Algorithms for Sensor Networks

GRAAL/AEOLUS School on Hot Topics in Network Algorithms

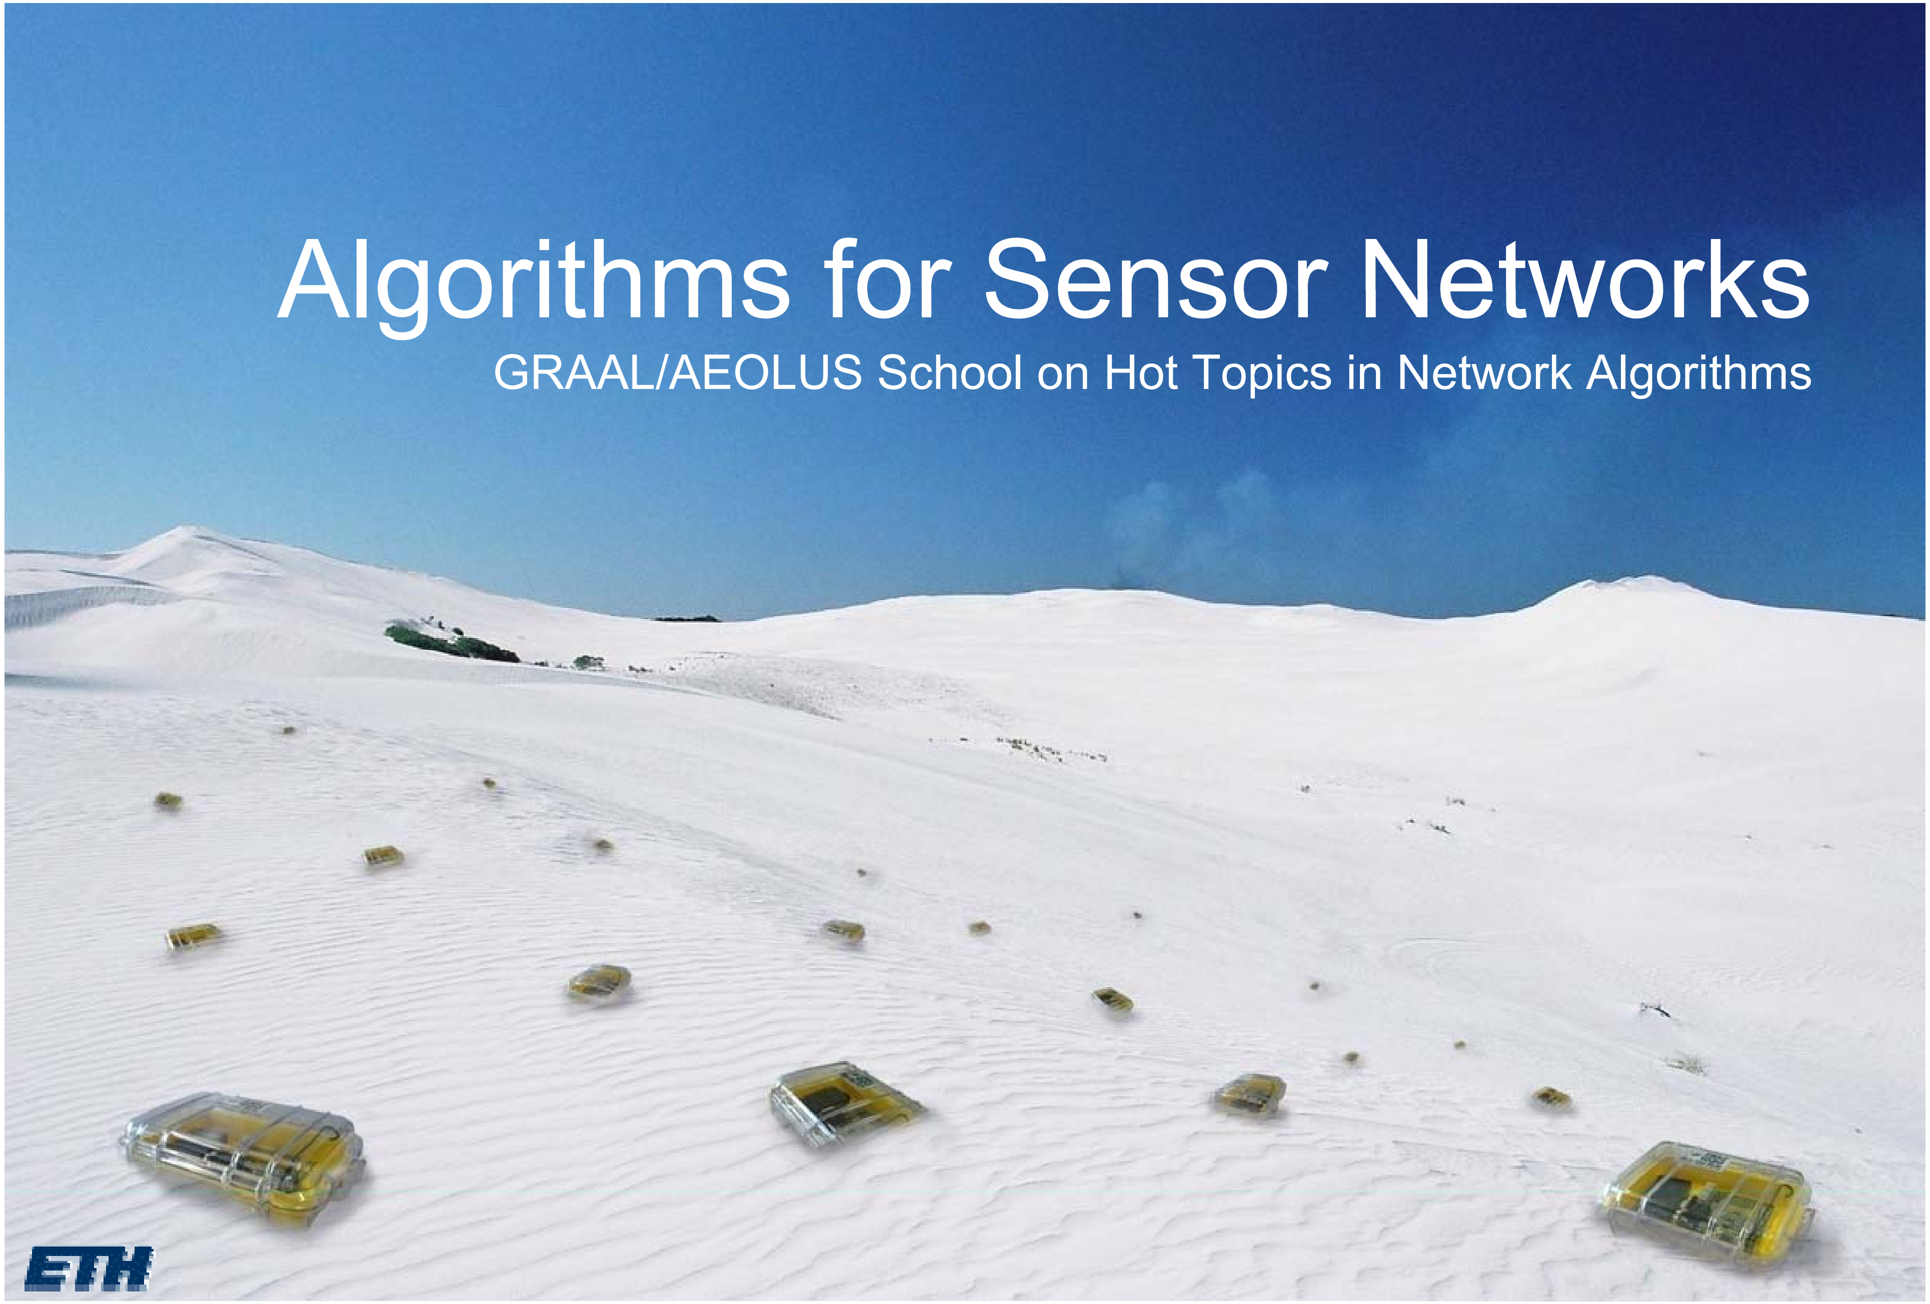

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Goals

- What do **YOU** want to learn?
  - How much do you know already?

- Problem: Huge area
  - with hundreds of workshops, literally!
  - At ETH Zurich, I teach a 28h course on this topic

- What I can (hopefully) offer
  - Learn some of the basic models and ideas
  - Learn some cool algorithms and techniques

- But mostly
  - Try to figure out what is really hot (research ideas)
  - Hybrid of really short lecture and really long marketing talk

Some topics

- **Application**: 1 Applications, 7 Positioning, 8 Time Synchronization
- **Transport**: 12 Transport
- **Network**: 5 Mobility, 3 Geo-Routing, 11 Routing, 6 Data Gathering
- **Link**: 10 MAC, 4 Topology Control, 7 Clustering
- **Physical**: 2 Basics, 2 Models, 13 Capacity

Literature

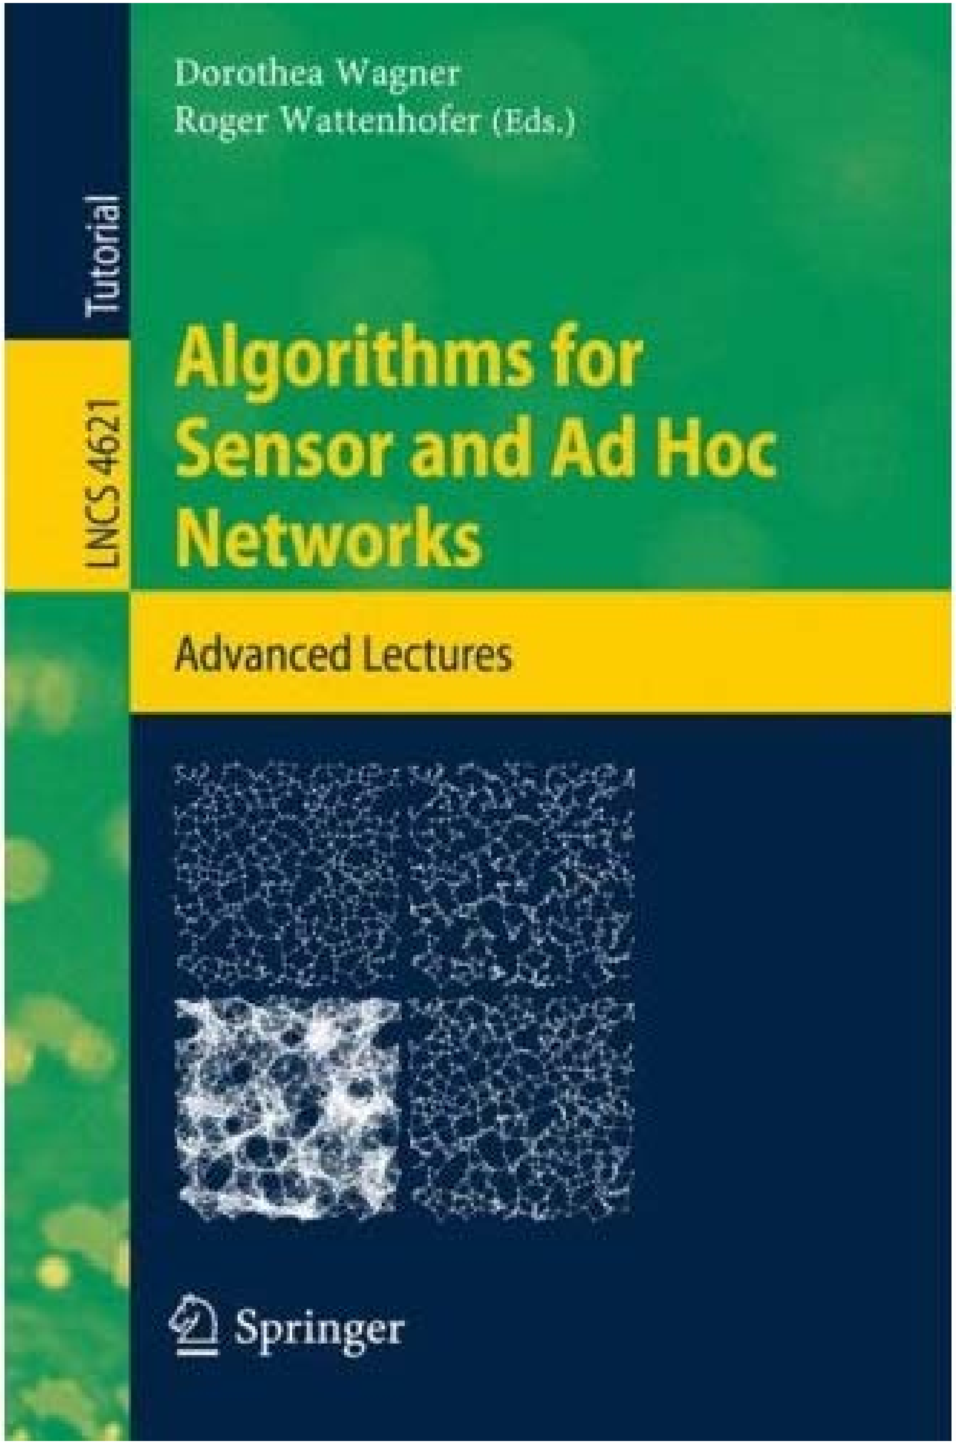
Dorothea Wagner
Roger Wattenhofer (Eds.)
Tutorial
LNCS 4621
Algorithms for Sensor and Ad Hoc Networks
Advanced Lectures
Springer

More Literature

- Bhaskar Krishnamachari – *Networking Wireless Sensors*
- Paolo Santi – *Topology Control in Wireless Ad Hoc and Sensor Networks*
- F. Zhao and L. Guibas – *Wireless Sensor Networks: An Information Processing Approach*
- Ivan Stojmeniovic – *Handbook of Wireless Networks and Mobile Computing*
- C. Siva Murthy and B. S. Manoj – *Ad Hoc Wireless Networks*
- Jochen Schiller – *Mobile Communications*
- Charles E. Perkins – *Ad-hoc Networking*
- Andrew Tanenbaum – *Computer Networks*

- *Plus tons of other books/articles*
- *Papers, papers, papers, …*

Introduction

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

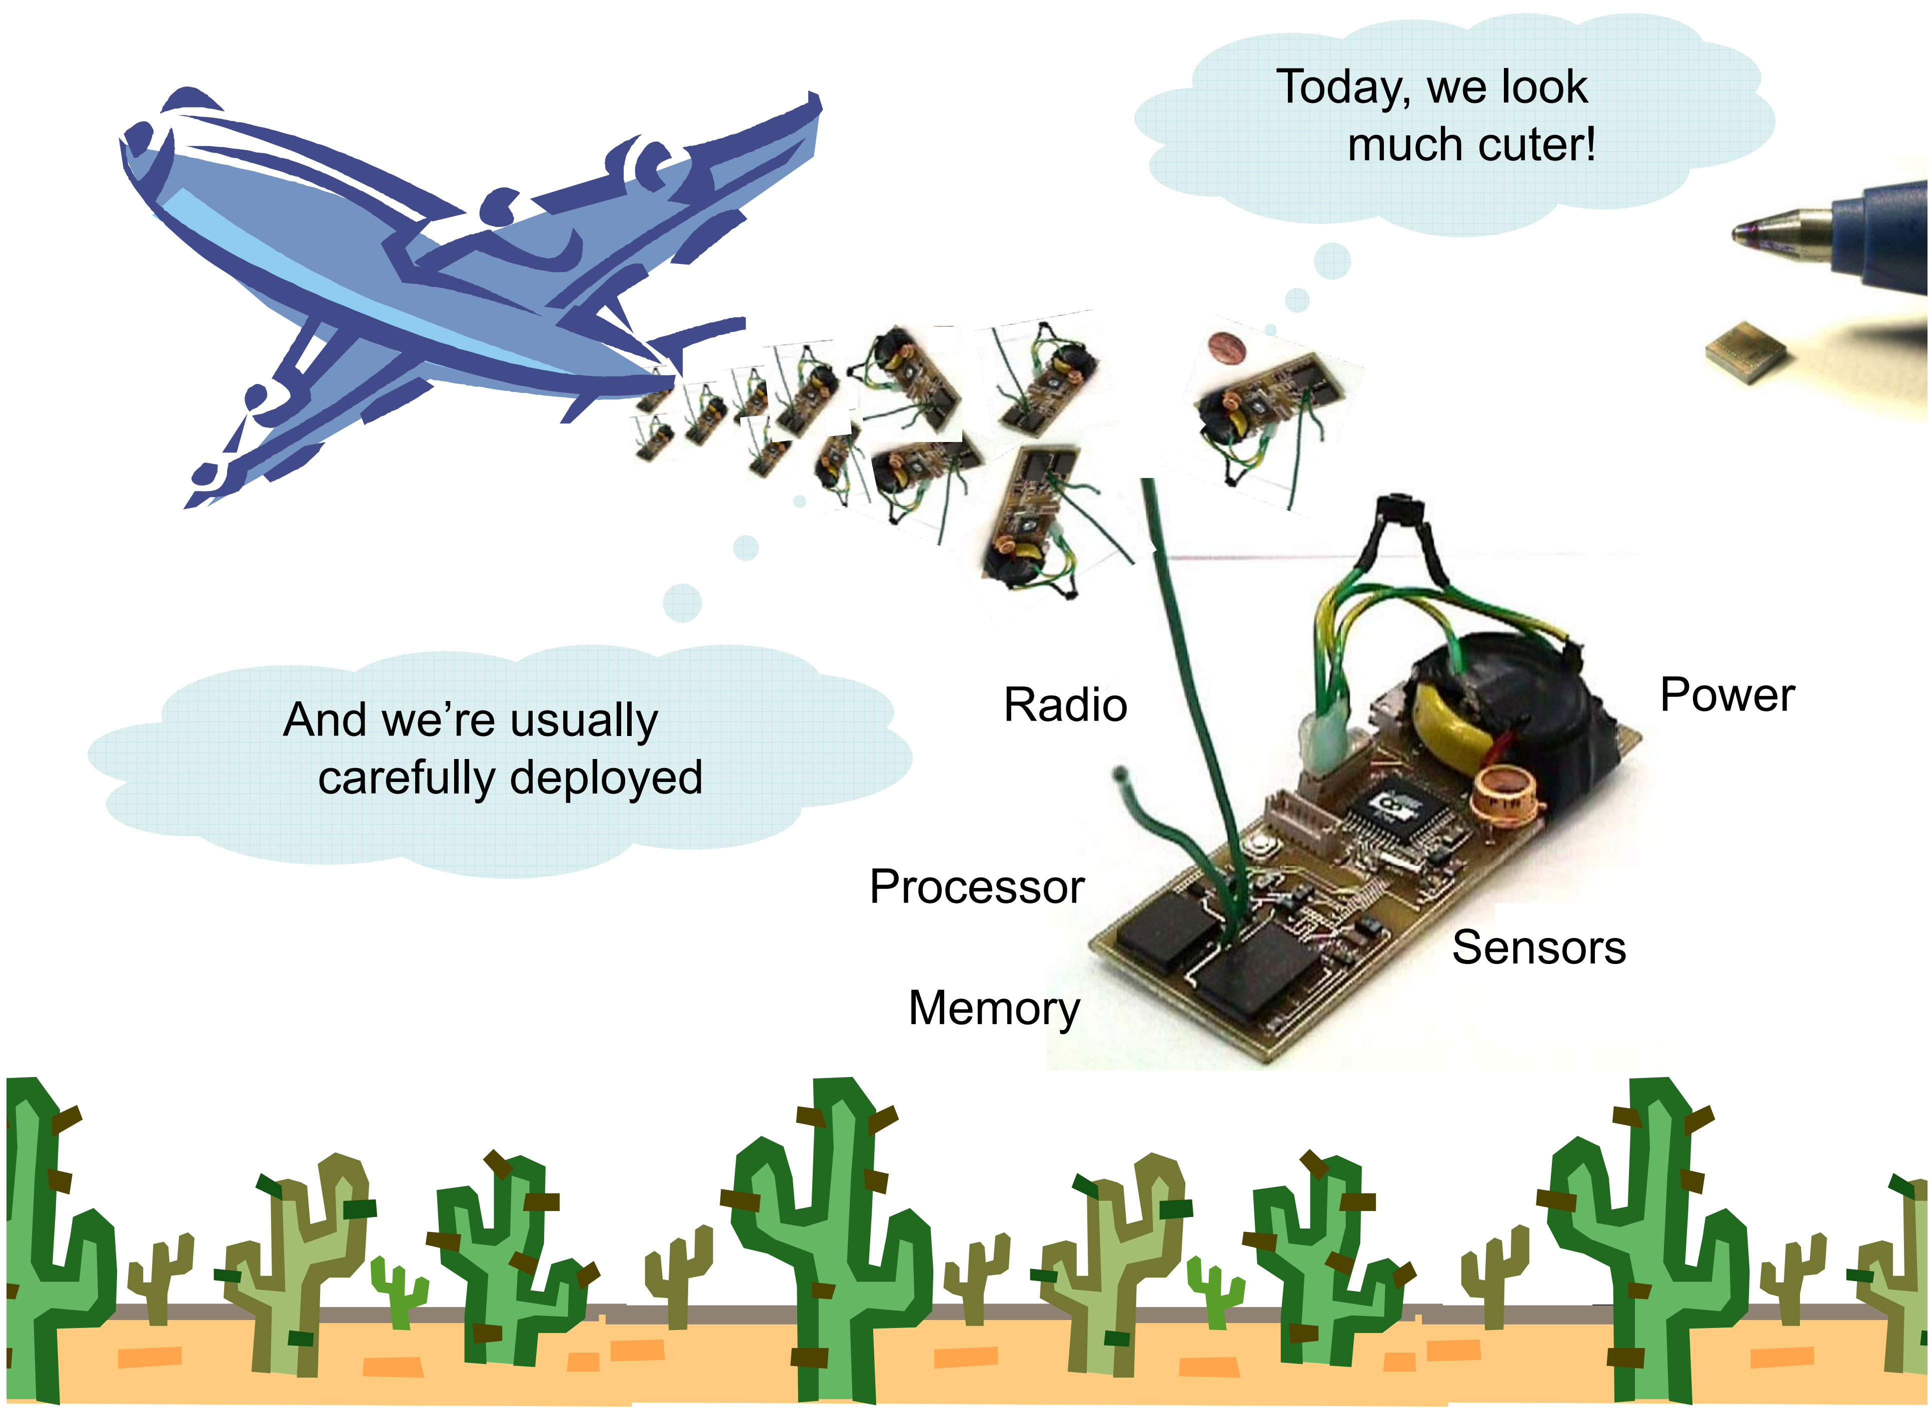
Today, we look much cuter!
And we're usually carefully deployed
Radio
Power
Processor
Sensors
Memory

A Typical Sensor Node: TinyNode 584

[Shockfish SA, The Sensor Network Museum]

- TI MSP430F1611 microcontroller @ 8 MHz
- 10k SRAM, 48k flash (code), 512k serial storage
- 868 MHz Xemics XE1205 multi channel radio
- Up to 115 kbps data rate, 200m outdoor range

	Current Draw	Power Consumption
uC sleep with timer on	6.5 uA	0.0195 mW
uC active, radio off	2.1 mA	6.3 mW
uC active, radio idle listening	16 mA	48 mW
uC active, radio TX/RX at +12dBm	62 mA	186 mW
Max. Power (uC active, radio TX/RX at +12dBm + flash write)	76.9 mA	230.7mW

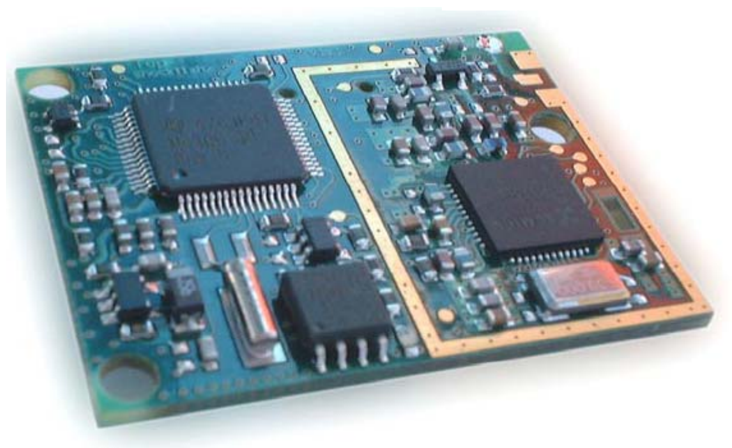

After Deployment

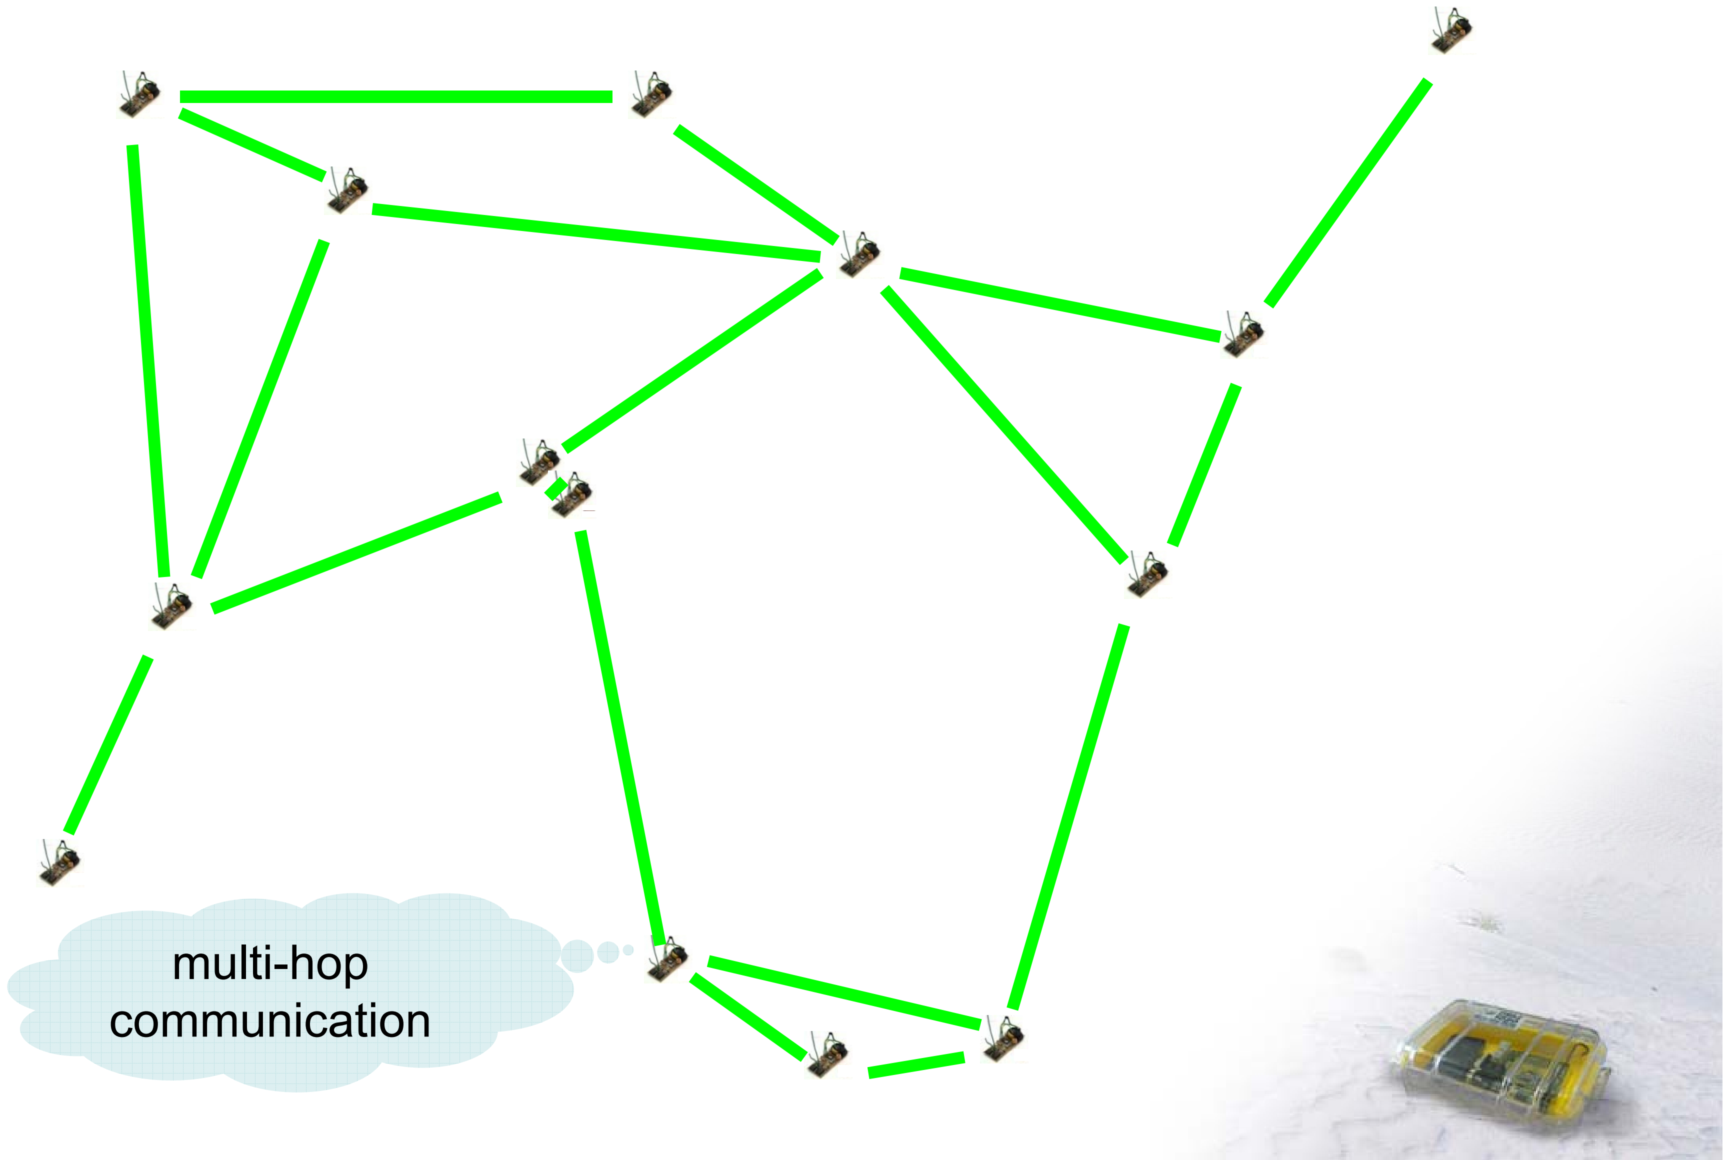

Even more visuals?!? No problem…

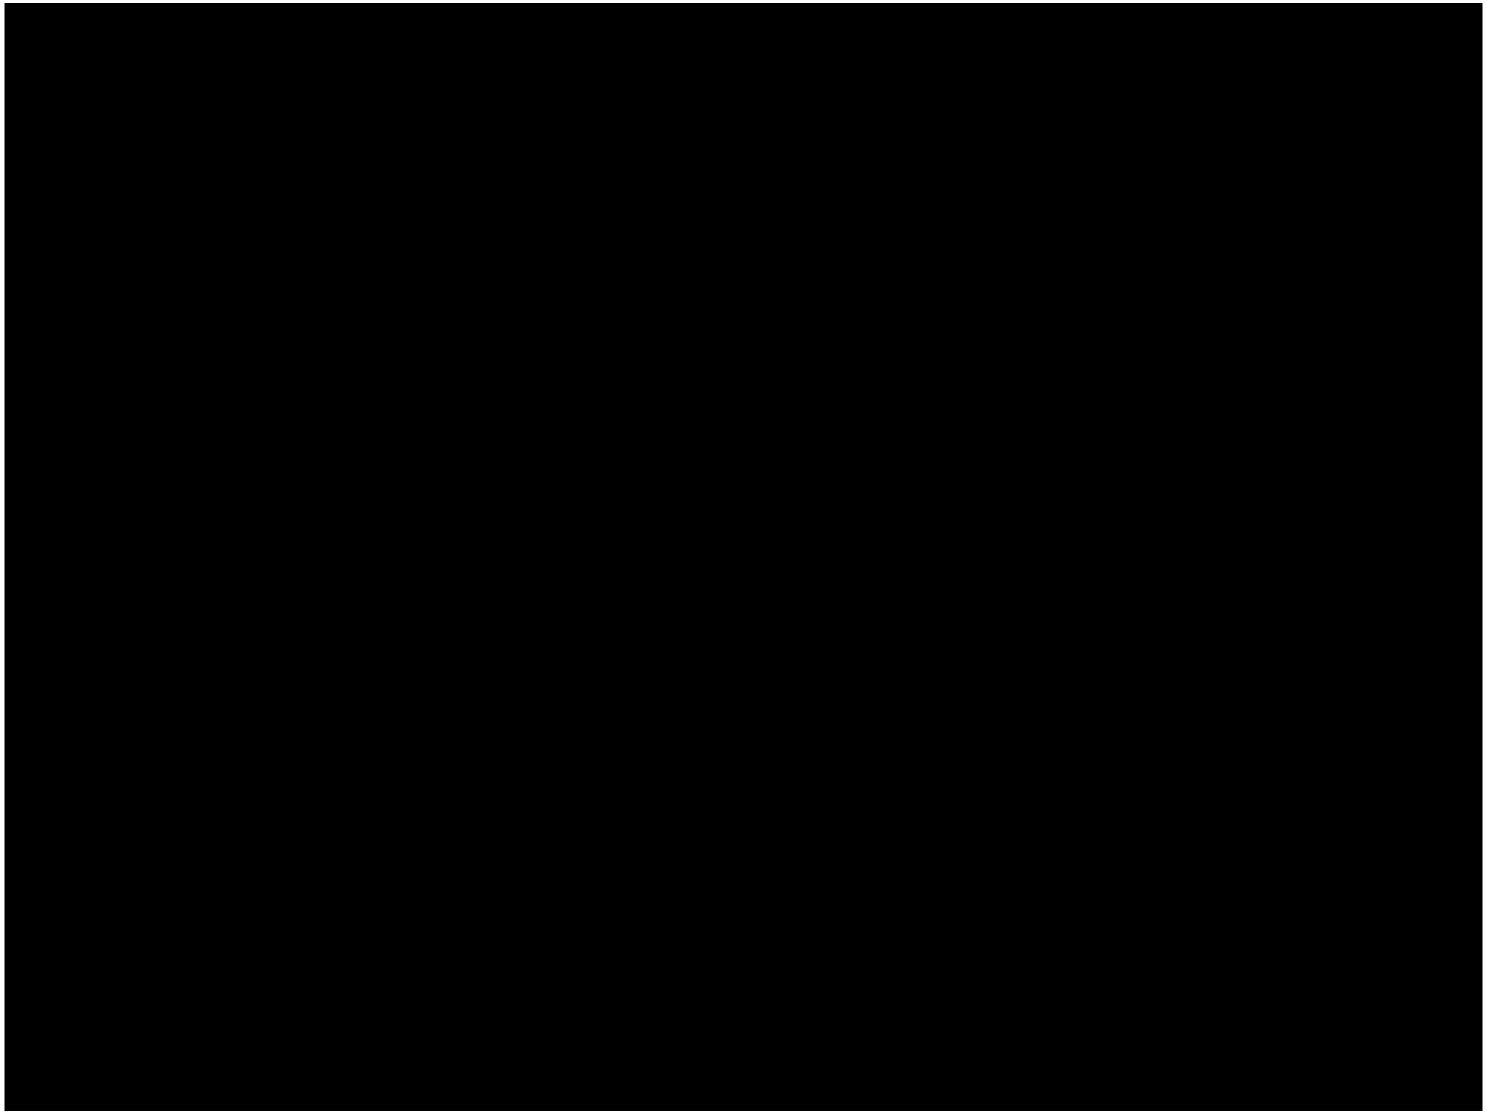

Ad Hoc Networks vs. Sensor Networks

Ad Hoc Networks

- Laptops, PDA's, cars, soldiers
- All-to-all routing
- Often with mobility (MANET's)
- Trust/Security an issue
  - No central coordinator
- Maybe high bandwidth

Sensor Networks

- Tiny nodes: 4 MHz, 32 kB, …
- Broadcast/Echo from/to sink
- Usually no mobility
  - but link failures
- One administrative control
- Long lifetime → Energy

There is no strict separation; more variants such as mesh or sensor/actor networks exist

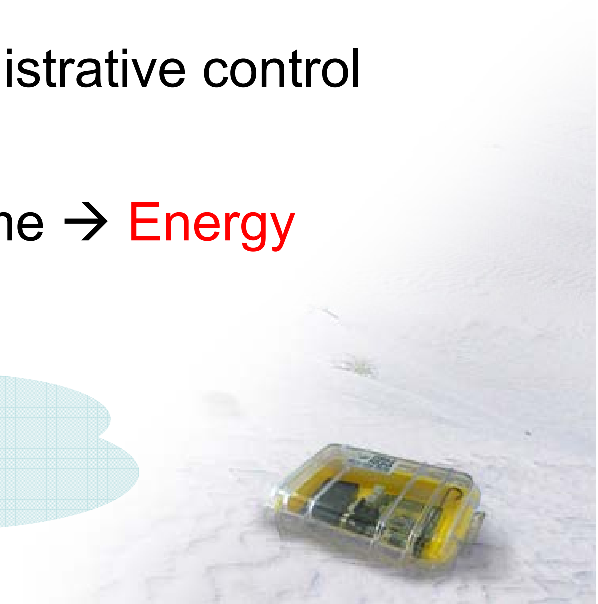

Overview

- Introduction
- Applications
- Case study “Worst-Case Capacity”

Animal Monitoring (Great Duck Island)

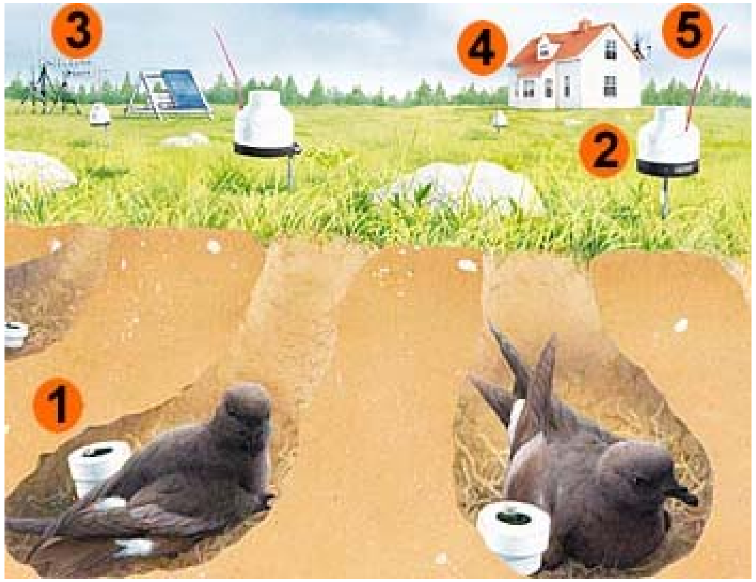

1. Biologists put sensors in underground nests of storm petrel
2. And on 10cm stilts
3. Devices record data about birds
4. Transmit to research station
5. And from there via satellite to lab

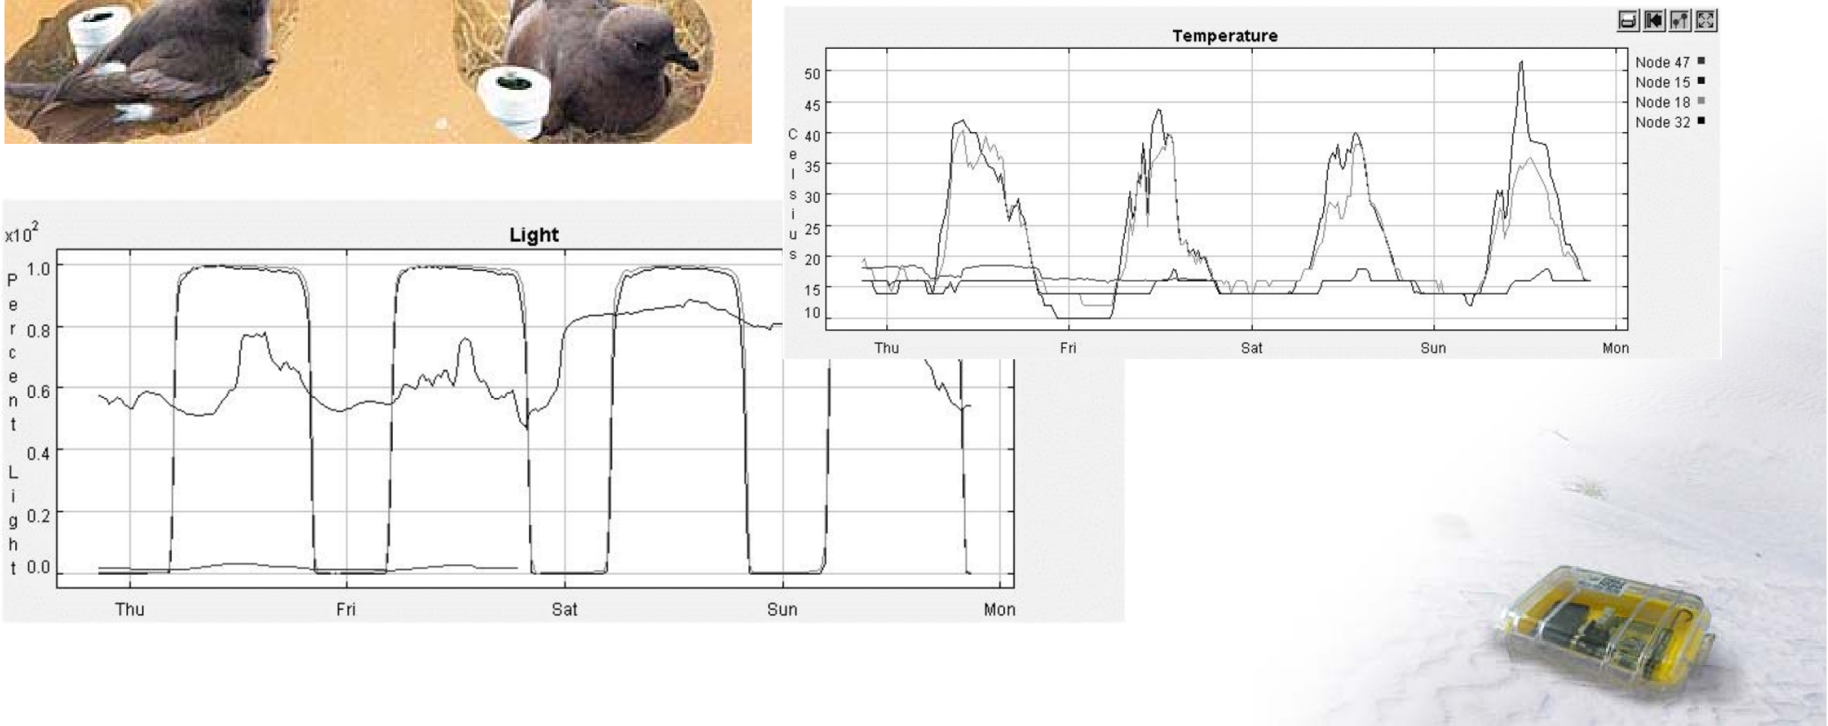

Environmental Monitoring (Redwood Tree)

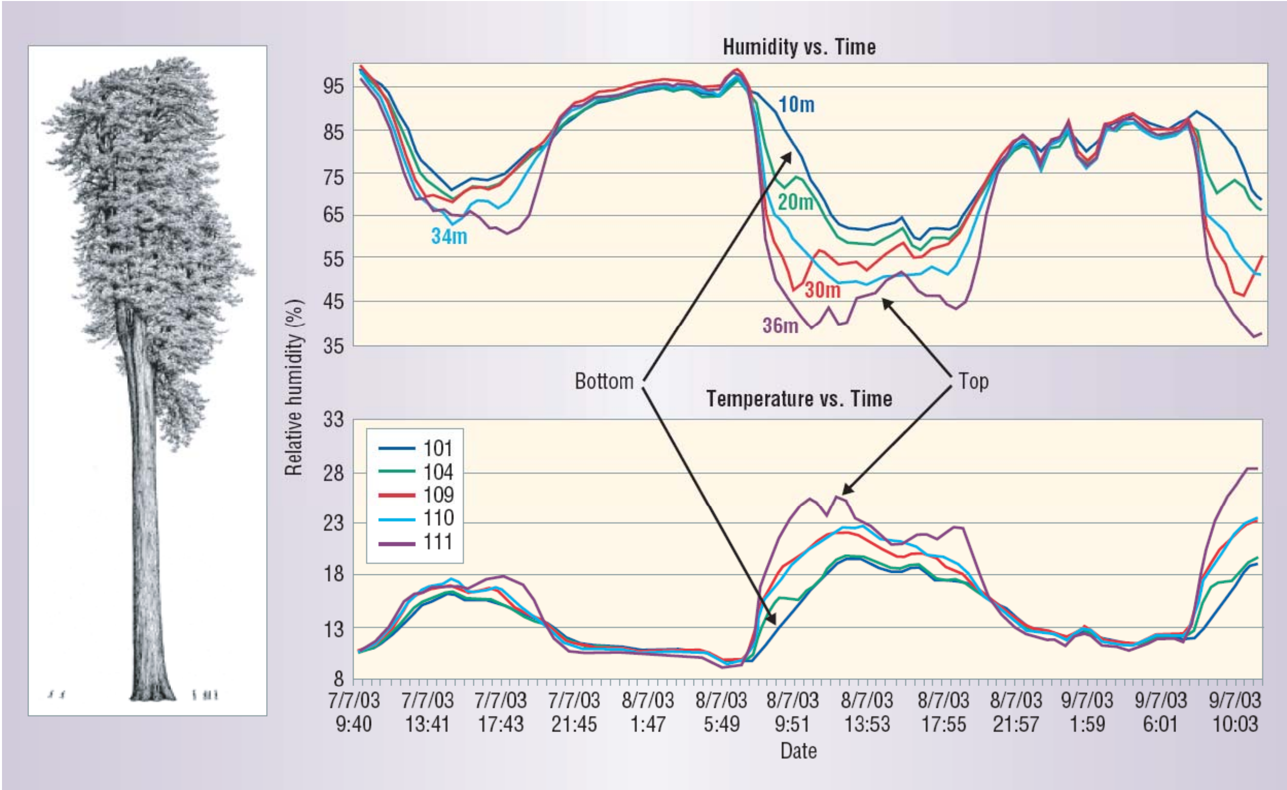

- Microclimate in a tree
- “10km less cables on a tree; easier to set up”

Environmental Monitoring (SensorScope)

- Comfortable access with web interface
- Swiss made (EPFL)
- Various deployments (campus, glacier, etc.)

Environmental Monitoring (Volcanic monitoring)

- Old hardware vs. new hardware
- Sensors: infrasonic mic (for pressure trace) and seismometer (for seismic velocity)
- Equivalent: Earthquake, Tsunami, etc.

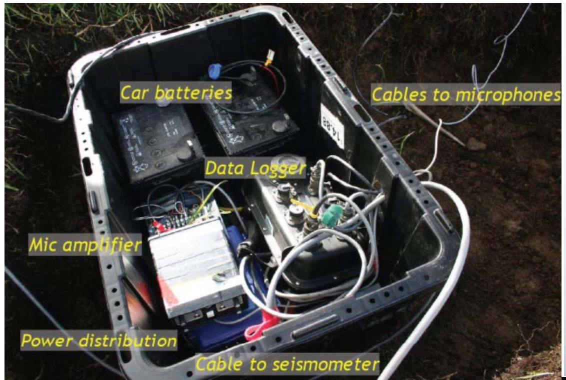

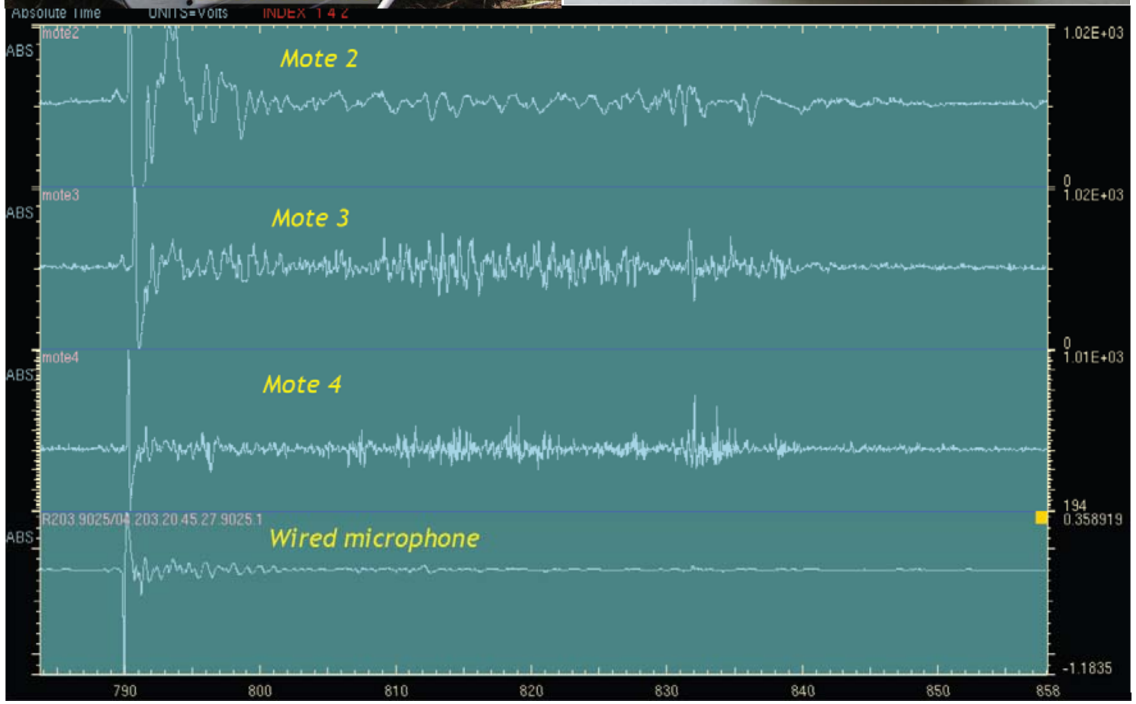

Environmental Monitoring (PermaSense)

- Understand global warming in alpine environment
- Harsh environmental conditions
- Swiss made (Basel, Zurich)

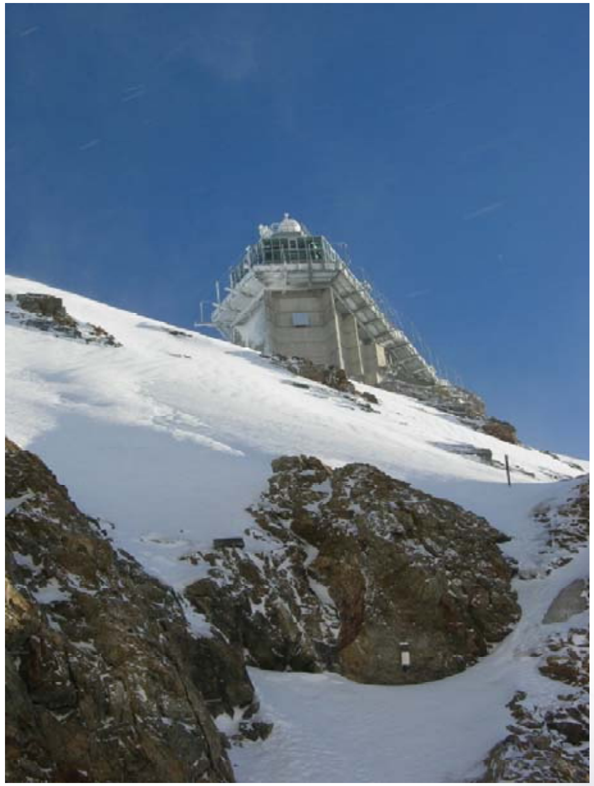

Underwater Sensor Networks

- Static sensor nodes plus mobile robots
- Dually networked
  - optical point-to-point transmission at 300kb/s
  - acoustical broadcast communication at 300b/s, over hundreds of meters range.

- Project AMOUR [MIT, CSIRO]

- Experiments
  - ocean
  - rivers
  - lakes

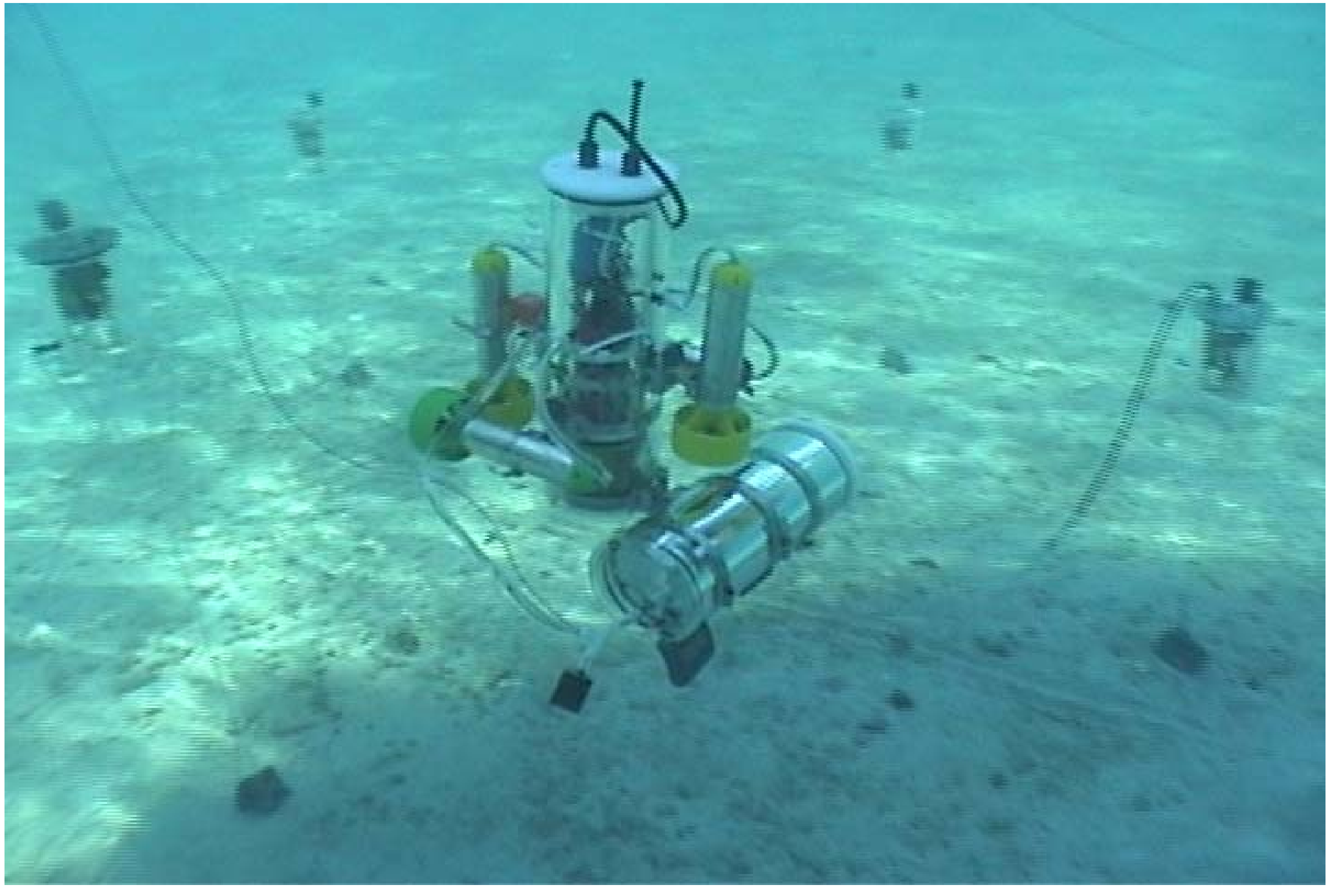

Vehicle Tracking

- Sensor nodes (equipped with magnetometers) are packaged, and dropped from fully autonomous GPS controlled “toy” air plane

- Nodes know dropping order, and use that for initial position guess

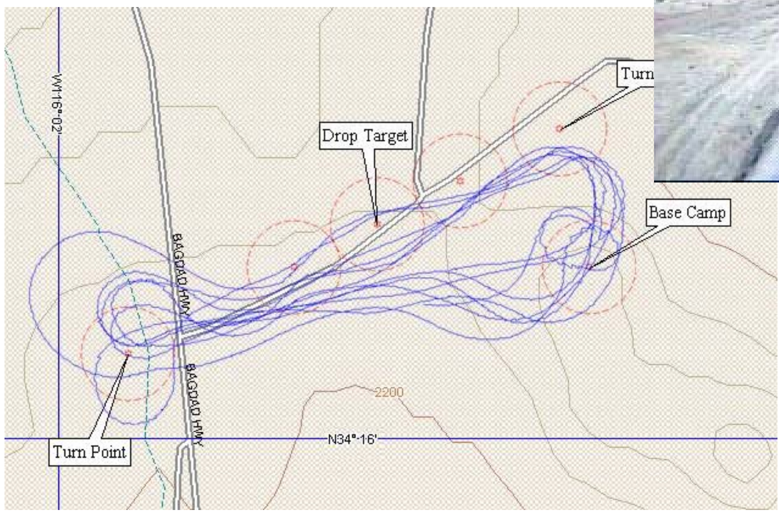

- Nodes then track vehicles (trucks mostly)

Smart Spaces (Car Parking)

- The good: Guide cars towards empty spots
- The bad: Check which cars do not have any time remaining
- The ugly: Meter running out: take picture and send fine

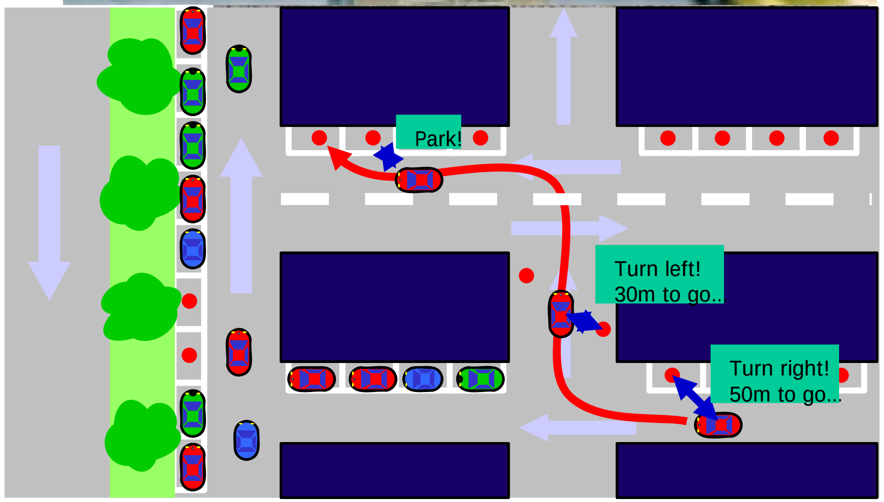

[Matthias Grossglauser, EPFL & Nokia Research]

Traffic Monitoring and Routing Planning (CarTel)

- GPS equipped cars for optimal route predictions, not necessarily “shortest” or “fastest” but also “most likely to get me to target by 9am”

- Various other applications
  e.g. Pothole Patrol

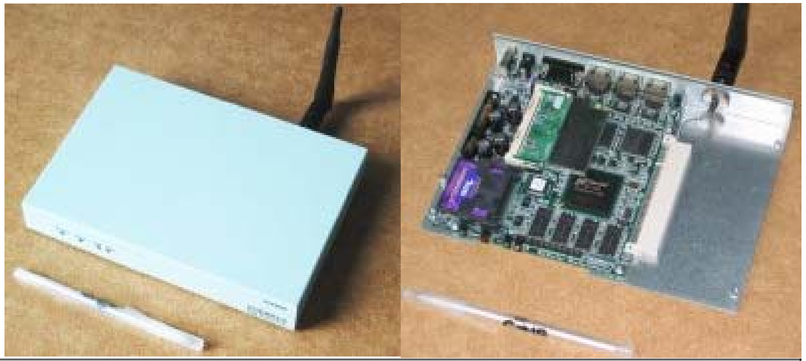

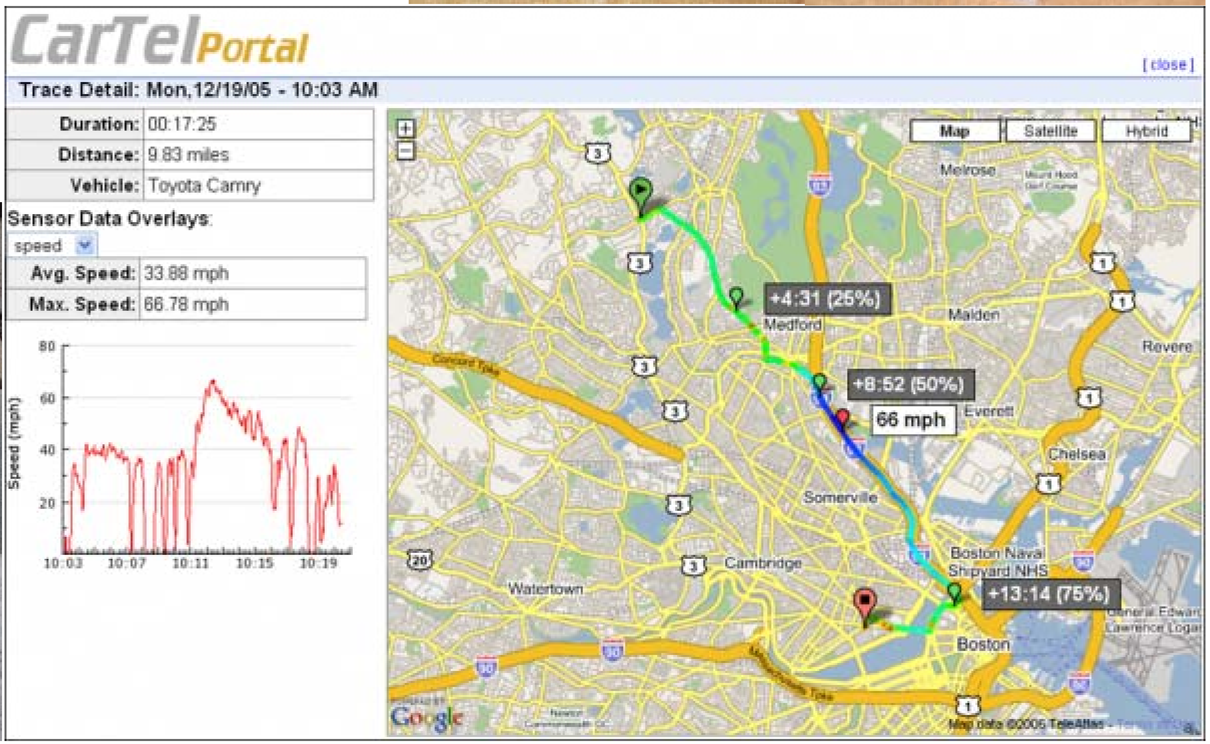

More Car Network Ideas

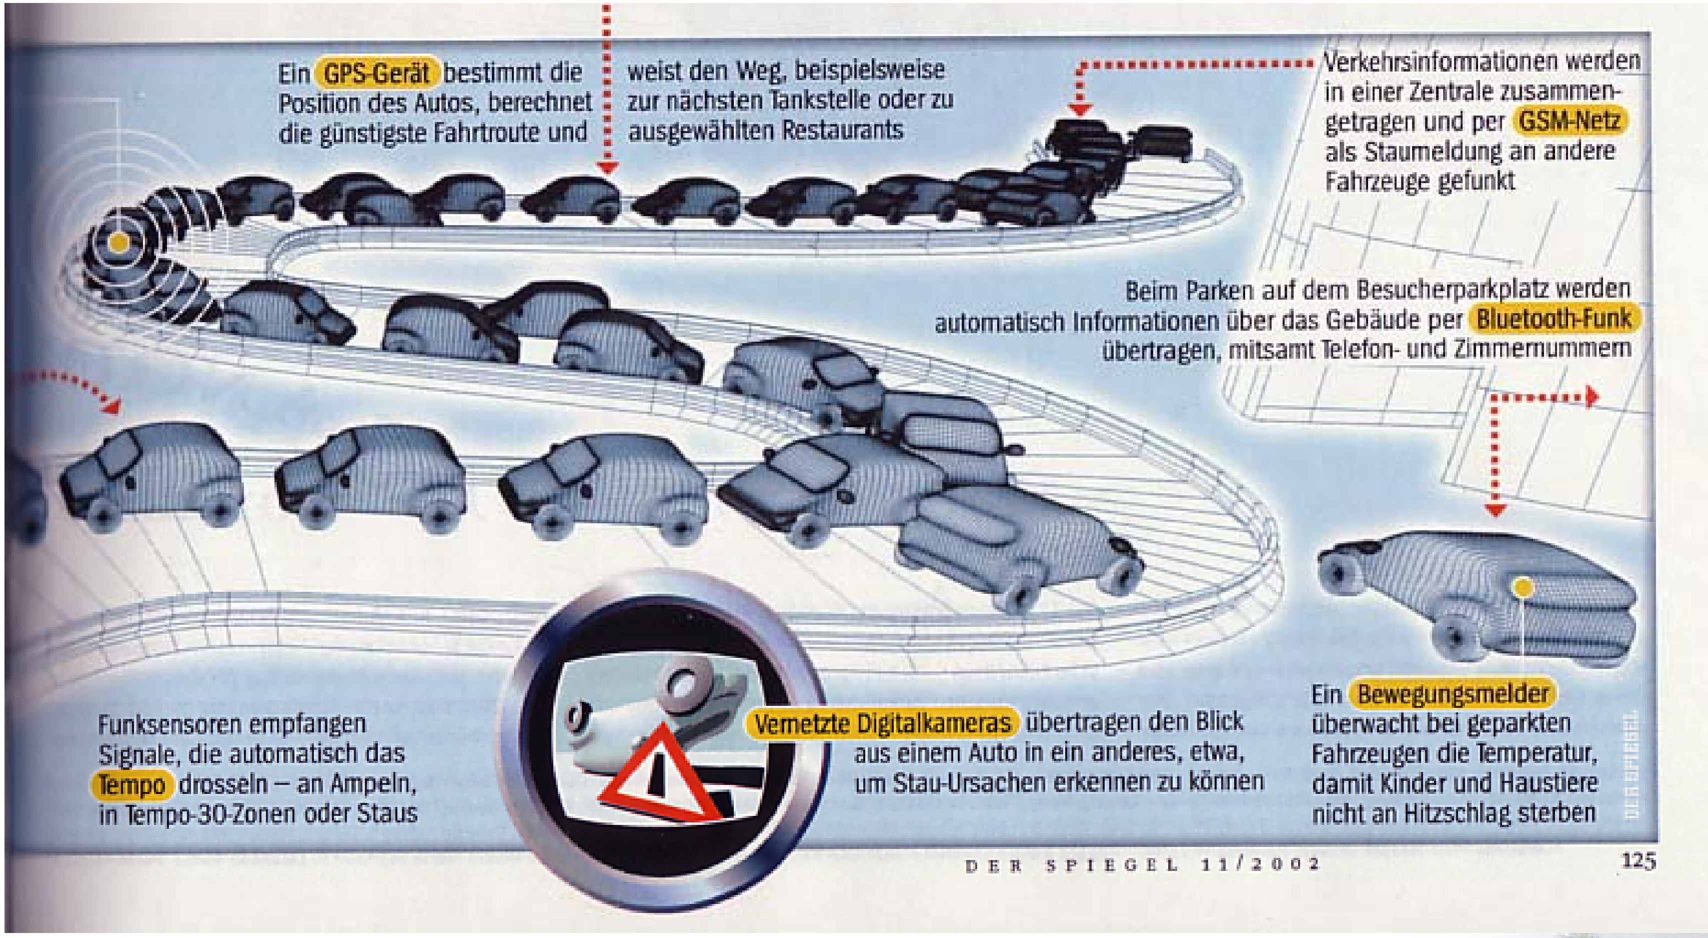

- CAR2CAR Consortium: Audi, BMW, Daimler, Fiat, GM, Honda, Renault, VW

Animal networks (e.g. DeerNet)

- Cars are not the only mobile objects…
- Objective: next-generation wildlife monitoring technology for behavior analysis, interaction modeling, disease tracking and control
- Two-tier system
- Including video data
- Other animals are available: ZebraNet, etc.

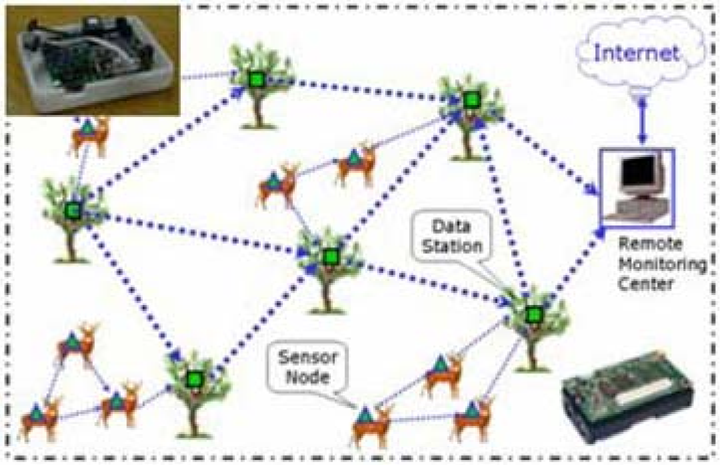

[U. Alberta]

Acoustic Detection (Shooter Detection)

- Sound travels much slower than radio signal (331 m/s)
- This allows for quite accurate distance estimation (cm)
- Main challenge is to deal with reflections and multiple events

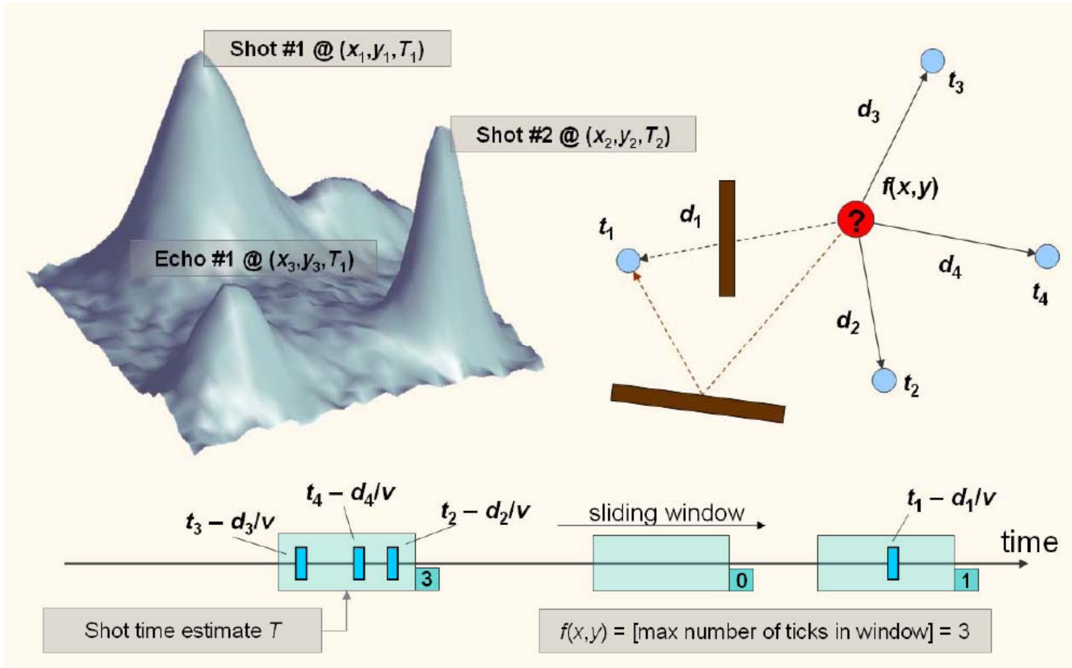

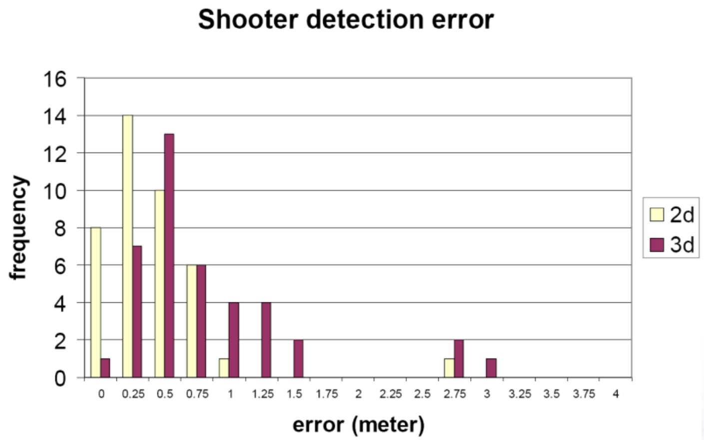

Structural Health Monitoring (Bridge)

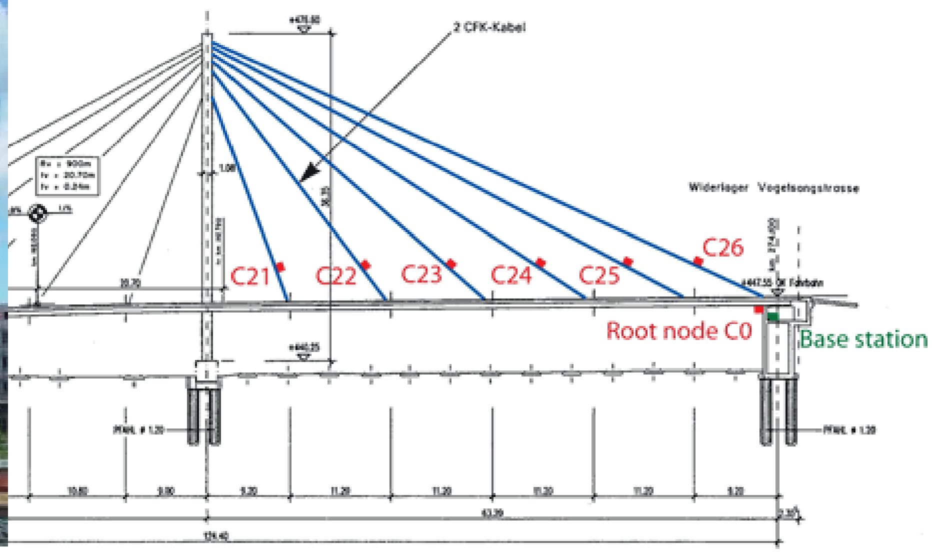

Swiss Made
[EMPA]

Detect structural defects, measuring temperature, humidity, vibration, etc.

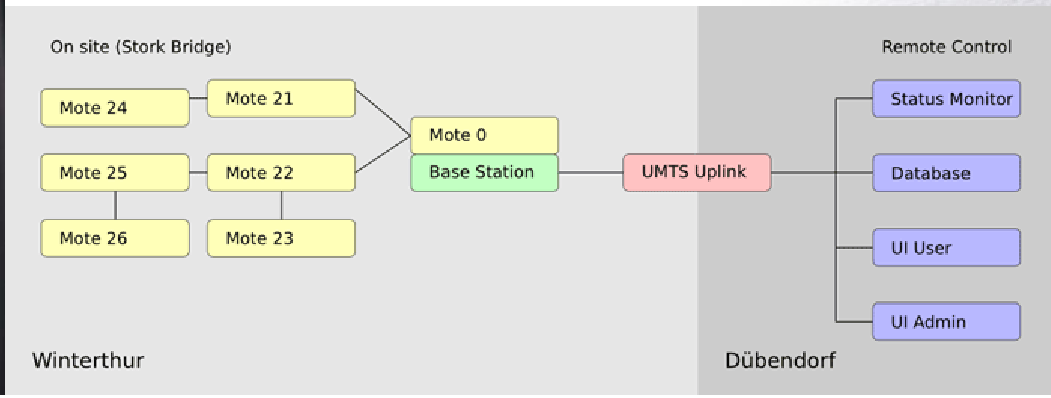

Home Automation

- Light
- Temperature
- Sun-Blinds
- Fans

- Energy Monitoring
- Audio/Video
- Security
  - Intrusion Detection
  - Fire Alarm

Standby Energy [digitalSTROM.org]

- 10 billion electrical devices in Europe
- 9.5 billion are not networked
- 6 billion euro per year energy lost

- Make electricity smart
  - cheap networking (over power)
  - true standby
  - remote control
  - electricity rates
  - universal serial number
  - …

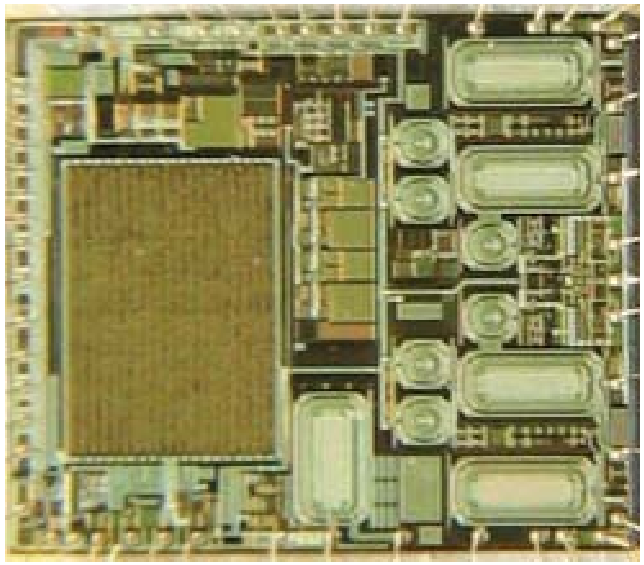

Inventory Tracking (Cargo Tracking)

- Current tracking systems require line-of-sight to satellite.

- Count and locate containers
- Search containers for specific item
- Monitor accelerometer for sudden motion
- Monitor light sensor for unauthorized entry into container

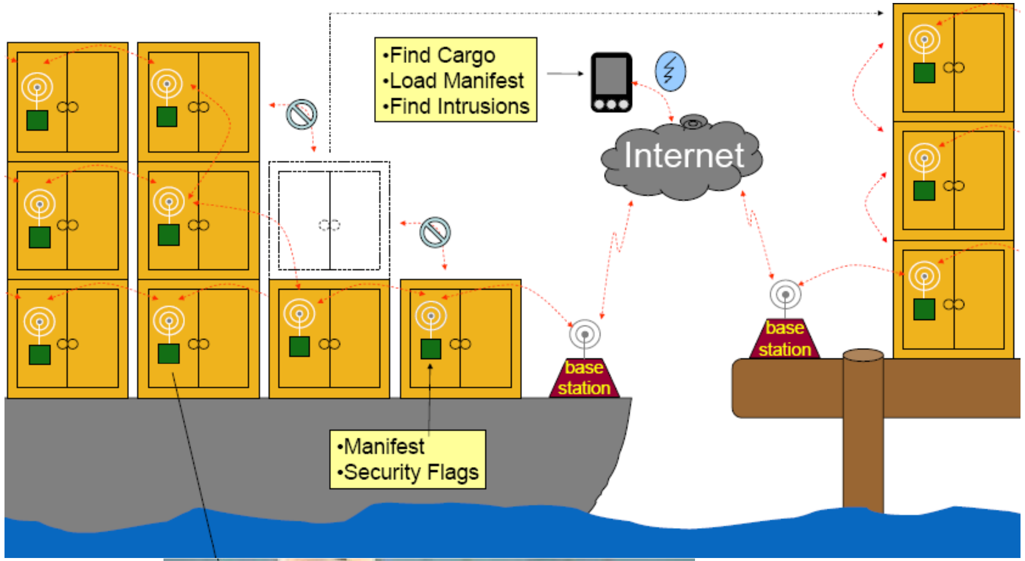

Agriculture (COMMONSense)

- Idea: Farming decision support system based on recent local environmental data.
- Irrigation, fertilization, pest control, etc. are output of function of sunlight, temperature, humidity, soil moisture, etc.
- (Actual sensors are mostly underground)

[EPFL & IIT]

Virtual Fence (CSIRO Australia)

- Download the fence to the cows. Today stay here, tomorrow go somewhere else.
- When a cow strays towards the co-ordinates, software running on the collar triggers a stimulus chosen to scare the cow away, a sound followed by an electric shock; this is the “virtual” fence. The software also "herds" the cows when the position of the virtual fence is moved.
- If you just want to make sure that cows stay together, GPS is not really needed…

Cows learn and need not to be shocked later… Moo!

Mesh Networking (Roofnet)

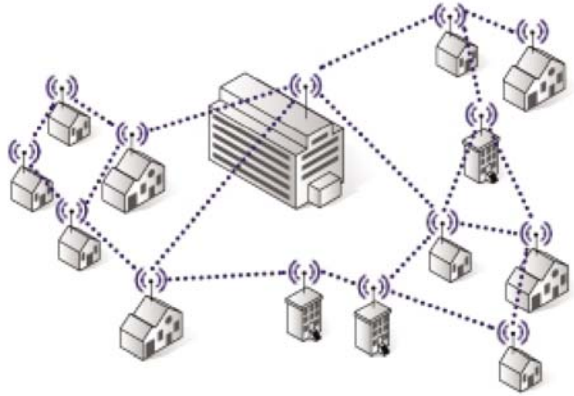

- Sharing Internet access
- Cheaper for everybody
- Several gateways → fault-tolerance
- Possible data backup
- Community add-ons
  - I borrow your hammer, you copy my homework
  - Get to know your neighbors

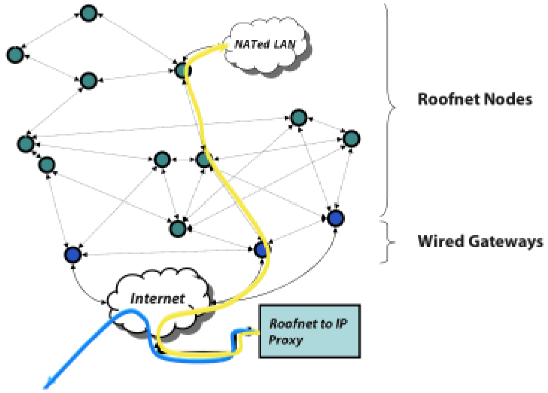

Games / Art

- Uncountable possibilities, below,e.g. a beer coaster that can interact with other coasters…

[sentilla]

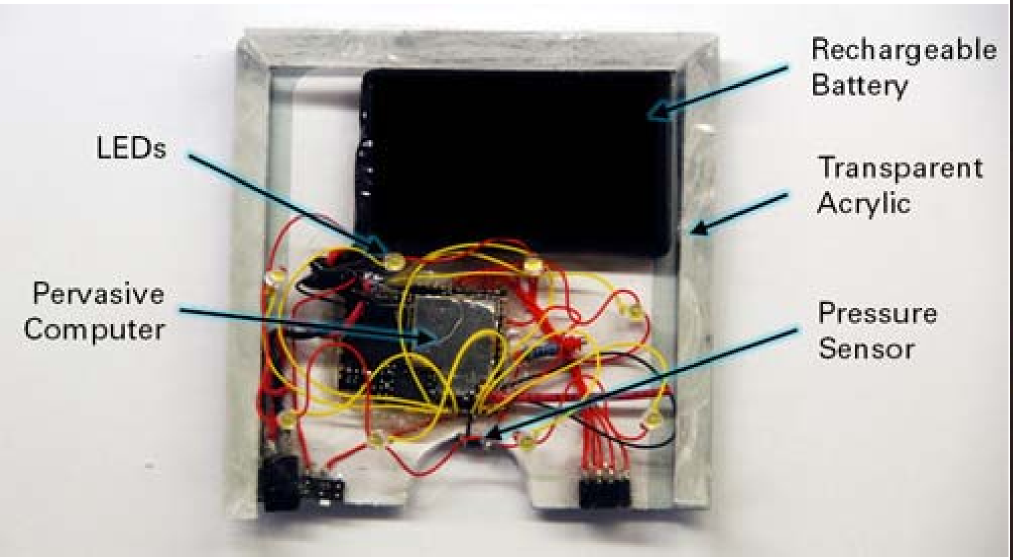

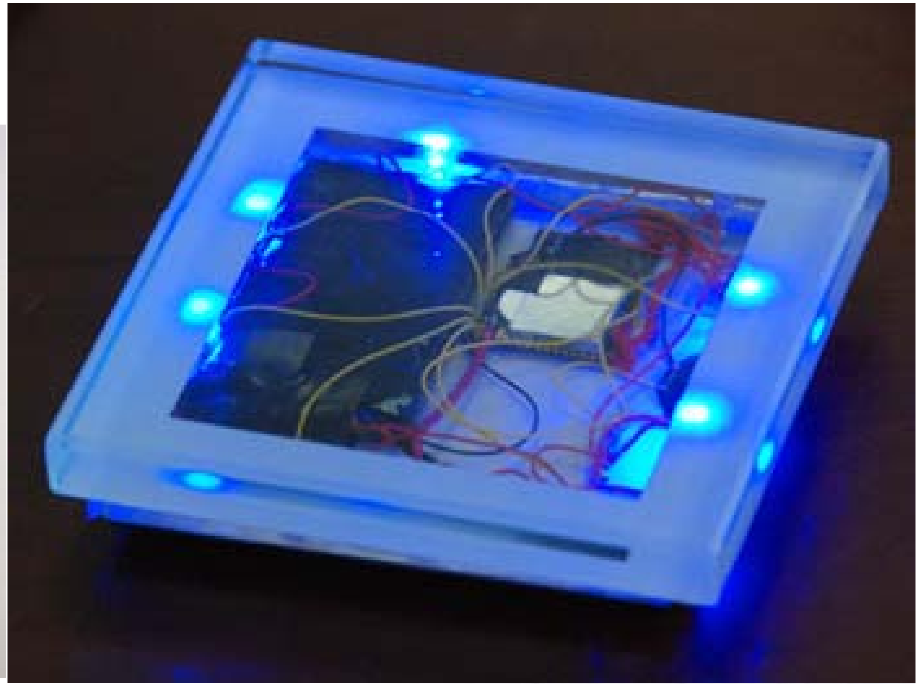

Economic Forecast

[Jean-Pierre Hubaux, EPFL]

- Industrial Monitoring (35% – 45%)
    - Monitor and control production chain
    - Storage management
    - Monitor and control distribution
- Building Monitoring and Control (20 – 30%)
    - Alarms (fire, intrusion etc.)
    - Access control
- Home Automation (15 – 25%)
    - Energy management (light, heating, AC etc.)
    - Remote control of appliances
- Automated Meter Reading (10-20%)
    - Water meter, electricity meter, etc.
- Environmental Monitoring (5%)
    - Agriculture
    - Wildlife monitoring

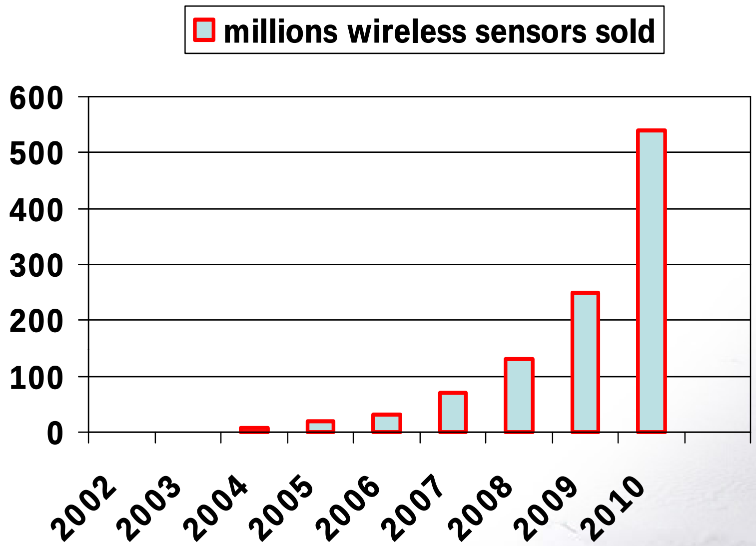

Related Areas

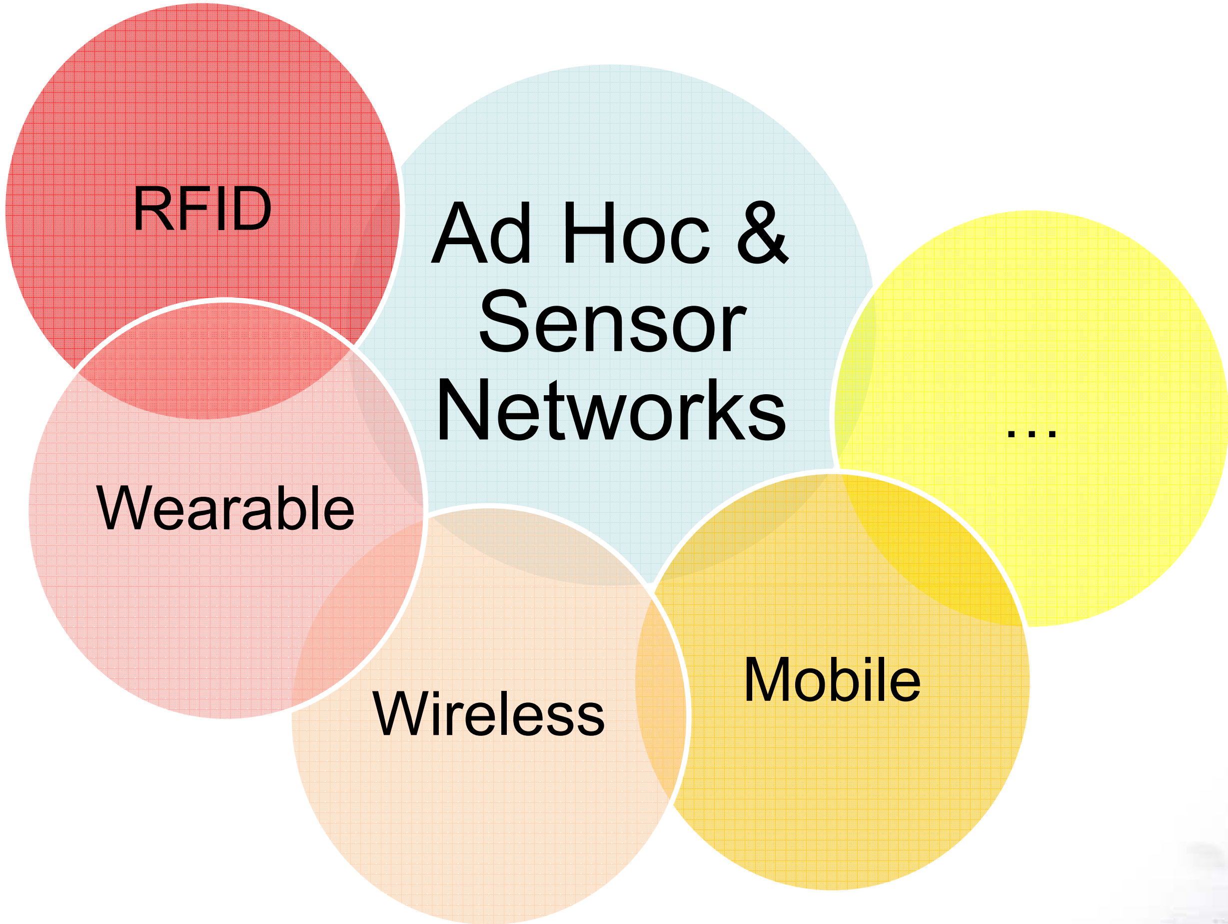

RFID Systems

- Fundamental difference between ad hoc/sensor networks and RFID: In RFID there is always the distinction between the passive tags/transponders (tiny/flat), and the reader (bulky/big).

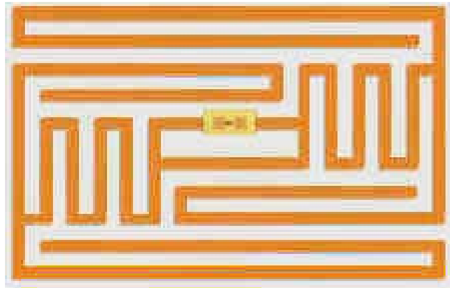

- There is another form of tag, the so-called active tag, which has its own internal power source that is used to power the integrated circuits and to broadcast the signal to the reader. An active tag is similar to a sensor node.

- More types are available, e.g. the semi-passive tag, where the battery is not used for transmission (but only for computing)

Wearable Computing / Ubiquitous Computing

- Tiny embedded “computers”
- UbiComp: Microsoft’s Doll

- I refer to my colleague Gerhard Troester and his lectures & seminars

[Schiele, Troester]

Wireless and/or Mobile

- Aspects of mobility
  - User mobility: users communicate “anytime, anywhere, with anyone” (example: read/write email on web browser)
  - Device portability: devices can be connected anytime, anywhere to the network
- Wireless vs. mobile Examples

Wireless	mobile	Examples
✗	✗	Stationary computer
✗	✓	Notebook in a hotel
✓	✗	Historic buildings; last mile
✓	✓	Personal Digital Assistant (PDA)

- The demand for mobile communication creates the need for integration of wireless networks and existing fixed networks
  - Local area networks: standardization of IEEE 802.11 or HIPERLAN
  - Wide area networks: GSM and ISDN
  - Internet: Mobile IP extension of the Internet protocol IP

Wireless & Mobile Examples

- Up-to-date localized information
  - Map
  - Pull/Push
- Ticketing
- Etc.

[Asus PDA, iPhone, Blackberry, Cybiko]

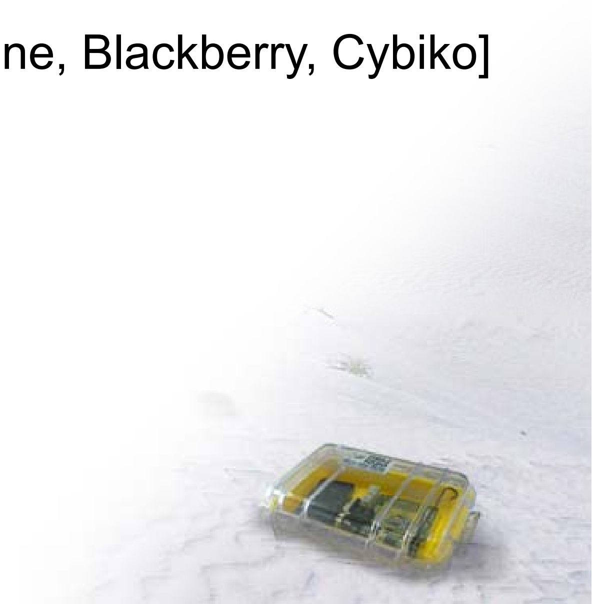

General Trend: A computer in 10 years?

- Advances in technology
  - More computing power in smaller devices
  - Flat, lightweight displays with low power consumption
  - New user interfaces due to small dimensions
  - More bandwidth (per second? per space?)
  - Multiple wireless techniques
- Technology in the background
  - Device location awareness: computers adapt to their environment
  - User location awareness: computers recognize the location of the user and react appropriately (call forwarding)
- “Computers” evolve
  - Small, cheap, portable, replaceable
  - Integration or disintegration?

Rating (of Applications)

- Area maturity

First steps — Text book

- Practical importance

No apps — Mission critical

- Theoretical importance

Not really — Must have

Open Problem

- Well, the open problem for this chapter is obvious:
- Find the killer application! Get rich and famous!!

Worst-Case Capacity

Rating

- Area maturity

First steps — Text book

- Practical importance

No apps — Mission critical

- Theoretical importance

Not really — Must have

Disclaimer…

- Work is about wireless networking in general
  - This presentation focusing on wireless sensor networks

Periodic data gathering in sensor networks

- All nodes produce relevant information about their vicinity periodically.
- Data is conveyed to an information sink for further processing.
- Data may or may not be aggregated.
- Variations
  - Sense event (e.g. fire, burglar)
  - SQL-like queries (e.g. TinyDB)

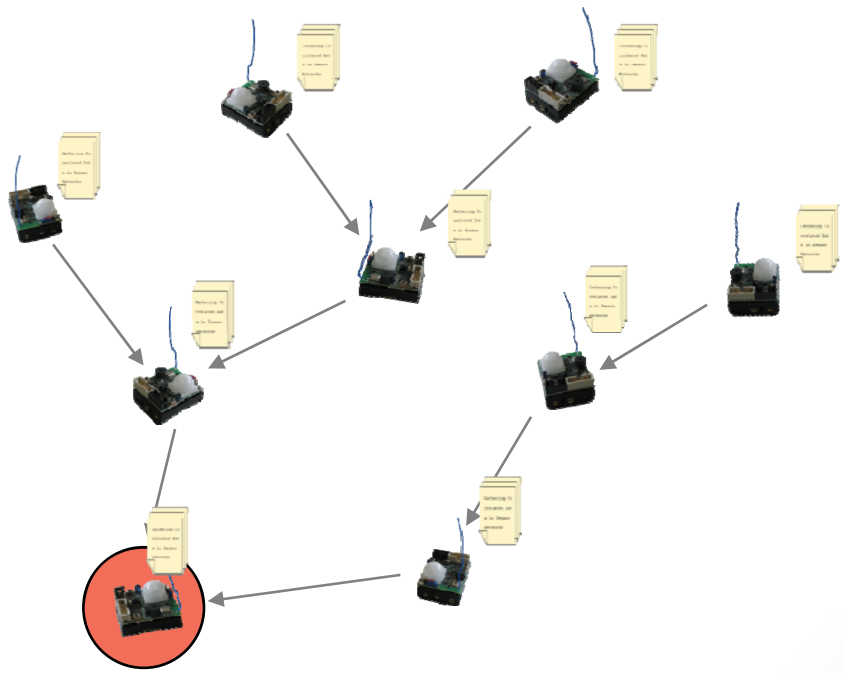

Data Gathering in Wireless Sensor Networks

- Data gathering & aggregation
  - Classic application of sensor networks
  - Sensor nodes periodically sense environment
  - Relevant information needs to be transmitted to sink
- Functional Capacity of Sensor Networks
  - Sink peridically wants to compute a function f_n of sensor data
  - At what rate can this function be computed?

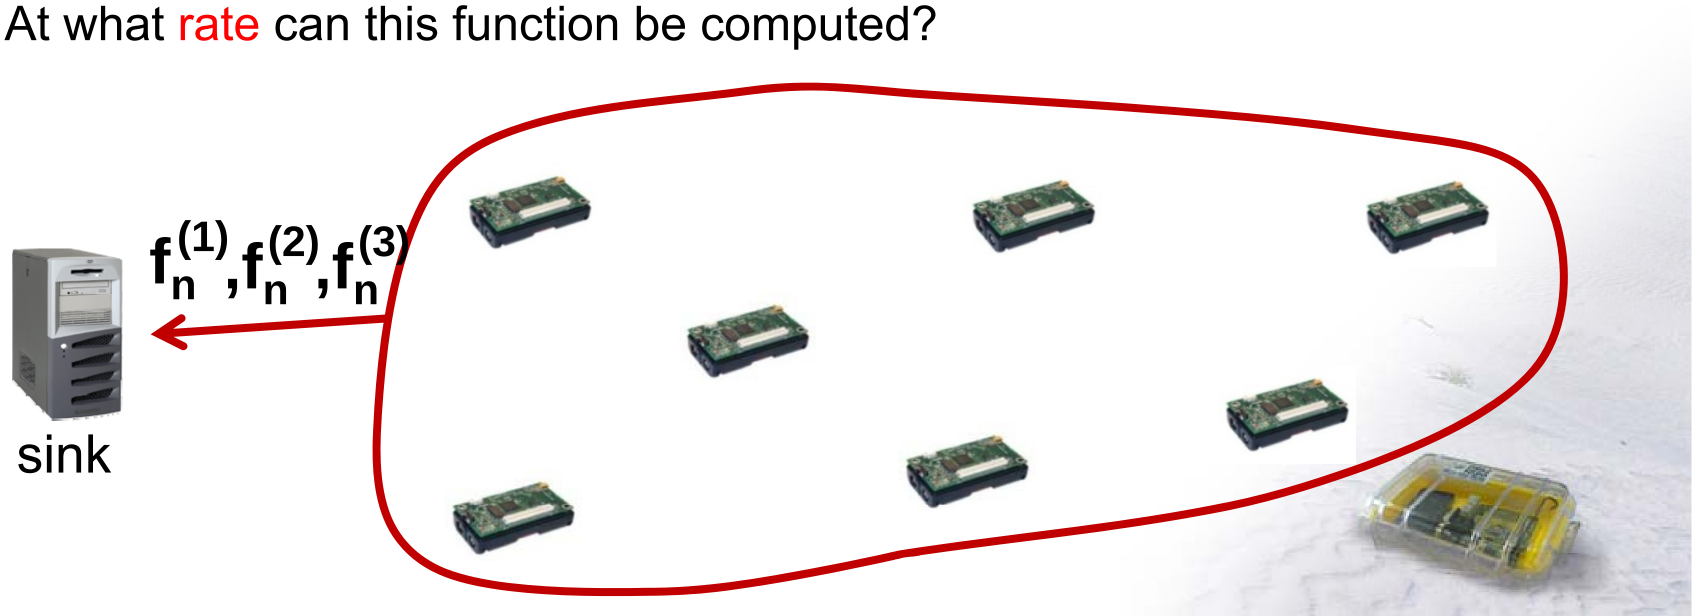

Data Gathering in Wireless Sensor Networks

Example: simple round-robin scheme
→ Each sensor reports its results directly to the root one after another

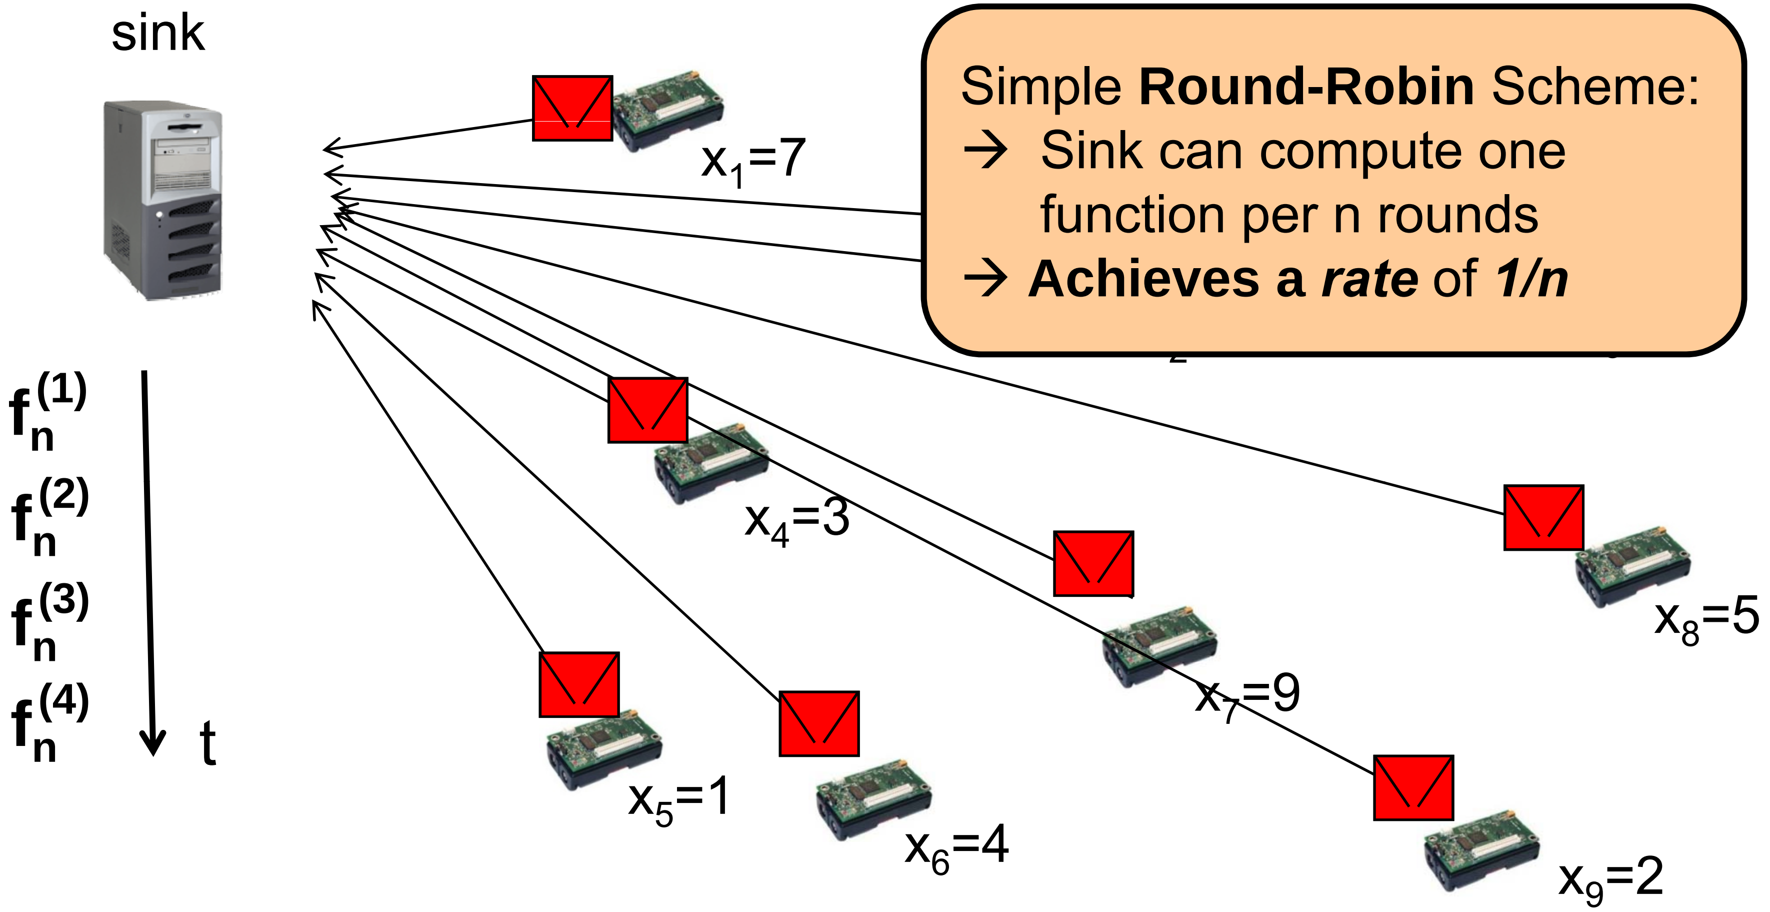

Data Gathering in Wireless Sensor Networks

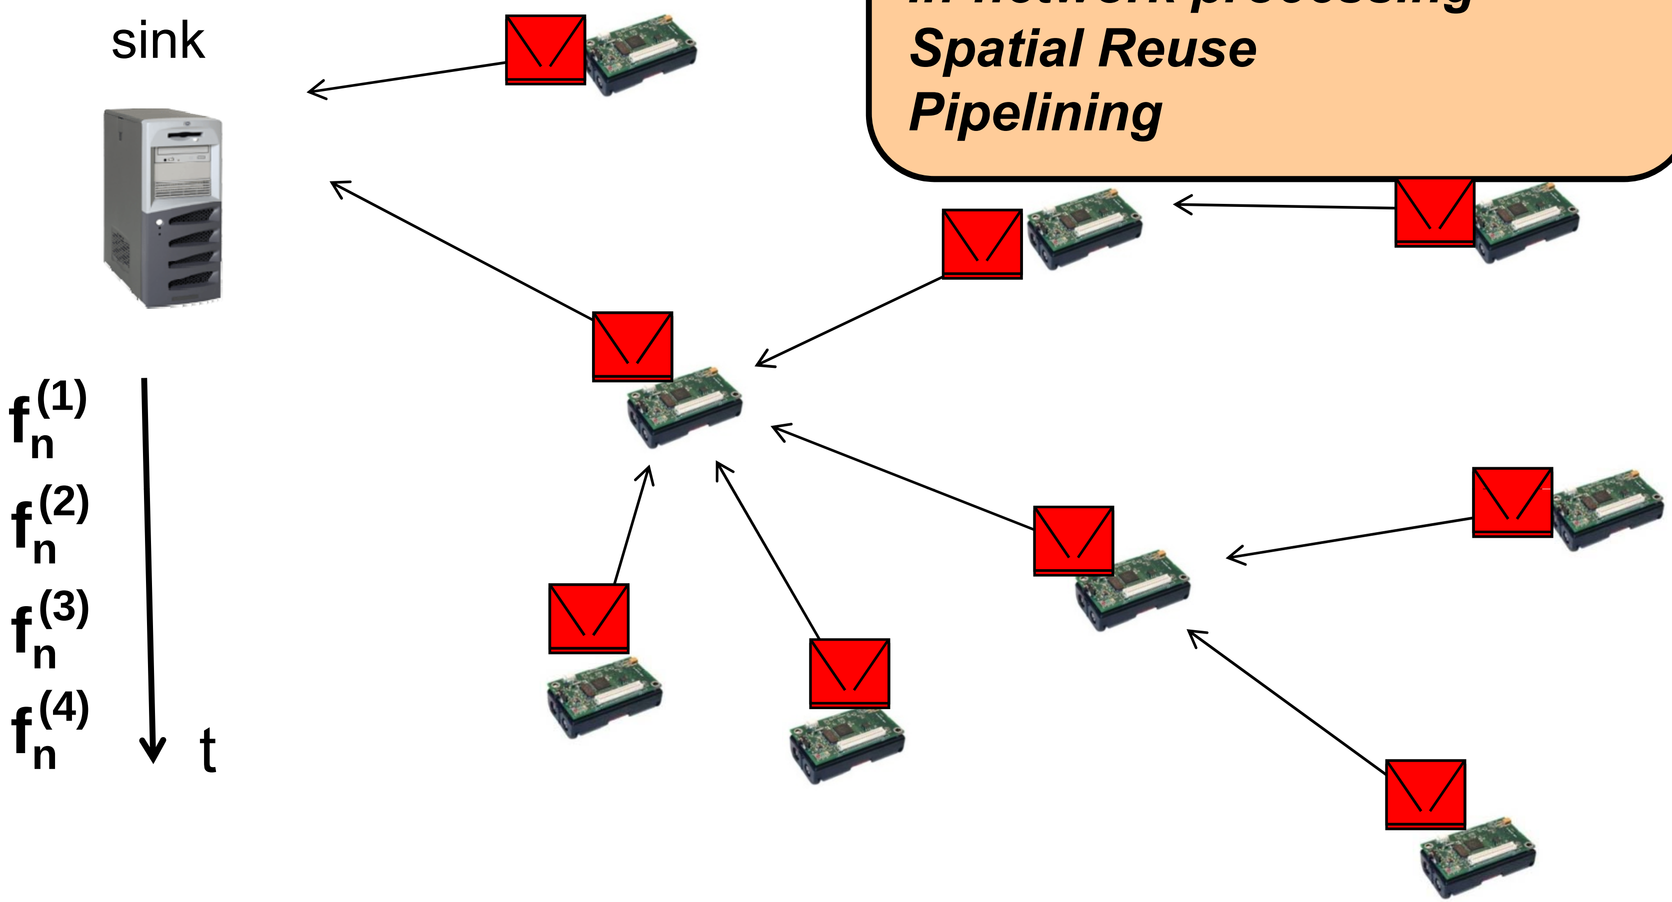

There are better schemes using
Multi-hop relaying
In-network processing
Spatial Reuse
Pipelining

Capacity in Wireless Sensor Networks

At what ***rate*** can sensors transmit data to the sink?
Scaling-laws → how does rate decrease as *n* increases…?

$\Theta(1/n)$ $\Theta(1/\sqrt{n})$ $\Theta(1/\log n)$ $\Theta(1)$

Answer depends on:
- Function to be computed → Only perfectly compressible functions (max, min, avg,…)
- Coding techniques → No fancy coding techniques
- Network topology
- Wireless communication model

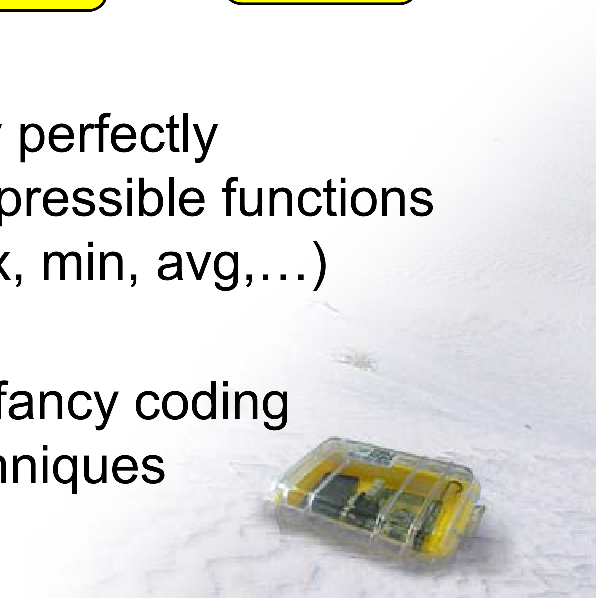

“Classic” Capacity…

The Capacity of Wireless Networks
Gupta, Kumar, 2000

[Arpacioglu et al, IPSN’04]
[Giridhar et al, JSAC’05]
[Barrenechea et al, IPSN’04]
[Grossglauser et al, INFOCOM’01]
[Liu et al, INFOCOM’03]
[Toumpis, TWC’03]
[Gamal et al, INFOCOM’04]
[Kyasanur et al, MOBICOM’05]
[Kodialam et al, MOBICOM’05]
[Gastpar et al, INFOCOM’02]
[Li et al, MOBICOM’01]
[Mitra et al, IPSN’04]
[Zhang et al, INFOCOM’05]
[Bansal et al, INFOCOM’03]
[Dousse et al, INFOCOM’04]
[Yi et al, MOBIHOC’03]
[Perevalov et al, INFOCOM’03]
etc…

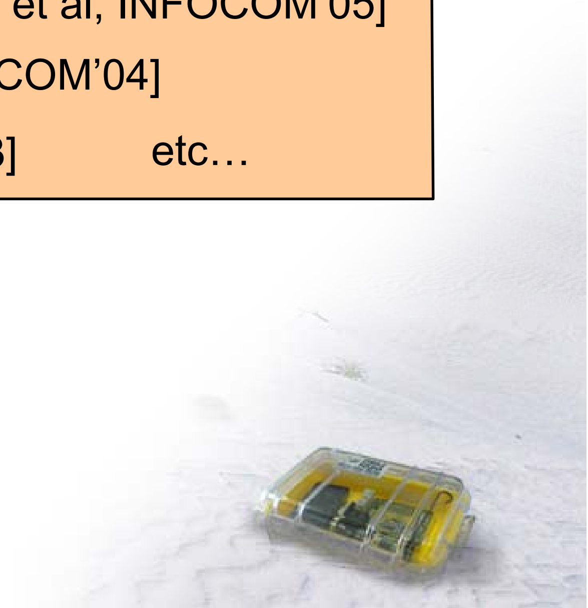

Worst-Case Capacity

- Capacity studies so far make very strong assumptions on node deployment, topologies
  - randomly, uniformly distributed nodes
  - nodes placed on a grid
  - etc...

What if a network looks differently…?

Like this?

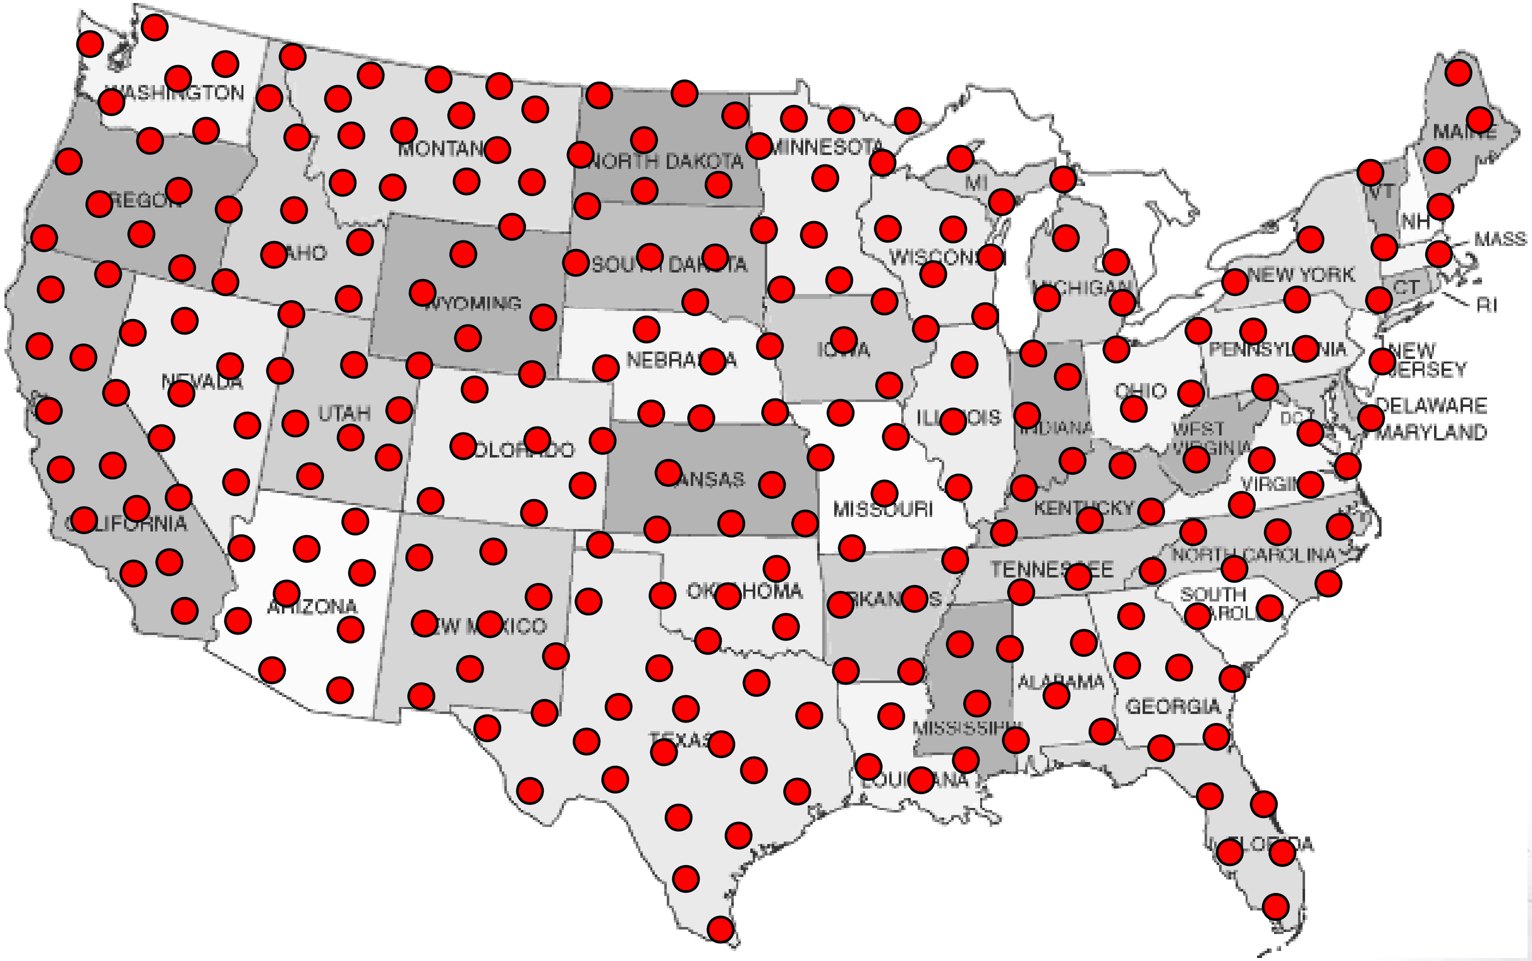
NORTH DAKOTA
MINNESOTA
MI
WYOMING
UTAH
NEVADA
OHIO
NEW YORK
VT
NH
MASS
CT
RI
NEW JERSEY
DELAWARE
MARYLAND
DC
KENTUCKY
NORTH CAROLINA
GEORGIA
MISSOURI
TEXAS

Or rather like this?

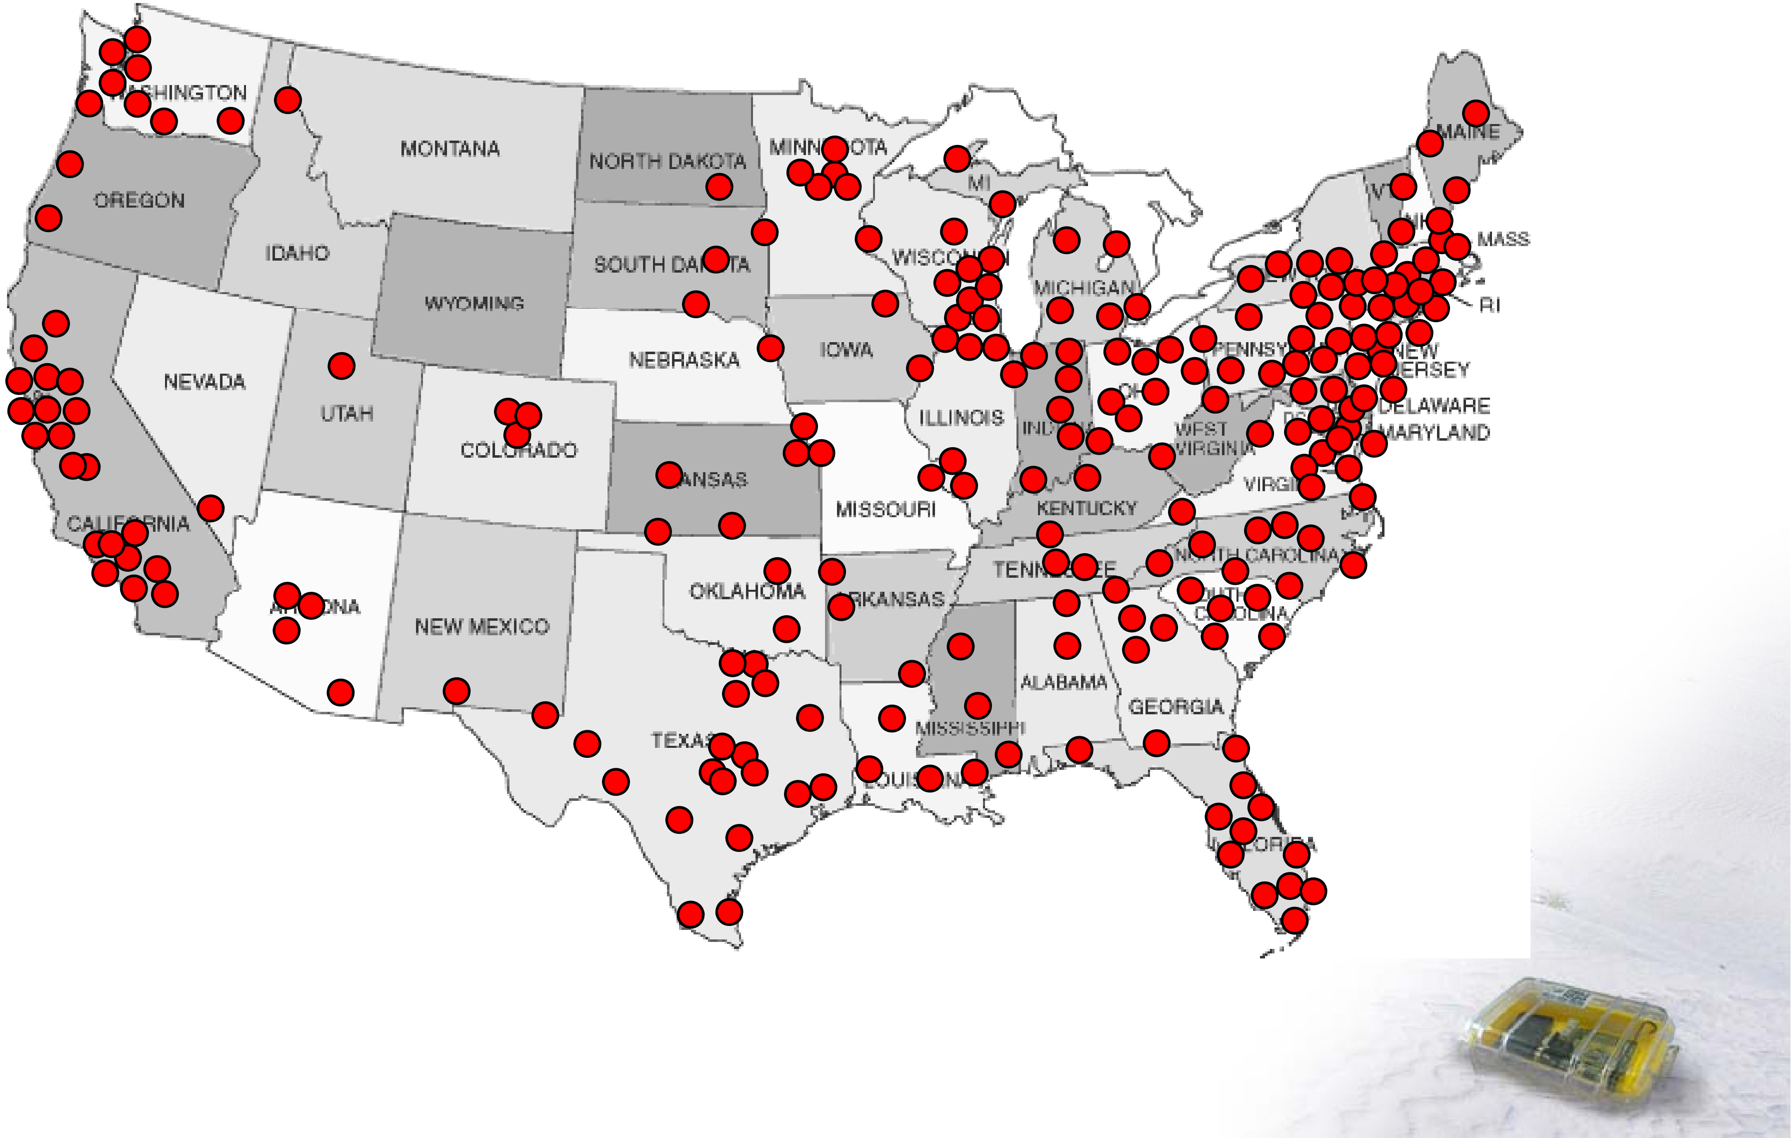

Worst-Case Capacity

- Capacity studies so far have made very strong assumptions on node deployment, topologies
  - randomly, uniformly distributed nodes
  - nodes placed on a grid
  - etc...

What if a network looks differently…?

We assume arbitrary node distribution

worst-case topologies

Classic Capacity ⟷ **Worst-Case Capacity**

How much information can be transmitted in nice, well-behaving networks

How much information can be Transmitted in any network

Models

- Two standard models in wireless networking

Protocol Model (graph-based, simpler) ⟷ Physical Model (SINR-based, more realistic)

Protocol Model

- Based on graph-based notion of interference
- Transmission range and interference range

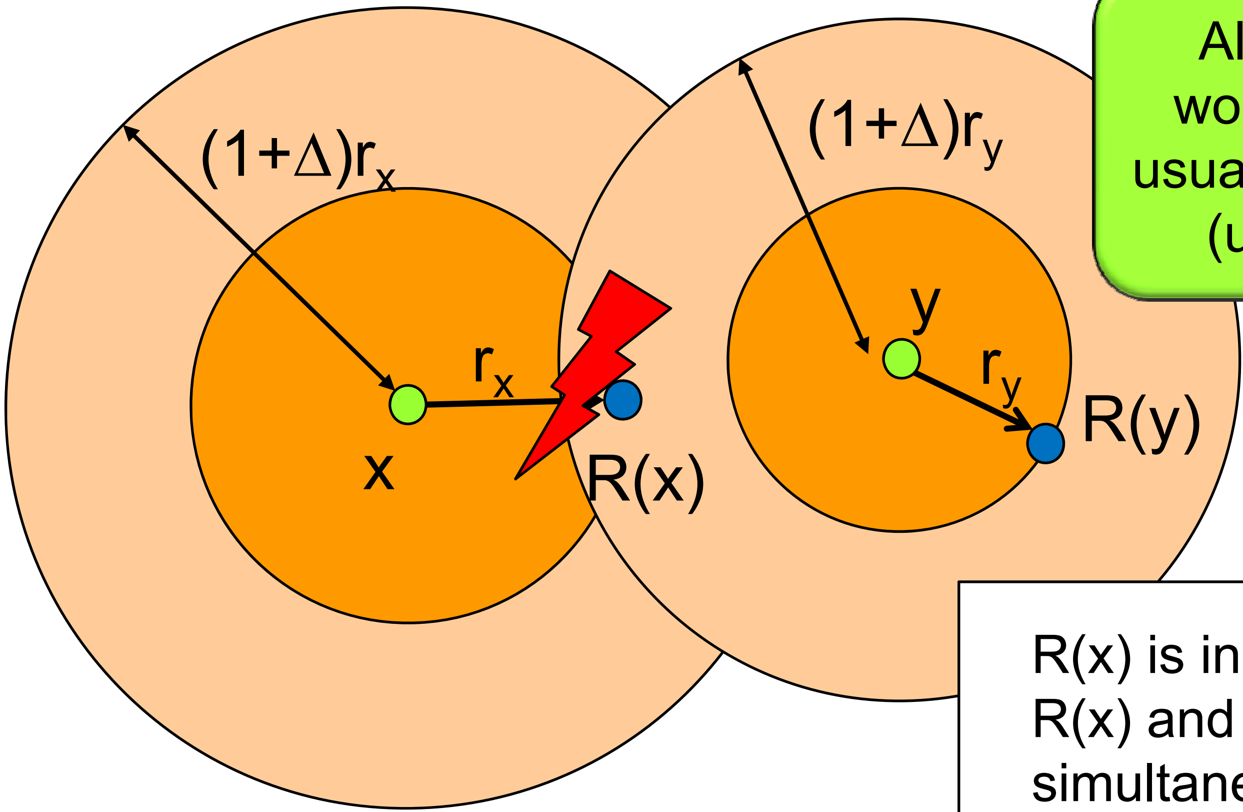

Algorithmic work on worst-case topologies usually in protocol models (unit disk graph,…)

R(x) is in interference range of y
R(x) and R(y) cannot simultaneously receive!

Physical Model

- Based on signal-to-noise-plus-interference (SINR)
- Simplest case:
  → packets can be decoded if SINR is larger than β at receiver

$$\frac{\frac{P_u}{d(u,v)^\alpha}}{N + \sum_{w \in V\setminus\{u\}} \frac{P_w}{d(w,v)^\alpha}} \geq \beta$$

Received signal power from sender

Power level of sender u

Path-loss exponent

Minimum signal-to-interference ratio

Noise

Received signal power from all other nodes (=interference)

Distance between two nodes

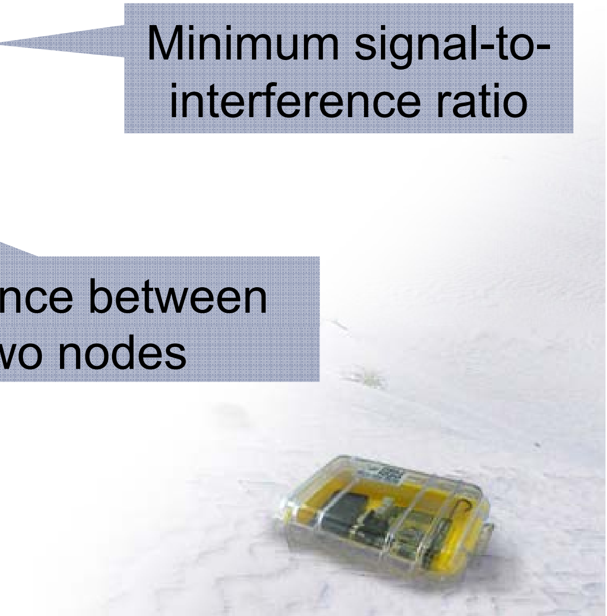

Models

- Two standard models of wireless communication

Protocol Model (graph-based, simpler) ⟷ Physical Model (SINR-based, more realistic)

- Algorithms typically designed and analyzed in protocol model

Premise: Results obtained in protocol model do not divert too much from more realistic model!

Justification:

Capacity results are typically (almost) the same in both models (e.g., Gupta, Kumar, etc...)

Example: Protocol vs. Physical Model

A sends to D, B sends to C

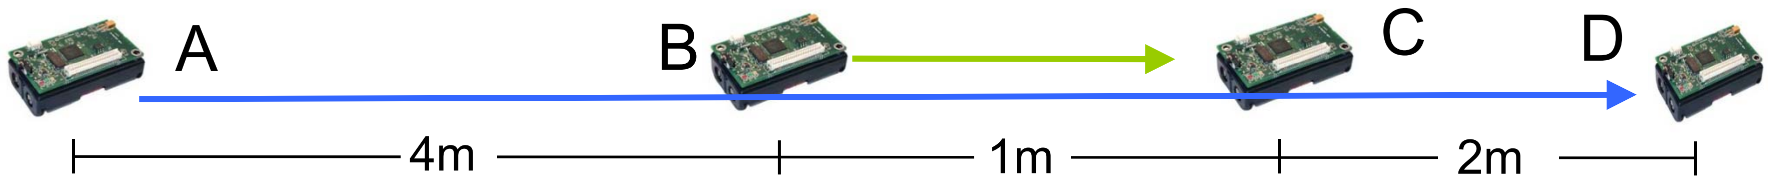

Assume a single frequency (and no fancy decoding techniques!)

Is spatial reuse possible?
- NO — Protocol Model
- YES — Physical Model
- YES — In Reality!

Let α=3, β=3, and N=10nW

Transmission powers: P_B= -15 dBm and P_A= 1 dBm

SINR of A at D: $$\frac{1.26mW/(7m)^3}{0.01\mu W + 31.6\mu W/(3m)^3} \approx 3.11 \geq \beta$$

SINR of B at C: $$\frac{31.6\mu W/(1m)^3}{0.01\mu W + 1.26mW/(5m)^3} \approx 3.13 \geq \beta$$

This works in practice!

- We did measurements using standard mica2 nodes!
- Replaced standard MAC protocol by a (tailor-made) „SINR-MAC"
- Measured for instance the following deployment...

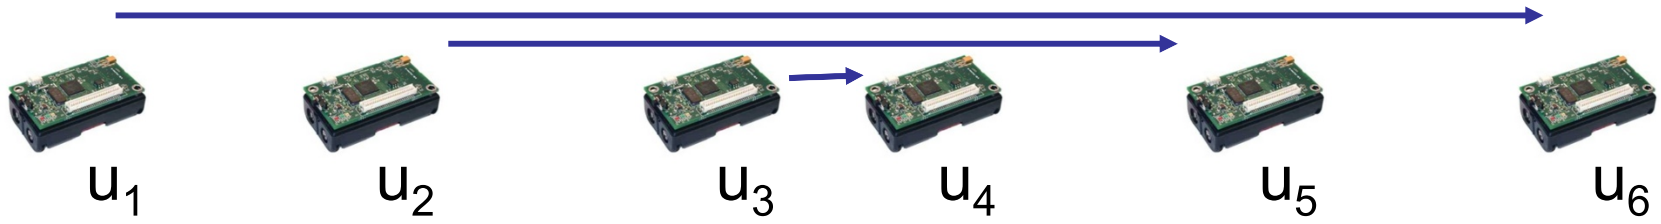

- Time for successfully transmitting 20'000 packets:

	Time required	
	standard MAC	"SINR-MAC"
Node u_1	$721s$	$267s$
Node u_2	$778s$	$268s$
Node u_3	$780s$	$270s$

	Messages received	
	standard MAC	"SINR-MAC"
Node u_4	19999	19773
Node u_5	18784	18488
Node u_6	16519	19498

Speed-up is almost a factor 3

[Moscibroda, Wattenhofer, Weber, Hotnets'06]

Upper Bound Protocol Model

- There are networks, in which at most one node can transmit!
  → like round-robin
- Consider exponential node chain
- Assume nodes can choose arbitrary transmission power

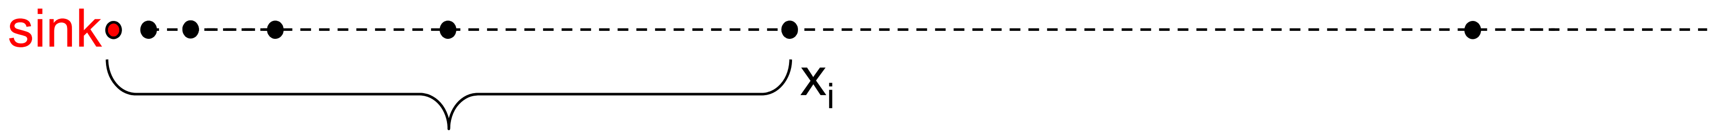

$d(sink, x_i) = (1+1/\Delta)^{i-1}$

- Whenever a node transmits to another node
  → All nodes to its left are in its interference range!
  → Network behaves like a single-hop network

In the **protocol model**, the achievable rate is $\Theta(1/n)$.

Lower Bound Physical Model

- Much better bounds in SINR-based physical model are possible (exponential gap)
- Paper presents a scheduling algorithm that achieves a rate of $\Omega(1/\log^3 n)$

In the **physical model**, the achievable rate is $\Omega(1/\text{polylog } n)$.

- Algorithm is centralized, highly complex → not practical
- But it shows that high rates are possible even in worst-case networks

- Basic idea: Enable spatial reuse by exploiting SINR effects.

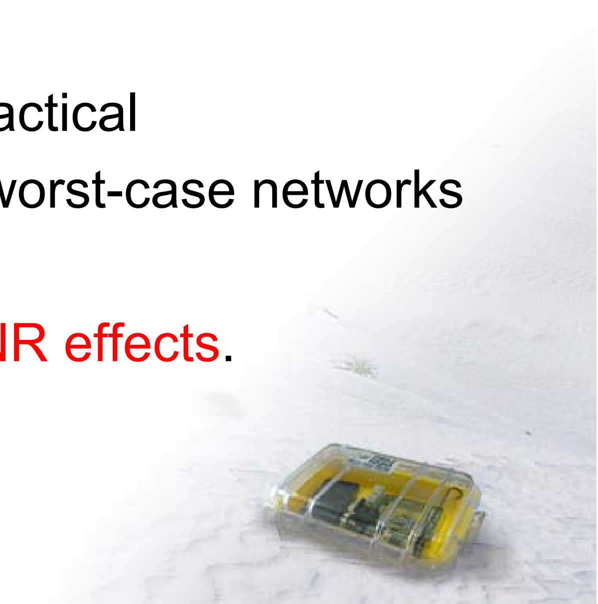

Scheduling Algorithm – High Level Procedure

- High-level idea is simple
- Construct a hierarchical tree T(X) that has desirable properties

1) Initially, each node is active
2) Each node connects to closest active node
3) Break cycles → yields forest
4) Only root of each tree remains active

loop until no active nodes

Can be adjusted if transmission power limited

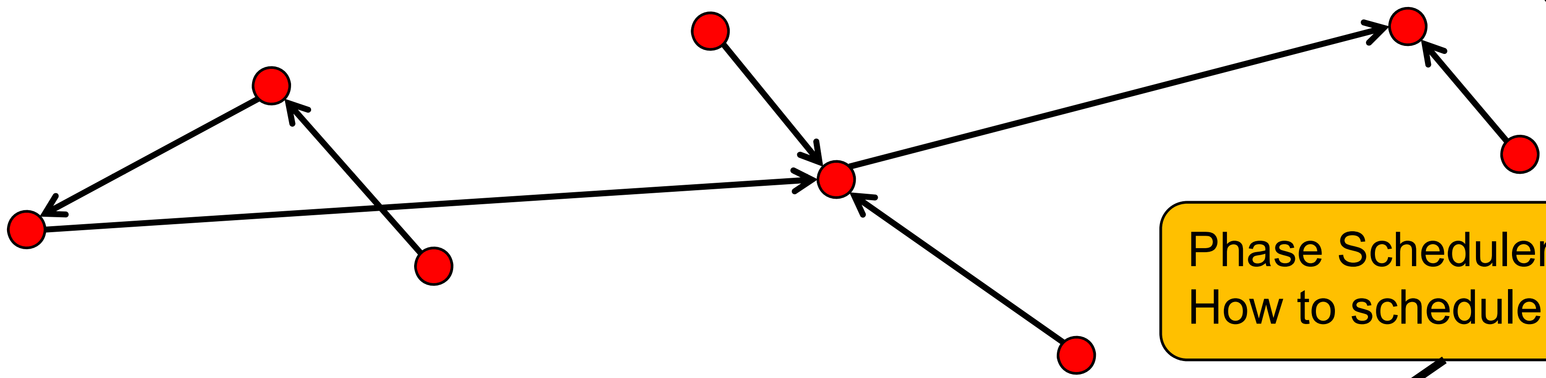

Phase Scheduler:
How to schedule T(X)?

The resulting structure has some nice properties
→ If each link of T(X) can be scheduled at least once in L(X) time-slots
→ Then, a rate of 1/L(X) can be achieved

Scheduling Algorithm – Phase Scheduler

- How to schedule T(X) efficiently
- We need to schedule links of different magnitude simultaneously!
- Only possibility:
  senders of small links must overpower their receiver!

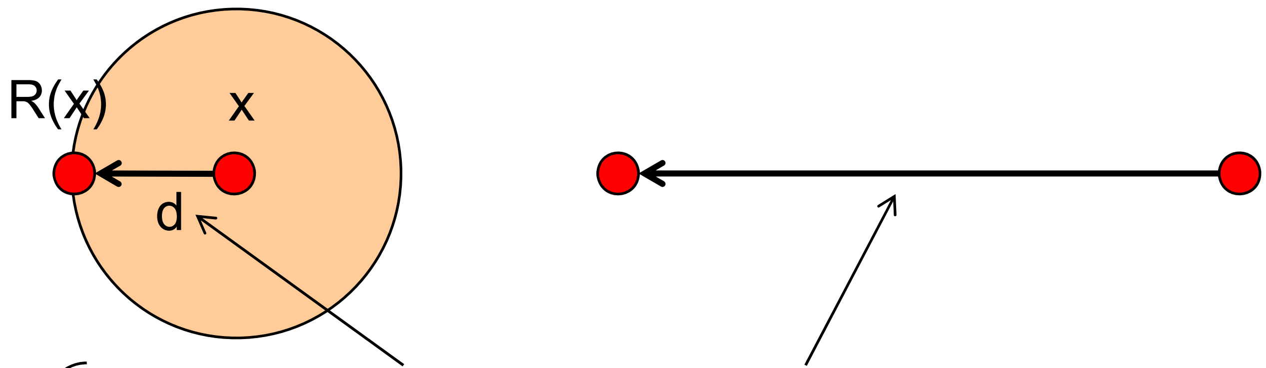

If we want to schedule both links…

1) … R(x) must be overpowered
→ Must transmit at power more than ~d^{α}

2) If senders of small links overpower their receiver…
… their "safety radius" increases (spatial reuse smaller)

Subtle balance is needed!

Scheduling Algorithm – Phase Scheduler

1) Partition links into sets of similar length

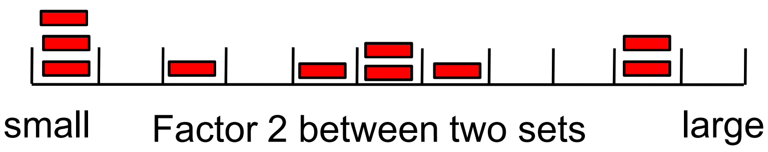

2) Group sets such that links a and b in two sets in the same group have at least $d_a \geq (\xi\beta)^{\xi(\tau a - \tau b)} \cdot d_b$

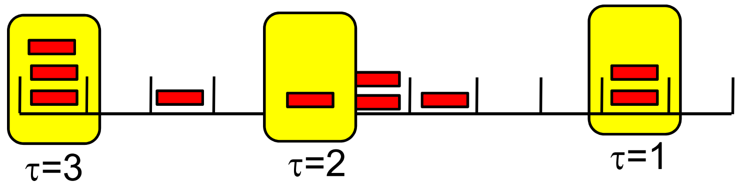

→ Each link gets a τ_{ij} value → Small links have large τ_{ij} and vice versa
→ Schedule links in these sets in one outer-loop iteration
→ Intuition: Schedule links of similar length or very different length

3) Schedule links in a group → Consider in order of decreasing length
(I will not show details because of time constraints.)

Together with structure of T(x) → $\Omega(1/\log^3 n)$ bound

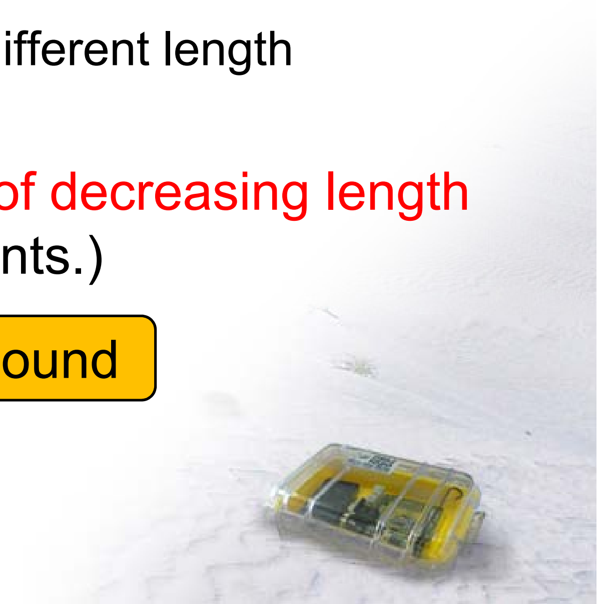

Worst-Case Capacity in Wireless Networks

Worst-Case Capacity

Traditional Capacity

Networks / Model	Max. rate in arbitrary, worst-case deployment	Max. rate in random, uniform deployment
Protocol Model	$\Theta(1/n)$	$\Theta(1/\log n)$
Physical Model	$\Omega(1/\log^3 n)$	$\Omega(1/\log n)$

[Giridhar, Kumar, 2005]

Exponential gap between protocol and physical model!

<u>The Price of Worst-Case Node Placement</u>

- Exponential in protocol model
- Polylogarithmic in physical model (almost no worst-case penalty!)

Possible Applications – Improved "Channel Capacity"

- Consider a channel consisting of wireless sensor nodes
- What is the throughput-capacity of this channel...?

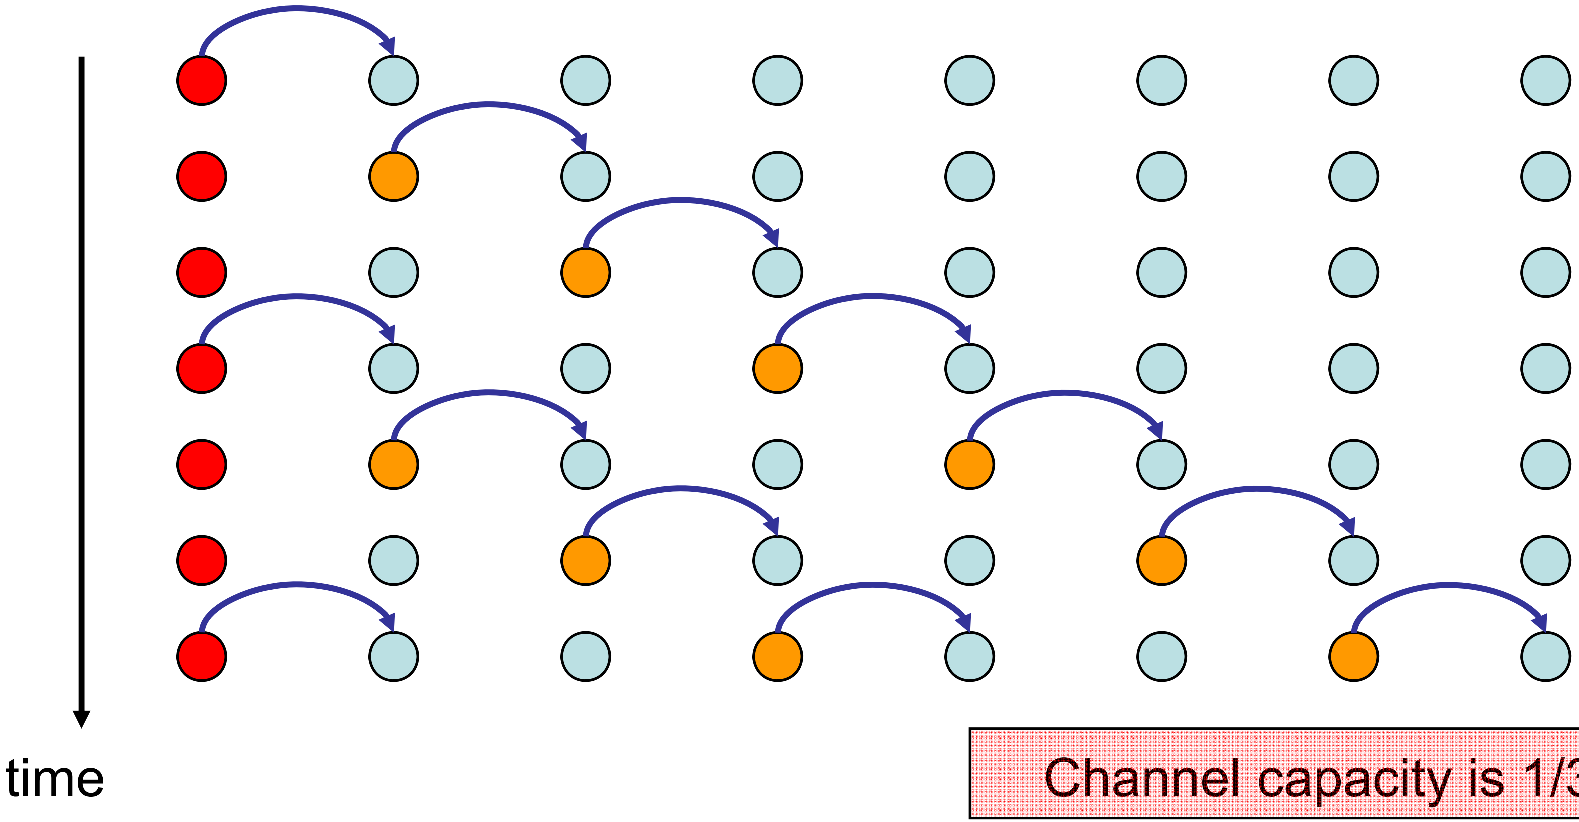

Channel capacity is 1/3

Possible Applications – Improved “Channel Capacity”

- A better strategy...
- Assume node can reach 3-hop neighbor

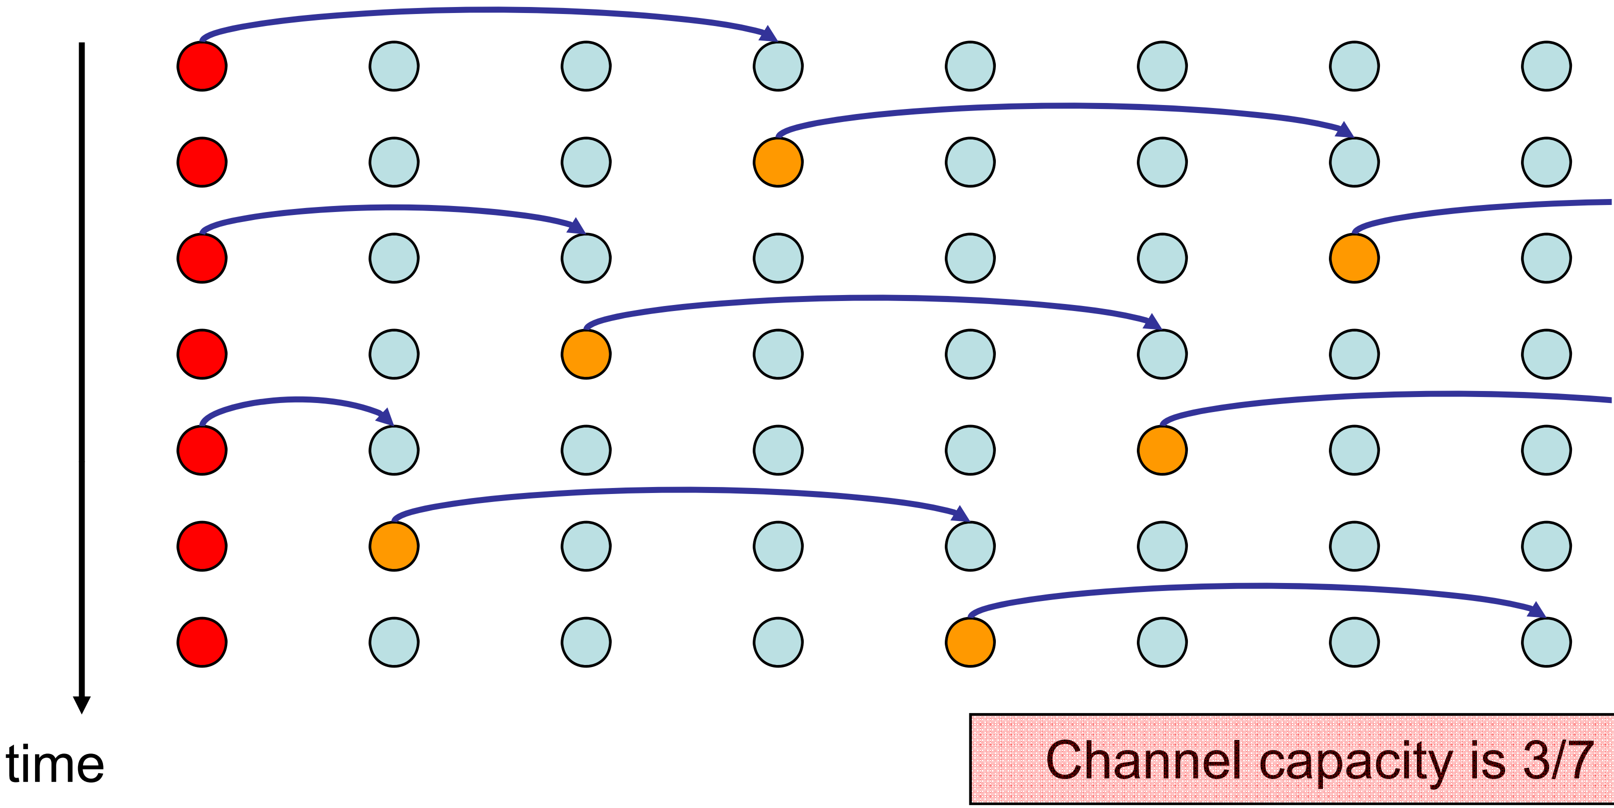

Channel capacity is 3/7

Possible Applications – Improved "Channel Capacity"

- All such (graph-based) strategies have capacity strictly less than 1/2!
- For certain α and β, the following strategy is better!

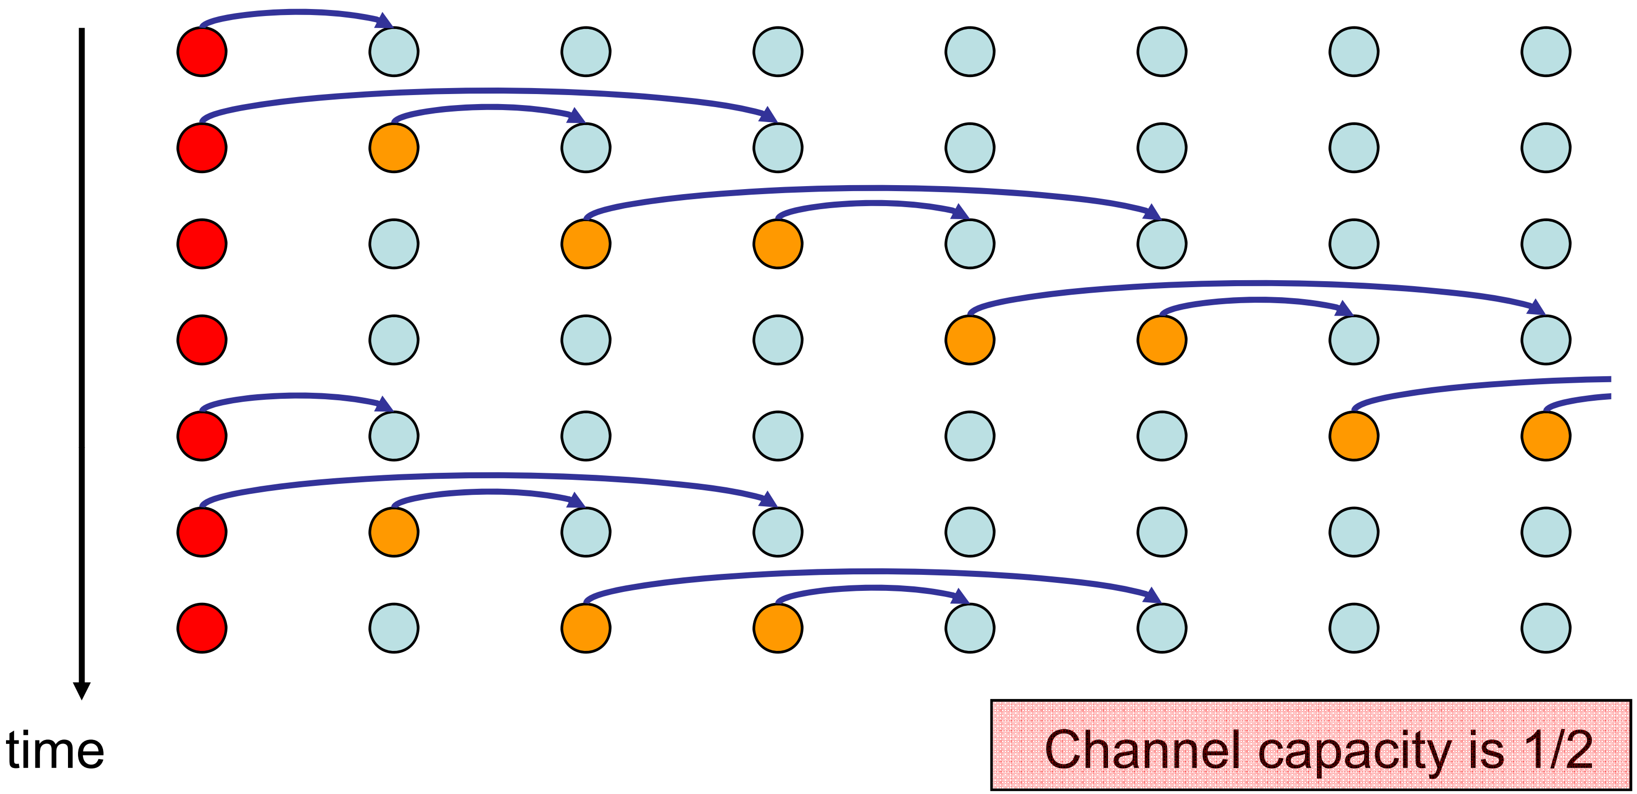

Possible Application – Hotspots in WLAN

- Traditionally: clients assigned to (more or less) closest access point → far-terminal problem → hotspots have less throughput

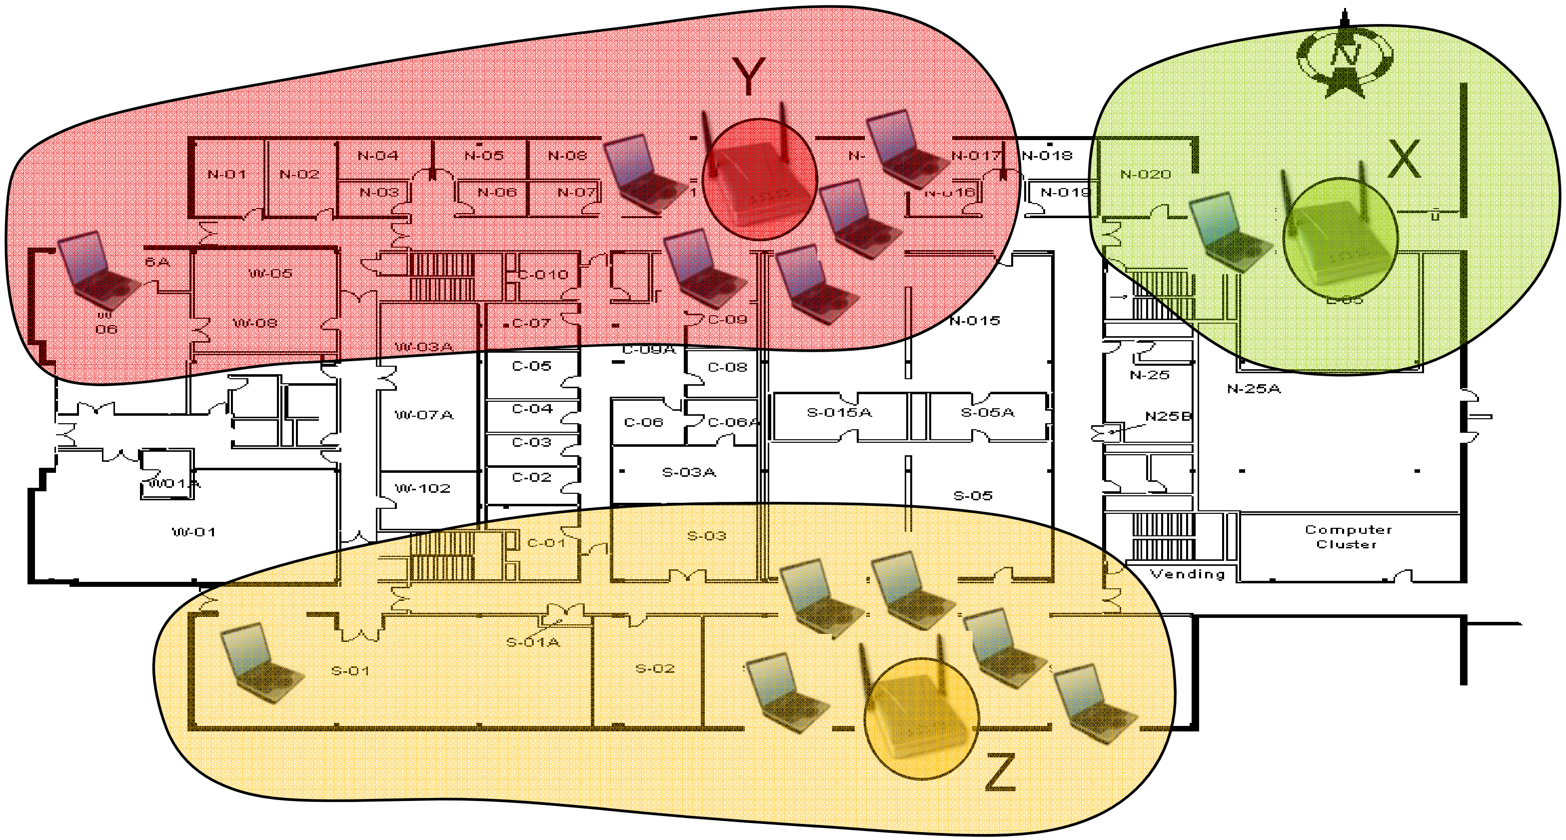

Possible Application – Hotspots in WLAN

- Potentially better: create hotspots with very high throughput
- Every client outside a hotspot is served by one base station

→ Better overall throughput – increase in capacity!

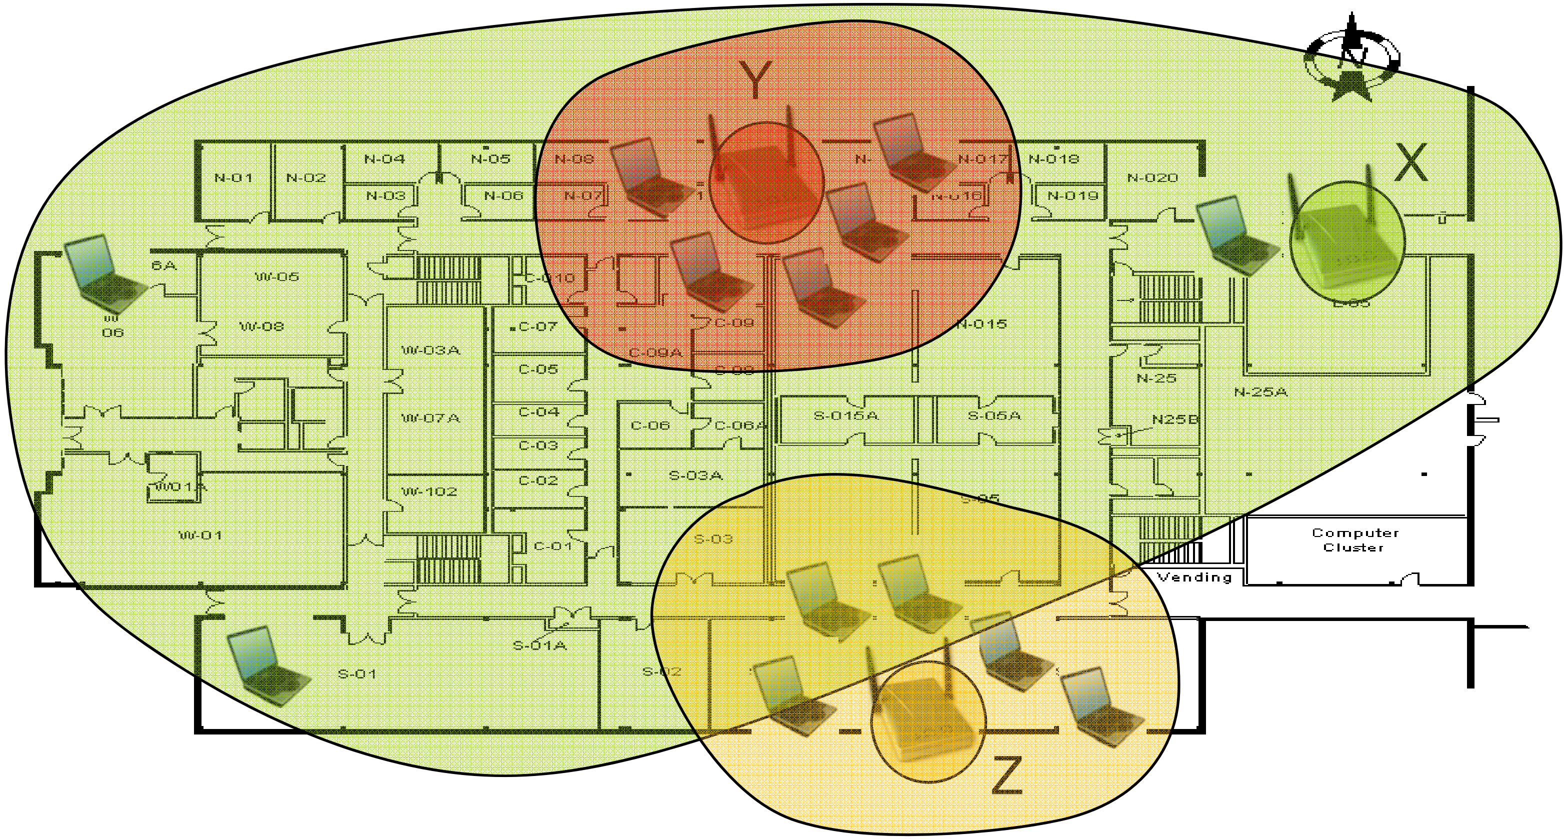

Possible Applications – Data Gathering

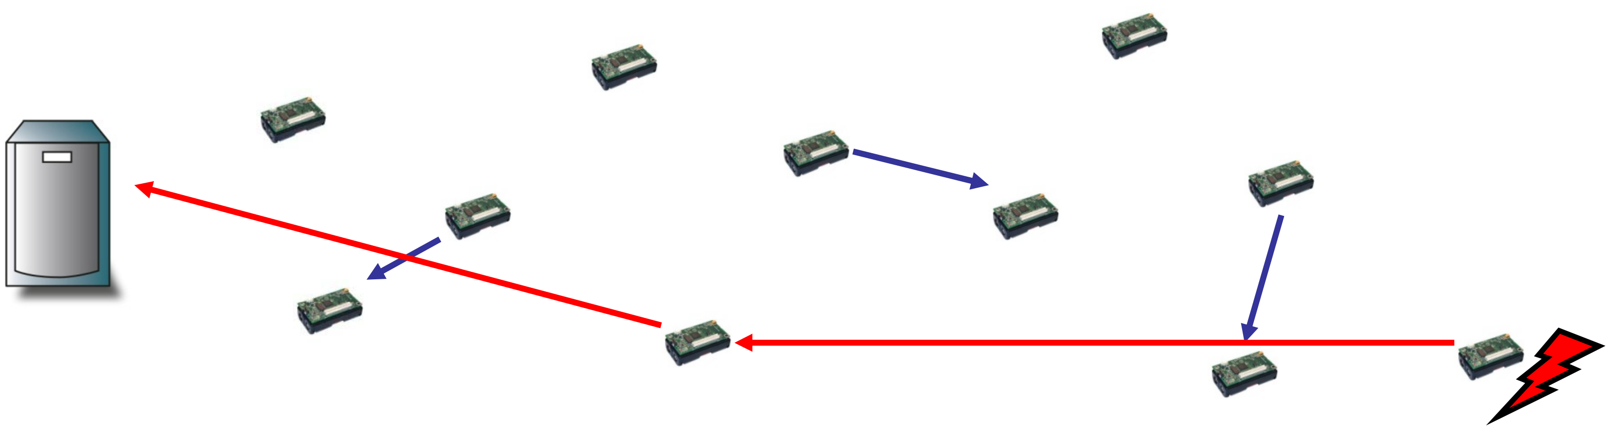

- Neighboring nodes must communicate periodically (for time synchronisation, neighborhood detection, etc…)
- Sending data to base station may be time critical → use long links
- Employing clever power control may reduce delay & reduce coordination overhead!

→ From theory (scheduling) to practice (protocol design)…?

Summary

- Introduce worst-case capacity of sensor networks
  → How much data can periodically be sent to data sink

- Complements existing capacity studies
- Many novel insights

1) Possibilities and limitations of wireless communication
2) Fundamentals of wireless communication models
3) How to devise efficient scheduling algorithms, protocols

Sensor Networks Scale!
Efficient data gathering is possible in every (even worst-case) network!

Protocol Model Poor!
Exponential gap between protocol and physical model!

Efficient Protocols!
Must use SINR-effects and power control to achieve high rate!

Remaining Questions…?

- My talk so far was based on the paper Moscibroda & W, The Complexity of Connectivity in Wireless Networks, Infocom 2006

- The paper was more general than my presentation
  - It was not about data gathering rate, but rather…
  1. Given an arbitrary network
  2. Connect the nodes in a meaningful way by links
  3. Schedule the links such that the network becomes strongly connected

- Question: Given *n* communication requests, assign a color (time slot) to each request, such that all requests sharing the same color can be handled correctly, i.e., the SINR condition is met at all destinations (the source powers are constant). The goal is to minimize the number of colors.

Is this a difficult problem?

Scheduling Wireless Links: How hard is it?

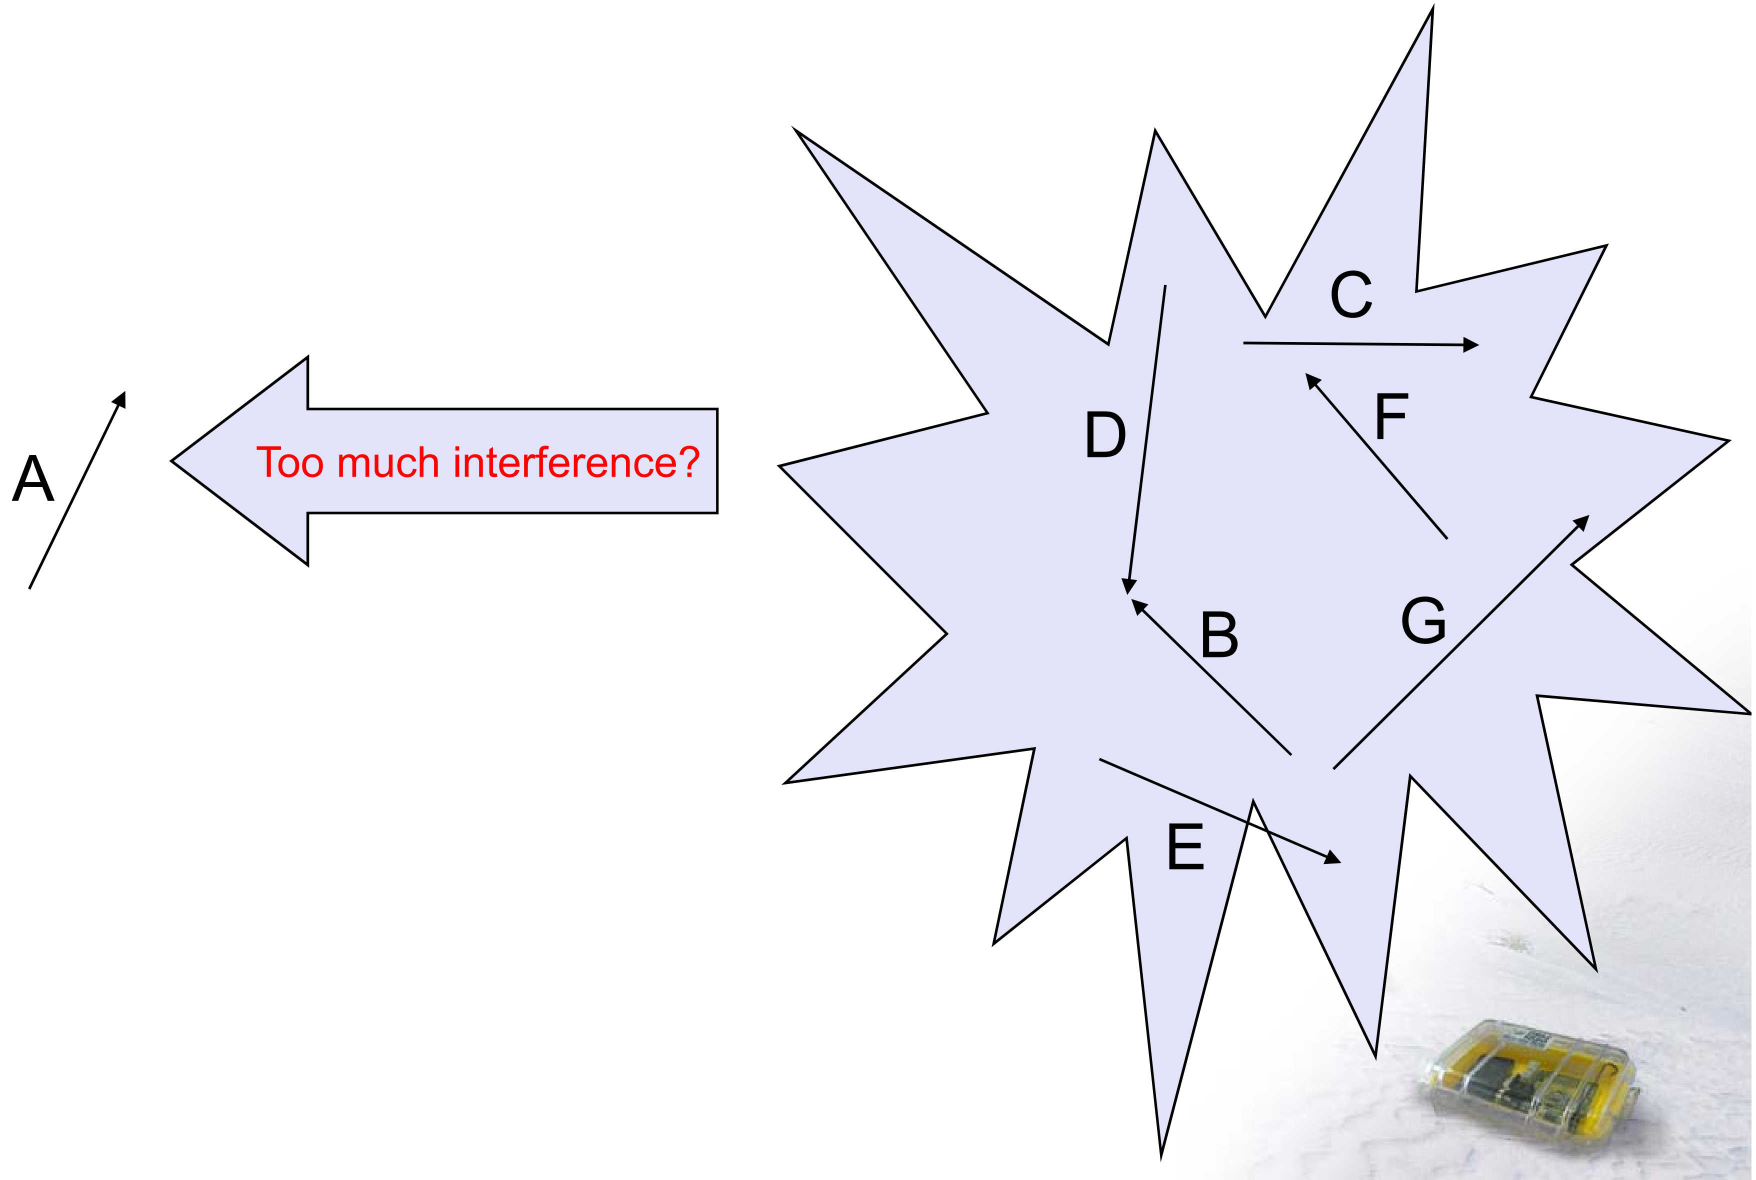

Scheduling: Problem Definition

- P: constant power level
- L: set of communication requests
- S: schedule $S = \{S_1, S_2, \ldots, S_T\}$

- Interference Model: *SINR*
  - A: path-loss matrix, defined for every pair of nodes

$$SINR(s,r) = \frac{\frac{P}{A_{sr}}}{N + \sum_{v \in V, v \neq s} \frac{P}{A_{vr}}} \geq \beta$$

Received signal power from sender

Min. SINR threshold

Ambient noise

Received signal power from all other nodes (Interference!)

- **Problem statement:**

> *Find a minimum-length schedule **S**, s.t. every link in **L** is scheduled in at least one time slot **t, 1≤t ≤T**, and all concurrently scheduled receivers in **S_t** satisfy the **SINR** constraints.*

“Scheduling as hard as coloring” … not really!

“The Wall Model”: Now only adjacent links interfere! (Has been shown to be as hard as coloring [Bjoerklund 2003])

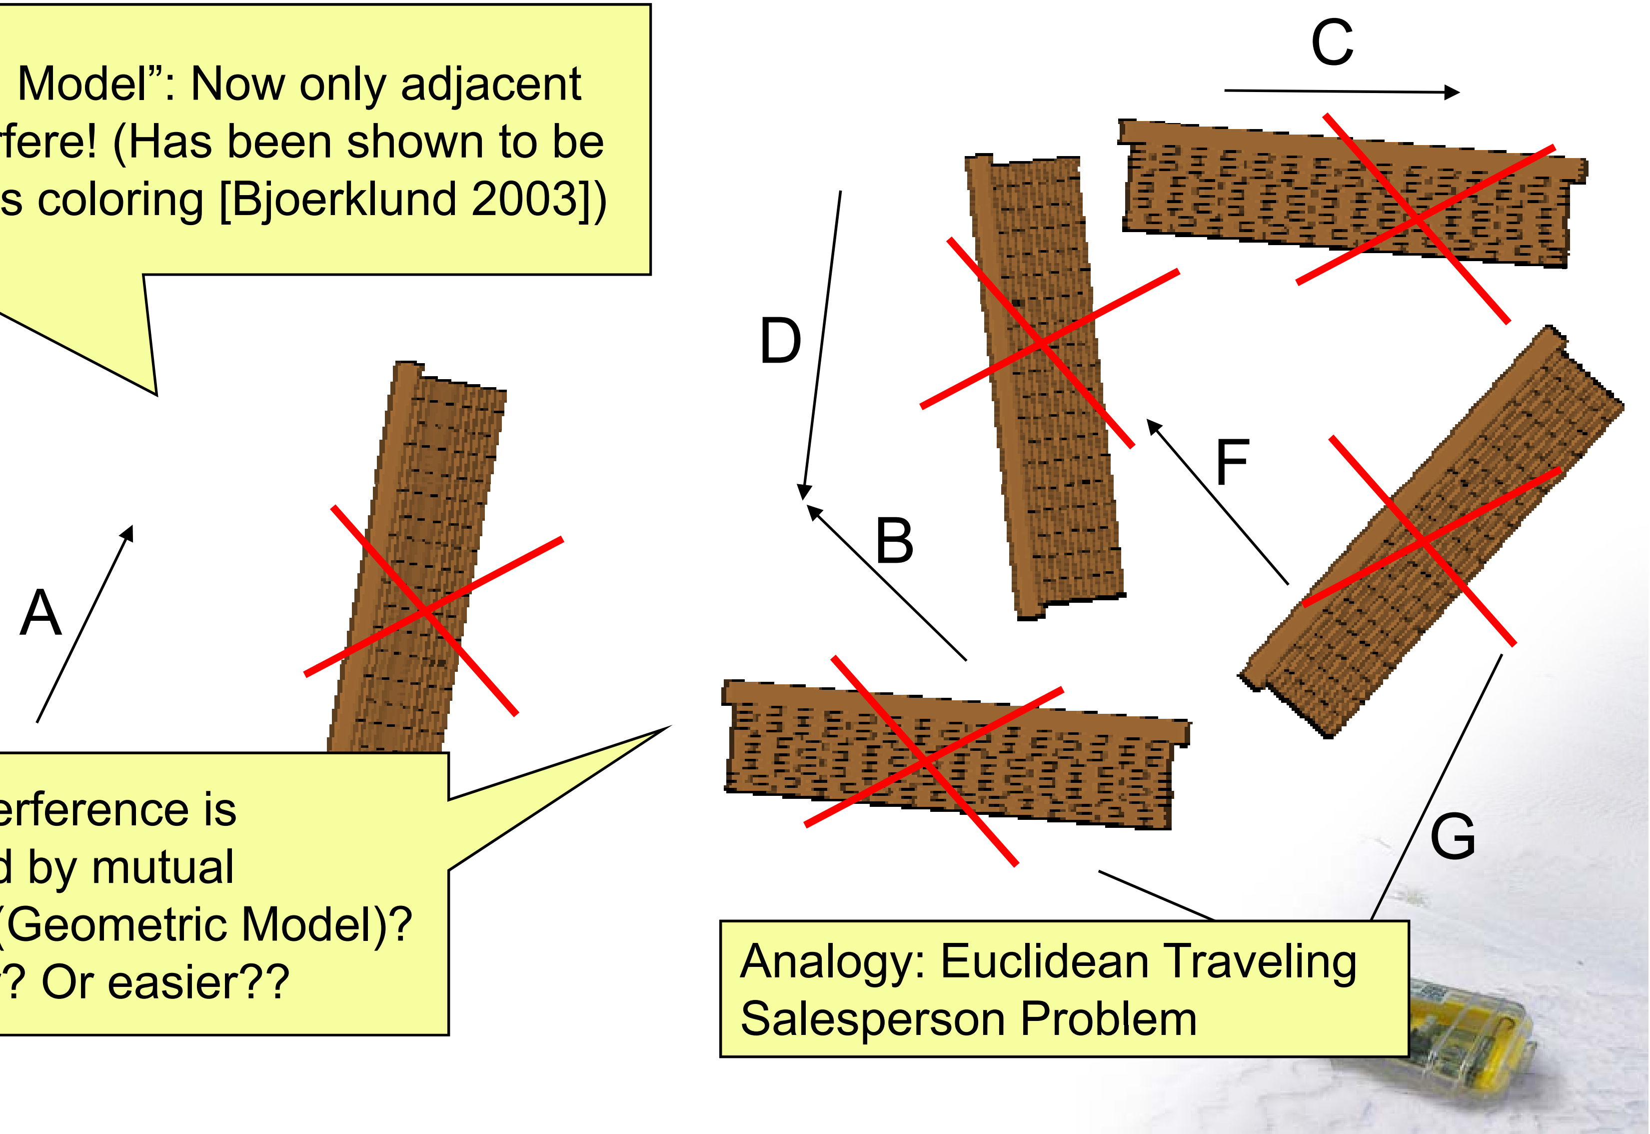

What if interference is determined by mutual distances (Geometric Model)? Is it harder? Or easier??

Analogy: Euclidean Traveling Salesperson Problem

Scheduling: Reduction from Partition

- Partition problem (NP-Complete [Karp 1972]):
  - Given a set of integers *I*, find two subsets of integers I_1, I_2, s.t.:
- Decision version of Scheduling: $T \leq 2$:
  - Consider a set of integers *I*, whose elements sum up to σ:

$$I_1, I_2 \subset I = \{i_1, \ldots, i_n\}$$
$$I_1 \cap I_2 = \emptyset,$$
$$I_1 \cup I_2 = I,$$
$$\sum_{i_j \in I_1} i_j = \sum_{i_j \in I_2} i_j = \frac{1}{2} \sum_{i_j \in I} i_j.$$

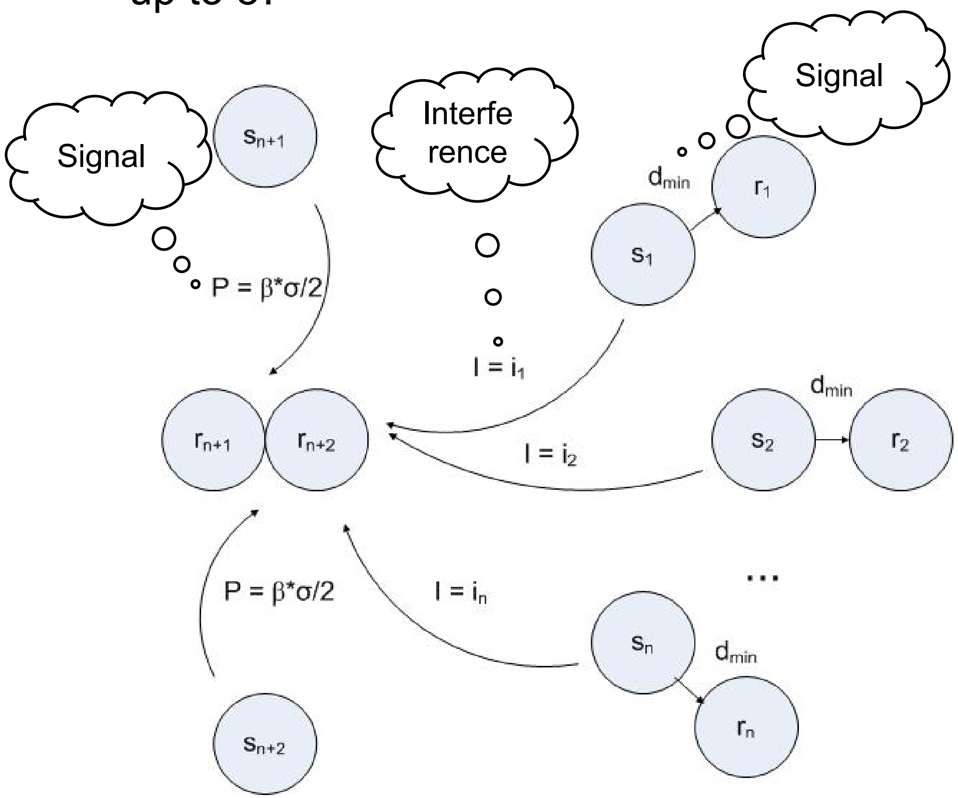

$$SINR_{r_{n+1}} = \frac{\beta \cdot \frac{\sigma}{2}}{\sum_{i_j \in I_1} i_j} = \beta$$

Schedule with time $T \leq 2 \leftrightarrow$ Partition

SINR Models

- Abstract SINR
  - Arbitrary path loss matrix
  - No notion of triangle inequality
  - If an algorithm works here, it works everywhere!
  - Best model for upper bounds

- Geometric SINR
  - Nodes are points in plane
  - Path loss is function of distance
  - If an impossibility result holds here, it holds everywhere!
  - Best model for lower bounds

too pessimistic ⟷ **too optimistic**

- Reality is here
  - Path loss roughly follows geometric constraints, but there are exceptions
  - Open field networks are closer to Geometric SINR
  - With more walls, you get more and more Abstract SINR

Overview of results so far

- Moscibroda, W, Infocom 2006
  - First paper in this area, O($\log^3 n$) bound for connectivity, and more
- Moscibroda, W, Weber, HotNets 2006
  - Practical experiments, ideas for capacity-improving protocol
- Goussevskaia, Oswald, W, MobiHoc 2007
  - Hardness results & constant approximation for constant power

- Moscibroda, W, Zollinger, MobiHoc 2006
  - First results beyond connectivity, namely in the topology control domain
- Moscibroda, Oswald, W, Infocom 2007
  - Generalizion of Infocom 2006, proof that known algorithms perform poorly
- Chafekar, Kumar, Marathe, Parthasarathy, Srinivasan, MobiHoc 2007
  - Cross layer analysis for scheduling and routing
- Moscibroda, IPSN 2007
  - Connection to data gathering, improved O($\log^2 n$) result
- Goussevskaia, W, FOWANC 2008
  - Hardness results for analog network coding
- Locher, von Rickenbach, W, ICDCN 2008
  - Still some major open problems

Main open question in this area

- Most papers so far deal with special cases, essentially scheduling a number of links with special properties. The general problem is still wide open:

- A communication request consists of a source and a destination, which are arbitrary points in the Euclidean plane. Given *n* communication requests, assign a color (time slot) to each request. For all requests sharing the same color specify power levels such that each request can be handled correctly, i.e., the SINR condition is met at all destinations. The goal is to minimize the number of colors.

- E.g., for arbitrary power levels not even hardness is known…

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich