(i)} \sim f_1$

- **The Setting:** Discrete time, $k \geq 0$. The change occurs at a random time $T \geq 0$.
  - Need to detect the change
  - We have a prior on T, and the observations $X_k^{(i)}$, $1 \leq i \leq N, k \geq 1$
    - Conditionally i.i.d. in time and space
- **The Problem:** Determine a stopping time τ, such that

$$\begin{aligned} \text{min} \quad & \mathsf{E}\Big[(\tau - T)^+\Big] \\ \text{s.t.} \quad & \mathsf{P}(\tau < T) \leqslant \alpha \end{aligned}$$

- **The Solution:** The optimal τ is given by [Shiryaev '63]

$$\begin{aligned} \tau &= \inf\left\{k : \Pi_k \geqslant \Gamma_\alpha\right\} \\ \text{where } \Pi_k &= \mathsf{P}\left(T \leqslant k \mid \mathbf{X}_{[1:k]}\right) \end{aligned}$$

The Classical Change Detection Problem (CDP)

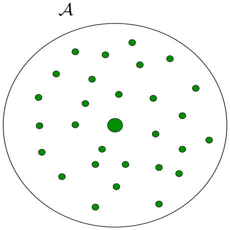

Before Change $X_k^{(i)} \sim f_0$

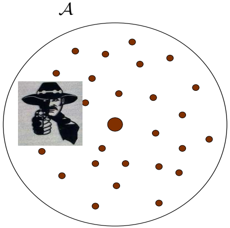

After Change $X_k^{(i)} \sim f_1$

- **The Setting:** Discrete time, $k \geq 0$. The change occurs at a random time $T \geq 0$.
  - Need to detect the change
  - We have a prior on T, and the observations $X_k^{(i)}$, $1 \leq i \leq N, k \geq 1$
    - Conditionally i.i.d. in time and space
- **The Problem:** Determine a stopping time τ, such that

$$\begin{aligned} \min \quad & \mathsf{E}\left[(\tau - T)^+\right] \\ \text{s.t.} \quad & \mathsf{P}(\tau < T) \leqslant \alpha \end{aligned}$$

- **The Solution:** The optimal τ is given by [Shiryaev '63]

$$\begin{aligned} \tau &= \inf\left\{k : \Pi_k \geqslant \Gamma_\alpha\right\} \\ \text{where } \Pi_k &= \mathsf{P}\left(T \leqslant k \mid \mathbf{X}_{[1:k]}\right) \end{aligned}$$

Limitations of the Classical CDP

- Recent motivation: event detection in wireless sensor networks
  - The classical model does not account for many practical aspects
- Sleep-wake scheduling of sensors (cost for making observations)
  - Optimal detection with sleep-wake scheduling: [Infocom 2008]
- Transient change: The event occurs and then disappears in a finite time
  - Quickest transient change detection: [IWAP 2010]
- Networked sensors
  - Random access packet network between the sensors and the fusion centre [SECON 2007, ACM TOSN 2010]
- Large extent network: spatial variation in the way the change (event) affects the sensors
  - Distributed detection and localisation in large wireless sensor networks: [Allerton 2009]

Limitations of the Classical CDP

- Recent motivation: event detection in wireless sensor networks
  - The classical model does not account for many practical aspects
- Sleep-wake scheduling of sensors (cost for making observations)
  - Optimal detection with sleep-wake scheduling: [Infocom 2008]
- Transient change: The event occurs and then disappears in a finite time
  - Quickest transient change detection: [IWAP 2010]
- Networked sensors
  - Random access packet network between the sensors and the fusion centre [SECON 2007, ACM TOSN 2010]
- Large extent network: spatial variation in the way the change (event) affects the sensors
  - Distributed detection and localisation in large wireless sensor networks: [Allerton 2009]

Limitations of the Classical CDP

- Recent motivation: event detection in wireless sensor networks
    - The classical model does not account for many practical aspects
- Sleep-wake scheduling of sensors (cost for making observations)
    - Optimal detection with sleep-wake scheduling: [Infocom 2008]
- Transient change: The event occurs and then disappears in a finite time
    - Quickest transient change detection: [IWAP 2010]
- Networked sensors
    - Random access packet network between the sensors and the fusion centre [SECON 2007, ACM TOSN 2010]
- Large extent network: spatial variation in the way the change (event) affects the sensors
    - Distributed detection and localisation in large wireless sensor networks: [Allerton 2009]

Limitations of the Classical CDP

- Recent motivation: event detection in wireless sensor networks
  - The classical model does not account for many practical aspects
- Sleep-wake scheduling of sensors (cost for making observations)
  - Optimal detection with sleep-wake scheduling: [Infocom 2008]
- Transient change: The event occurs and then disappears in a finite time
  - Quickest transient change detection: [IWAP 2010]
- Networked sensors
  - Random access packet network between the sensors and the fusion centre [SECON 2007, ACM TOSN 2010]
- Large extent network: spatial variation in the way the change (event) affects the sensors
  - Distributed detection and localisation in large wireless sensor networks: [Allerton 2009]

Limitations of the Classical CDP

- Recent motivation: event detection in wireless sensor networks
  - The classical model does not account for many practical aspects
- Sleep-wake scheduling of sensors (cost for making observations)
  - Optimal detection with sleep-wake scheduling: [Infocom 2008]
- Transient change: The event occurs and then disappears in a finite time
  - Quickest transient change detection: [IWAP 2010]
- Networked sensors
  - Random access packet network between the sensors and the fusion centre [SECON 2007, ACM TOSN 2010]
- Large extent network: spatial variation in the way the change (event) affects the sensors
  - Distributed detection and localisation in large wireless sensor networks: [Allerton 2009]

- Nodes make periodic observations (e.g., temperature; gas concentration (e.g., LPG))
- Measurements (or local decisions) are sent to the "fusion centre"...
- ... over a random access network

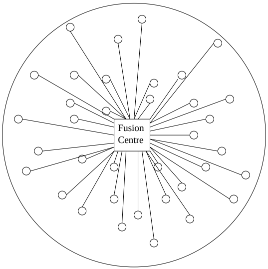

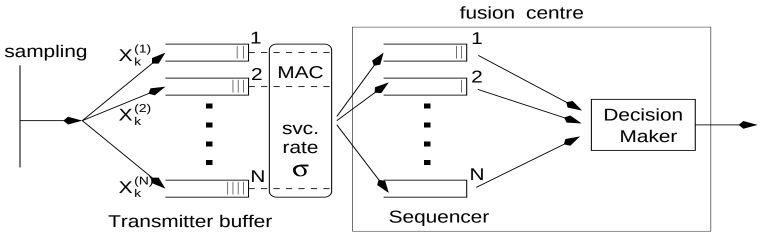

The Sequential Detection Problem with Network Delays

- The sequencer at the fusion centre releases complete batches of samples to the detection algorithm

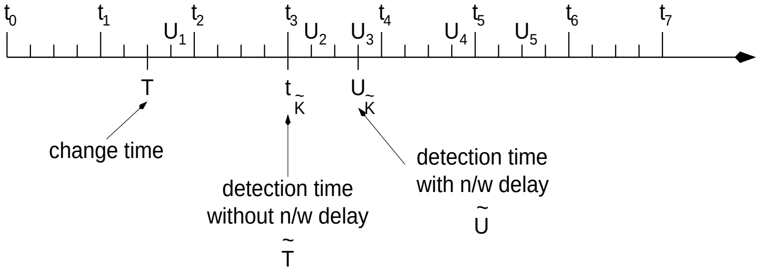

- Minimise detection delay s.t. a false alarm probability bound, i.e.,

$$\min \mathsf{E}\left((\tilde{U} - T) I_{\{\tilde{T} \geq T\}}\right)$$

subject to

$$\mathsf{Pr}\left(\tilde{T} < T\right) \leq \alpha$$

Collocated Nodes: Change Detection over a Network

Decision and Network Delays Decouple

Theorem

If the sampling is periodic at rate r and the batch sojourn time process D_b, $b \geqslant 1$, is stationary with mean $d(r)$, then

$$\min_{\Pi_\alpha} \mathsf{E}\Big((\tilde{U} - T) I_{\{\tilde{T} \geq T\}}\Big)$$

$$= \quad l(r)\,(1-\alpha) + d(r)\,(1-\alpha) + \frac{1}{r} \min_{\Pi_\alpha} \mathsf{E}\Big(\tilde{K} - K\Big)^+$$

where $l(r)$ is the delay due to (coarse) sampling, and K is the index of the first sample after the change.

- A cross-layer result: combines *network performance* & *statistical performance* for the first time
- The problem $\min_{\Pi_\alpha} \mathsf{E}\left(\tilde{K} - K\right)^+$ is the classical one
  - The threshold will depend on the sampling rate r

Theorem

If the sampling is periodic at rate r and the batch sojourn time process D_b, $b \geqslant 1$, is stationary with mean $d(r)$, then

$$\min_{\Pi_\alpha} \mathsf{E}\Big((\tilde{U} - T) I_{\{\tilde{T} \geq T\}}\Big) = l(r)\,(1-\alpha) + d(r)\,(1-\alpha) + \frac{1}{r} \min_{\Pi_\alpha} \mathsf{E}\Big(\tilde{K} - K\Big)^{+}$$

where $l(r)$ is the delay due to (coarse) sampling, and K is the index of the first sample after the change.

- A cross-layer result: combines *network performance* & *statistical performance* for the first time
- The problem $\min_{\Pi_\alpha} \mathsf{E}\left(\tilde{K} - K\right)^{+}$ is the classical one
  - The threshold will depend on the sampling rate r

Theorem

If the sampling is periodic at rate r and the batch sojourn time process D_b, $b \geqslant 1$, is stationary with mean $d(r)$, then

$$\min_{\Pi_\alpha} \mathsf{E}\Big((\tilde{U} - T) I_{\{\tilde{T} \geq T\}}\Big) = l(r)\,(1-\alpha) + d(r)\,(1-\alpha) + \frac{1}{r} \min_{\Pi_\alpha} \mathsf{E}\Big(\tilde{K} - K\Big)^{+}$$

where $l(r)$ is the delay due to (coarse) sampling, and K is the index of the first sample after the change.

- A cross-layer result: combines *network performance* & *statistical performance* for the first time
- The problem $\min_{\Pi_\alpha} \mathsf{E}\Big(\tilde{K} - K\Big)^{+}$ is the classical one
  - The threshold will depend on the sampling rate r

- Example: $\rho = 0, p = 0.0005, \alpha = 0.01, \sigma = 0.3636, N = 10$
- Stability requirement $N \times r < \sigma$, the success rate
- Mean detection delay vs. sampling rate r

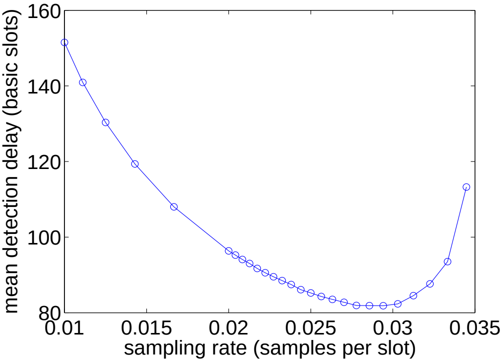

- The optimal sampling rate is $r = 0.029$ (sampling interval = 34 slots)

Collocated Nodes: Optimal Node Energy Management

- Detect an event at a random time, T
  - n sensors in the observation field
- Sensor i makes the measurements $X_k^{(i)}, k \geq 1$
- How many sensors should we turn on?
- Tradeoffs: detection delay, false alarm rate, battery life

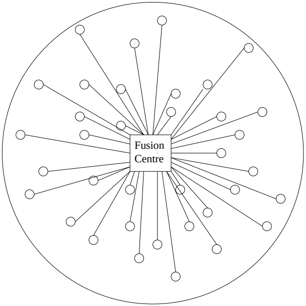

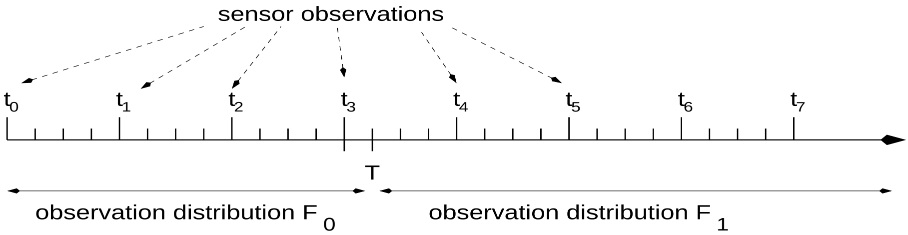

Collocated Nodes: Optimal Node Energy Management

- Detect an event at a random time, T
  - n sensors in the observation field
- Sensor i makes the measurements $X_k^{(i)}, k \geq 1$
- How many sensors should we turn on?
- Tradeoffs: detection delay, false alarm rate, battery life

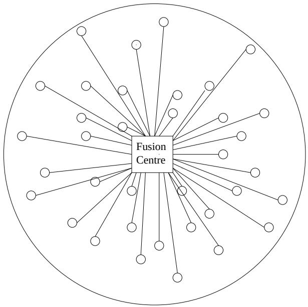

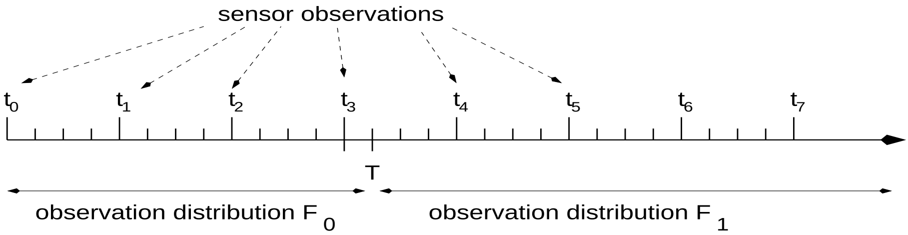

- Detect an event at a random time, T
  - n sensors in the observation field
- Sensor i makes the measurements $X_k^{(i)}, k \geq 1$
- How many sensors should we turn on?
- Tradeoffs: detection delay, false alarm rate, battery life

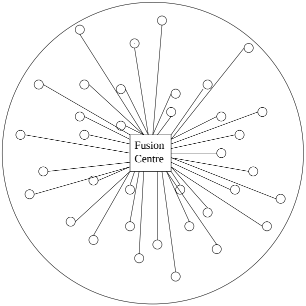

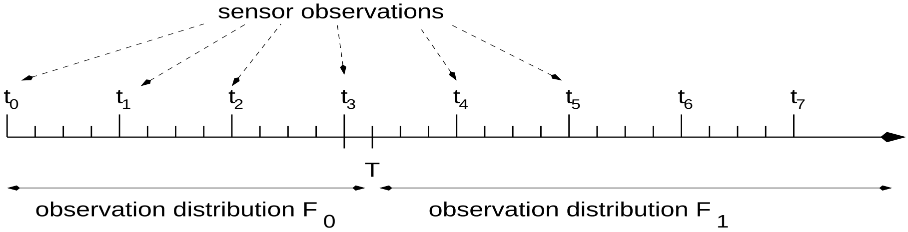

Collocated Nodes: Optimal Node Energy Management

Dynamic Control of Mote Sleep-Wake Cycling

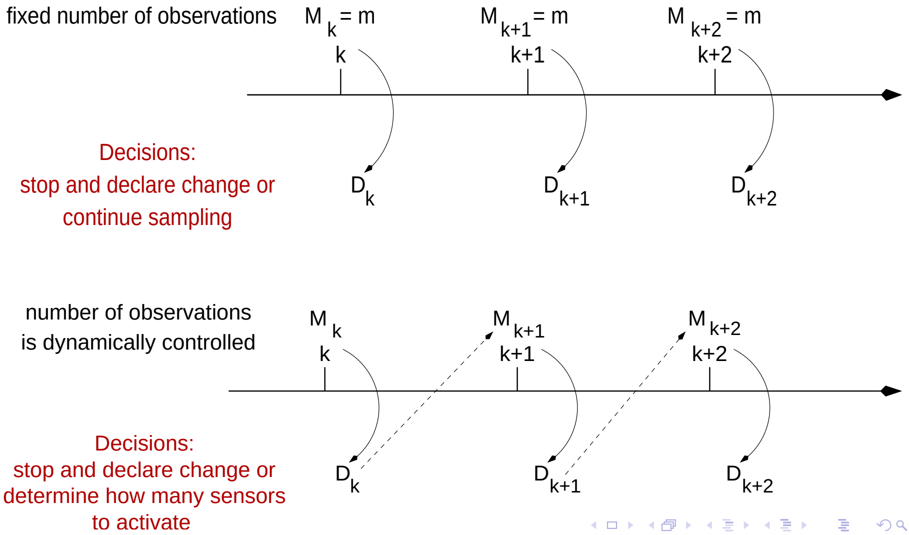

Collocated Nodes: Optimal Node Energy Management

The Decision Rule

- At sample time k, the observation vector is $\mathbf{X}_k^{(M_k)}$
  - the vector of observations from M_k sensors
- Compute $\Pi_k = \mathsf{P}\left\{T \le k \,\middle|\, \mathbf{X}_1^{(M_1)}, \cdots, \mathbf{X}_k^{(M_k)}\right\}$
  - $\Pi_k =$ posterior probability of the event having occurred at or before k
- An example
  - $n = 10$ sensors
  - False alarm cost $= 100.0$
  - Sensor activation cost $= 0.5$
  - Delay cost $= 1$ per sample delay
  - $f_0 \sim \mathcal{N}(0, 1)$, $f_1 \sim \mathcal{N}(1, 1)$
  - Optimal algorithm
    - Stop if $\Pi_k \ge 0.9$
    - Else use M_k as in the plot

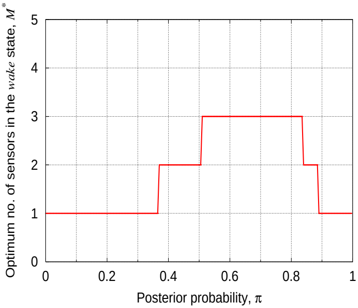

- At sample time k, the observation vector is $\mathbf{X}_k^{(M_k)}$
  - the vector of observations from M_k sensors
- Compute $\Pi_k = \mathsf{P}\left\{T \le k \,\middle|\, \mathbf{X}_1^{(M_1)}, \cdots, \mathbf{X}_k^{(M_k)}\right\}$
  - Π_k = posterior probability of the event having occurred at or before k

- An example
  - $n = 10$ sensors
  - False alarm cost $= 100.0$
  - Sensor activation cost $= 0.5$
  - Delay cost $= 1$ per sample delay
  - $f_0 \sim \mathcal{N}(0,1)$, $f_1 \sim \mathcal{N}(1,1)$
  - Optimal algorithm
    - Stop if $\Pi_k \ge 0.9$
    - Else use M_k as in the plot

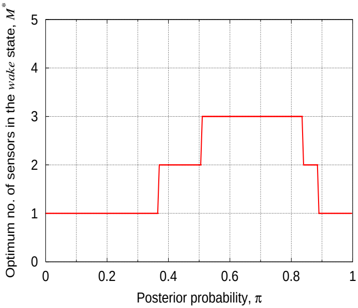

- At sample time k, the observation vector is $\mathbf{X}_k^{(M_k)}$
  - the vector of observations from M_k sensors
- Compute $\Pi_k = \mathsf{P}\left\{ T \le k \,\middle|\, \mathbf{X}_1^{(M_1)}, \cdots, \mathbf{X}_k^{(M_k)} \right\}$
  - Π_k = posterior probability of the event having occurred at or before k
- An example
  - $n = 10$ sensors
  - False alarm cost $= 100.0$
  - Sensor activation cost $= 0.5$
  - Delay cost $= 1$ per sample delay
  - $f_0 \sim \mathcal{N}(0,1)$, $f_1 \sim \mathcal{N}(1,1)$
  - Optimal algorithm
    - Stop if $\Pi_k \ge 0.9$
    - Else use M_k as in the plot

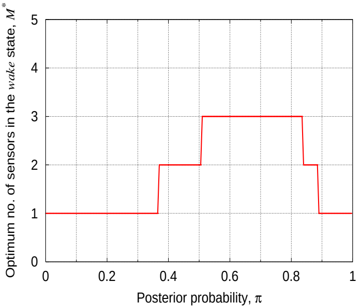

Collocated Nodes: Optimal Node Energy Management

Comparison with a Fixed Number of Active Sensors

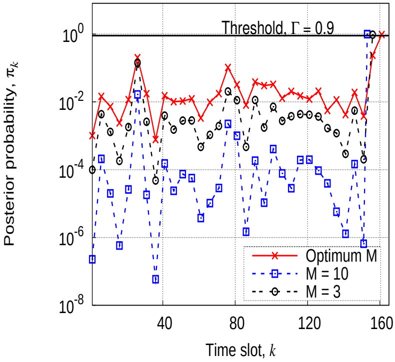

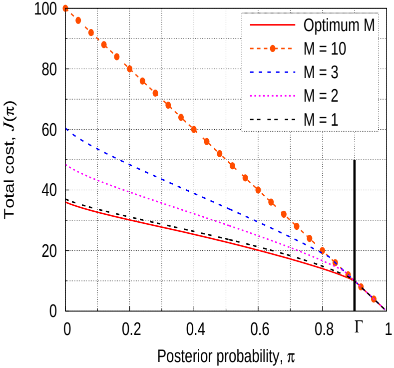

- The optimum rule initially keeps just one sensor active
- Rapidly increases the number of active sensors when the event occurs
- Plots on the right show the average cost comparison with various fixed number of sensors

Distributed Detection: Event Detection in a Large Network

Before the Event

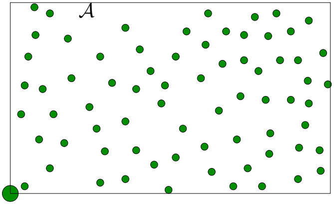

$$X_k^{(i)} \sim f_0.$$

After the Event

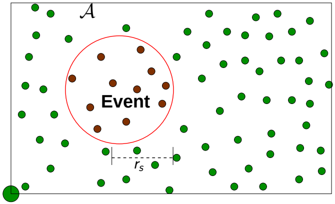

$$X_k^{(i)} \sim \begin{cases} f_1, & \text{if } i \in \mathcal{N}_e(r_s), \\ f_0, & \text{if } i \notin \mathcal{N}_e(r_s). \end{cases}$$

Problem: An event happens at an unknown time T and at an unknown location ℓ_e. **Detect** and **locate** the event

Objective: Raise an alarm with a **small detection delay** subject to **Time to False Alarm (TFA)** $\geqslant \gamma$ and **Time to False Isolation (TFI)** $\geqslant \gamma$.

Distributed Detection: Event Detection in a Large Network

Before the Event

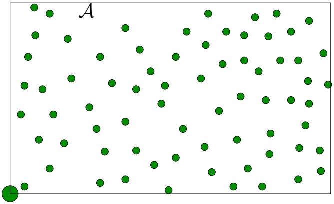

$$X_k^{(i)} \sim f_0.$$

After the Event

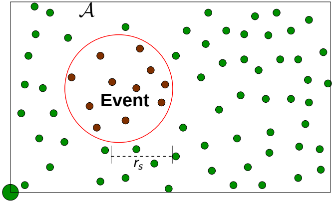

$$X_k^{(i)} \sim \begin{cases} f_1, & \text{if } i \in \mathcal{N}_e(r_s), \\ f_0, & \text{if } i \notin \mathcal{N}_e(r_s). \end{cases}$$

Problem: An event happens at an unknown time T and at an unknown location ℓ_e. **Detect** and **locate** the event

Objective: Raise an alarm with a **small detection delay** subject to **Time to False Alarm (TFA)** $\geqslant \gamma$ and **Time to False Isolation (TFI)** $\geqslant \gamma$.

We deployed 10 nodes in the ROI in a deterministic manner such that we get $N = 9$ regions and $\mathcal{N}_1 = \{1,2,4\}$, $\mathcal{N}_2 = \{2,4,5\}$, $\mathcal{N}_3 = \{2,3,5\}$, $\mathcal{N}_4 = \{4,6,7\}$, $\mathcal{N}_5 = \{4,5,7\}$, $\mathcal{N}_6 = \{5,7,8\}$, $\mathcal{N}_7 = \{6,7,9\}$, $\mathcal{N}_8 = \{7,9,10\}$, and $\mathcal{N}_9 = \{7,8,10\}$.

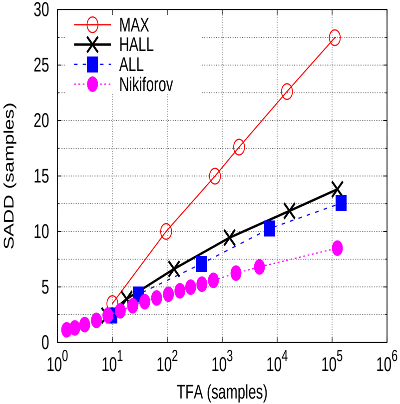

- Nikiforov (1995) provides the optimal centralised algorithm
- Distributed algorithms are based on fusion of local CUSUM statistics
- We also provide an approach for a power-law signal decay model

Outline of Talk

1. Wireless Sensor Networks (WSNs): Quick Overview
2. Optimum Event Detection
3. Other Algorithm Design Work

Other WSN Algorithm Design Work

- WSN node placement
  - Minimum cost
  - Performance objectives (delay and network reliability)
- Node localisation without GPS
  - Hop count based localisation
  - Theory based on random graph models
- Geographical forwarding of alarm packets in WSN
  - Optimal forwarding when nodes sleep-wake cycle
  - Algorithms for "hole" avoidance