


Wireless Sensor Networks for Emergency Medical Care

Matt Welsh, David Malan, Breanne Duncan, and
Thaddeus Fulford-Jones

Harvard University

Steve Moulton


Boston University School of Medicine



Introduction: Sensor Networks



WeC (1999)



Rene (2000)



Dot (2001)

Exciting emerging domain of deeply networked systems


- Low-power, wireless “motes” with tiny amount of CPU/memory
- Large federated networks for high-resolution sensing of environment

Drive towards miniaturization and low power

- Eventual goal - complete systems in 1 mm^3 , MEMS sensors
- Family of Berkeley motes as COTS experimental platform




MICA (2002)



Speck (2003)

The Berkeley MICA Mote



- ATMEGA 128L (7.3 Mhz 8-bit CPU)
- 128 KB code, 4 KB data SRAM
- 512 KB flash for logging
- 433 or 916 Mhz, 76.8 Kbps radio (100m max)
- Sandwich-on sensor boards
- Powered by 2AA batteries

Several thousand produced, used by 100s of research groups

- Get yours at www.xbow.com, about \$150 a pop

Great platform for experimentation (though not particularly small)

- Easy to integrate new sensors & actuators
- 15-20 mA active (5-6 days on 2 AAs)
- 15 μ A sleeping (21 years, but limited by shelf life of battery!)

Typical Applications



Vehicle tracking

- Sensors take magnetometer readings, localize object
- Communicate using geographic routing to base station
- Robust against node and radio link failures



Habitat monitoring – Great Duck Island

- Gather temp, IR, humidity, and other readings from bird nests on island
- Determining occupancy of nests to understand breeding & migration behavior
- Live readings at www.greatduckisland.net



Potential Medical Applications

Real-time, continuous patient monitoring

- Pre-hospital, in-hospital, and ambulatory monitoring possible
- Replace expensive and cumbersome wired telemetry systems

Home monitoring for chronic and elderly patients

- Collect periodic or continuous data and upload to physician
- Allows long-term care and trend analysis
- Reduce length of hospital stay

Collection of long-term databases of clinical data

- Correlation of biosensor readings with other patient information
- Longitudinal studies across populations
- Study effects of interventions and data mining



Mass Casualty Events

Large accidents, fires, terrorist attacks

- Normal organized community support may be damaged or destroyed
- Large numbers of patients, severe load on emergency personnel

Manual tracking of patient status is difficult

- Current systems are paper, phone, radio based

Sensor nets have potential for large impact

- Real-time, continuous vital monitoring
- “Electronic triage tag” to store patient data
- Immediate alerts of changes in patient status
- Relay data to hospital, correlate with pt. records



CONTAMINATED

Personal Property Receipt/ Evidence Tag
 Destination _____
 Via _____

TRIAGE TAG

S L U D G E M
Calculator Landmeter Urinometer Defibrillator G.I. Distress Emesis Mics

AUTO INJECTOR 1 2 3 4 5

Yes No Gases Decon
 Yes No Secondary Decon

Solution _____

Blunt Trauma
 Burn
 D. Saine
 Cardiac
 Crushing
 Fracture
 Laceration
 Penetrating Injury

Other: _____

VITAL SIGNS

Time	B/P	Pulse	Respiration

Time	Drug	Solution	Dose

CONTAMINATED

Comments/Information

Patient's Name _____

RESPIRATIONS Yes + 2 Sec. Can Do
 No - 2 Sec. Can't Do

MENTAL STATUS M Can Do Can't Do

Move the Walking Wounded ▶ **MINOR**

No Respirations After Head Tilt ▶ **MORGUE**

Resorations - Over 30 ▶ **IMMEDIATE**

Perfusion - Capillary Refill Over 2 Seconds ▶ **IMMEDIATE**

Mental Status - Unable to Follow Simply Commands ▶ **IMMEDIATE**

Otherwise ▶ **DELAYED**

POWERED BY: 1996 Disaster Management Systems, Inc. • Pomona, CA 92568-0400 • www.1996tag.com

PERSONAL INFORMATION

NAME _____
 ADDRESS _____
 CITY _____ ST _____ ZIP _____
 PHONE _____
 COMMENTS _____ RELIGIOUS PREF _____



MORGUE
 Pulseless/Non-Breathing

IMMEDIATE Life Threatening Injury	IMMEDIATE Life Threatening Injury
DELAYED Serious Non Life Threatening	DELAYED Serious Non Life Threatening
MINOR Walking Wounded	MINOR Walking Wounded

EVIDENCE

EVIDENCE

VitalDust: Wireless vital sign monitoring



Mica2-based pulse oximeter using BCI, Inc. OEM board

- Measures heart rate, blood oxygen saturation
- Mote-based ECG currently under development
- PDA- and PC-based applications for multi-patient triage
- Integration with iRevive, PDA-based patient care record system for EMTs

Sensor Network Challenges

Low computational power

- Current mote processors run at < 10 MIPS
- Not enough horsepower to do real signal processing
 - *DSP integration may be possible*
- 4 KB of memory not enough to store significant data

Poor communication bandwidth

- Current radios achieve about 10 Kbps per mote
- Note that raw channel capacity is much greater
 - *Overhead due to CSMA backoff, noise floor detection, start symbol, etc.*
- 802.15.4 (Zigbee) radios now available at 250 Kbps
 - *Not clear how much bandwidth available to applications*

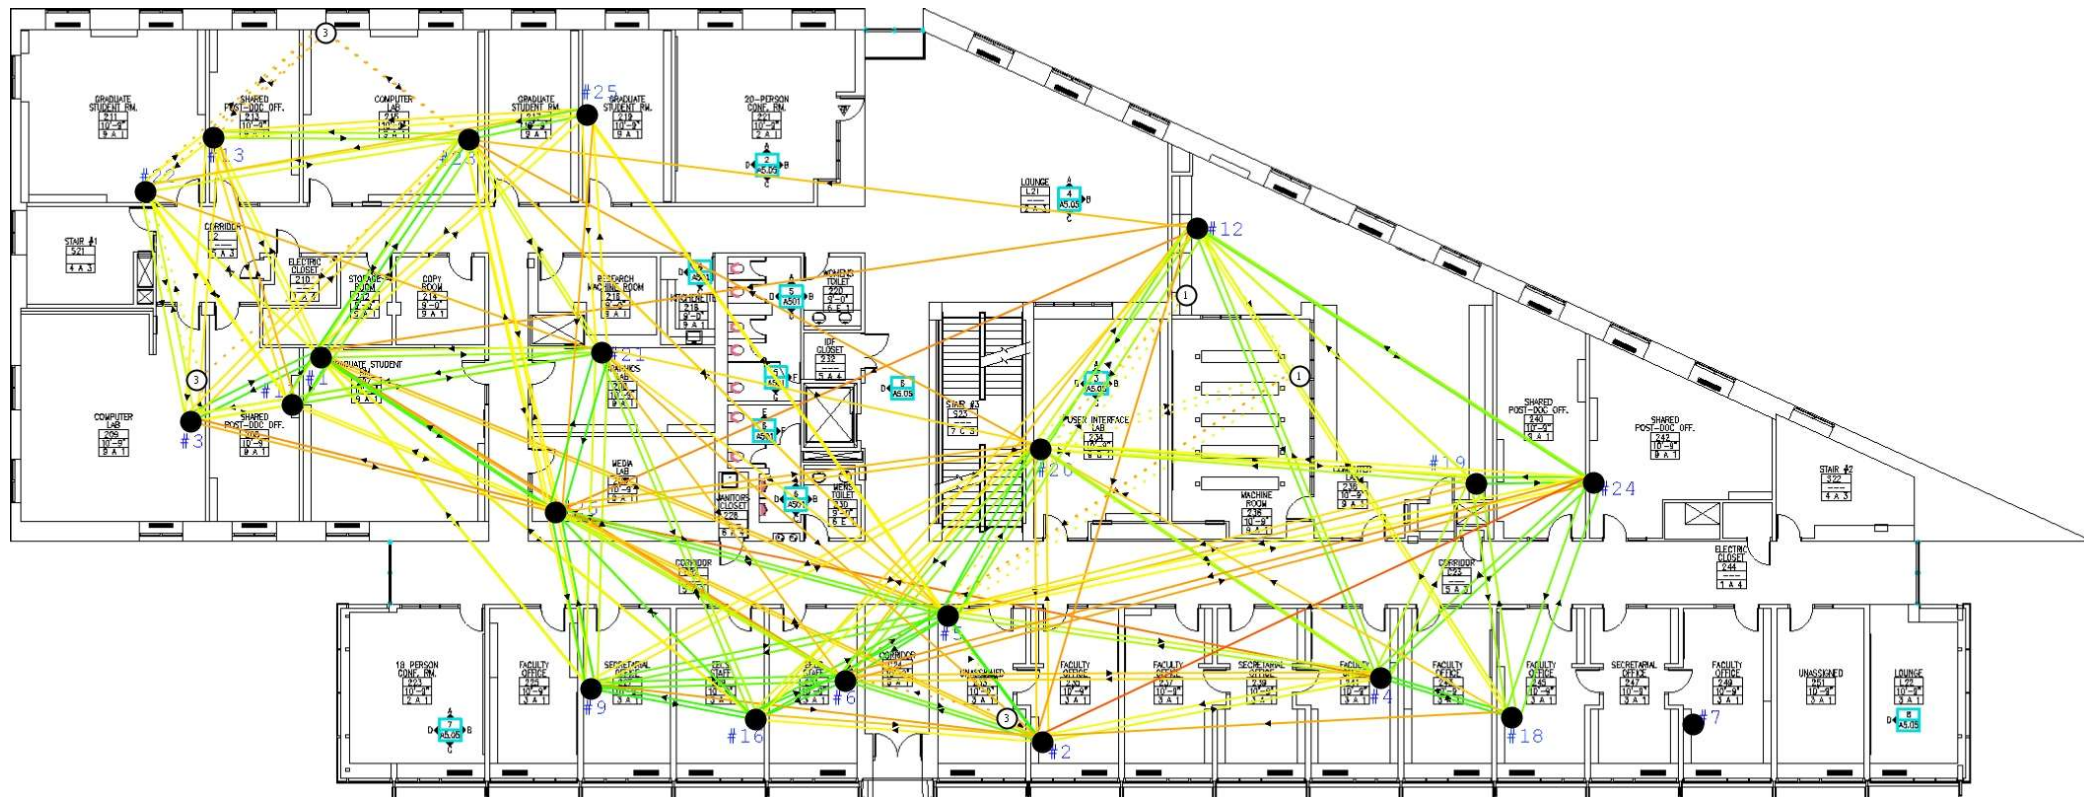
Limited energy budget

- 2 AA motes provide about 2850 mAh
- Coin-cell Li-Ion batteries provide around 800 mAh
- Solar cells can generate around 5 mA/cm² in direct sunlight
- Must use low duty cycle operation to extend lifetime beyond a few days

Sensor Net Challenges

Radio connectivity is highly volatile!

- Packet loss not well correlated with distance
- Affected by receiver sensitivity, wall attenuation, antenna orientation, etc.
- Many links are asymmetric!



motelab.eecs.harvard.edu

Sensor Net Challenges

Multihop routing is extremely dynamic


- Nodes must continuously evaluate link quality of neighbors
- Collect packet loss statistics over time
- Periodically broadcast “link report”

Time synchronization

- Complex link dynamics make this difficult
- Scaling and overhead of time sync traffic is an issue

Localization

- Commonly use ultrasound time-of-flight ranging
- Send RF and ultrasound pulse simultaneously
- Lots of difficulties arise due to lack of receiver calibration
- RF-signature based localization achieves indoor accuracy of 2-3 meters



CodeBlue: Motivation

Medical applications have a unique set of demands!

Must be extremely robust

- System cannot break down under stress
- Communication paths must have low latencies

Must scale to very dense networks


- Many nodes, high degree of mobility
- Complex radio link dynamics

Must preserve privacy of medical data (HIPAA Privacy Rule)


- Cannot disclose medical data or patient identification indiscriminately
- Different health care providers have different access rights

CodeBlue Architecture

Vital sign sensors and active tags



Location beacons



CodeBlue Information Plane

Naming <i>Discovery</i>	Authentication <i>Encryption</i>	Event Delivery	Filtering <i>Aggregation</i>	Handoff
----------------------------	-------------------------------------	----------------	---------------------------------	---------



Other hospital information systems



Wireless PDAs and fixed terminals



Ambulance MDTs

CodeBlue Architecture

Scalable, robust “information plane” for critical care

- Ad hoc, any-to-any routing with dynamic discovery of routes
- Runs across a range of devices, from motes to PDAs to PCs

Publish/subscribe data delivery model

- Sensor nodes publish vital signs, location, identity
- Rescue/medical personnel subscribe to data of interest
- In-network filtering and aggregation of data to limit bandwidth and information overload

Reliable delivery of critical data

- Content-based prioritization
 - *e.g., Patient stops breathing or loss of network connectivity*
- Scale transmit power to limit interference or issue “SOS” messages

Decentralized authentication and security

- Handoff of credentials across rescue personnel
- Seamless access control across patient transfers

Emergency Response Setting

Notes attached to patients
collect vital signs (pulse ox, heart rate, etc.)



Ambulance system makes
triage decisions, relays to EMTs



PDAs carried by EMTs
receive vital signs and enter
into field report



Correlate with patient records
at hospital

Research Challenges

Scalable, flexible routing infrastructure

- Many existing ad hoc routing schemes are connection-oriented
- Much work in sensor networks focused on many-to-one data collection
- We require more flexible naming and pub/sub semantics

Rapid, robust, ad hoc deployment

- Must operate without external network or computational infrastructure
- Decision making must be fully distributed
- Zero administrative overhead for setup and configuration

Coping with enormous ranges of density and node volatility

- Must scale to very high node densities
- Communication must adapt to widely varying network conditions

Lightweight, decentralized security mechanisms

- Sensor nodes are too primitive for expensive public key approaches
- Still must support flexible security policies

ECC-based Public Key Encryption

TinySec: Skipjack-based symmetric key encryption

- Extremely lightweight and works well on motes

How to perform secure key exchange?

- Traditional public-key systems are extremely expensive
 - *No floating point processor on the motes*
- Diffie-Hellman key exchange with 160-bit exponent:
 - *Requires 31 sec (768-bit modulus) or 57.9 sec (1024-bit modulus)*

Elliptic Curve Cryptography (ECC) to the rescue?

- Possible to implement with integer arithmetic only
- Smaller modulus size (163 bits)
 - *First implementation: Fast but memory intensive: 1.7 sec for 33-bit keys*
 - *Second implementation: Lean but slow: 466 sec for 163-bit keys*
- Currently retooling code, significant optimizations possible

Power-Aware Routing

Dynamically adjust transmit power of individual nodes

- Attempt to reduce energy consumption and avoid interference
- Goal: Get “good enough” connectivity to next hop in route

CC1000 supports wide range of transmission power levels

- 1 μW (4.65 mA) to 10 mW (22.41 mA)
- Can adjust Tx power level in software

Spanning-tree-based multihop routing

- Based on TinyOS “Surge” protocol
- Nodes collect link quality assessments from neighbors, periodically broadcast

Dynamic power scaling algorithm

- Reduce transmission power if link quality to parent is above threshold
- Increase Tx power if below low-water mark

MoteTrack: RF-Based Localization

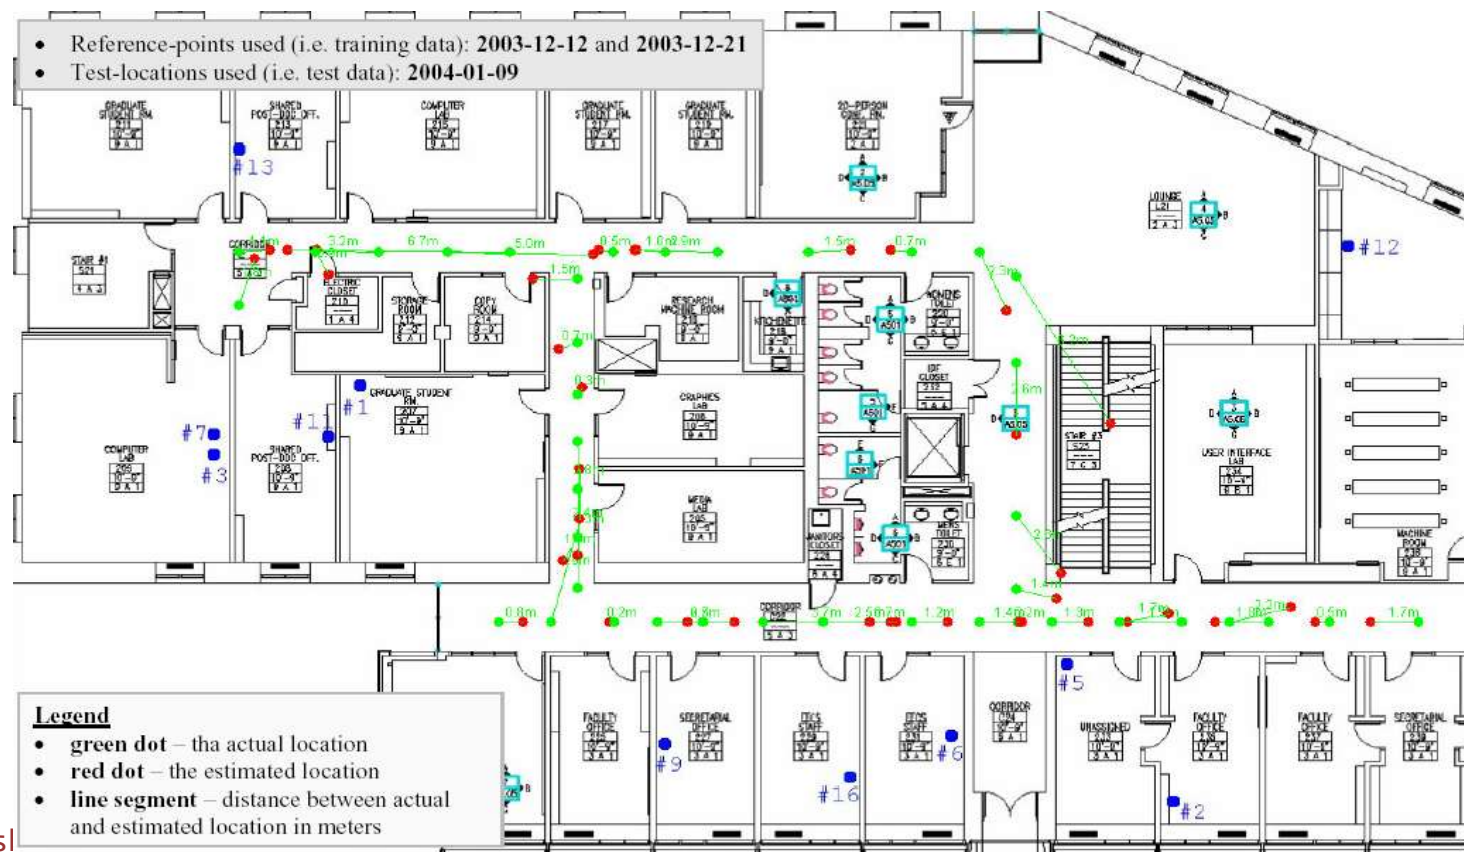
Collect RF signal “signatures” from various points in building

- Use MoteLab testbed with 30 beacon nodes
- Similar to RADAR scheme for 802.11 networks, with much higher density

Nodes compute location by comparing to stored signatures

- Centroid of weighted signature distance from known points

Good results: 80th percentile error of 2 meters



The Hourglass Data Collection Network

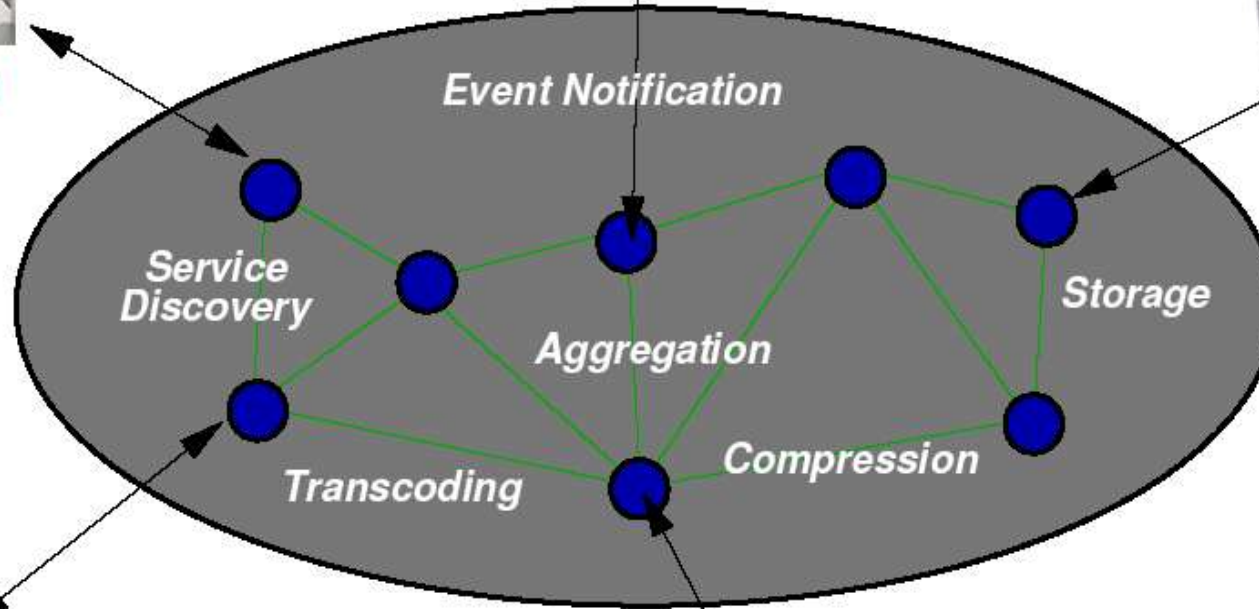


911 dispatch



EMS dispatch

Hospital staff



Hospital information systems

CodeBlue Network at disaster site



Hourglass Architecture

Internet-based event delivery and aggregation network

- Based on peer-to-peer overlay architecture: robust to node and link failures
- Hospitals, EMS, 911 dispatch, ambulances act as publishers and subscribers
 - *Coordinate disaster response across multiple entities*

In-network services

- Ability to push filtering, compression, aggregation, and storage into network
- Dynamically optimize network paths based on bandwidth, latency, and load
- High degree of node heterogeneity

Intermittent connectivity at the edge

- Data sources and sinks may be temporarily disconnected
 - *e.g., Ambulance using cellular link*
- Must provide seamless operation when rejoining network

Conclusions

Medical applications pose many challenges for wireless sensors

- Robustness
- Scalability
- Security

CodeBlue infrastructure for emergency response

- An ad hoc, wireless infrastructure for vital sign monitoring and triage

Current status

- Initial design of CodeBlue complete
- Prototypes of several major components completed
- Forming collaborations with several Boston area hospitals

For more information:

- <http://www.eecs.harvard.edu/~mdw> or mdw@eecs.harvard.edu