CIS 3990

Mobile and IoT Computing

https://penn-waves-lab.github.io/cis3990-24spring

Lecture 10: Network Connectivity

Instructor: Mingmin Zhao (mingminz@cis.upenn.edu)

TA: Haowen Lai (hwlai@cis.upenn.edu)

Some material adapted from Hari Balakrishnan (MIT) and Fadel Adib (MIT)

Mobile and IoT Computing

The convergence of sensing, communication, and computation that allows us to:

Acquire data from the environment

Aggregate data from multiple sources

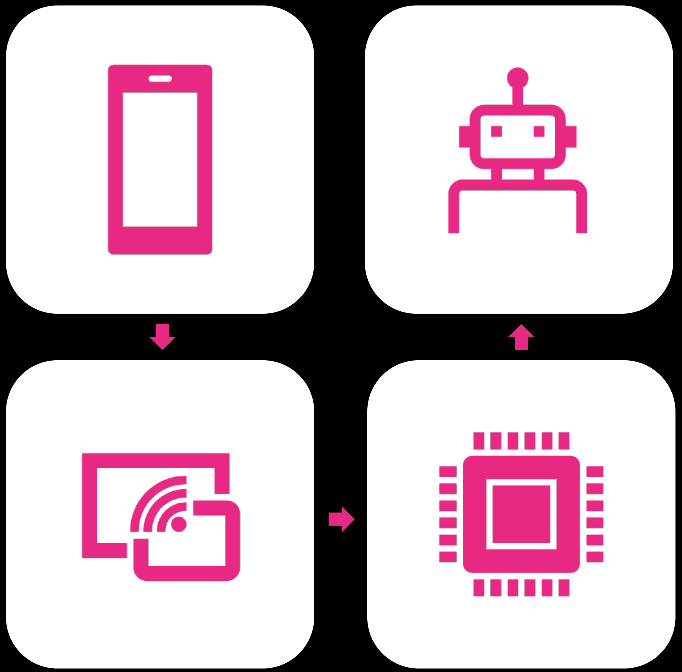

Act based on the data

Analyze data (cloud/edge) and provide insights about the world

Mobile and IoT Computing

The convergence of sensing, communication, and computation that allows us to:

Acquire data from the environment

Act based on the data

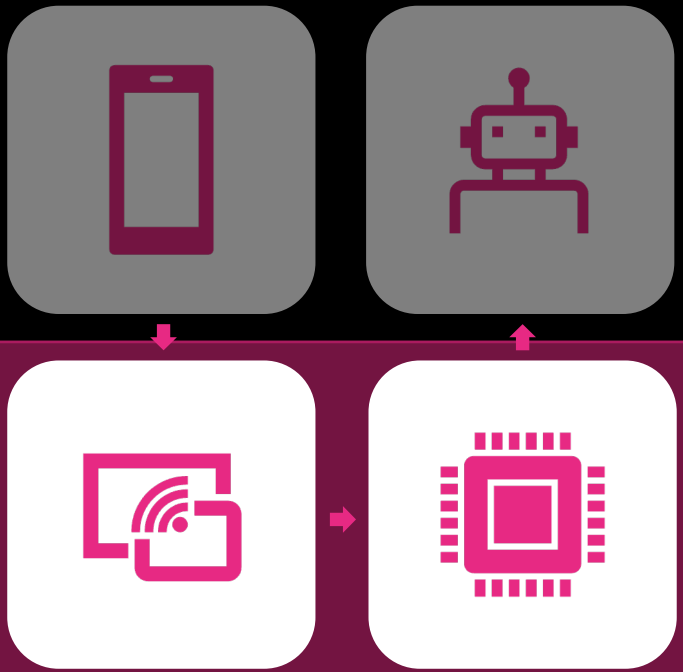

Aggregate data from multiple sources

Analyze data (cloud/edge) and provide insights about the world

This course

Main Component of IoT Systems

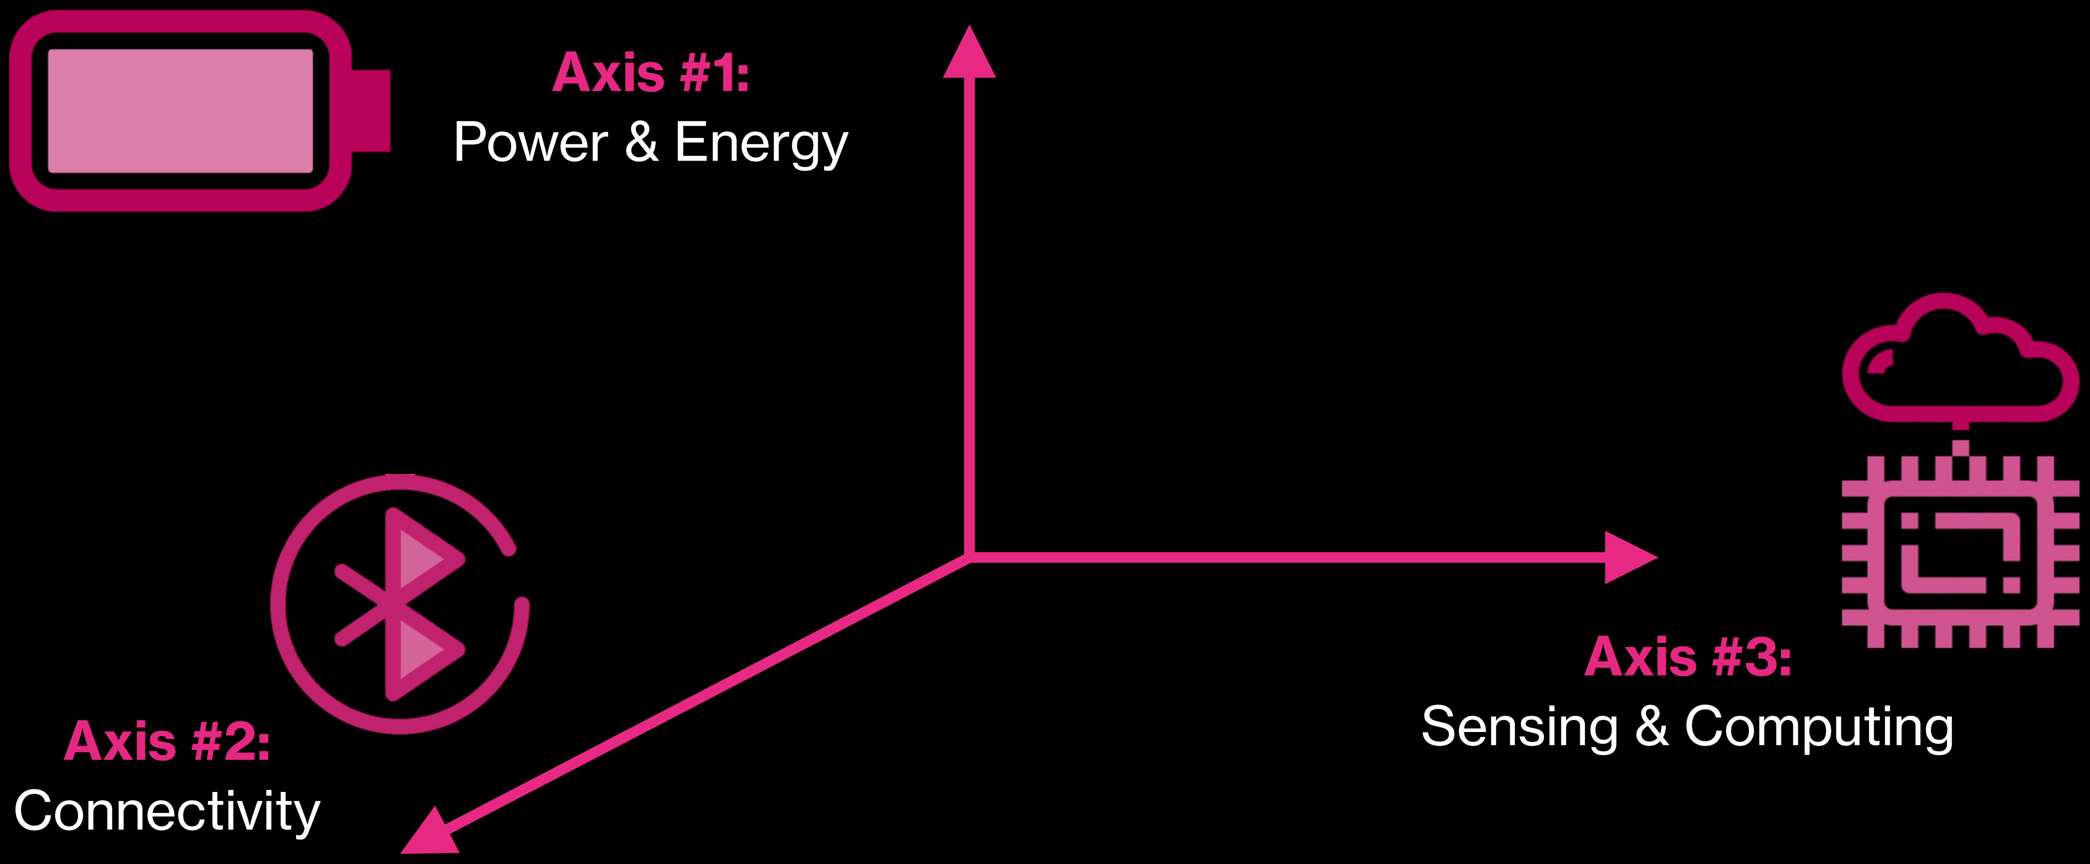

Course Organization

Module 1: Localization and Motion Tracking

Module 2: Sensing

Module 3: Connectivity

Module 4: Low-power IoT & Efficient Computing

Module 5: Emerging Topics

Objectives of This Module

Learn the fundamentals, applications, and implications of Network technologies for Mobile and IoT Systems

1. What are the various classes of network technologies?
2. How do we choose the right technology for a given application?
3. What are the different routing architectures?
4. How does energy impact the system design?
5. How does batteryless connection work?

Networking: “GLUE” for the Mobile and IoT

Mobile and IoT systems “technology push” from the convergence of:

- Increasing powerful nodes with embedded computing + miniaturized sensing
- Wireless network connectivity among the nodes

The past, present, and future

2G

Voice

The past, present, and future

2G

Voice

3G

Broadband Data

Low cost Voice &
Voice capacity

The past, present, and future

2G	3G	4G
Voice	Broadband Data	MTC, Higher data rates
	Low cost Voice & Voice capacity	Low cost data & data capacity
		Zero cost voice

The past, present, and future

2G	3G	4G	5G
Voice	Broadband Data	MTC, Higher data rates	Industrial IoT, Higher data rates
	Low cost Voice & Voice capacity	Low cost data & data capacity	Zero cost data & zero cost voice
		Zero cost voice	MTC Optimized?

The past, present, and future

2G

Voice

3G

Broadband Data

Low cost Voice & Voice capacity

4G

MTC, Higher data rates

Low cost data & data capacity

Zero cost voice

5G

Industrial IoT, Higher data rates

Zero cost data & zero cost voice

MTC Optimized?

6G

RF Sensing, 100 Gbits/sec @ short range Gbps @ high speed

Industrial IoT optimized for low cost & perf.

Zero cost data, voice, massive MTC

Each new generation is about optimizing the new use case of the previous generation to reduce cost and introduction of new use cases

2030


Ubiquitous Compute


Precision Sensing & Actuation

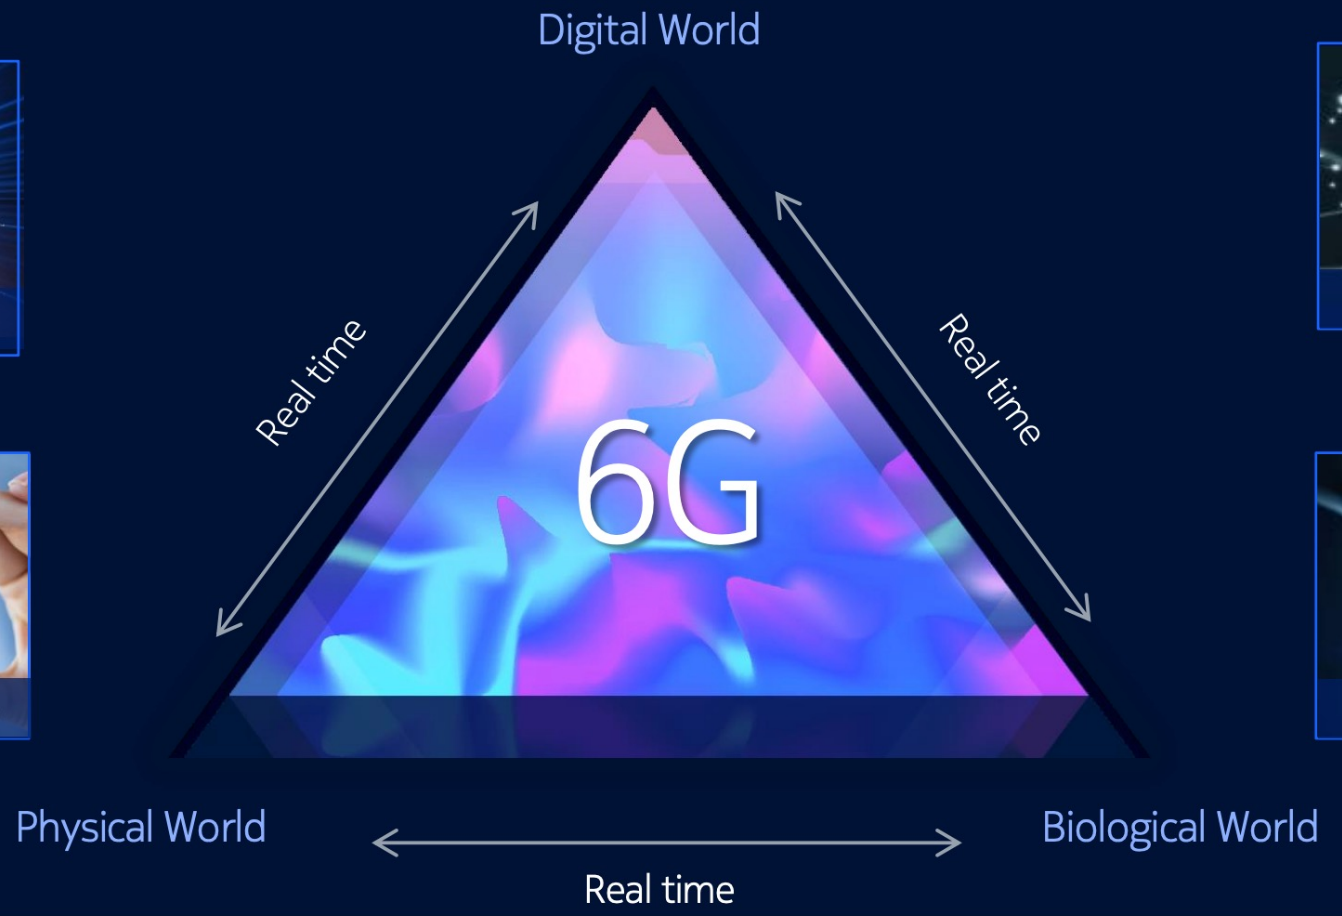


Knowledge systems


Human Machine Interface

6G to unify the experience across physical, digital and biological worlds

NOKIA Bell Labs

The IoT Connectivity Arsenal

Bluetooth® SMART

ZigBee®

6LoWPAN
IPv6-based Low-power Wireless Personal Area Networks

IEEE 802.15.4

WiFi diving into 900 MHz
802.11ah
IEEE

LoRa™

3GPP™
A GLOBAL INITIATIVE

a b g WiFi n ac

Lte™ 4DVANCED

4G LTE

5G

Z WAVE

SIGFOX
One network A billion dreams

ingenu

WEIGHTLESS™

INSTEON®

enocean

ALLSEEN ALLIANCE
AllJoyn® Framework

THREAD

OPEN INTERCONNECT
CONSORTIUM℠

industrial internet CONSORTIUM®

Networking: “GLUE” for the Mobile and IoT

Many different approaches, many different proposed standards. Much confusion.

- One size does not fit all: best technology depends on application
- What are the key organizing principles and ideas?

Architecture: Direct, Gateways & Edges

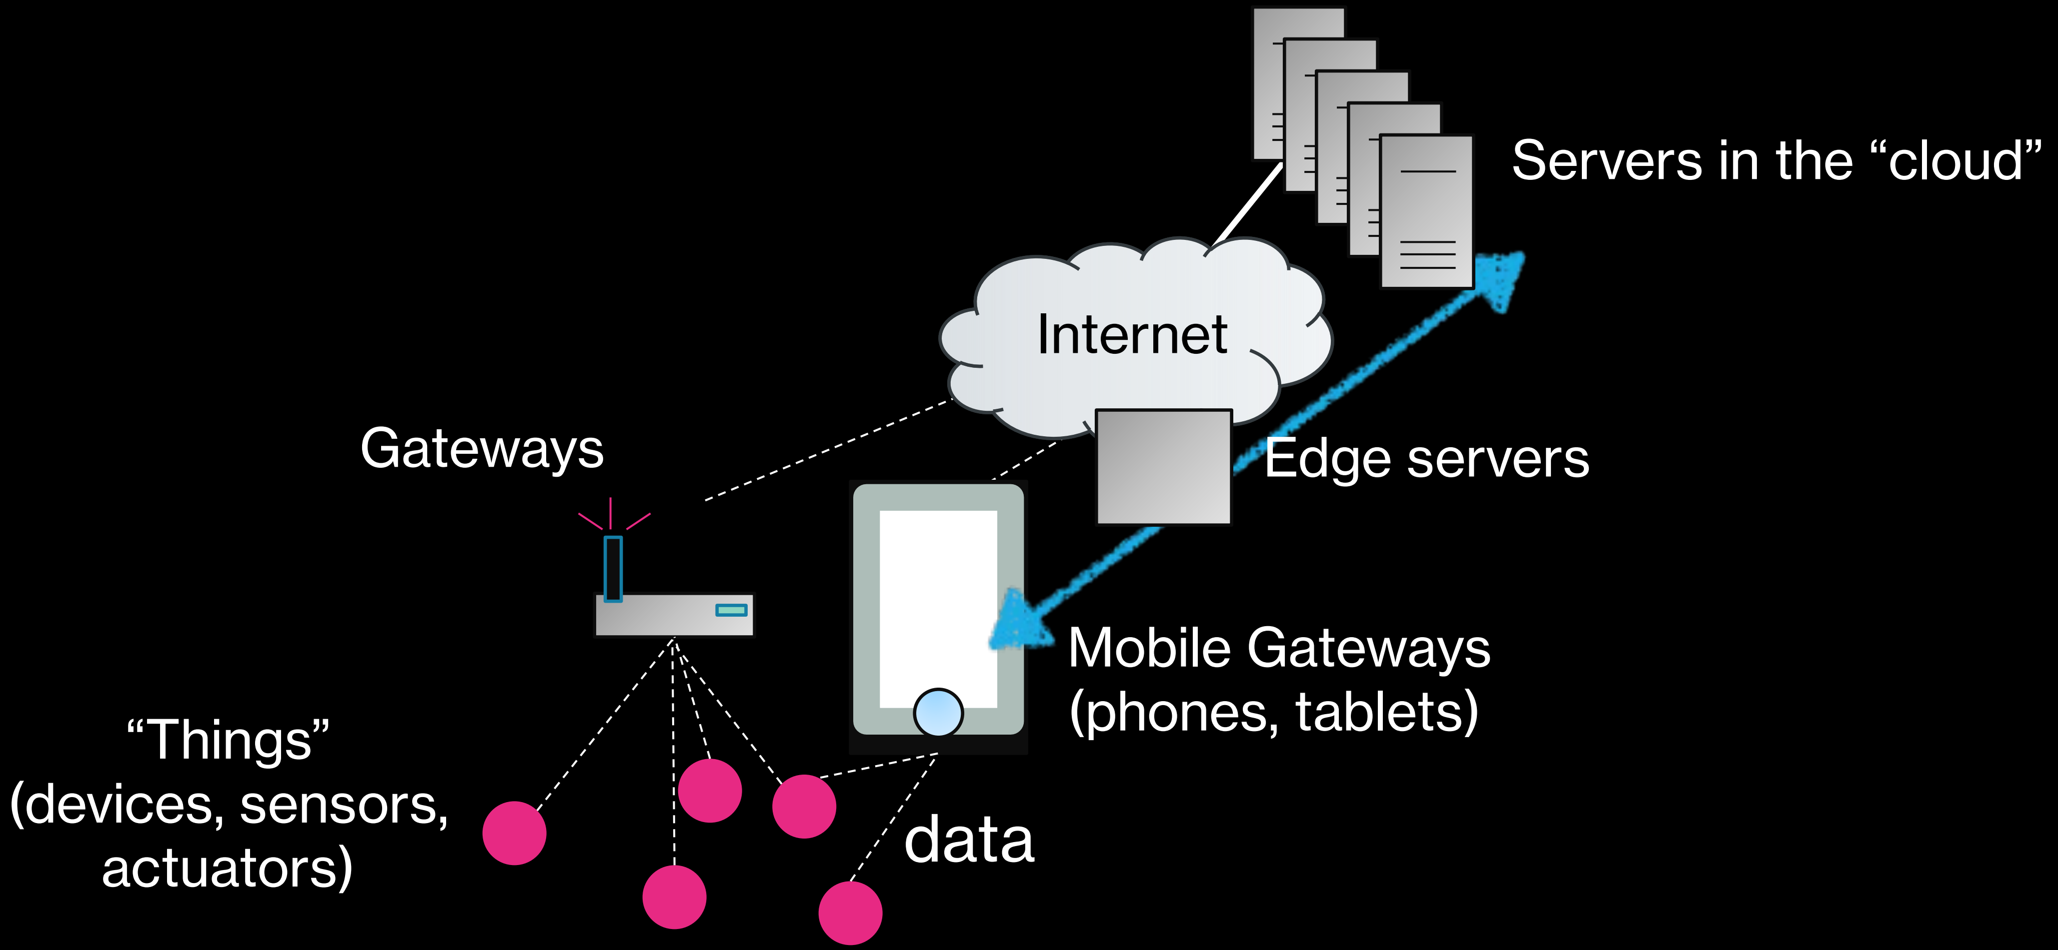

But, in fact, a rich design space

How should gateways and things communicate?

Many answers, many approaches

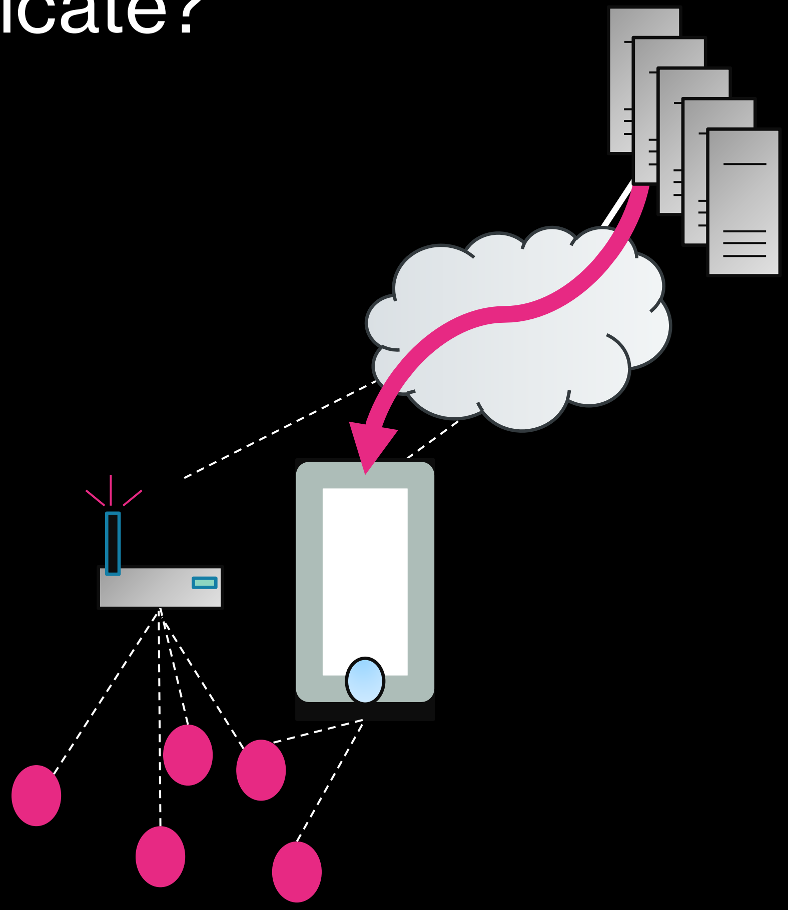

Can’t we just use the wireless internet?

Cellular and Wi-Fi?

Yes, we can…

except when we can’t!

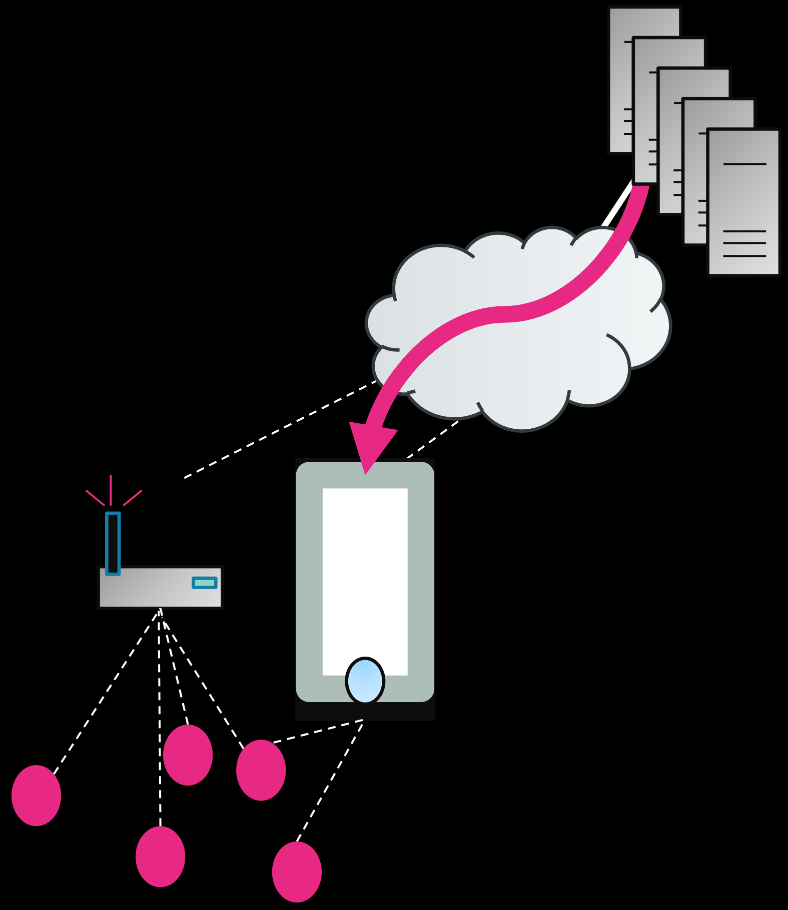

Wireless internet for IoT?

Cellular (5G, LTE/4G, 3G, 2G) and Wi-Fi are:

- Widely available (cellular in the wide-area and Wi-Fi for static uses)
- High bandwidth (for most purposes), so can support high-rate apps

But, they have big drawbacks:

- High power: not ideal for battery-operated scenarios
- Cellular: often high cost (esp. per byte if usage-per-thing is low)
- Wi-Fi: OK in most buildings, but not for longer range

Wi-Fi: In-building powered things (speakers, washers, refrigerators, …)

Cellular: High-valued powered things (e.g., “connected car”)

Connectivity Design Space

What are the metrics that we care about?

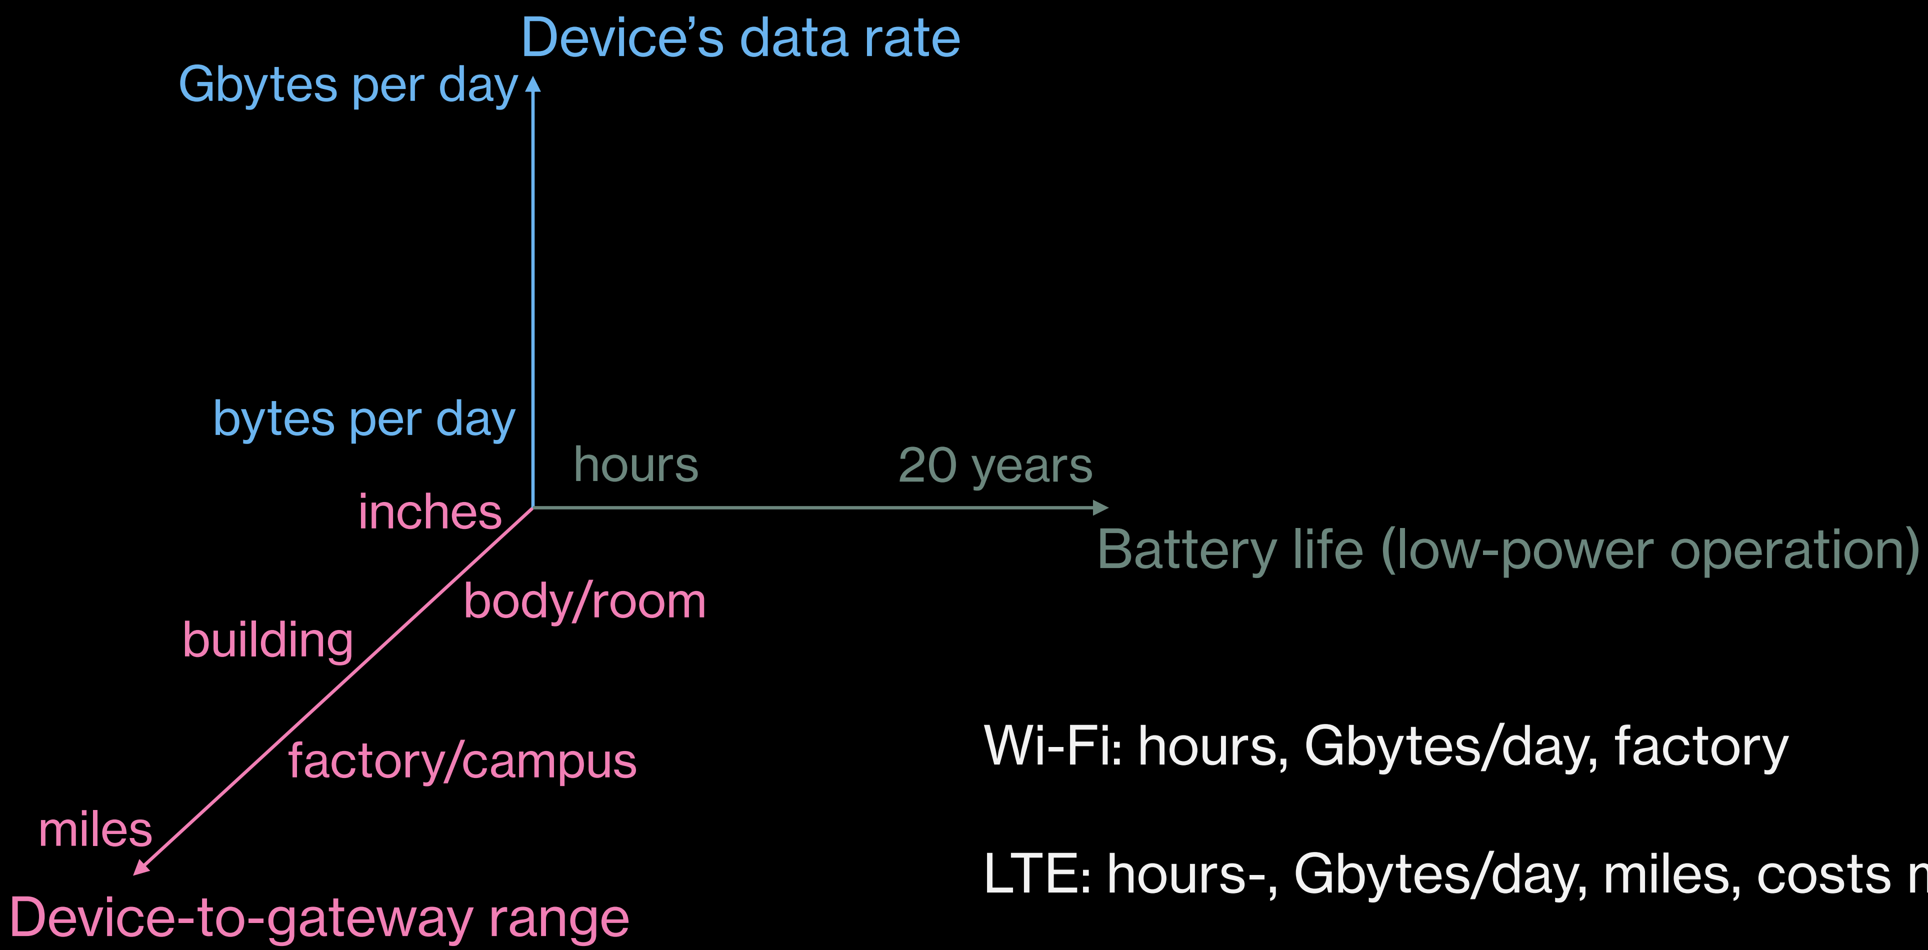

Wi-Fi: hours, Gbytes/day, factory

LTE: hours-, Gbytes/day, miles, costs more

Why so many solutions?

- Axes are not independent
- Technology evolves fast
- Bundling into popular devices speeds-up adoption, change the economics
  + cf. Wi-Fi & laptops
  + Bluetooth classic & cell phones & wireless headsets
  + Bluetooth Low Energy (BLE) & smartphones & “body/room” connectivity with months-to-years with low duty cycles

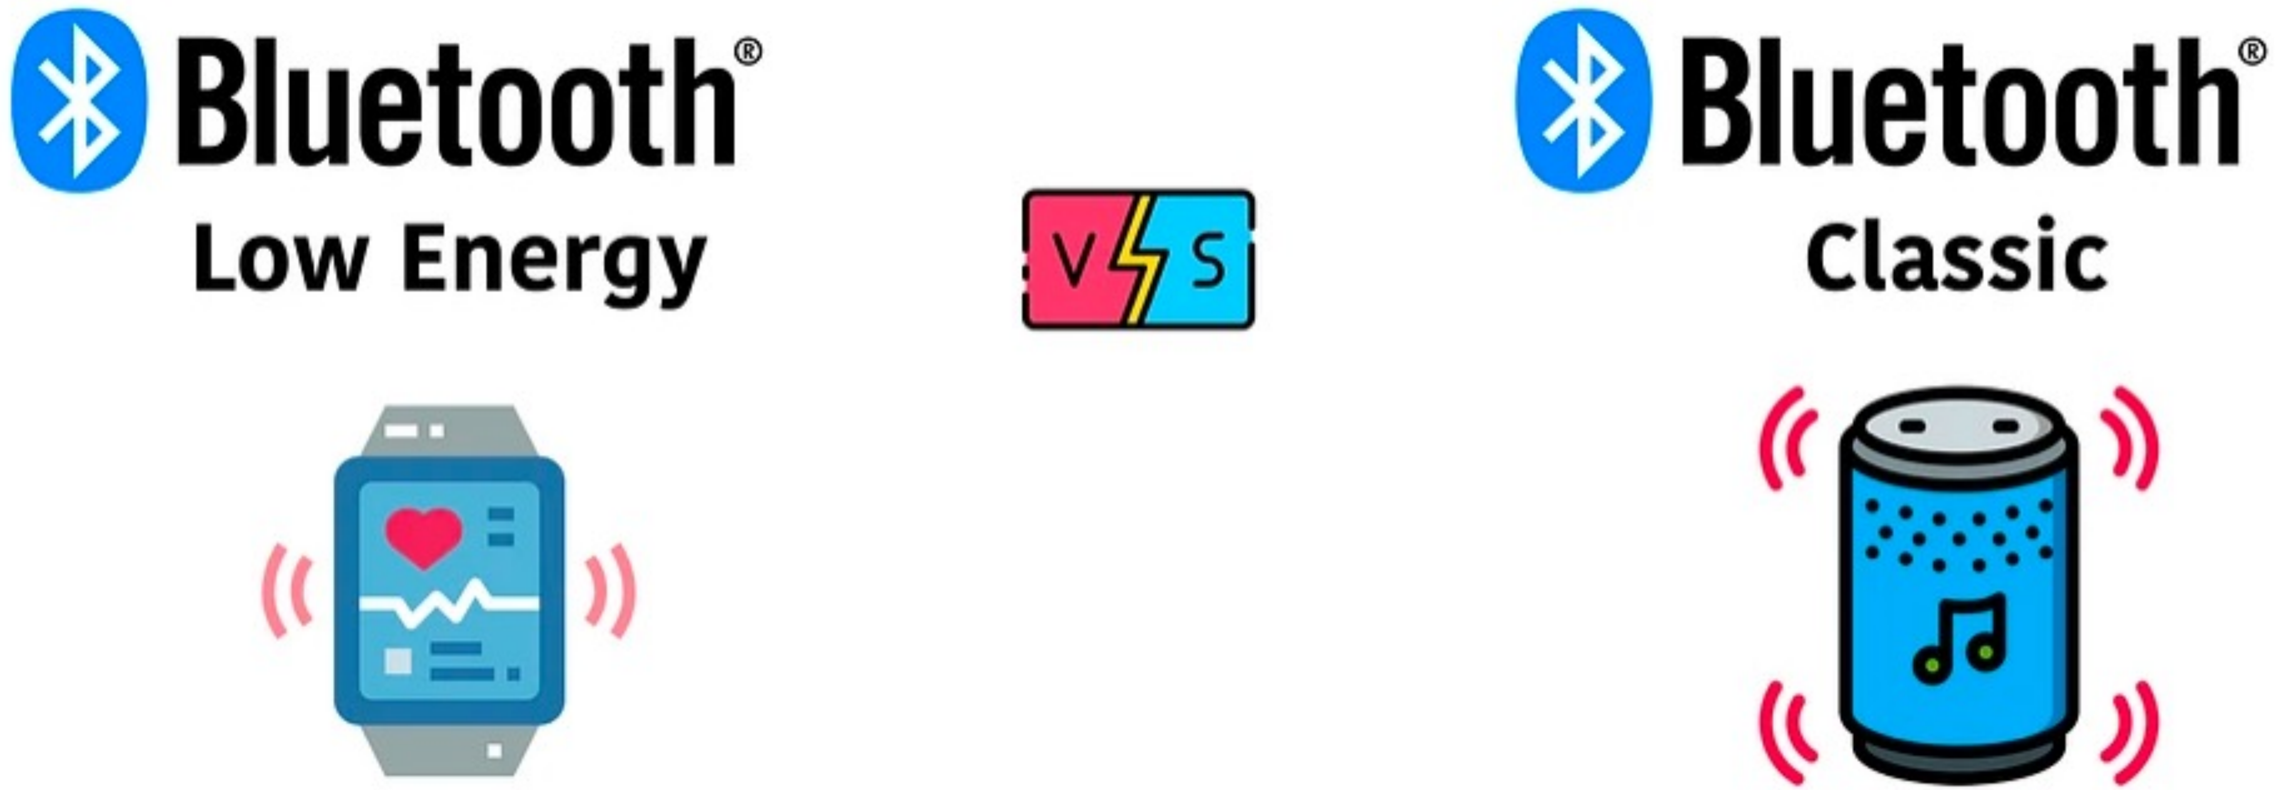

Bluetooth Classic and Bluetooth Low Energy (BLE)

How does BLE work?

Two parts:

1. Advertisements (aka beaconing) for device discovery
2. Connection phase for data exchange

BLE Advertisements are periodic

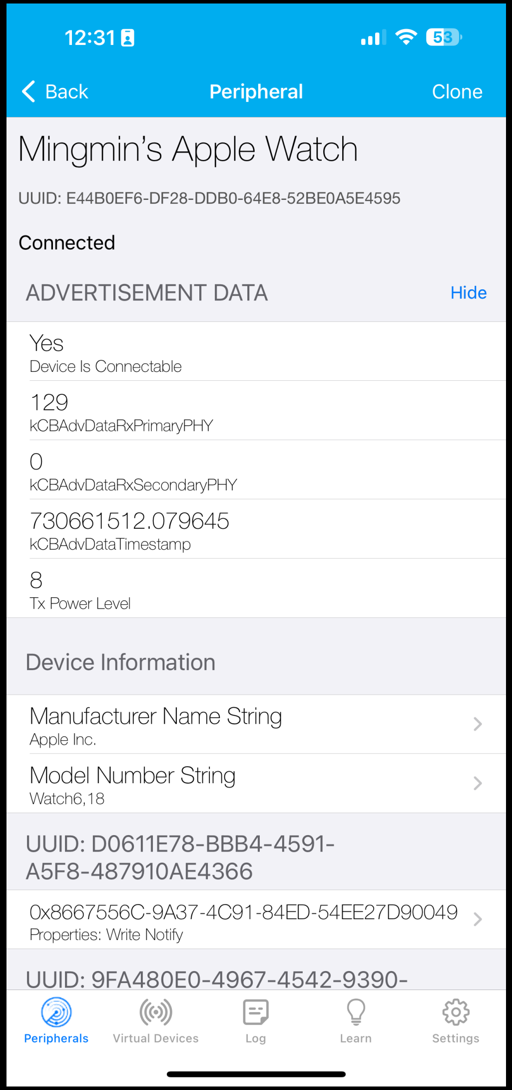

- Typical period: 100 ms
- Less frequent is fine
- Triggered advertisement are often a good idea

- Trade-off between energy consumption and discovery latency

On connection

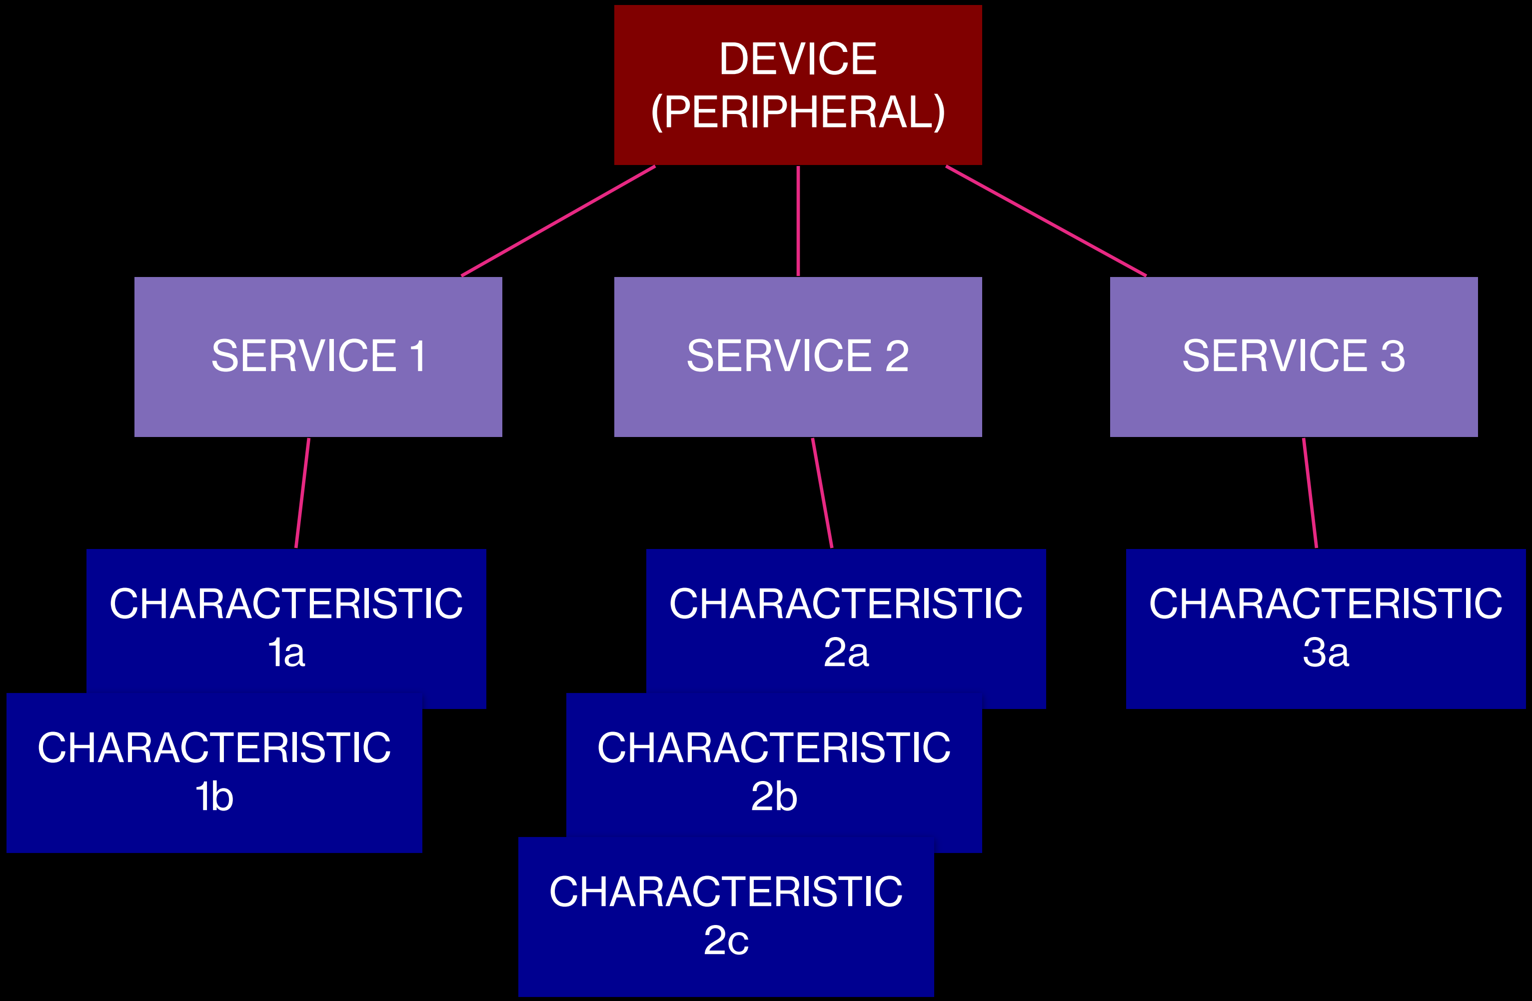

READABLE
READ/WRITE
NOTIFICATIONS

Usually support OTA (over-the-air upgrades)

On Connection: MAC protocol

Central orchestrates data communication

Key idea: time-schedule to reduce energy consumption

On connect: exchange parameters

- Frequency hopping sequence
- Connection interval, i.e., periodicity of data exchange (T milliseconds)

Every T milliseconds, Central and Peripheral exchange up to 4 packets, alternating turns

Then Peripheral can go back to sleep until next interval

Extending communication range

What are the metrics that we care about?

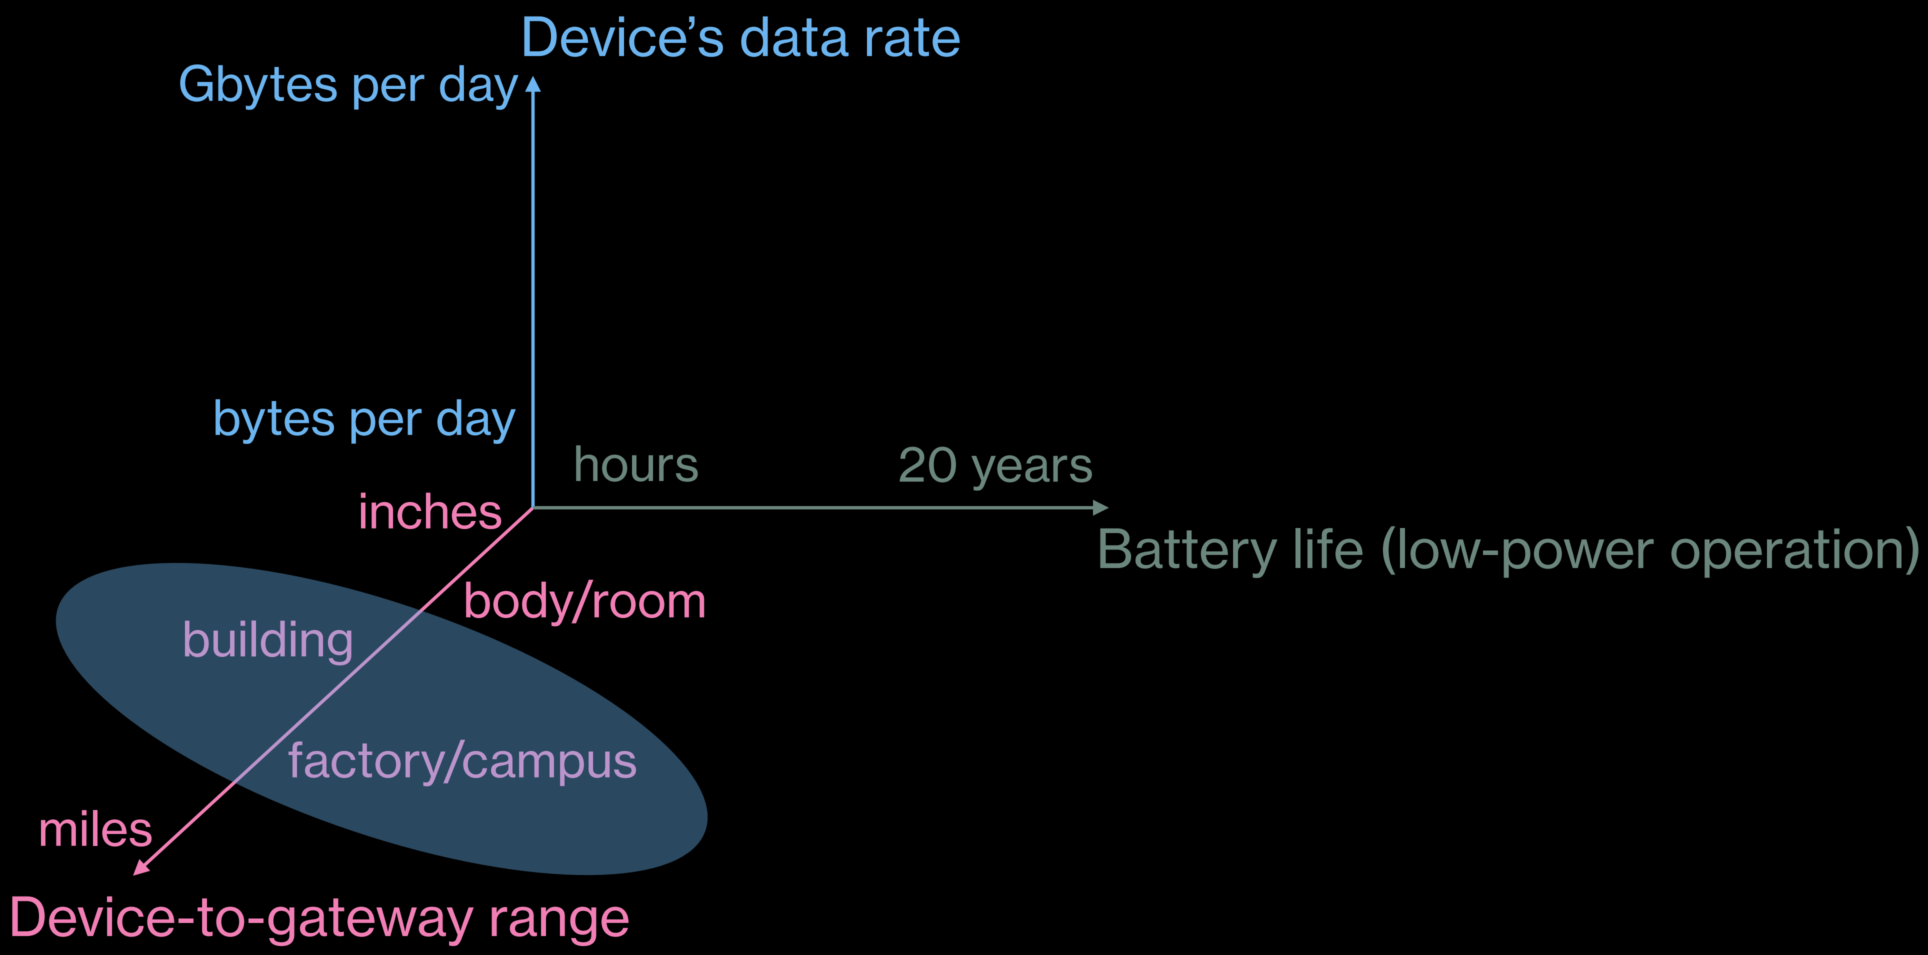

Extending communication range: mesh networks

1980s: DARPA packet radio networks

The DARPA Packet Radio Network Protocols

JOHN JUBIN AND JANET D. TORNOW, ASSOCIATE, IEEE

Invited Paper

In this paper we describe the current state of the DARPA packet radio network. Fully automated algorithms and protocols to organize, control, maintain, and move traffic through the packet radio network have been designed, implemented, and tested. By means of protocols, networks of about 50 packet radios with some degree of nodal mobility can be organized and maintained under a fully distributed mode of control. We have described the algorithms and illustrated how the PRNET provides highly reliable network transport and datagram service, by dynamically determining optimal routes, effectively controlling congestion, and fairly allocating the channel in the face of changing link conditions, mobility, and varying traffic loads.

I. INTRODUCTION

In 1973, the Defense Advanced Research Projects Agency (DARPA) initiated research on the feasibility of using packet-switched, store-and-forward radio communications to provide reliable computer communications [1]. This development was motivated by the need to provide computer network access to mobile hosts and terminals, and to provide computer communications in a mobile environment. Packet radio networking offers a highly efficient way of using a multiple-access channel, particularly with bursty traffic [2]. The DARPA Packet Radio Network (PRNET) has evolved through the years to be a robust, reliable, operational experimental network [3]. The development process has been of an incremental, evolutionary nature [4]; as algorithms were designed and implemented, new versions of the PRNET with increased capabilities were demonstrated. The PRNET has been in daily operation for experimental purposes for nearly ten years. In this paper we describe the current state of the DARPA PRNET.

We begin by providing a synopsis of the PRNET system concepts, attributes, and physical components in Section II. In Section III, we illustrate the mechanisms by which a packet radio automatically keeps track of a potentially continuously changing network topology. In Section IV, we describe the algorithms used to route a packet through the packet radio communications subnet. In Section V, we examine the protocols for transmitting packets. In Section VI, we describe some of the hardware capabilities of the packet radio that strongly influence the design and characteristics of the PRNET protocols. We conclude by looking briefly at some applications of packet radio networks and by summarizing the state of the current technology.

II. DESCRIPTION OF THE PACKET RADIO SYSTEM

A. Broadcast Radio

The PRNET provides, via a common radio channel, the exchange of data between computers that are geographically separated. As a communications medium, broadcast radio (as opposed to wires and antenna-directed radio) provides important advantages to the user of the network. One of the benefits is mobility; a packet radio (PR) can operate while in motion. Second, the network can be installed or deployed quickly; there are no wires to set up. A third advantage is the ease of reconfiguration and redeployment. The PRNET protocols take advantage of broadcasting and common-channel properties to allow the PRNET to be expanded or contracted automatically and dynamically. A group of packet radios leaving the original area simply departs. Having done so, it can function as an autonomous group and may later rejoin the original network or join another group.

The broadcasting and common channel properties of radio have disadvantages too. These properties, for all practical purposes, prohibit the building of a radio that is able to transmit and receive at the same time. Therefore, the PRNET protocols must attempt to schedule each transmission when the intended PR is not itself transmitting. Also, transmissions often reach unintended PRs and interfere with intended receptions. Therefore, the protocols must attempt to schedule each transmission when the intended PR is not receiving another PR's transmission.

B. Automated Network Management

The PRNET features fully automated network management. It is self-configuring upon network initialization, reconfigures upon gain or loss of packet radios, and has dy-

Manuscript received February 1, 1986; revised July 30, 1986. The work of J. Jubin was supported by the Defense Advanced Research Projects Agency of the Department of Defense under Contract MDA903-85-C-0205. The work of J. D. Tornow was supported by the Defense Advanced Research Projects Agency of the Department of Defense under Contract MDA903-85-C-0254.

J. Jubin is with Collins Defense Communications, Rockwell International, Richardson, TX 75081, USA.

J. D. Tornow is with SRI International, Menlo Park, CA 94025, USA.

0018-9219/87/0100-0021$01.00 © 1987 IEEE

1990s: mobile ad hoc networks (MANET)

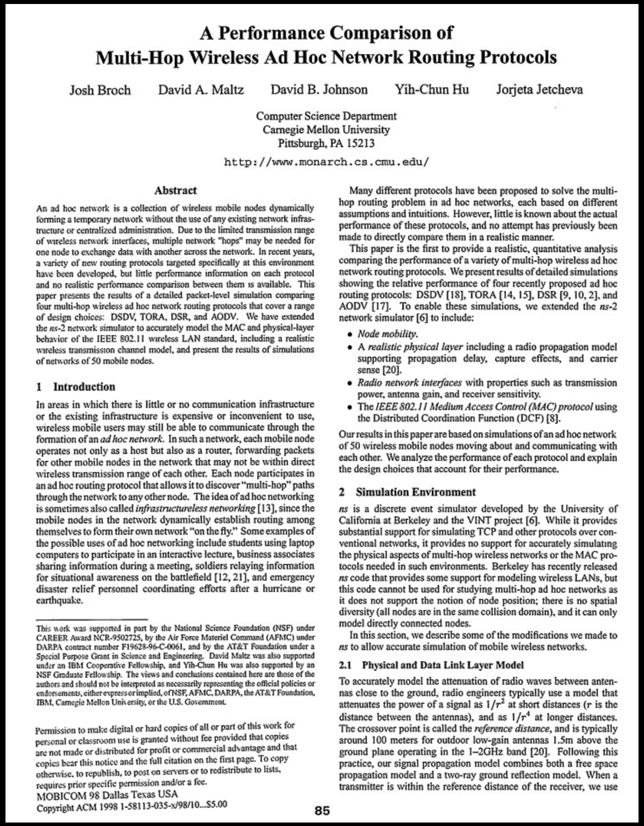

A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols

Josh Broch David A. Maltz David B. Johnson Yih-Chun Hu Jorjeta Jetcheva

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

http://www.monarch.cs.cmu.edu/

Abstract

An ad hoc network is a collection of wireless mobile nodes dynamically forming a temporary network without the use of any existing network infrastructure or centralized administration. Due to the limited transmission range of wireless network interfaces, multiple network "hops" may be needed for one node to exchange data with another across the network. In recent years, a variety of new routing protocols targeted specifically at this environment have been developed, but little performance information on each protocol and no realistic performance comparison between them is available. This paper presents the results of a detailed packet-level simulation comparing four multi-hop wireless ad hoc network routing protocols that cover a range of design choices: DSDV, TORA, DSR, and AODV. We have extended the *ns*-2 network simulator to accurately model the MAC and physical-layer behavior of the IEEE 802.11 wireless LAN standard, including a realistic wireless transmission channel model, and present the results of simulations of networks of 50 mobile nodes.

1 Introduction

In areas in which there is little or no communication infrastructure or the existing infrastructure is expensive or inconvenient to use, wireless mobile users may still be able to communicate through the formation of an *ad hoc network*. In such a network, each mobile node operates not only as a host but also as a router, forwarding packets for other mobile nodes in the network that may not be within direct wireless transmission range of each other. Each node participates in an ad hoc routing protocol that allows it to discover "multi-hop" paths through the network to any other node. The idea of ad hoc networking is sometimes also called *infrastructureless networking* [13], since the mobile nodes in the network dynamically establish routing among themselves to form their own network "on the fly." Some examples of the possible uses of ad hoc networking include students using laptop computers to participate in an interactive lecture, business associates sharing information during a meeting, soldiers relaying information for situational awareness on the battlefield [12, 21], and emergency disaster relief personnel coordinating efforts after a hurricane or earthquake.

Many different protocols have been proposed to solve the multi-hop routing problem in ad hoc networks, each based on different assumptions and intuitions. However, little is known about the actual performance of these protocols, and no attempt has previously been made to directly compare them in a realistic manner.

This paper is the first to provide a realistic, quantitative analysis comparing the performance of a variety of multi-hop wireless ad hoc network routing protocols. We present results of detailed simulations showing the relative performance of four recently proposed ad hoc routing protocols: DSDV [18], TORA [14, 15], DSR [9, 10, 2], and AODV [17]. To enable these simulations, we extended the *ns*-2 network simulator [6] to include:

- *Node mobility*.
- A *realistic physical layer* including a radio propagation model supporting propagation delay, capture effects, and carrier sense [20].
- *Radio network interfaces* with properties such as transmission power, antenna gain, and receiver sensitivity.
- The *IEEE 802.11 Medium Access Control (MAC) protocol* using the Distributed Coordination Function (DCF) [8].

Our results in this paper are based on simulations of an ad hoc network of 50 wireless mobile nodes moving about and communicating with each other. We analyze the performance of each protocol and explain the design choices that account for their performance.

2 Simulation Environment

ns is a discrete event simulator developed by the University of California at Berkeley and the VINT project [6]. While it provides substantial support for simulating TCP and other protocols over conventional networks, it provides no support for accurately simulating the physical aspects of multi-hop wireless networks or the MAC protocols needed in such environments. Berkeley has recently released *ns* code that provides some support for modeling wireless LANs, but this code cannot be used for studying multi-hop ad hoc networks as it does not support the notion of node position; there is no spatial diversity (all nodes are in the same collision domain), and it can only model directly connected nodes.

In this section, we describe some of the modifications we made to *ns* to allow accurate simulation of mobile wireless networks.

2.1 Physical and Data Link Layer Model

To accurately model the attenuation of radio waves between antennas close to the ground, radio engineers typically use a model that attenuates the power of a signal as $1/r^2$ at short distances (r is the distance between the antennas), and as $1/r^4$ at longer distances. The crossover point is called the *reference distance*, and is typically around 100 meters for outdoor low-gain antennas 1.5m above the ground plane operating in the 1–2GHz band [20]. Following this practice, our signal propagation model combines both a free space propagation model and a two-ray ground reflection model. When a transmitter is within the reference distance of the receiver, we use

This work was supported in part by the National Science Foundation (NSF) under CAREER Award NCR-9502725, by the Air Force Materiel Command (AFMC) under DARPA contract number F19628-96-C-0061, and by the AT&T Foundation under a Special Purpose Grant in Science and Engineering. David Maltz was also supported under an IBM Cooperative Fellowship, and Yih-Chun Hu was also supported by an NSF Graduate Fellowship. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of NSF, AFMC, DARPA, the AT&T Foundation, IBM, Carnegie Mellon University, or the U.S. Government.

Extending communication range: mesh networks

Late 90s, 2000s: Sensor networks

Next Century Challenges: Scalable Coordination in Sensor Networks

Deborah Estrin
Ramesh Govindan
John Heidemann
Satish Kumar

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292, USA
{estrin,govindan,johnh,kkumar}@isi.edu

Abstract

Networked sensors—those that coordinate amongst themselves to achieve a larger sensing task—will revolutionize information gathering and processing both in urban environments and in inhospitable terrain. The sheer numbers of these sensors and the expected dynamics in these environments present unique challenges in the design of unattended autonomous sensor networks. These challenges lead us to hypothesize that sensor network coordination applications may need to be structured differently from traditional network applications. In particular, we believe that *localized algorithms* (in which simple local node behavior achieves a desired global objective) may be necessary for sensor network coordination. In this paper, we describe localized algorithms, and then discuss *directed diffusion*, a simple communication model for describing localized algorithms.

1 Introduction

Integrated low-power sensing devices will permit remote object monitoring and tracking in many different contexts: in the field (vehicles, equipment, personnel), the office building (projectors, furniture, books, people), the hospital ward (syringes, bandages, IVs) and the factory floor (motors, small robotic devices). Networking these sensors—empowering them with the ability to *coordinate* amongst themselves on a larger sensing task—will revolutionize information gathering and processing in many situations. Large scale, dynamically changing, and robust sensor colonies can be deployed in inhospitable physical environments such as remote geographic regions or toxic urban locations. They will also enable low maintenance sensing in more benign, but less accessible, environments: large industrial plants, aircraft interiors etc.

To motivate the challenges in designing these sensor networks, consider the following scenario. Several thousand sensors are rapidly deployed (*e.g.*, thrown from an aircraft) in remote terrain. The sensors coordinate to establish a communication network, divide the task of mapping and monitoring the terrain amongst themselves in an energy-efficient manner, adapt their overall sensing accuracy to the remaining total resources, and re-organize upon sensor failure. When additional sensors are added or old sensors fail, the sensors re-organize themselves to take advantage of the added system resources.

Several aspects of this scenario present systems design challenges different from those posed by existing computer networks (Section 2). The *sheer numbers* of these devices, and their unattended deployment, will preclude reliance on broadcast communication or the configuration currently needed to deploy and operate networked devices. Devices may be battery constrained or subject to hostile environments, so individual *device failure* will be a regular or common event. In addition, the configuration devices will *frequently change* in terms of position, reachability, power availability, and even task details. Finally, because these devices interact with the physical environment, they, and the network as a whole, will experience a significant range of *task dynamics*.

The WINS project [1] has considered device-level communication primitives needed to satisfy these requirements. However, these requirements potentially affect many other aspects of network design: routing and addressing mechanisms, naming and binding services, application architectures, security mechanisms, and so forth. This paper focuses on the principles underlying the design of services and applications in sensor networks. In particular, since the sensing is inherently distributed, we argue that sensor network applications will themselves be distributed.

Many of the lessons learned from Internet and mobile network design will be applicable to designing sensor network applications. However, this paper hypothesizes that sensor networks have different enough requirements to at least warrant re-considering the overall structure of applications and services. Specifically, we believe there are significant robustness and scalability advantages to designing applications using *localized* algorithms—where sensors only interact with other sensors in a restricted vicinity, but nevertheless collectively achieve a desired global objective (Section 3). We also describe *directed diffusion*, a promising model for describing localized algorithms (Section 4).

Our research project is starting to investigate the design of localized algorithms using the directed diffusion model. These ideas were developed in the context of a DARPA ISAT study, chaired by one of the authors (Estrin). The

Mobicom '99 Seattle Washington USA

An Application-Specific Protocol Architecture for Wireless Microsensor Networks

Wendi B. Heinzelman, *Member, IEEE*, Anantha P. Chandrakasan, *Senior Member, IEEE*, and Hari Balakrishnan, *Member, IEEE*

***Abstract*—Networking together hundreds or thousands of cheap microsensor nodes allows users to accurately monitor a remote environment by intelligently combining the data from the individual nodes. These networks require robust wireless communication protocols that are energy efficient and provide low latency. In this paper, we develop and analyze low-energy adaptive clustering hierarchy (LEACH), a protocol architecture for microsensor networks that combines the ideas of energy-efficient cluster-based routing and media access together with application-specific data aggregation to achieve good performance in terms of system lifetime, latency, and application-perceived quality. LEACH includes a new, distributed cluster formation technique that enables self-organization of large numbers of nodes, algorithms for adapting clusters and rotating cluster head positions to evenly distribute the energy load among all the nodes, and techniques to enable distributed signal processing to save communication resources. Our results show that LEACH can improve system lifetime by an order of magnitude compared with general-purpose multihop approaches.**

***Index Terms*—Data aggregation, protocol architecture, wireless microsensor networks.**

I. INTRODUCTION

ADVANCES iN sensor technology, low-power electronics, and low-power radio frequency (RF) design have enabled the development of small, relatively inexpensive and low-power sensors, called *microsensors*, that can be connected via a wireless network. These wireless microsensor networks represent a new paradigm for extracting data from the environment and enable the reliable monitoring of a variety of environments for applications that include surveillance, machine failure diagnosis, and chemical/biological detection. An important challenge in the design of these networks is that two key resources—communication bandwidth and energy—are significantly more limited than in a tethered network environment. These constraints require innovative design techniques to use the available bandwidth and energy efficiently.

Manuscript received January 9, 2001; revised July 1, 2001 and August 24, 2001; accepted August 24, 2001. The editor coordinating the review of this paper and approving it for publication is M. Zorzi. The work of W. B. Heinzelman was supported by a Kodak Fellowship. This work was supported in part by the Defense Advanced Research Project Agency (DARPA) Power Aware Computing/Communication Program and the U. S. Air Force Research Laboratory, Air Force Materiel Command, under Agreement F30602-00-2-0551.

W. B. Heinzelman was with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA. She is now with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627-0126 USA (e-mail: wheinzel@ece.rochester.edu).

A. P. Chandrakasan and H. Balakrishnan are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: anantha@mtl.mit.edu; hari@lcs.mit.edu).

Digital Object Identifier 10.1109/TWC.2002.804190

In order to design good protocols for wireless microsensor networks, it is important to understand the parameters that are relevant to the sensor applications. While there are many ways in which the properties of a sensor network protocol can be evaluated, we use the following metrics.

A. Ease of Deployment

Sensor networks may contain hundreds or thousands of nodes, and they may need to be deployed in remote or dangerous environments, allowing users to extract information in ways that would not have been possible otherwise. This requires that nodes be able to communicate with each other even in the absence of an established network infrastructure and predefined node locations.

B. System Lifetime

These networks should function for as long as possible. It may be inconvenient or impossible to recharge node batteries. Therefore, all aspects of the node, from the hardware to the protocols, must be designed to be extremely energy efficient.

C. Latency

Data from sensor networks are typically time sensitive, so it is important to receive the data in a timely manner.

D. Quality

The notion of "quality" in a microsensor network is very different than in traditional wireless data networks. For sensor networks, the end user does not require all the data in the network because 1) the data from neighboring nodes are highly correlated, making the data redundant and 2) the end user cares about a higher-level description of events occurring in the environment being monitored. The quality of the network is, therefore, based on the quality of the aggregate data set, so protocols should be designed to optimize for the unique, application- specific quality of a sensor network.

This paper builds on the work described in [11] by giving a detailed description and analysis of low-energy adaptive clustering hierarchy (LEACH), an application-specific protocol architecture for wireless microsensor networks. LEACH employs the following techniques to achieve the design goals stated: 1) randomized, adaptive, self-configuring cluster formation; 2) localized control for data transfers; 3) low-energy media access control (MAC); and 4) application-specific data processing, such as data aggregation or compression. Simulation results show that LEACH is able to achieve the desired properties of sensor networks.

2000s: Mesh networks for Internet

Architecture and Evaluation of an Unplanned 802.11b Mesh Network

John Bicket, Daniel Aguayo, Sanjit Biswas, Robert Morris
M.I.T. Computer Science and Artificial Intelligence Laboratory
jbicket, aguayo, biswas, rtm @csail.mit.edu

ABSTRACT

This paper evaluates the ability of a wireless mesh architecture to provide high performance Internet access while demanding little deployment planning or operational management. The architecture considered in this paper has unplanned node placement (rather than planned topology), omni-directional antennas (rather than directional links), and multi-hop routing (rather than single-hop base stations). These design decisions contribute to ease of deployment, an important requirement for community wireless networks. However, this architecture carries the risk that lack of planning might render the network's performance unusably low. For example, it might be necessary to place nodes carefully to ensure connectivity; the omni-directional antennas might provide uselessly short radio ranges; or the inefficiency of multi-hop forwarding might leave some users effectively disconnected.

The paper evaluates this unplanned mesh architecture with a case study of the Roofnet 802.11b mesh network. Roofnet consists of 37 nodes spread over four square kilometers of an urban area. The network provides users with usable performance despite lack of planning: the average inter-node throughput is 627 kbits/second, even though the average route has three hops.

The paper evaluates multiple aspects of the architecture: the effect of node density on connectivity and throughput; the characteristics of the links that the routing protocol elects to use; the usefulness of the highly connected mesh afforded by omni-directional antennas for robustness and throughput; and the potential performance of a single-hop network using the same nodes as Roofnet.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*; C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing protocols*

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Mesh networks, Multi-hop wireless networks, Ad hoc networks, Wireless routing, Route metrics

1. INTRODUCTION

Community wireless networks typically share a few wired Internet connections among many users spread over an urban area. Two approaches to constructing community networks are common. The first approach is to carefully construct a multi-hop network with nodes in chosen locations and directional antennas aimed to engineer high-quality radio links [31, 8, 29]; these networks require well-coordinated groups with technical expertise, but result in high throughput and good connectivity. The second approach consists of individuals operating "hot-spot" access points to which clients directly connect [5, 4]. These access points often operate independently and are loosely connected, if at all. Access-point networks do not require much coordination to deploy and operate, but usually do not provide as much coverage per wired connection as multi-hop networks.

A more ambitious vision for community networks would combine the best characteristics of both network types, operating without extensive planning or central management but still providing wide coverage and acceptable performance. This paper provides an evaluation of such an architecture, consisting of the following design decisions:

1. Unconstrained node placement, rather than a topology planned for coverage or performance. The network should work well even if the topology is determined solely by where participants happen to live.
2. Omni-directional antennas, rather than directional antennas used to form particular high-quality links. Users should be able to install an antenna without knowing in advance what nodes the antenna might talk to. Nodes should be able to route data through whatever neighbors they happen to find.
3. Multi-hop routing, rather than single-hop base stations or access points. Multi-hop routing can improve coverage and performance despite lack of planning and lack of specifically engineered links.

Extending communication range: mesh networks

2010s: Mesh networks for IoT

Zigbee

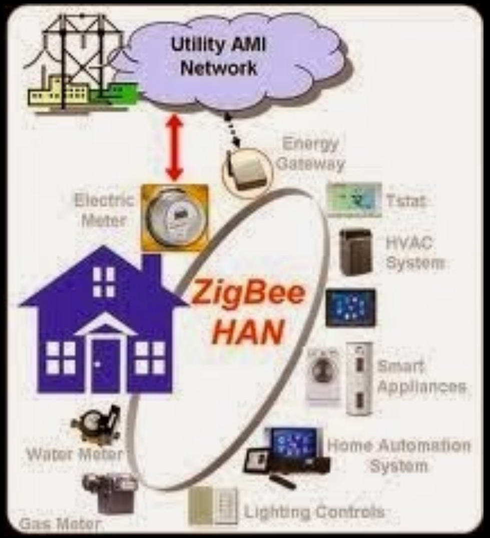

6LoWPAN: IPv6 over low-power wireless personal area networks

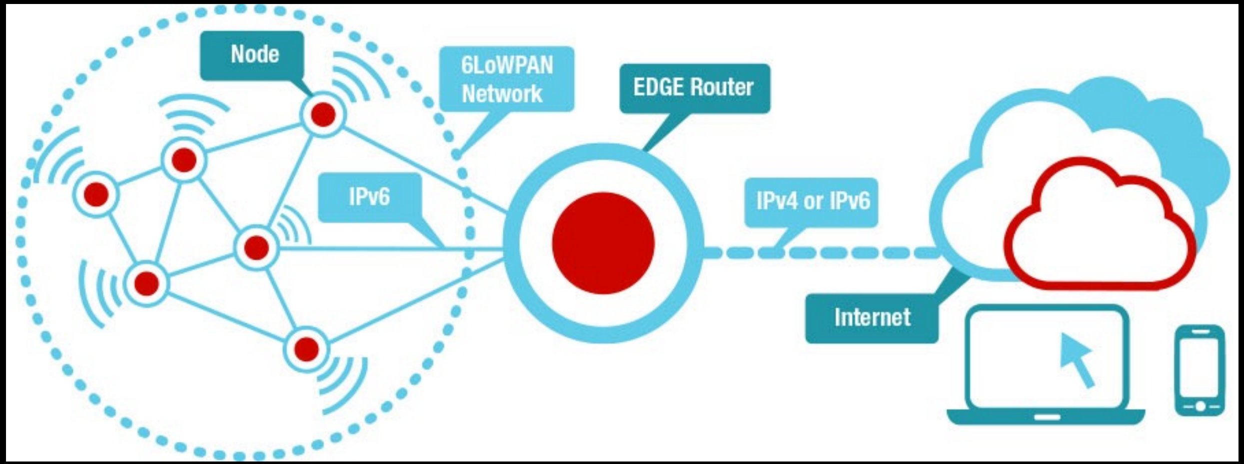

Both (typically) run over the 802.15.4 MAC standard
Routing protocol with different metrics, such as “expected transmission time”
Use case: devices communicating with gateway across multiple hops
Node duty cycles higher, some nodes do much more work

Wireless Network Architectures

There are 3 kinds of wireless network architectures

Access Network

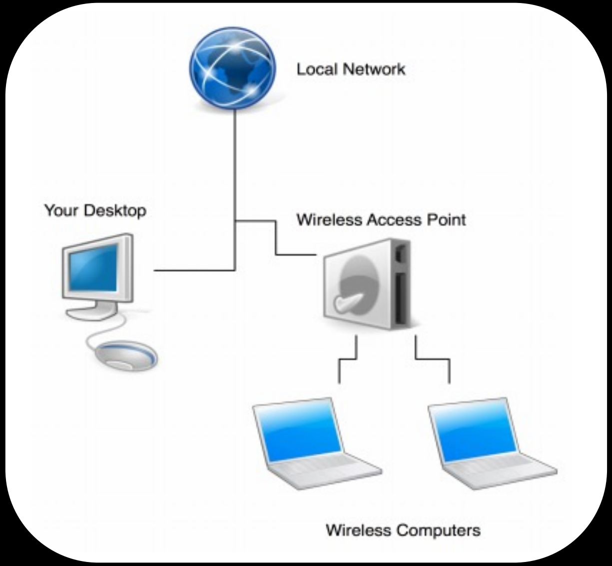

e.g., WiFi, cellular

Device-to-device

e.g., Bluetooth

Ad Hoc Network

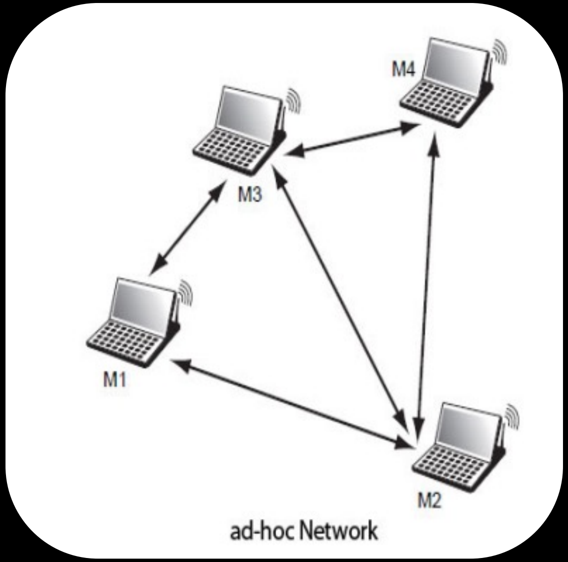

e.g., leverage P2P to reach internet (crises)

One-hop

Multi-hop

RoofNet

MIT Technology Review

Networking From the Rooftop

MIT researchers are developing new routing strategies for a wireless network that hops data in the roofs of the city.

by **Erico Guizzo** Aug 29, 2003

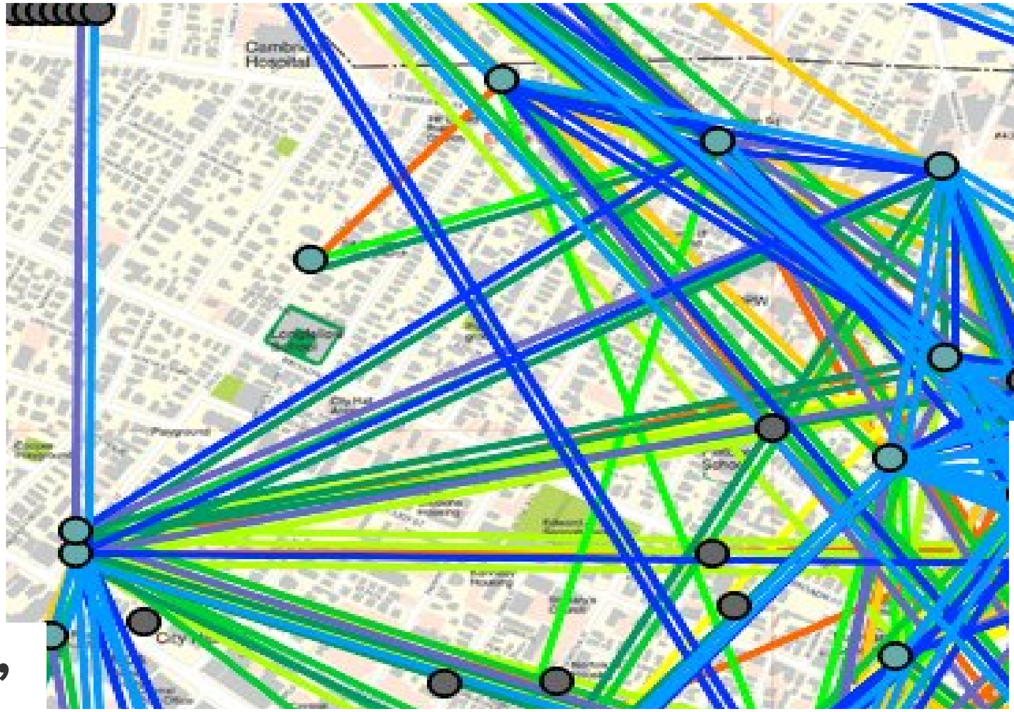

7 YEARS AFTER ROOFNET, MIT AND CSAIL CHOOSE MERAKI FOR WIRELESS LAN

February 17, 2010
Posted by: @merakisimon
Share

The Tech Established 1881 ONLINE EDITI

HOME | NEWS | OPINION | ARTS | SPORTS | CAMPUS LIFE | PHOTOS | BLO

Volume 125 >> Issue 65 : Wednesday, February 1, 2006

MIT and City Collaborate To Provide Free Wireless

TC

Cisco Acquires Enterprise Wi-Fi Startup Meraki For $1.2 Billion In Cash

Josh Constine @joshconstine / 6:36 pm EST • November 18, 2012 Comment

CISCO Buys meraki

Networking tech giant Cisco has just agreed to acquire cloud infrastructure startup Meraki, and my industry sources confirm the purchase price was $1.2 billion, all in cash. I've also gotten ahold of

Single Path Routing

Represent the wireless network as a graph

- Two nodes have an edge if they can communicate (i.e., are within radio range)
- Each edge is labeled with a weight (where a smaller weight indicates a preferred edge)

Run shortest path algorithm on the graph (e.g., Dijkstra)

- Produce the minimum weight path between every pair of nodes

How do you pick the edge weights?

- i.e., what metric should shortest path minimize?

Rest of this lecture

Approach 1:

Assign all edges the same weight → Minimize number of hops

Reasoning:

- Links in route share radio spectrum
- Extra hops reduce throughput

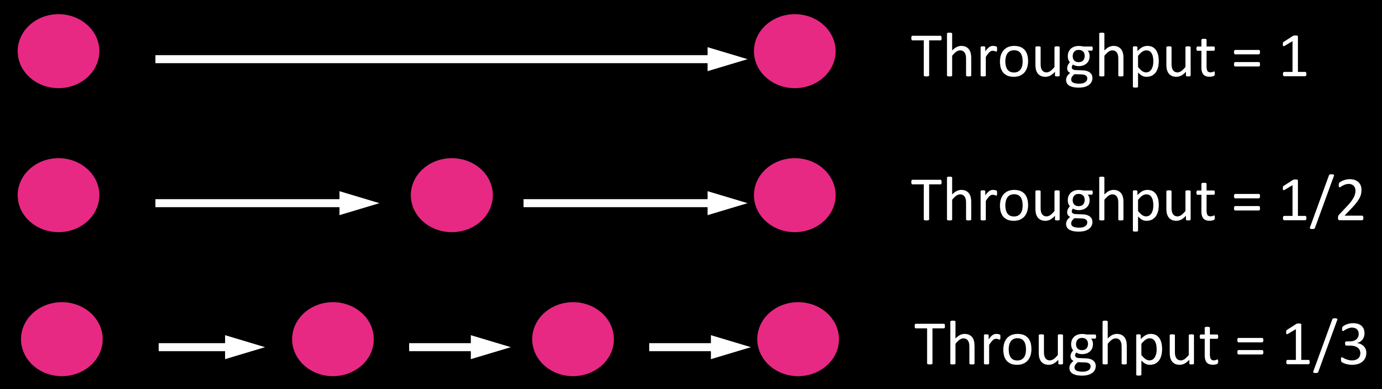

Pros? Cons?

Challenge: many links are lossy

One-hop broadcast delivery ratios

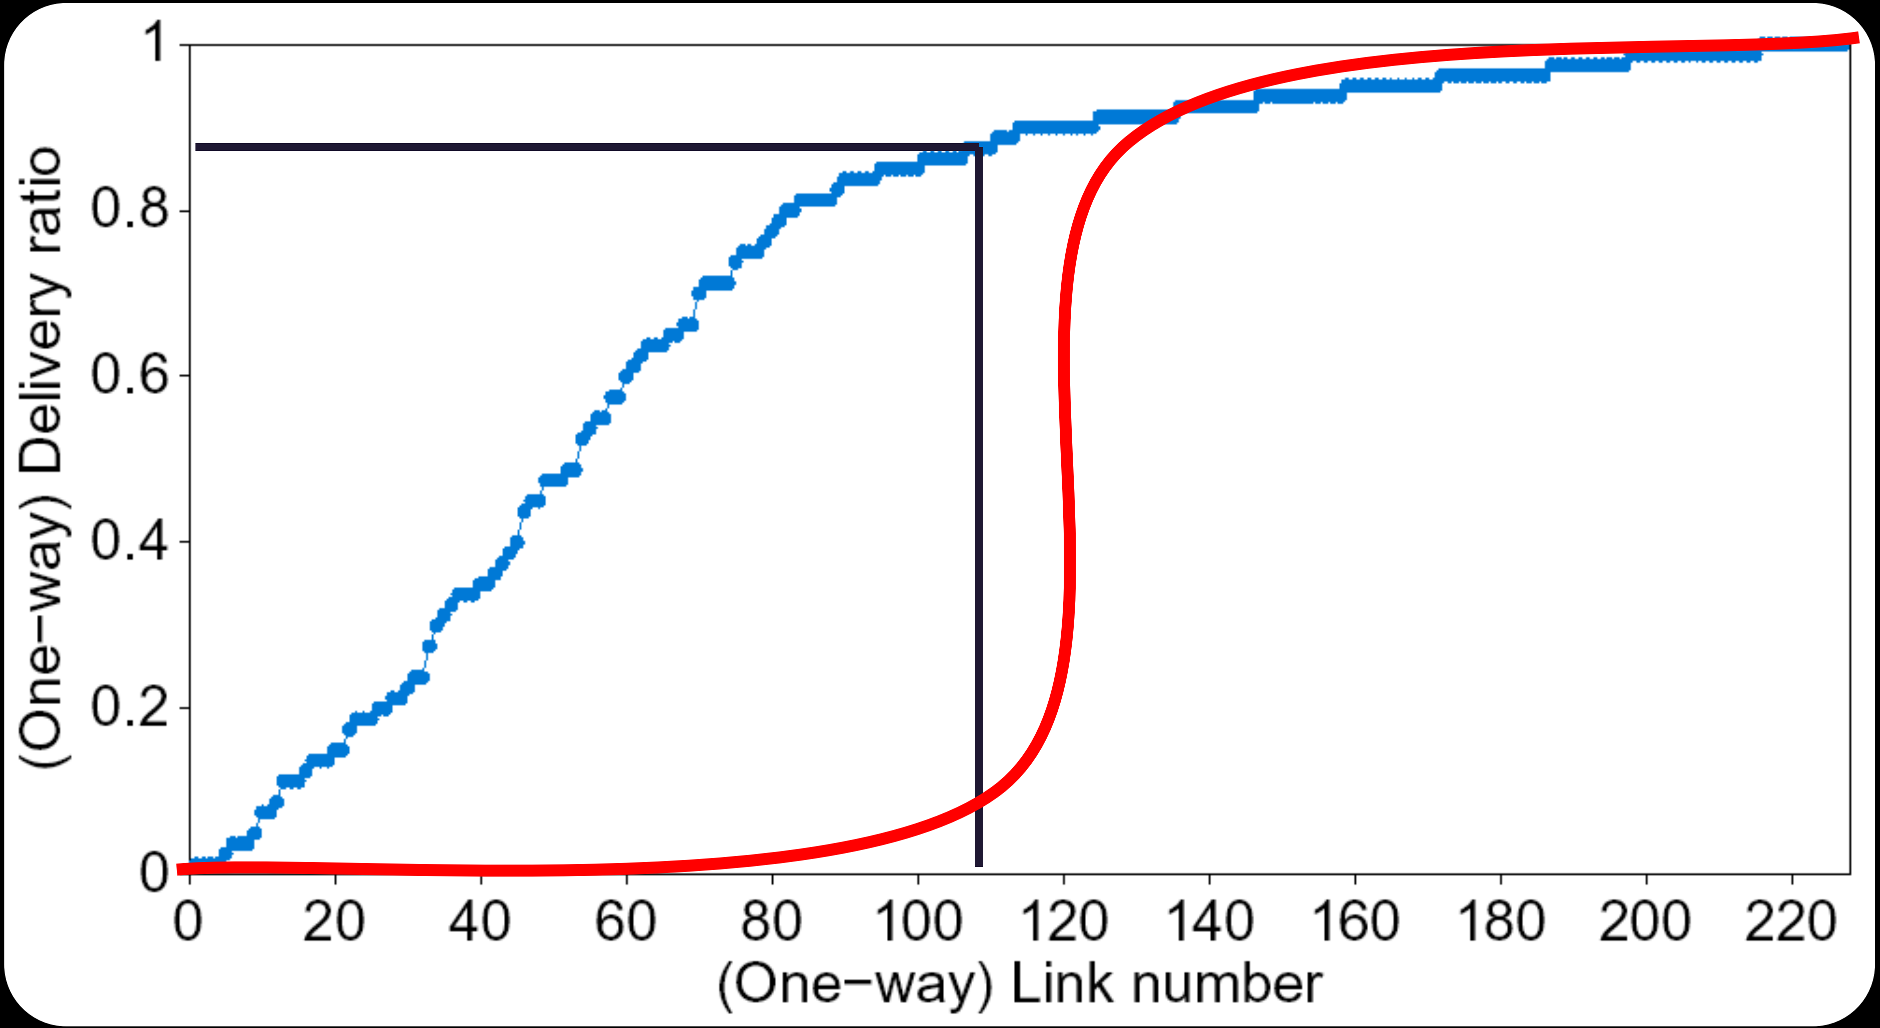

Smooth link distribution complicates link classification.

Approach 2:
Maximize bottleneck throughput

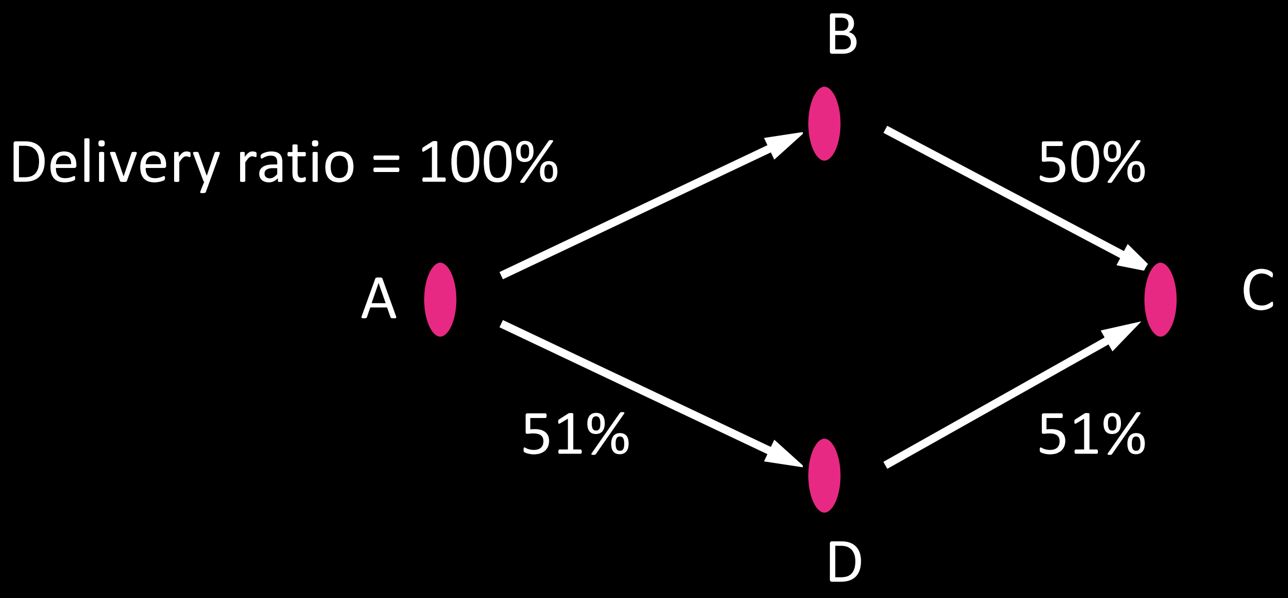

Bottleneck throughput:
- A-B-C = 50%
- A-D-C = <u>51%</u>

Actual throughput:
- A-B-C : ABBABBABB = <u>33%</u>
- A-D-C : AADDAADD = 25%

Pros? Cons?

Approach #3:
Maximize end-to-end delivery ratio

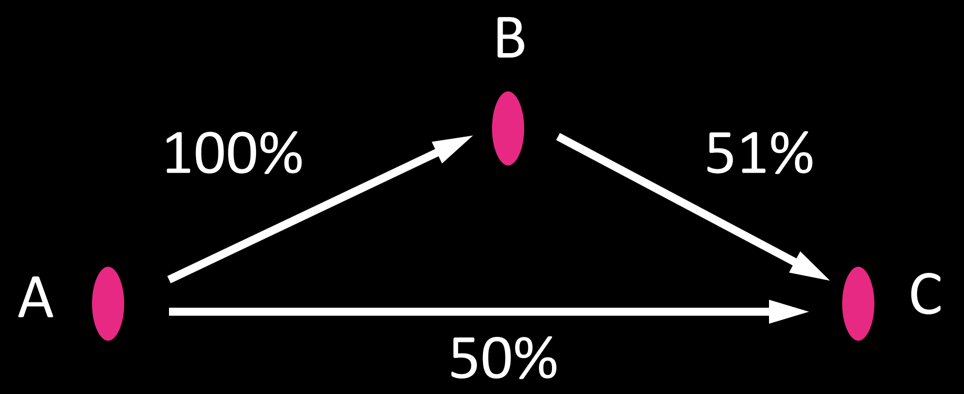

End-to-end delivery ratio:
- A-B-C = <u>51%</u>
- A-C = 50%

Actual throughput:
- A-B-C : ABBABBABB = 33%
- A-C : AAAAAAAA = <u>50%</u>

Pros? Cons?

Approach #4: Wireless routing metric: ETX

Minimize total transmissions per packet
(ETX, ‘Expected Transmission Count’)

Link throughput ≈ 1/ Link ETX

Delivery Ratio		Link ETX	Throughput
100%	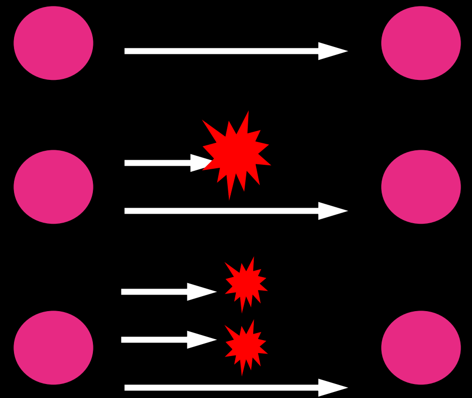	1	100%
50%		2	50%
33%		3	33%

Route ETX

Route ETX = Sum of link ETXs

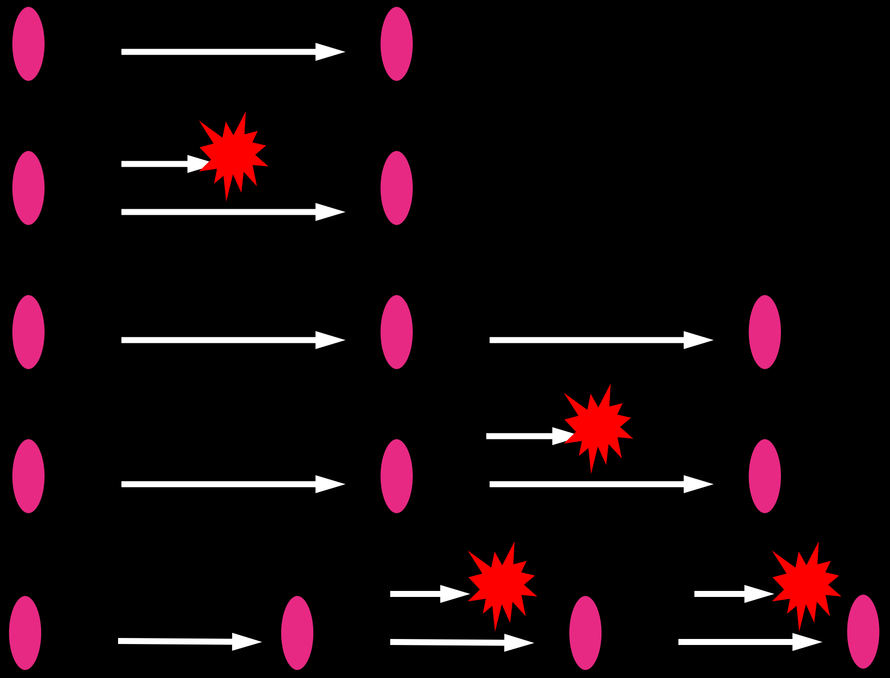

Route ETX	Throughput
1	100%
2	50%
2	50%
3	33%
5	20%

Calculating Link ETX

- Assuming 802.11 link-layer acknowledgments (ACKs) and retransmissions:
- P(TX success) = P(Data success) × P(ACK success)
- Link ETX = 1 / P(TX success)
  = 1 / [P(Data success) × P(ACK success)]
- Estimating link ETX:
- P(Data success) ≈ measured fwd delivery ratio r_{fwd}
- P(ACK success) ≈ measured rev delivery ratio r_{rev}
- Link ETX ≈ $1 / (r_{fwd} \times r_{rev})$

How can we measure delivery ratios?

- Each node broadcasts small link probes once per second
- Nodes remember probes received over past 10 seconds
- Reverse delivery ratios estimated as

  $$r_{\text{rev}} \approx \text{pkts received / pkts sent}$$

- Forward delivery ratios obtained from neighbors (piggybacked on probes)

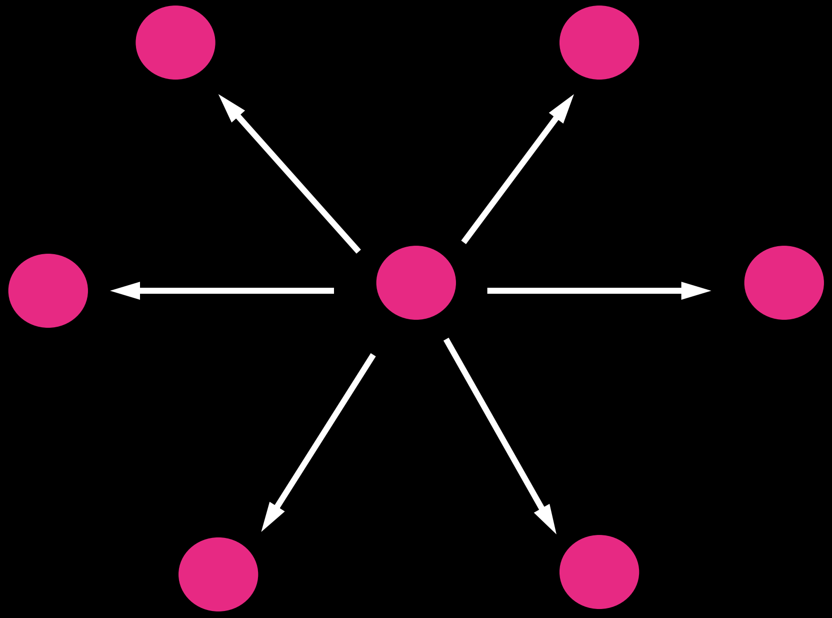

ETX Pros?

- ETX predicts throughput for short routes (1, 2, and 3 hops)
- ETX captures loss
- ETX captures asymmetry

ETX Caveats

- It is hard to measure link quality/loss
  - ➢ Changes as a function of load
  - ➢ Changes with time

- ETX ignores differences in bit-rate and packet size

- ETX ignores spatial re-use (i.e., assumes all links interfere)

IEEE SPECTRUM Topics Reports Blogs Multimedia

Why Wi-Fi Stinks—and How to Fix It

Neglected channels could add Wi-Fi capacity if router makers used them properly

By **Terry Ngo**

Global Market Insights
Insights to innovation.

Wireless Mesh Network Market revenue to hit USD 8 Bn by 2026, growing at around 15%: Global Market Insights, Inc.

Lebanon Protests: How To Communicate Securely in Case of a Network Disruption

October 18, 2019 News, Services

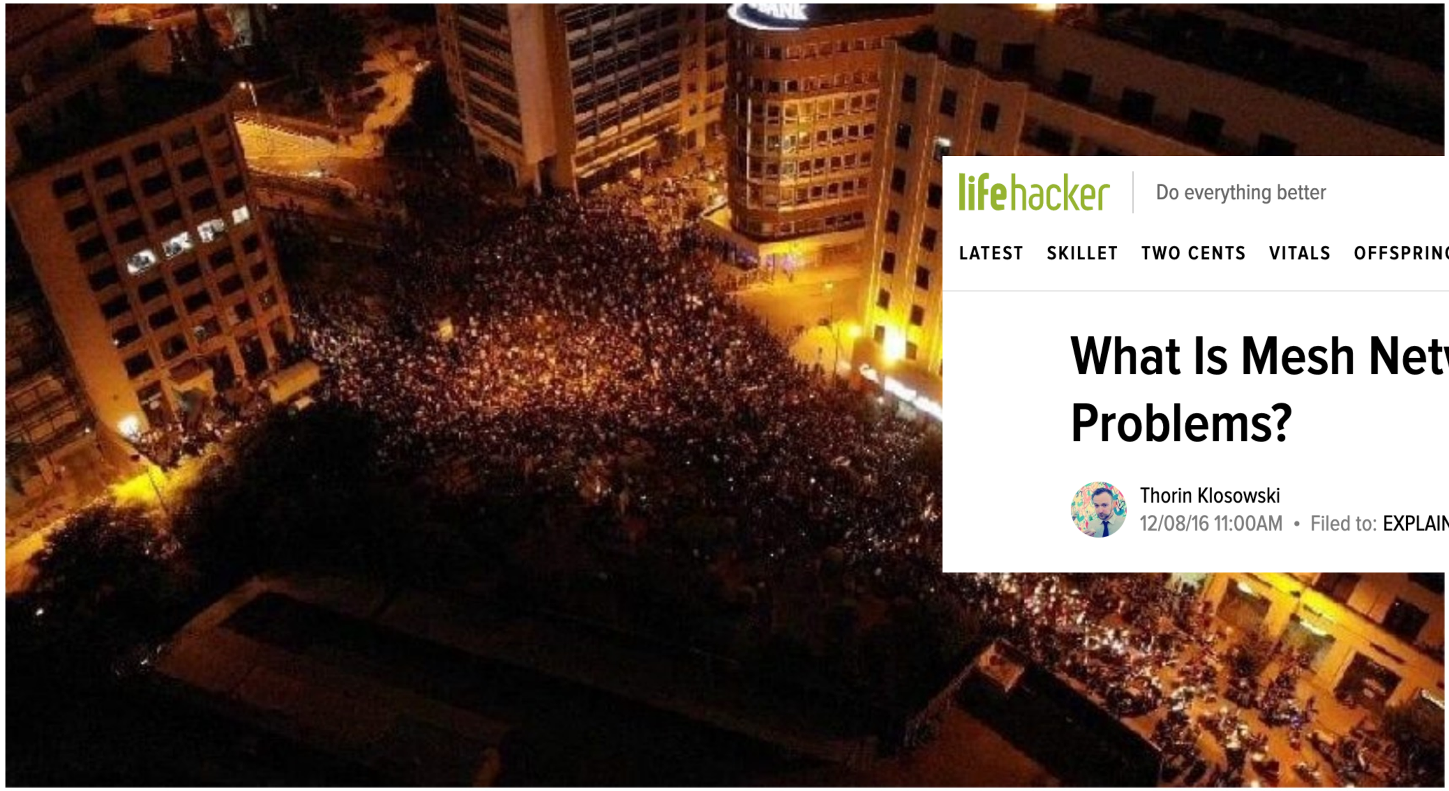

lifehacker Do everything better Subscribe

LATEST SKILLET TWO CENTS VITALS OFFSPRING THE UPGRADE APP DIRECTORY HOW I WORK VIDEO

What Is Mesh Networking and Will It Solve My Wi-Fi Problems?

Thorin Klosowski
12/08/16 11:00AM • Filed to: EXPLAINER

203.1K 69 6

Course Project

- Proposal discussion (+suggested ideas) after the spring break
- Proposal due one week after that
- Introductions?

Next Lecture

- Time: Wed Feb 28th
- Topic: Batteryless connectivity and RFID