Energy-Centric Theory for Wireless Networks

February 11, 2005

Abbas El Gamal
Department of Electrical Engineering
Stanford University

Acknowledgments: A. Ercan, J. Mammen, S. Zahedi

Wireless Imaging Sensor Networks

- Imaging sensors are all around us:
  - Cell phones
  - Digital cameras
  - Cars
  - Security cameras

- These sensors are:
  - Low cost
  - Low power
  - Digital
  - Integrated with processing
  - High resolution
  - Can operate at 100’s of frames/sec and achieve over 120dB of dynamic range

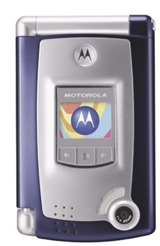

- My research group has been involved in the design of such sensors and we are now looking at networking them

Current and Future Networks

- There are already “networks” of surveillance cameras:
  - Analog and cabled
  - Fixed nodes
  - Data is shipped to monitors, observed by human operators
  - Expensive, non-scalable

- Future networks:
  - Digital and wireless
  - Fixed or mobile nodes (cell phones)
  - Automated detection
  - Low cost and scalable
  - More pervasive

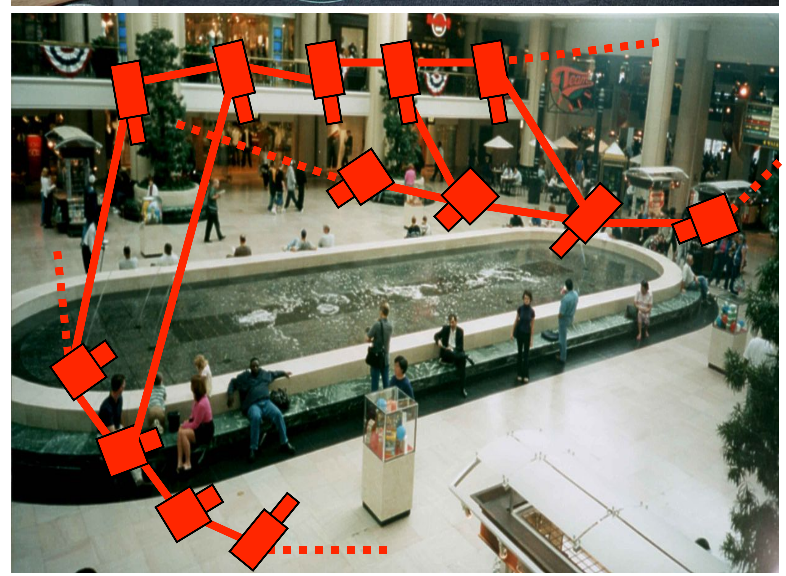

Energy is a Primary Concern

- Imaging sensors can produce *a lot* of data:
  - **Send everything**: Consume too much ***communication energy***
    - Compression may not be enough
  - **Local processing**: Implementing sophisticated vision algorithms require too much ***processing energy***
- Wireless networks today optimized for *throughput-delay:*
  - Energy is considered, but mainly at hardware level
- Imaging sensor networks must consider energy in *all* aspects of the design
  - Network architecture, communication, processing, hardware
- In some cases we have some theory to help us
- In others, we don’t have much theory

Outline

- Examples where we have some theory:
  - Does multi-hopping save energy?
  - Joint backlog and channel adaptation saves energy
  - Multi-user coding saves energy
- Example where we can save a lot of energy using local processing, but there is no theory:
  - Collaborative task-driven wireless imaging network
- General comments

Communication Energy

- Communication energy consists of:
  - RF transmission energy
  - Transceiver (circuit) energy
  - Standby energy
- Let’s ignore standby energy:
  - Transmission and transceiver energy likely to dominate in imaging sensor networks
  - Standby can be significantly reduced [Rabaey]
    - Avoid synchronous monitoring
    - Make radio wake up to external events
    - Reduce leakages

Communication Energy Model

- Power decays with distance as $r^{2\delta}, \delta \geq 1$
- AWGN channel model with noise power N
- Treat interference from other transmitting users as Gaussian noise with power I
- Rate at transmission power P is given by:

$$R(P,r) = \frac{1}{2}\log\left(1 + \frac{Pr^{-2\delta}}{N+I}\right)$$

- *Energy-per-bit* as function of $q = 1/R$ and r:

$$\mathcal{E}_b(q,r) = (N+I)r^{2\delta}q\left(2^{2/q} - 1\right) + qP_c\ ,$$

where P_c is circuit power dissipated at Tx and Rx

Energy-per-bit

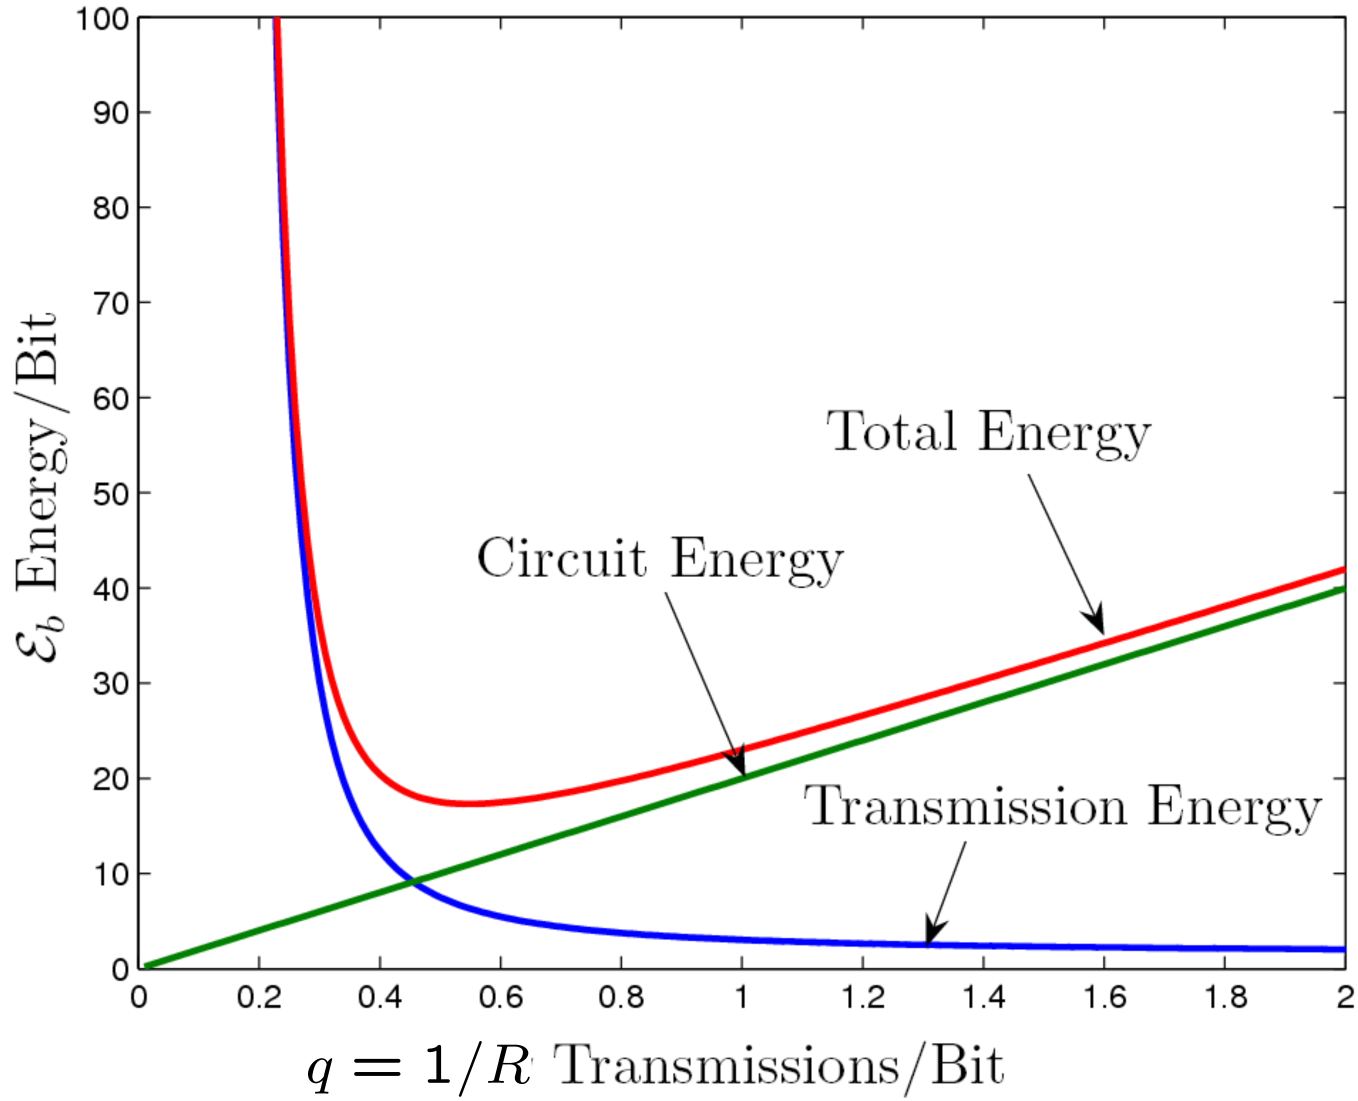

Example 1: To Hop or Not to Hop

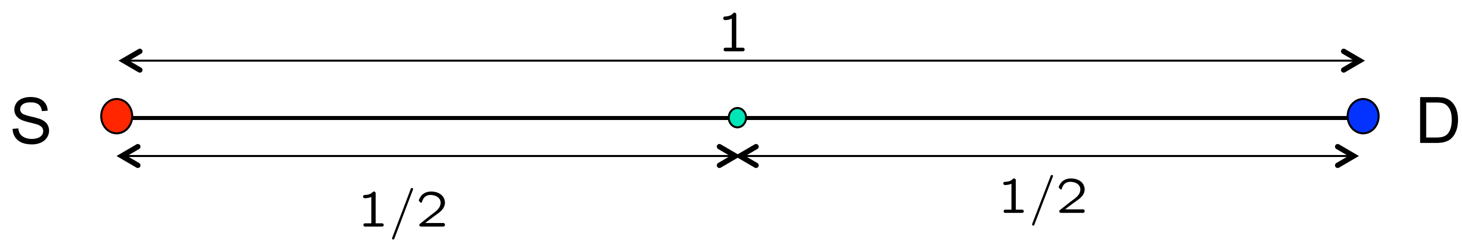

- Compare energy-per-bit for:
  - Direct transmission
  - 2 hops $(I \neq 0)$
- Conclusion: Hopping may or may not help
- How about multi-hopping ?

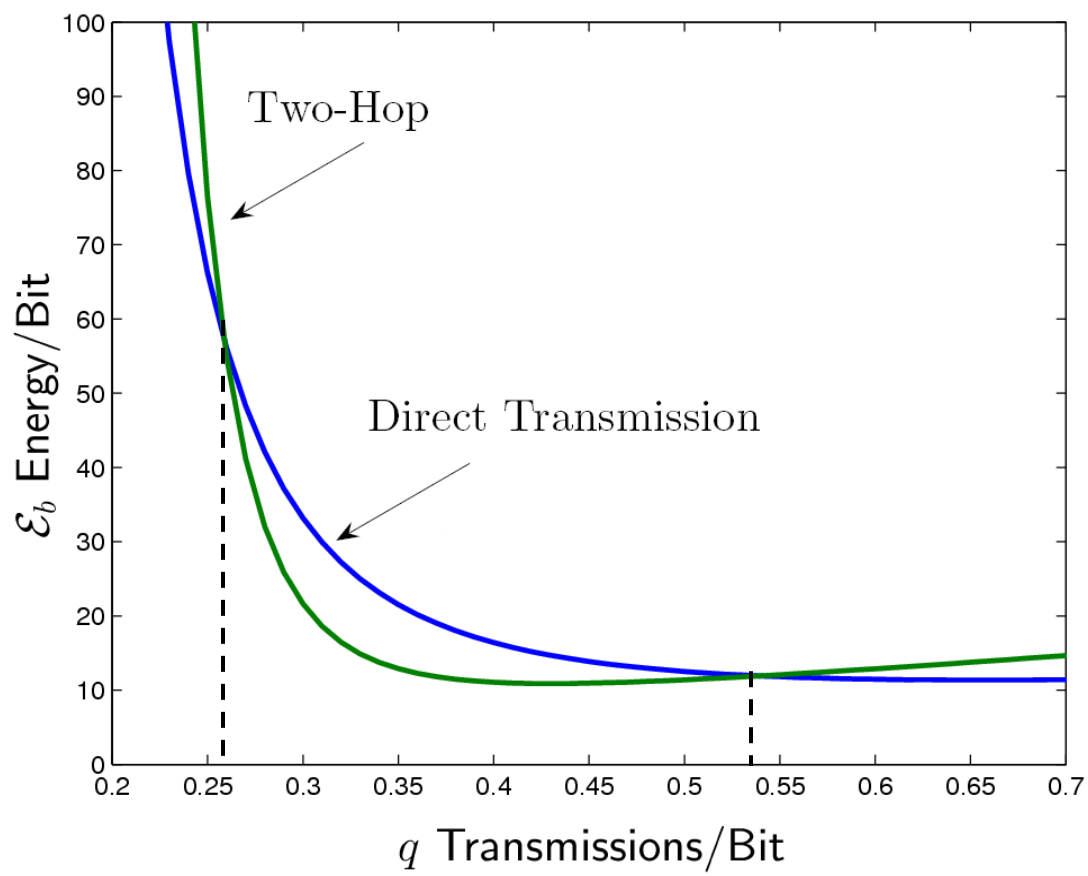

Example 2: Constant Distance

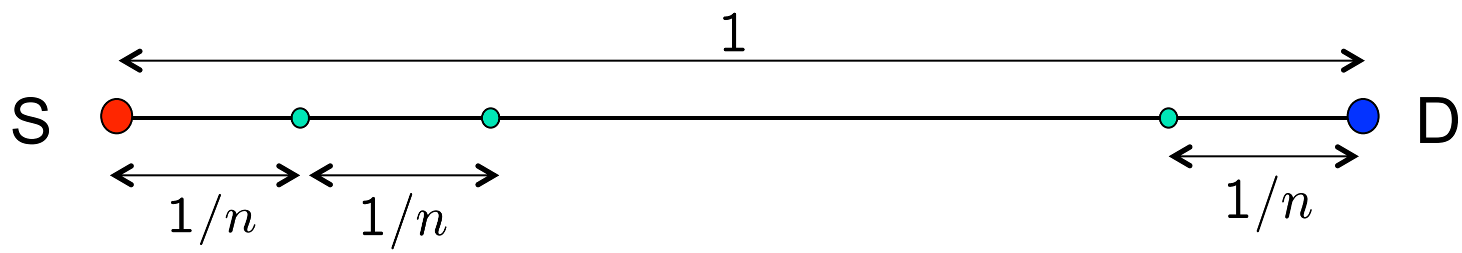

- Compare energy-per-bit for:
  - Direct transmission, $\mathcal{E}_0$ $(I = 0)$
  - n hops, $\mathcal{E}_n$ $(I \neq 0)$
- With n hops:

$$\mathcal{E}_n = \Theta\left(n\left(\left(\frac{1}{n}\right)^{2\delta} + 1\right)\right)$$

- Thus

$$\mathcal{E}_0/\mathcal{E}_n = \Theta(1/n)$$

- Suggests too much hopping is bad for energy

Example 3: Constant Density

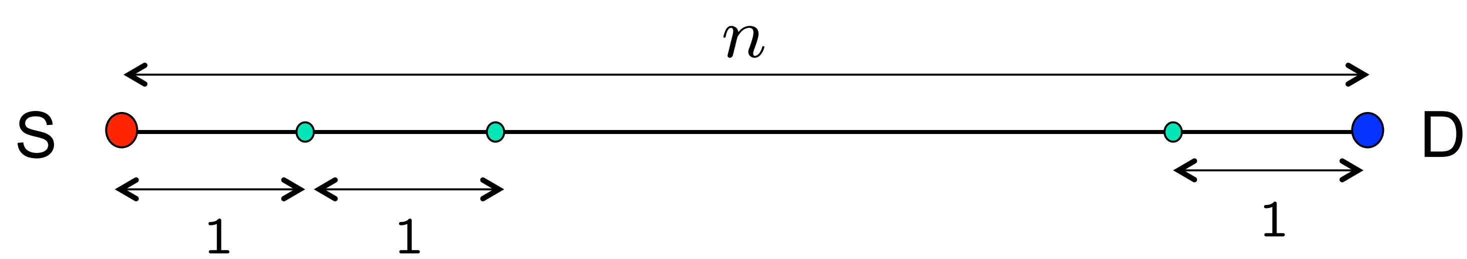

- Compare energy-per-bit for
  - Direct transmission, $\mathcal{E}_0 = \Theta(n^{2\delta})$
  - n hops, $\mathcal{E}_n = \Theta(n)$
- Result: $\mathcal{E}_0/\mathcal{E}_n = \Theta\left(n^{(2\delta-1)}\right)$
- Suggests a lot of hopping is good for energy
- Questions:
  - What is the conclusion for "real world" networks with arbitrary placements and many S-D pairs ?
  - What about corresponding throughput and delay ?

Energy Scaling in Networks

- Consider random network model [GK00]
- n nodes placed uniformly and at random in a unit square
- Nodes divided randomly into $n/2$ source-destination (S-D) pairs
- Each node can act as a relay to help other S-D pairs
- Discrete-time transmission model
- Protocol or physical model used to capture interference effects

Throughput and Delay [GK00,EMPS04]

- S-D pair throughput $T(n)$ scales as $\Theta\left(1/\sqrt{n \log n}\right)$
  - Achieved using cellular TDM scheme with cells of size $\Theta(\log n/n)$
  - *Highest* throughput at *largest* number of hops
- Varying cell-size results in the optimal throughput-delay trade-off: $D(n) = \Theta(nT(n))$

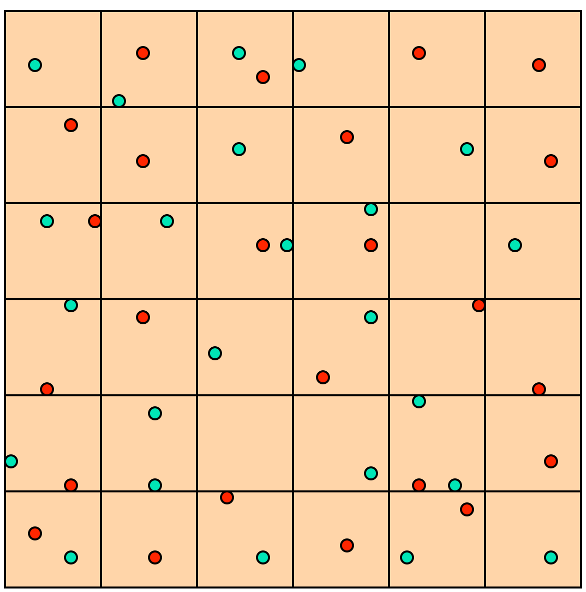

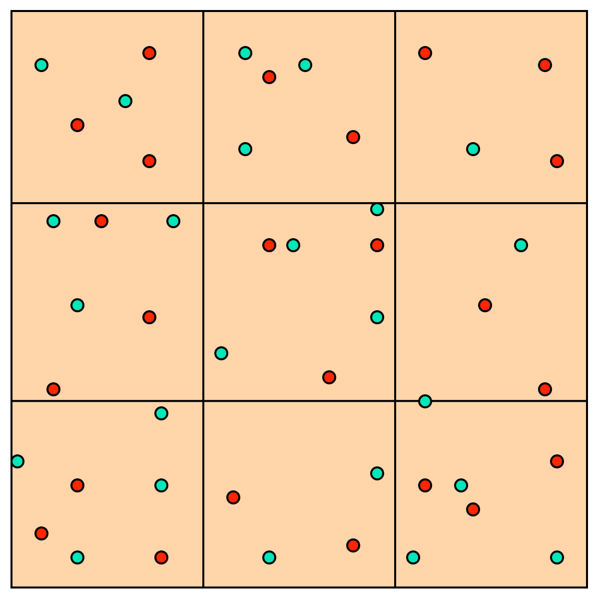

Energy Scaling

- Consider two random network models:
  - Constant area -- n nodes in unit square
  - Constant density -- n nodes in square of area n
- Use the energy-based rate model, where

$$\mathcal{E}_b(q,r) = (N+I)r^{2\delta} q \left(2^{2/q} - 1\right) + qP_c$$

- Under this energy-based model:
  - Throughput and delay scaling results are essentially unchanged for both network models
  - But energy-per-bit will depend on the network model

Energy for Constant Area Model

- Energy-per-bit $E(n)$ is determined by:
  - The average hop length $r(n)$ and
  - The number of hops $\Theta(1/r(n))$
- Total circuit energy increases linearly with the number of hops as $\Theta(1/r(n))$
- Total transmission energy scales down by hopping as $\Theta\left(r(n)^{2\delta-1}\right)$
- Result: Circuit energy dominates over transmission energy
  - Hopping is bad for energy

Energy-Constrained Throughput

- Many hops $\Rightarrow$ high $T(n)$, but also high $E(n), D(n)$
- Optimal trade-off: $T(n) = \Theta(E(n)/n)$

$$D(n) = \Theta(E(n))$$

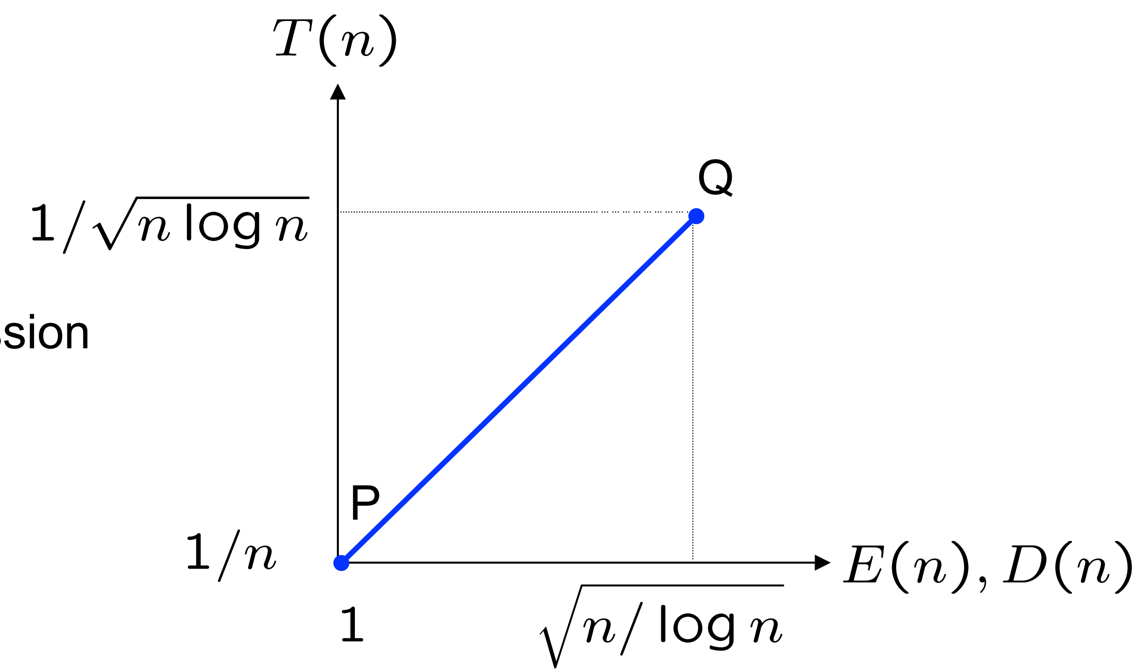

Point P: Direct transmission

Point Q: Max hopping

Energy Scaling for Constant Density

- Transmission energy dominates over circuit energy
- Smallest cells of area $\Theta(\log n)$ result in:
  - *Highest* throughput: $T(n) = \Theta\left(1/\sqrt{n}(\log n)^{\delta+1/2}\right)$
  - *Lowest* energy-per-bit: $E(n) = \Theta\left(\sqrt{n}(\log n)^{\delta+1/2}\right)$
- But at expense of *high* delay: $D(n) = \Theta\left(\sqrt{n/\log n}\right)$
- Without delay constraint, *any* lower throughput scaling can be achieved at minimum energy
- With delay constraint: Lower throughput and higher energy-per-bit

What do Scaling Laws Tell us?

- Hopping is good for throughput, bad for delay
- Hopping may or may not be good for energy:
  - Reduces RF transmission energy, but
  - Increases circuit energy
- **Constant area:** Transceiver circuit energy dominates in order; need hopping only to achieve good throughput
- **Constant density:** RF transmission energy dominates; hopping achieves maximum throughput at minimum energy/bit
- Conclusion: Some hopping is needed
- Details in an upcoming paper …

Outline

- Examples where we have some theory:
  - Does multi-hopping save energy?
  - **Joint backlog and channel adaptation saves energy**
  - Multi-user coding saves energy
- Example where we can save a lot of energy using local processing, but there is no theory:
  - Collaborative task-driven wireless imaging network
- General comments

Joint Adaptation

- Information theory tells us how to optimize throughput by adapting *rate* to *channel state* using *water-filling*
  - But assumes that data has *constant* rate
- In many applications (e.g., imaging sensors, web session) data rate is unknown and varies over a very wide range
  - Assuming highest rate and adapting only to channel can be very energy-inefficient
- Theory tells us that:
  - Adapting to data-rate (backlog) by itself can save a lot of energy
  - Jointly adapting to data rate and channel can save even more energy

Key Observation

- Transmission energy/bit is *monotonically decreasing* and *convex* in transmission time

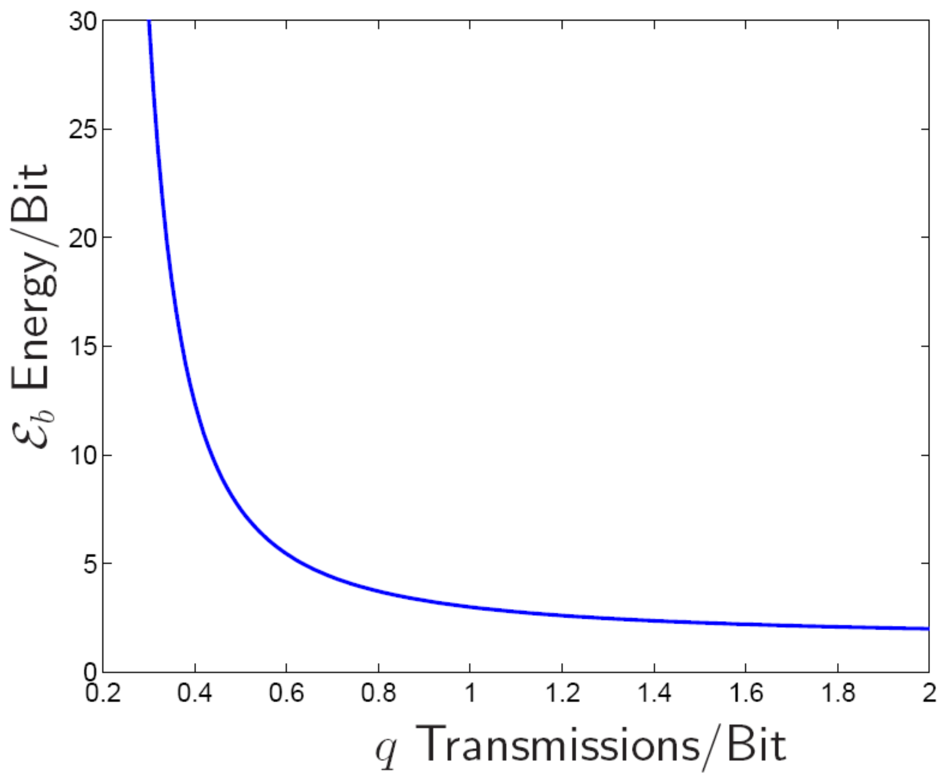

- To save energy transmit as slowly as you can, but no slower

Work on Joint Adaptation

- Berry and Gallager, "Communication Over Fading Channels with Delay Constraints," *IEEE Trans. on IT*, 2002.
- Berry, "Optimal Power-Delay Trade-offs in Fading Channels – Small Delay Asymptotics," Submitted.

  IID arrivals; Minimize energy subject to average delay constraint; Determined power-delay trade-off for asymptotically large or small delay

- Ata, "Dynamic Control of Stochastic Networks", Ph.D. Thesis, Graduate School of Business, Stanford University, 2003

  Poisson arrivals, finite buffer, exponential transmission time, rate is function of backlog; Minimize energy subject to drop rate constraint; Found optimal policy

- Prabhakar, Uysal-Biyikoglu and El Gamal, "Energy-efficient Transmission over a Wireless Link via Lazy Packet Scheduling," *IEEE Infocom*, 2001.
- Uysal-Biyikoglu and El Gamal, "On Adaptive Transmission for Energy Efficiency in Wireless Data Networks," *IEEE Trans. on IT*, 2004

Energy-Efficient Scheduling [PUE01]

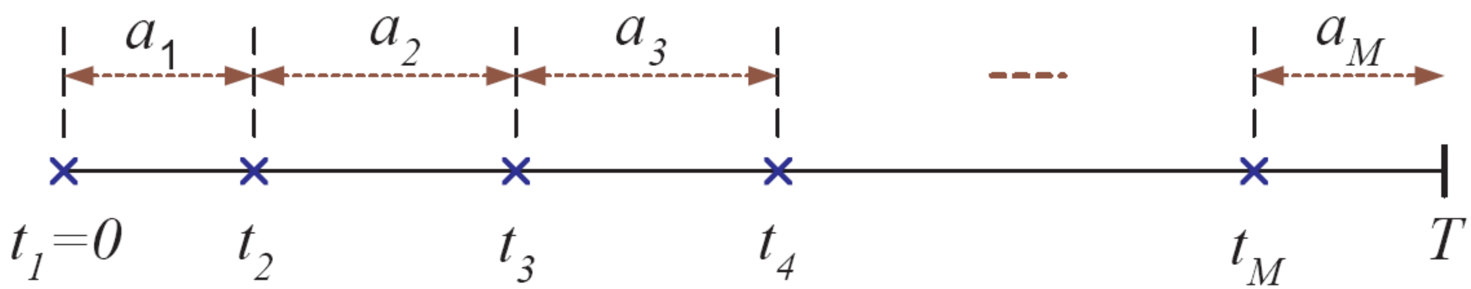

- Determine *offline* schedule that minimizes total transmission energy subject to deadline constraint T
- A convex optimization problem; has an explicit solution
- Lazy scheduling:
  - Use long transmission time when backlog is small
  - Use short transmission time when backlog is large
- Offline scheduling for uplink and downlink scenarios are also convex
- Results can be extended to total communication energy

Look-Ahead Online Heuristic

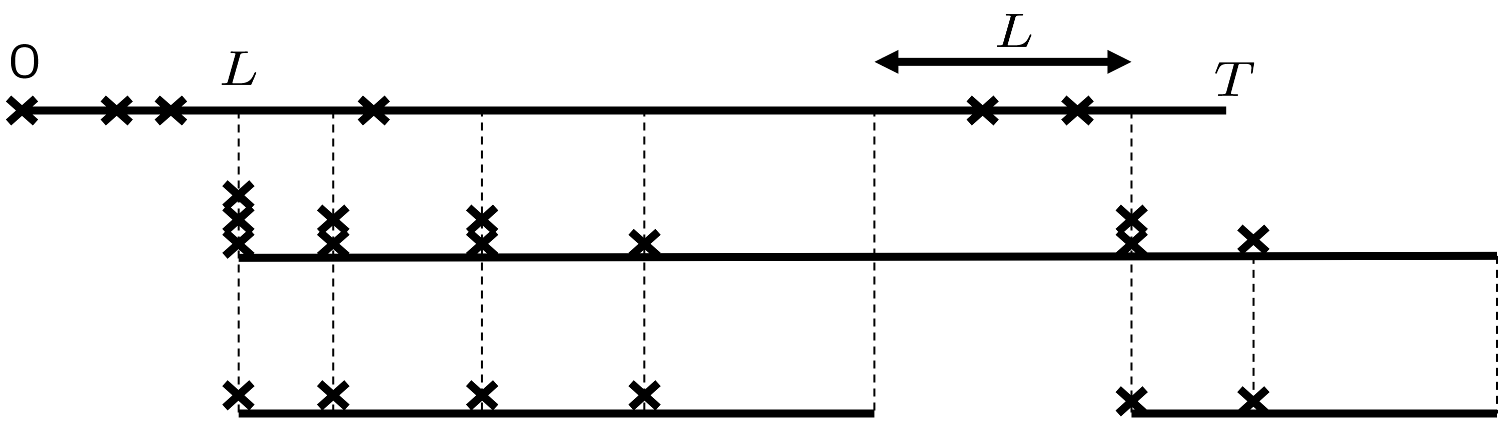

- Transmission time for first packet is L
- Transmit time for second packet is L/backlog at time L, etc.
- Performs close to optimal with small increase in delay

Joint Scheduling [UE04]

- Assume an AWGN block fading channel

$$Y = \sqrt{S}X + Z$$

- Assume both transmitter and receiver know the state of the channel before transmission
- Offline schedule is again a convex optimization problem
- Online heuristic: *Look-ahead water-filling*
  - Adapt to channel state by performing water-filling in each scheduling window of look-ahead algorithm

Joint Adaptation Example [UE04]

Packet size: 10^3 bits; Max rate: 1 packet/unit time; Data arriving at 0.2 packet/unit time

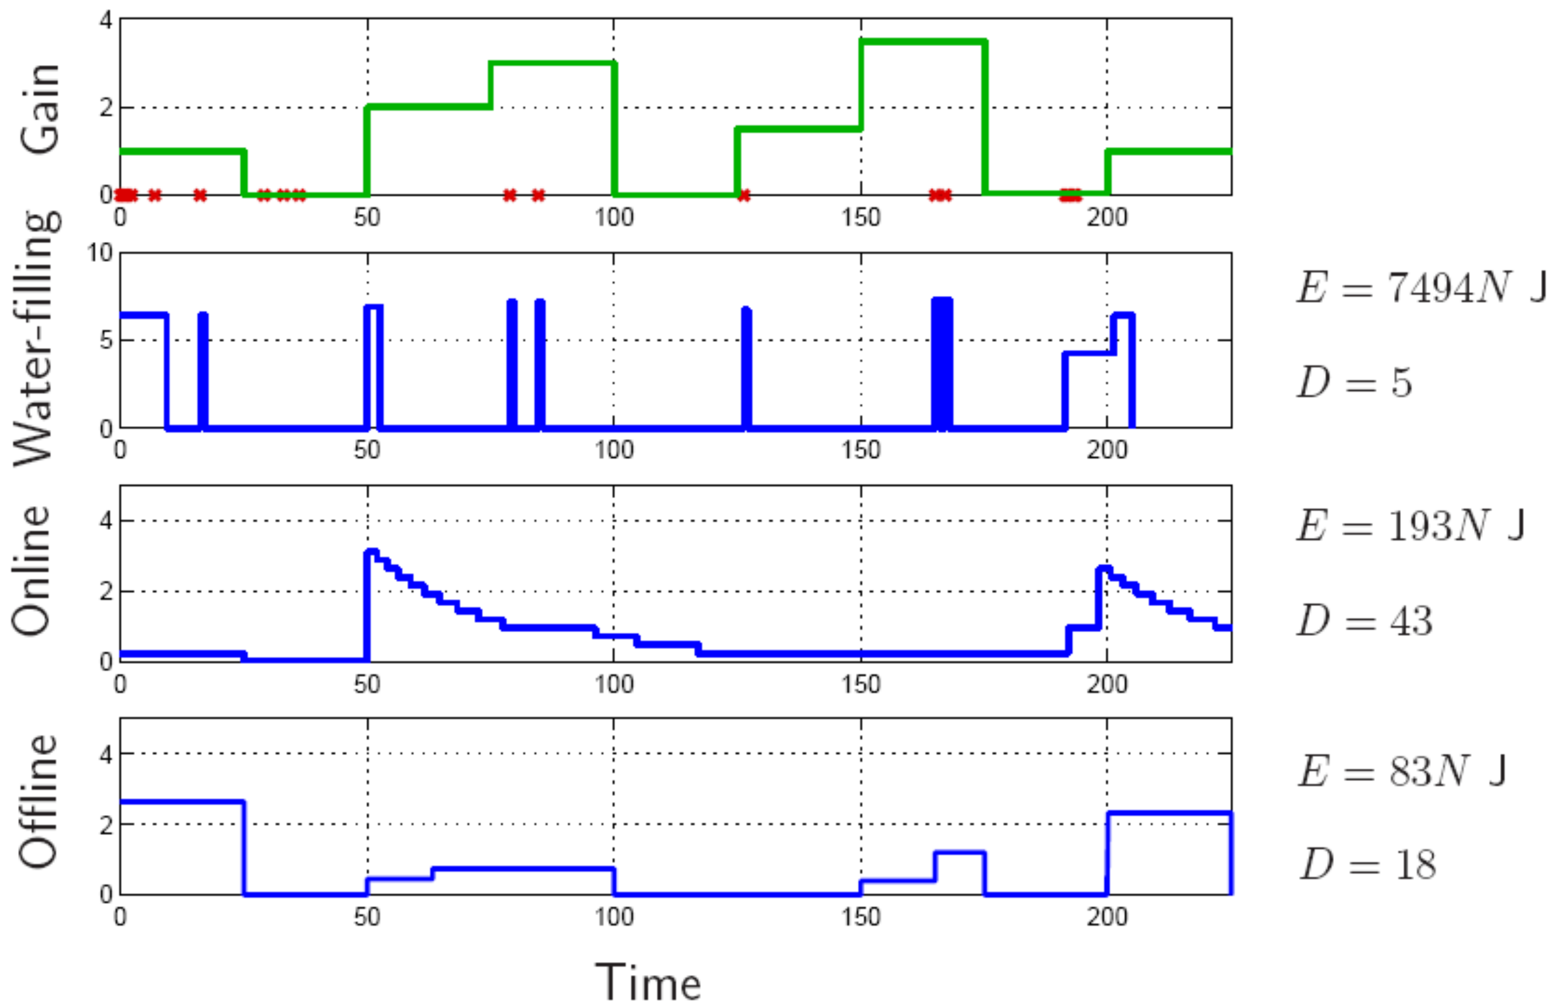

Outline

- Examples where we have some theory:
  - Does multi-hopping save energy?
  - Joint backlog and channel adaptation saves energy
  - **Multi-user coding saves energy**
- Example where we can save a lot of energy using local processing, but there is no theory:
  - Collaborative task-driven wireless imaging network
- General comments

Multi-Access Coding Saves Energy

- Consider energy-efficient scheduling problem over uplink channel [UE04]
- Only maximum data rate is known
- Compare schedules for:
  - **Time division**: Packet transmission from different users do not overlap
  - **Multi-access coding**: Packet transmissions can overlap; successive cancellation decoding used
- Optimal offline scheduling is easy to find
- Use look-ahead online scheduling

Multi-access vs. Time-division [UE04]

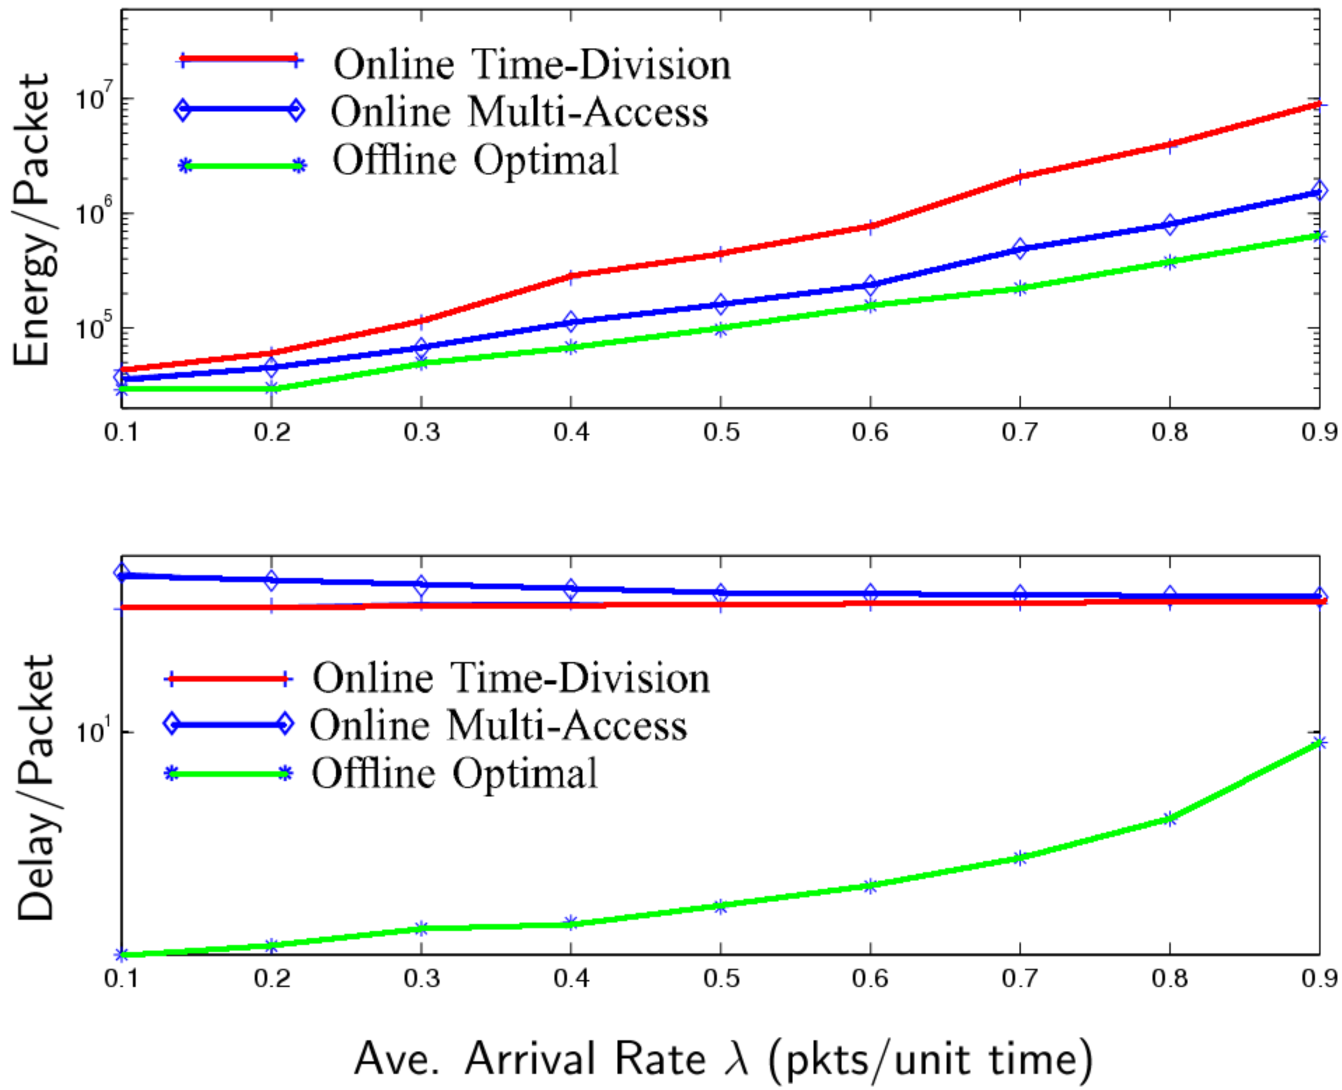

Relaying Can Save Energy [EZ03]

- Consider AWGN Relay channel

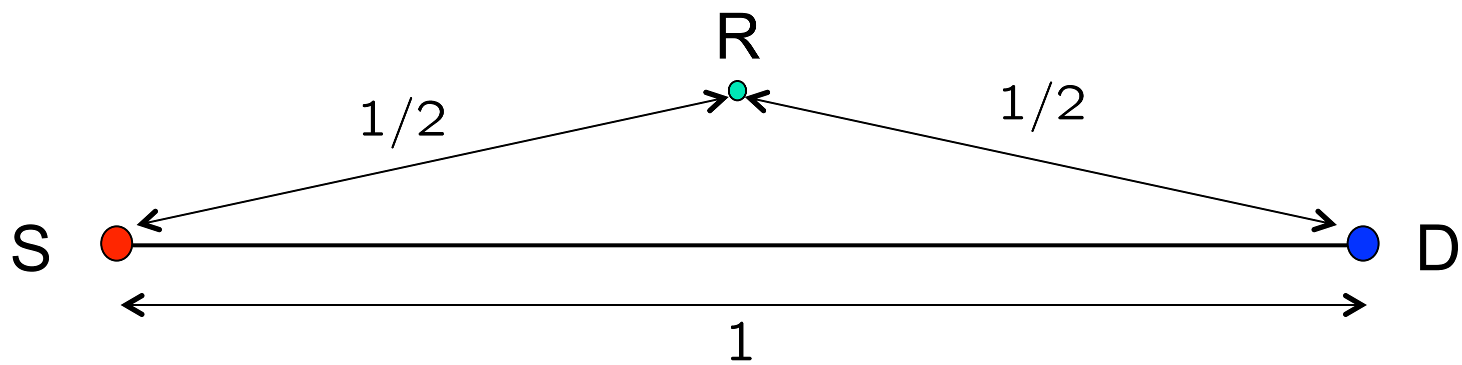

- Two-hop scheme
  - Packets are decoded by relay and retransmitted to final destination
  - Receiver treats direct transmission as noise
- Information theoretic relaying scheme [CE81]
  - Block-Markov encoding is used

Relaying Saves Transmission Energy

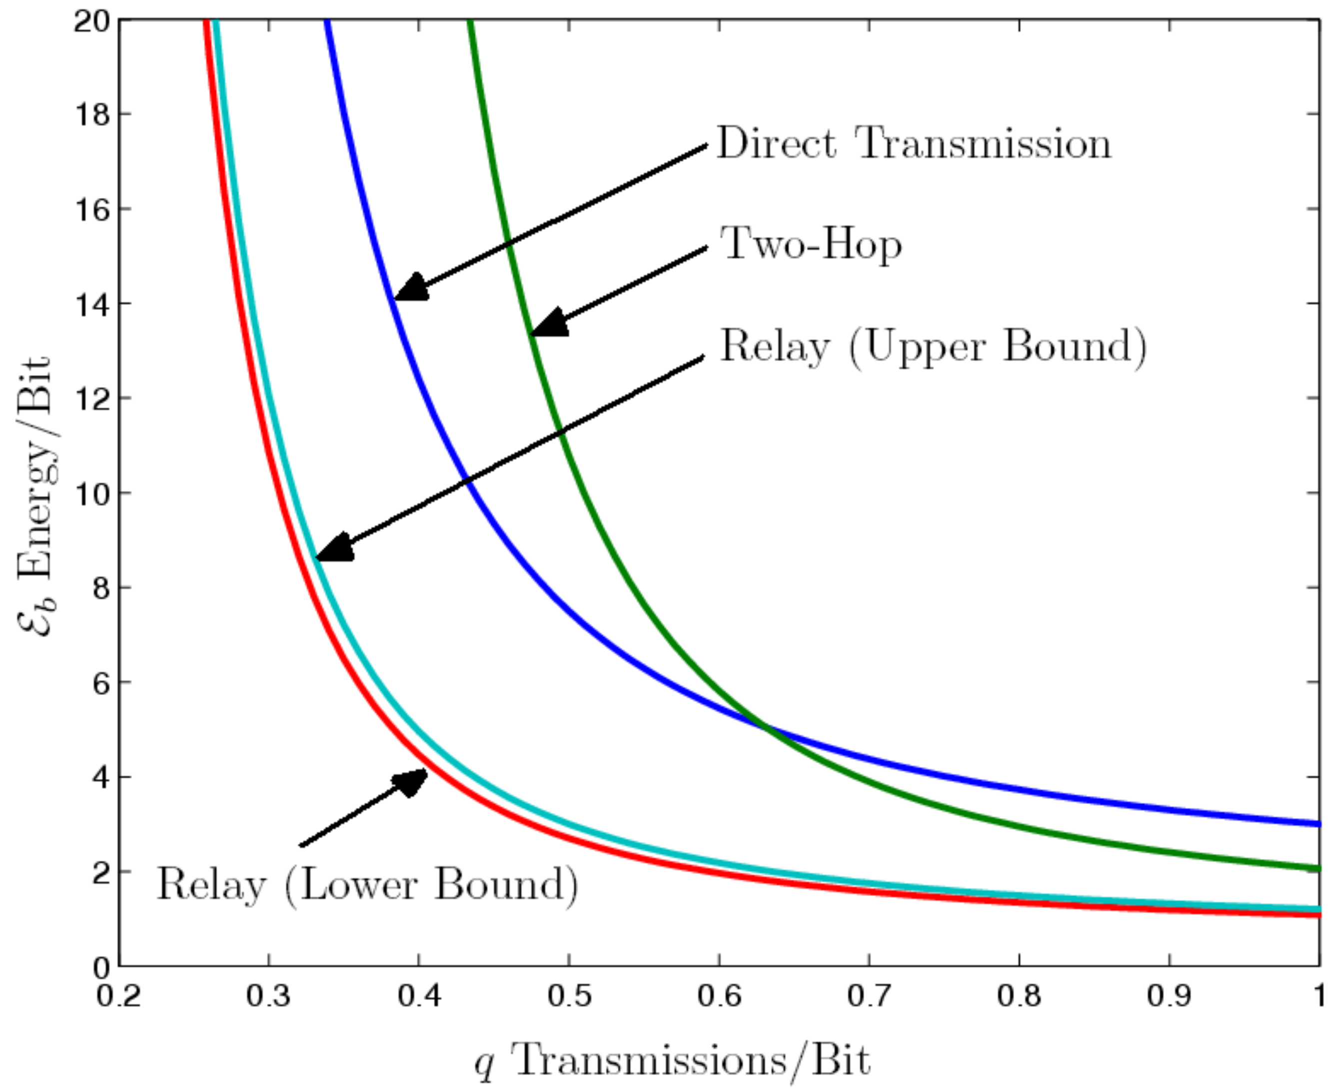

Outline

- Examples where we have some theory:
  - Does multi-hopping save energy?
  - Joint backlog and channel adaptation saves energy
  - Multi-user coding saves energy
- **Example where we can save a lot of energy using local processing, but there is no theory:**
  - Collaborative task-driven wireless imaging network
- General comments

Task: Occupancy Reasoning*

- Determine areas occupied by objects (e.g., people) in a room
- Occupancy information is valuable for many higher level tasks:
  - Monitoring, tracking, counting people, …
- Crowd causes multiple *occlusions* – no single camera by itself can perform task well
- Use a network of cameras to reason about occupancy [YGG03]:
  - Simple local processing
  - Collaborative computing

*Joint project with L. Guibas

Experimental Setup

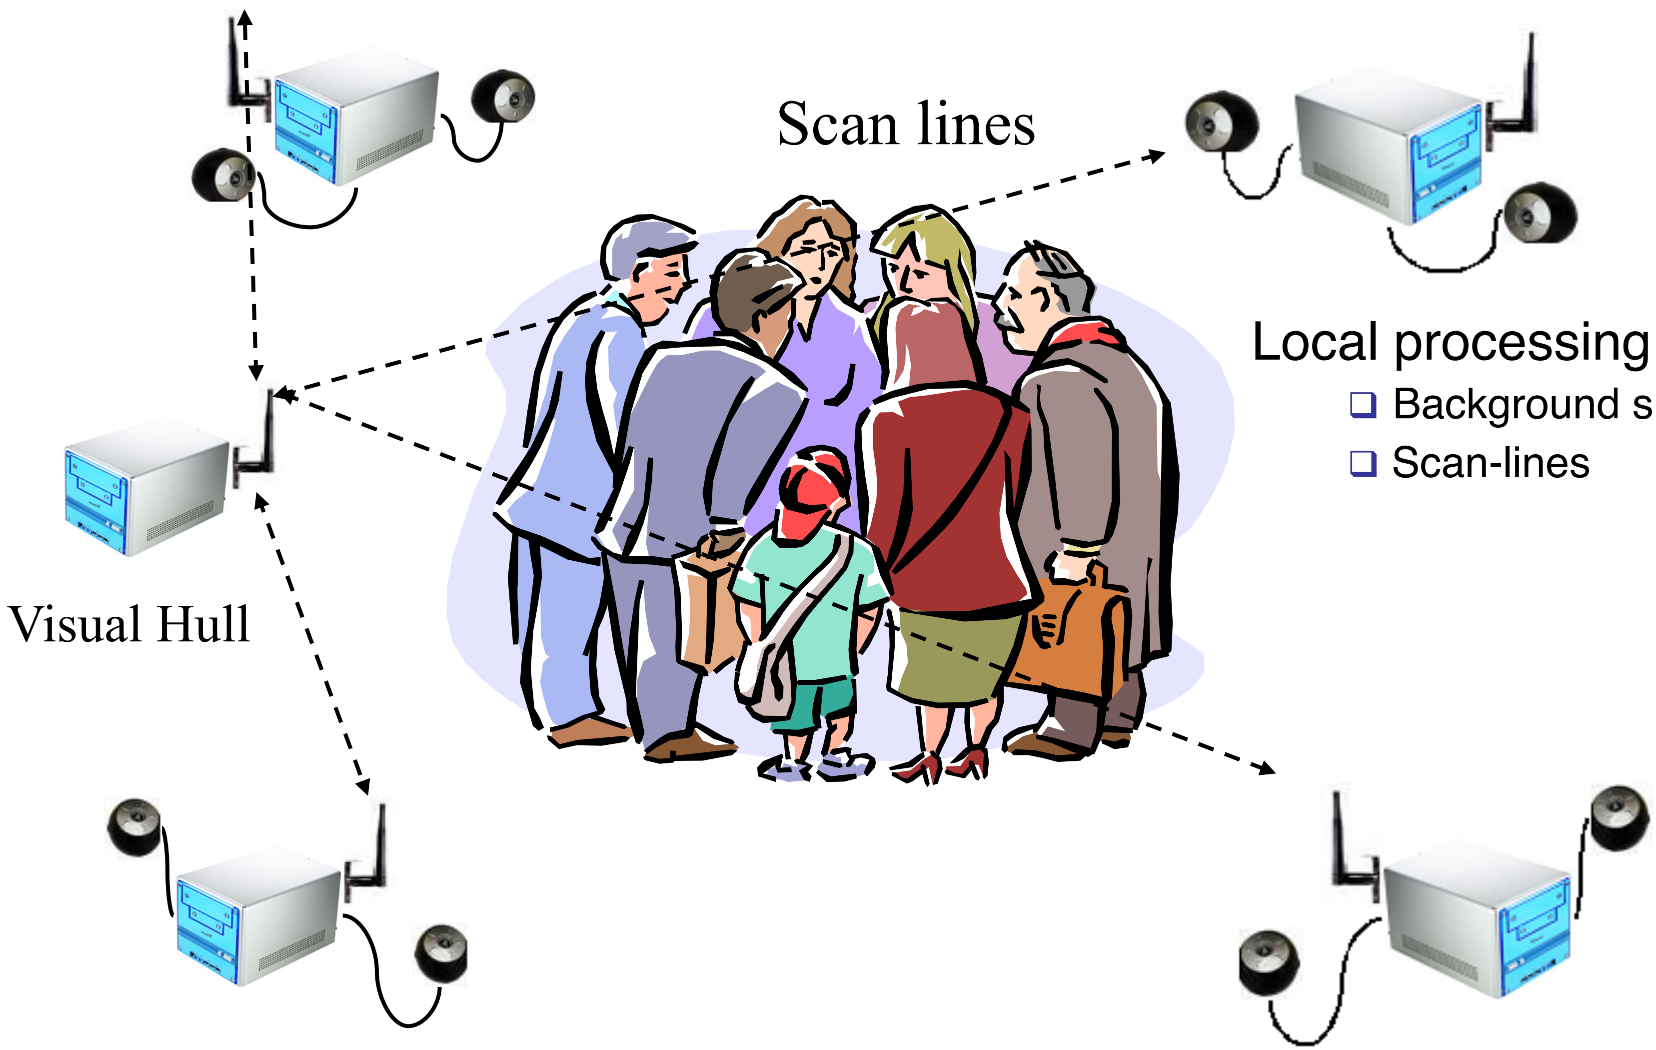

Local processing:

- Background subtract
- Scan-lines

Local Processing

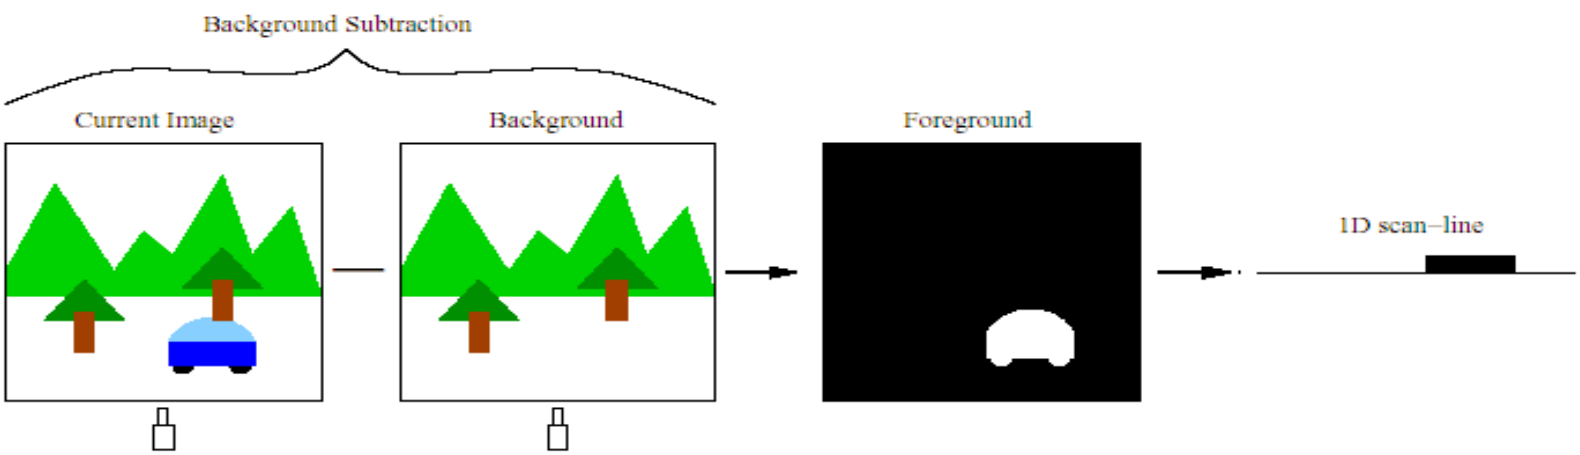

- Background subtraction
- Collapse to a scan-line
- 640x480 RGB image → 640 bit scan-line (can be compressed further)

Visual Hull

Combine scan-lines to generate visual hull – an *upper bound* on occupancy

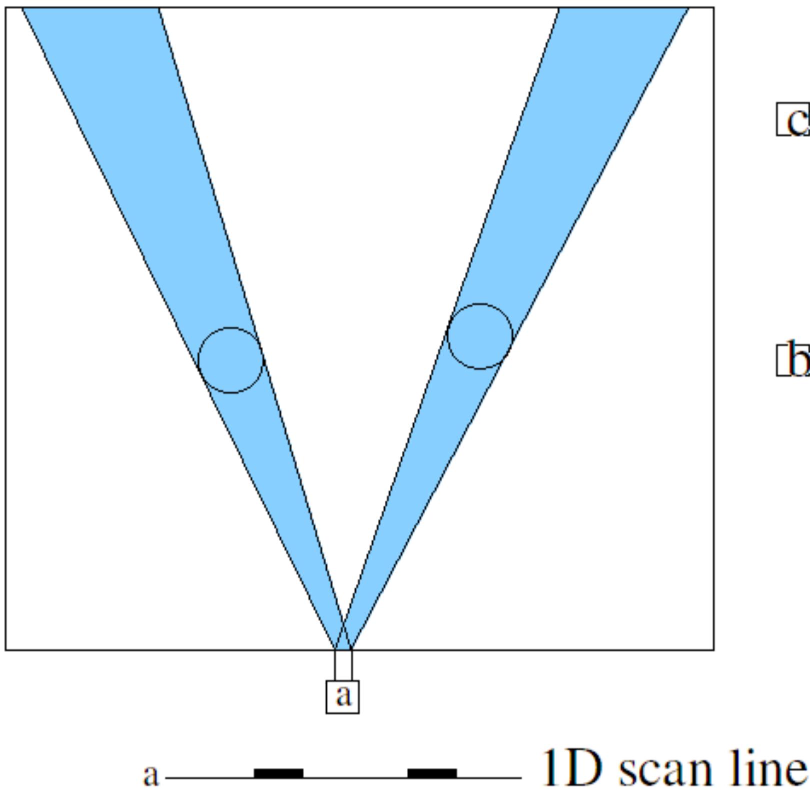

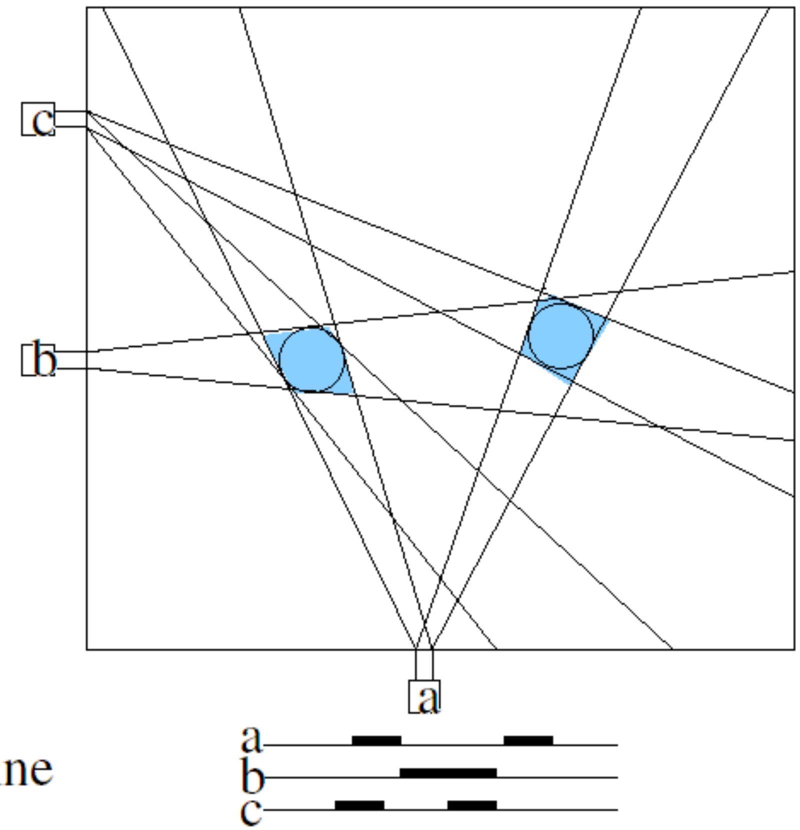

Demo

View of the setup

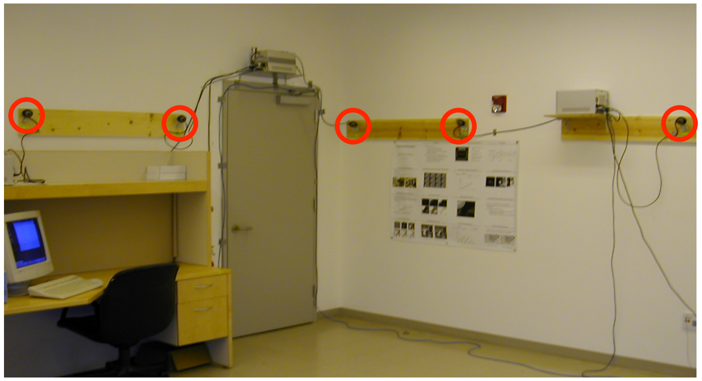

View from a camera

Top view of room

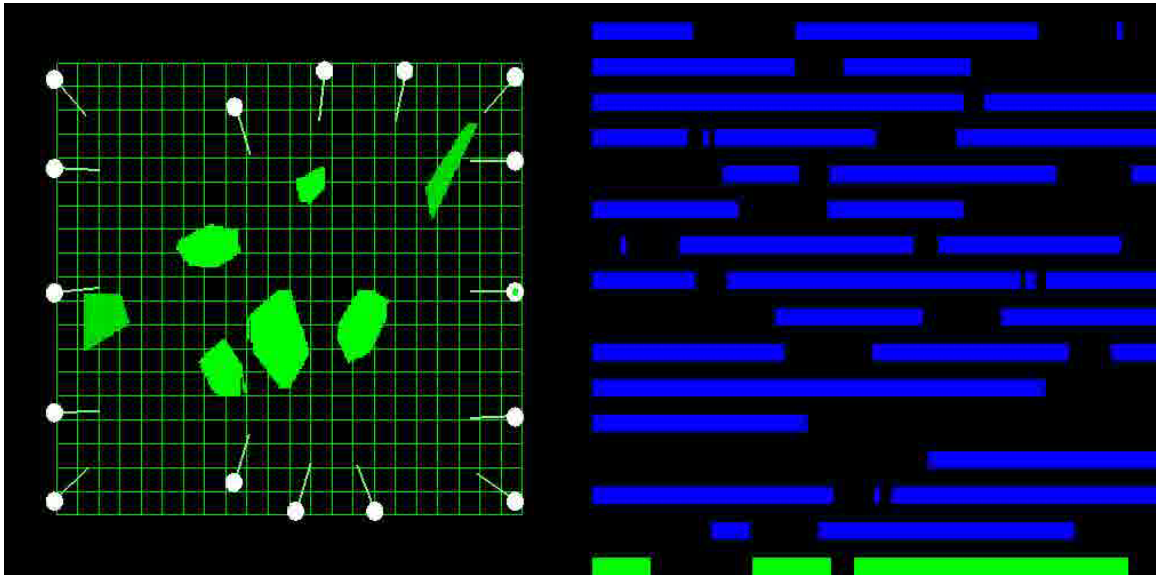

Scan-lines from 16 cameras

Tasking Best Subset of Cameras

- Having cameras operate all the time wastes energy
- Sensor tasking:
  - Choose best subset of nodes for given occupancy

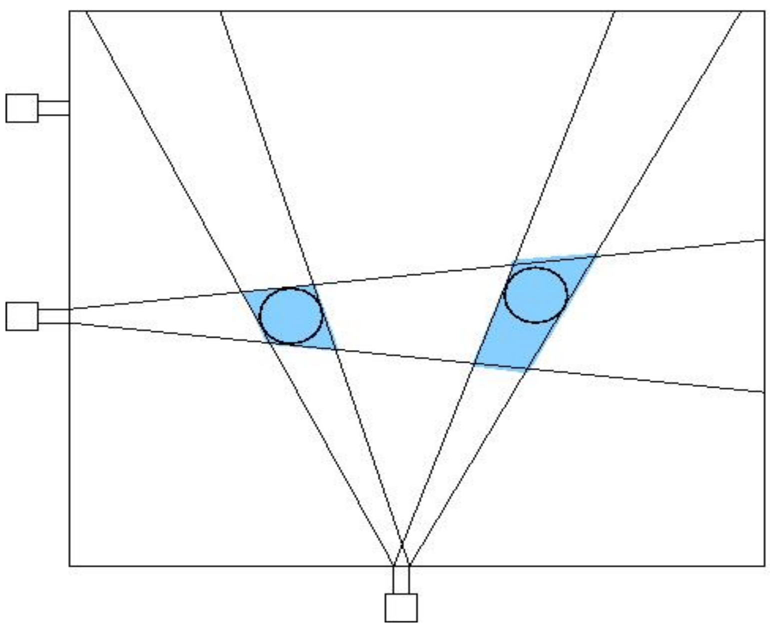

Good

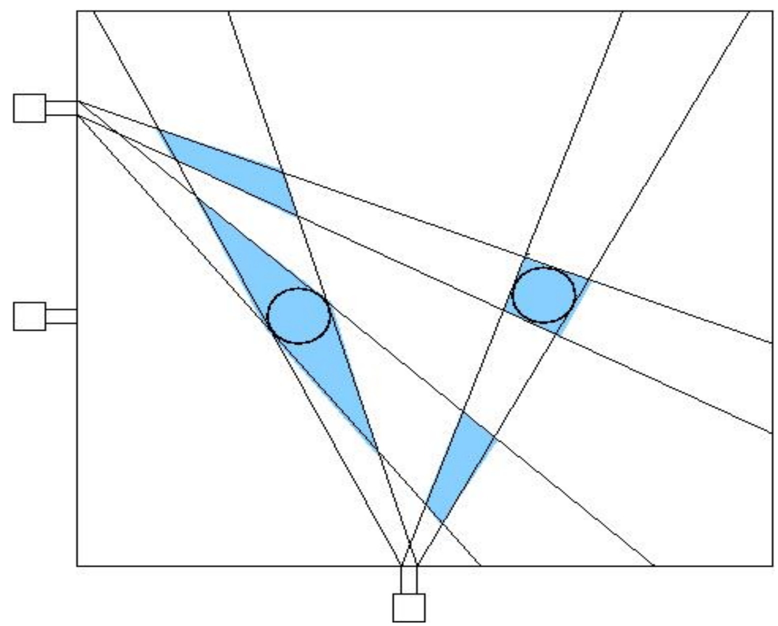

Bad

Camera Tasking Strategies

- **Uniform:** optimal for randomly placed objects
- **Greedy:** optimally select one camera at a time
- **Incremental greedy:** change one camera at a time:
  - Turning too many cameras on and off costs energy
- Results:
  - Greedy close to optimal
  - Incremental greedy almost as efficient as greedy

Uniform vs. Incremental Greedy

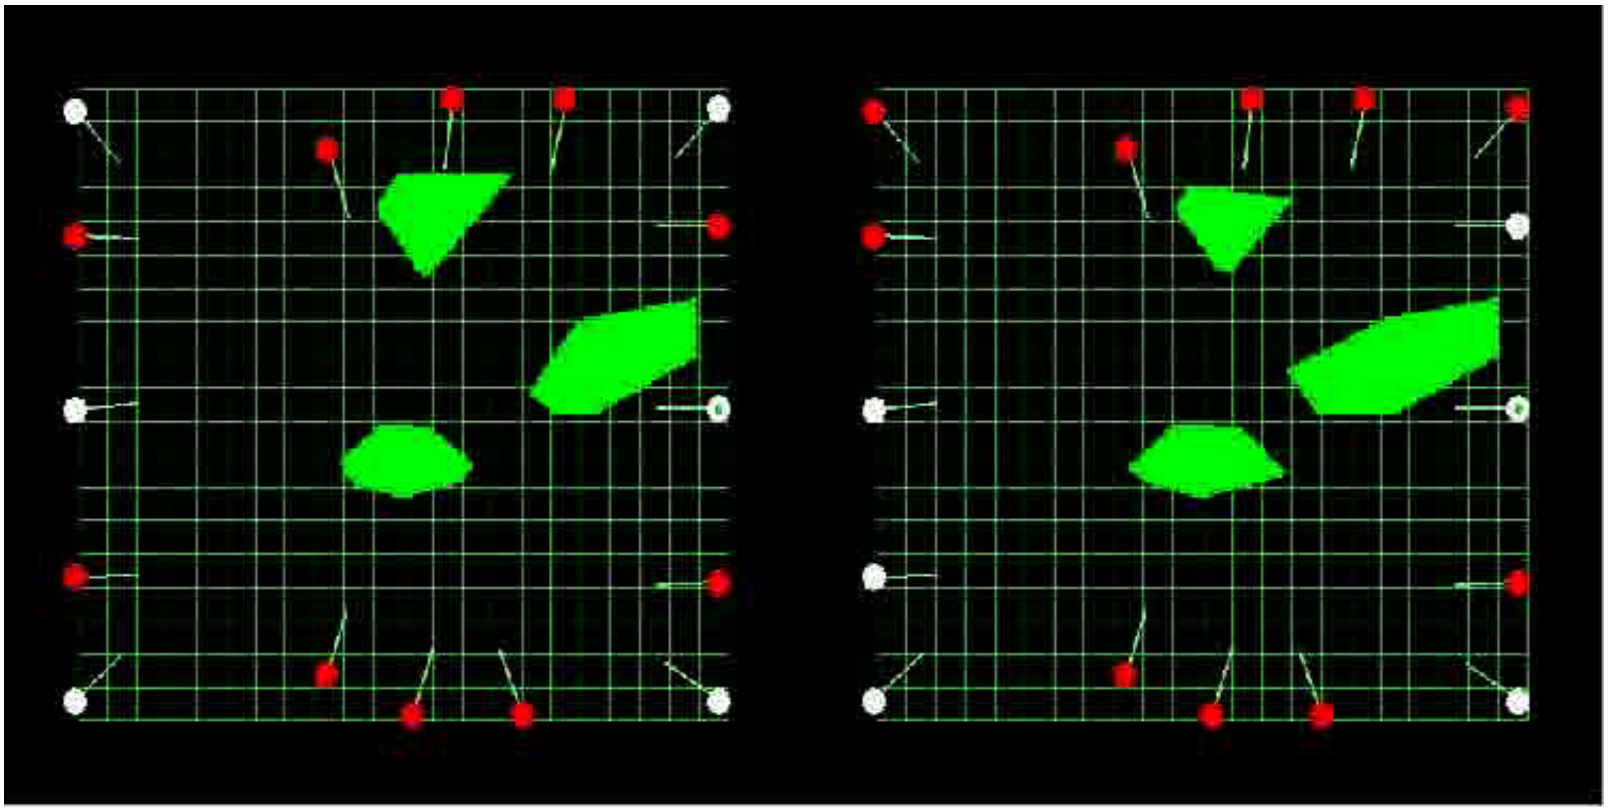

No Tasking　　　　With Tasking

Application: Tracking

- Example: Track 6 people moving in a room; one person track is shown

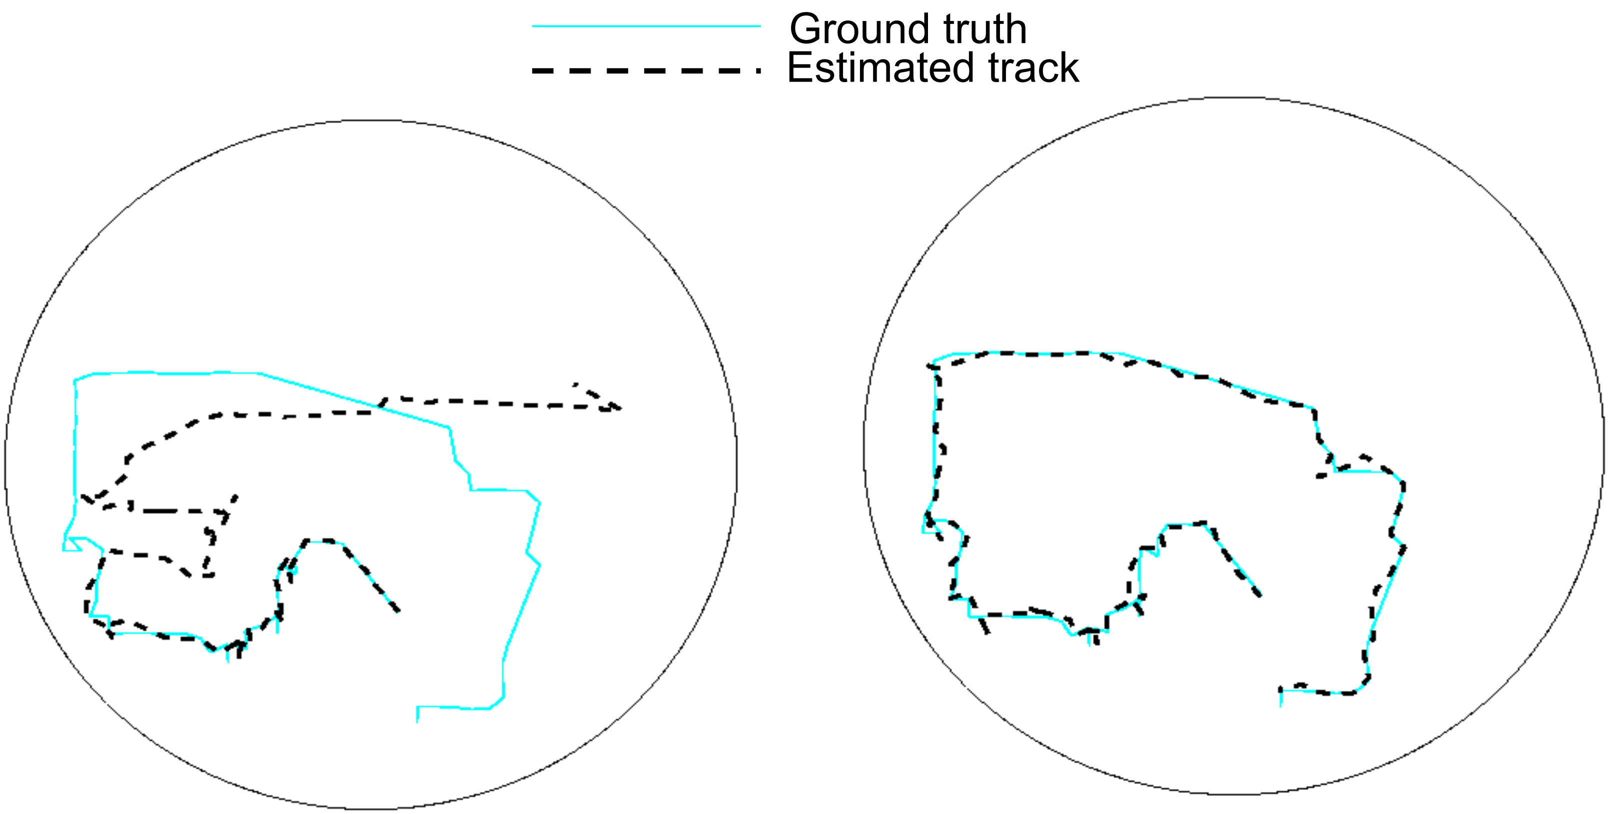

6 cameras: uniform placement   Tasking 5 cameras with inc. greedy

Summary

- Three examples where we have theory to help us save communication energy
  - Multi-hopping
  - Joint adaptation
  - Multi-user coding
- Can save a lot more communication energy by
  - Collaborative processing; simple local processing
  - Sensor tasking
  - But, there is no general theory to guide us

General Comments

- Energy is primary concern in future wireless networks
- We have some examples of energy-centric theory
- A lot more theory needs to be developed:
  - Many basic problems are still open, e.g., how to optimally deal with interference, relaying, etc.
  - We don't know how to incorporate processing energy in theory
  - Need theory to guide collaborative task-driven processing
  - Theory also needs to incorporate correlated sources
- A more comprehensive energy-centric theory will need
  - New insights
  - New models
  - New tools

But more importantly, the smartest and best trained people involved …