Microsoft Research

人·計算·现实世界

Computing, People, and the Physical World

第十二届“二十一世纪的计算”学术研讨会

The 12th Computing in the 21st Century

The Internet of Every Thing - a step toward sustainability

David E. Culler

Professor, Computer Science Chair

University of California, Berkeley

Microsoft Research

The Internet … a decade ago

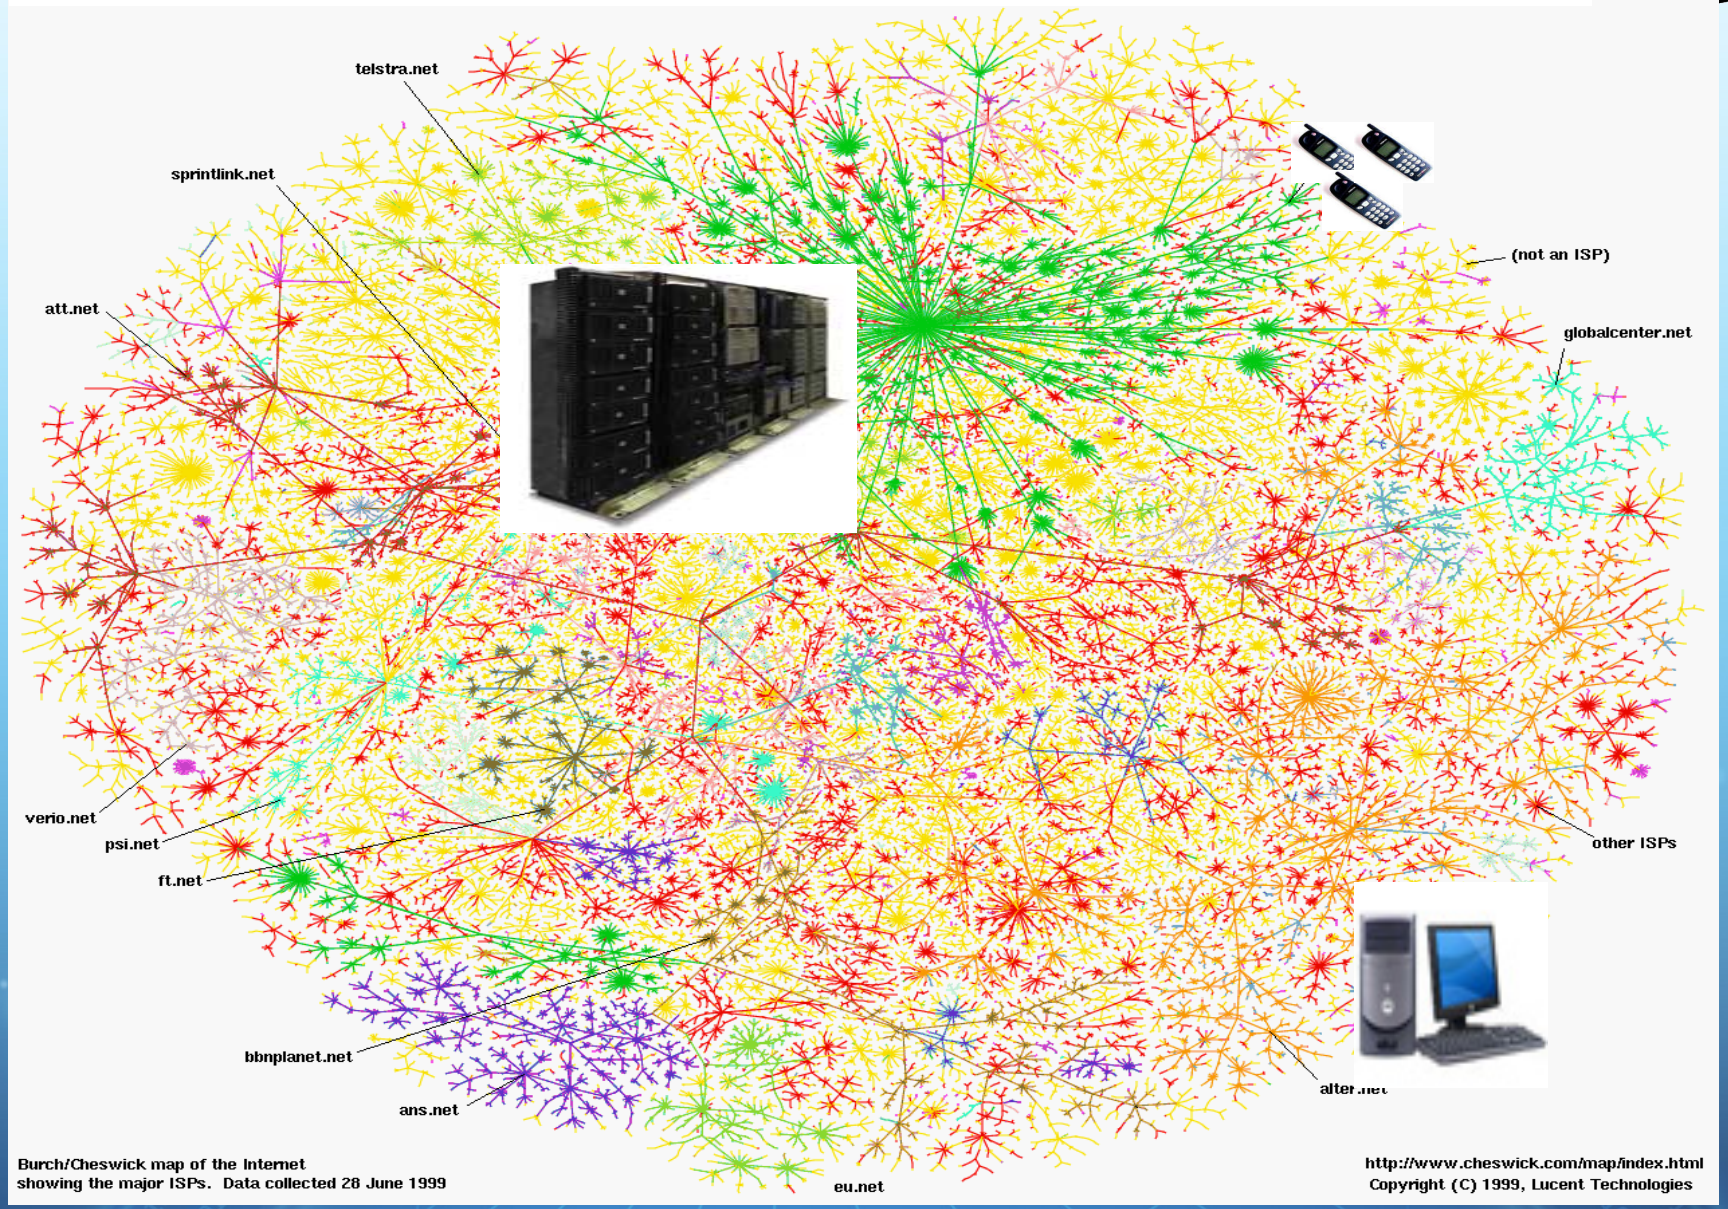

The Internet … that we envisioned

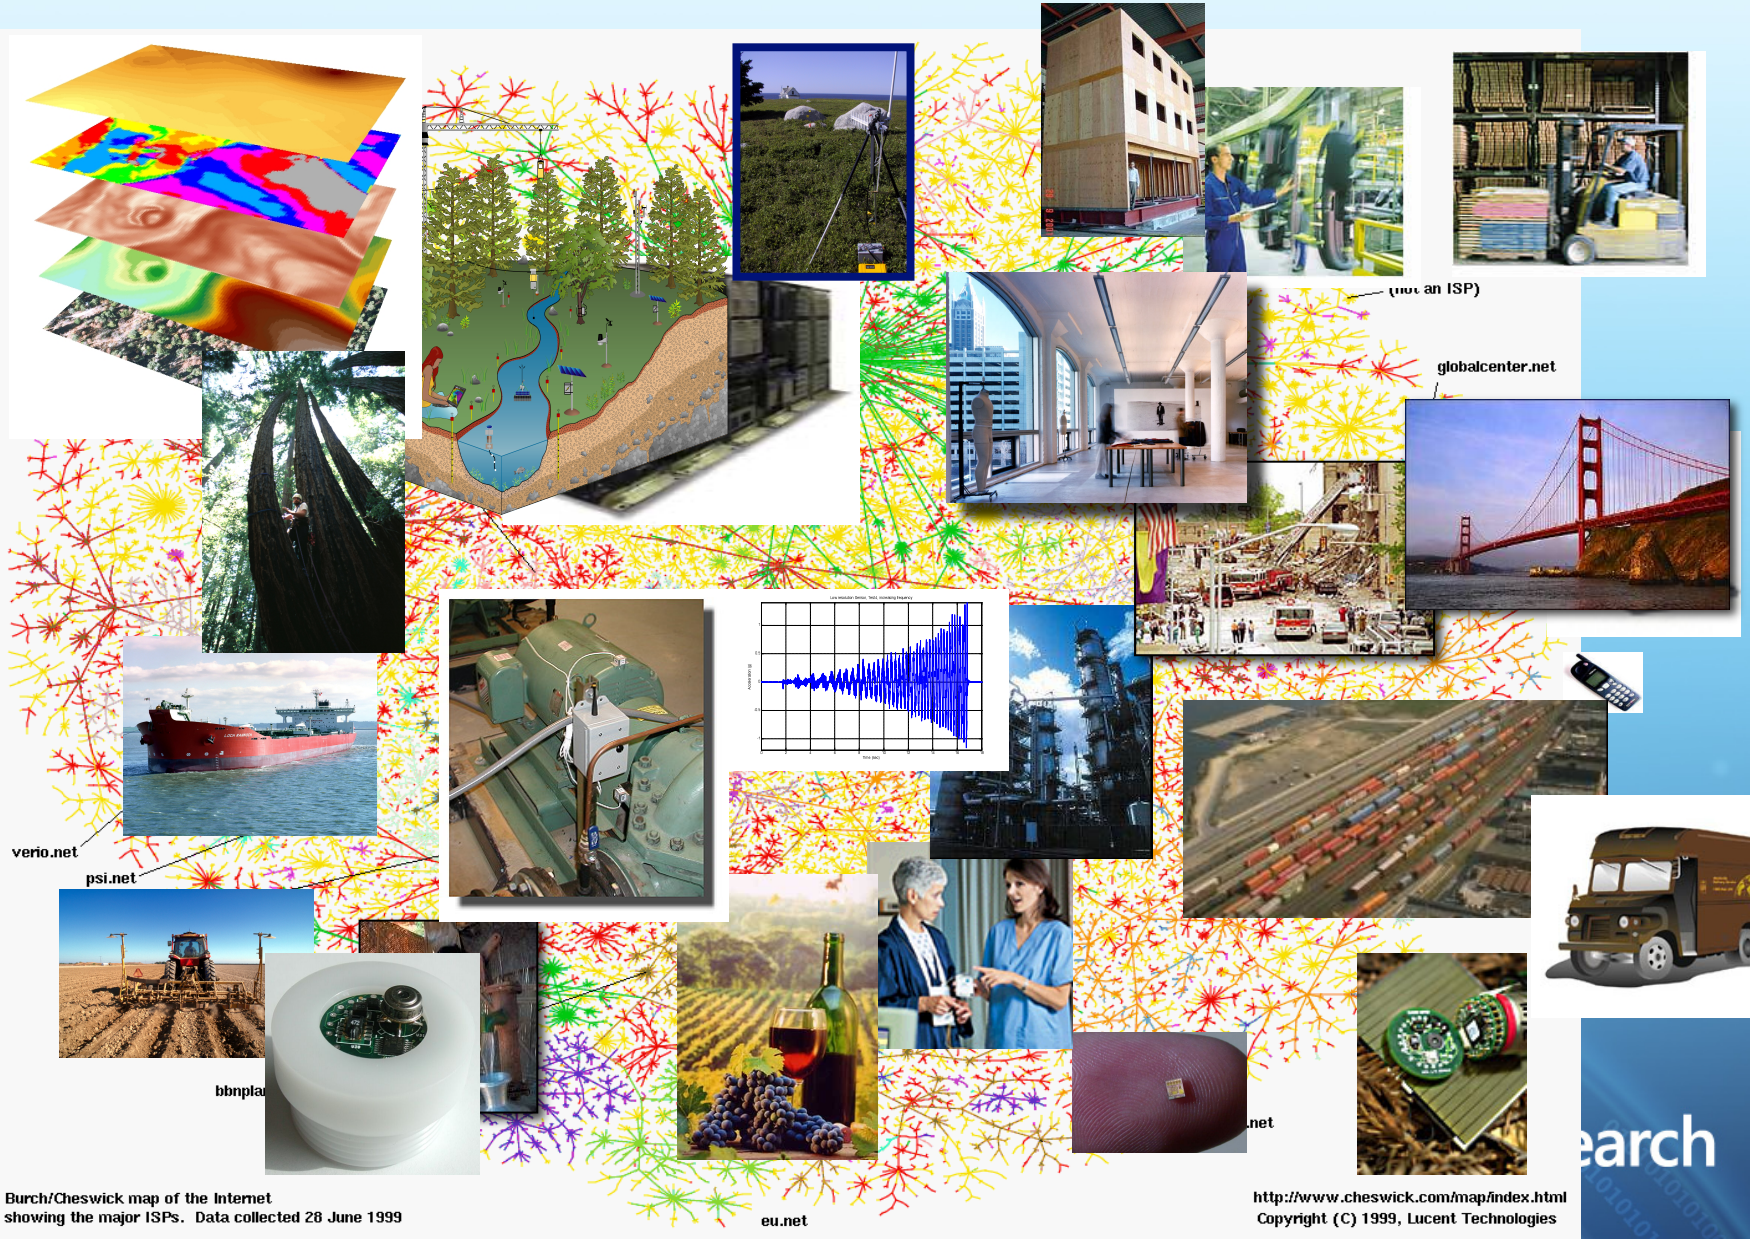

Why "Real" Information is so Important

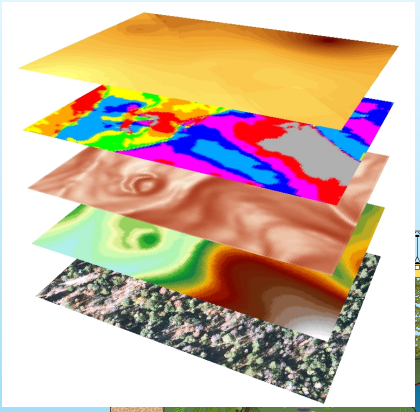

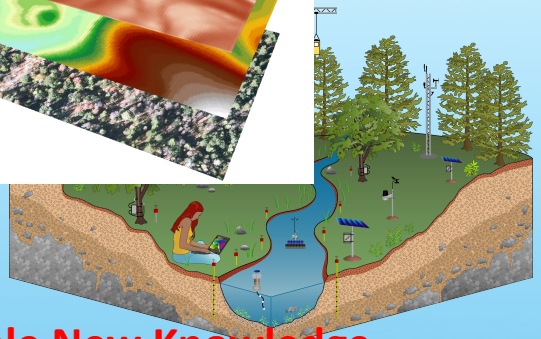

Enable New Knowledge

Save Resources

Improve Productivity

Increase Comfort

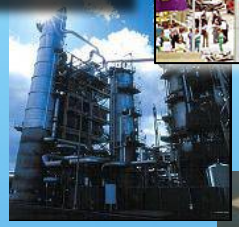

Enhance Safety & Security

High-Confidence Transport

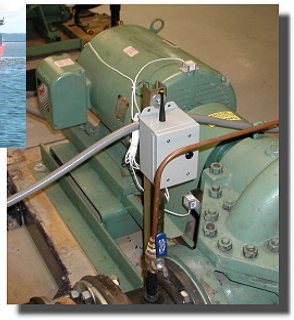

Preventing Failures

Improve Food & H20

Protect Health

Microsoft Research

Enabling Technology

Network

Microcontroller

Flash Storage

Radio Communication

Sensors

IEEE 802.15.4

Microsoft Research

Internet Concepts: Layering…

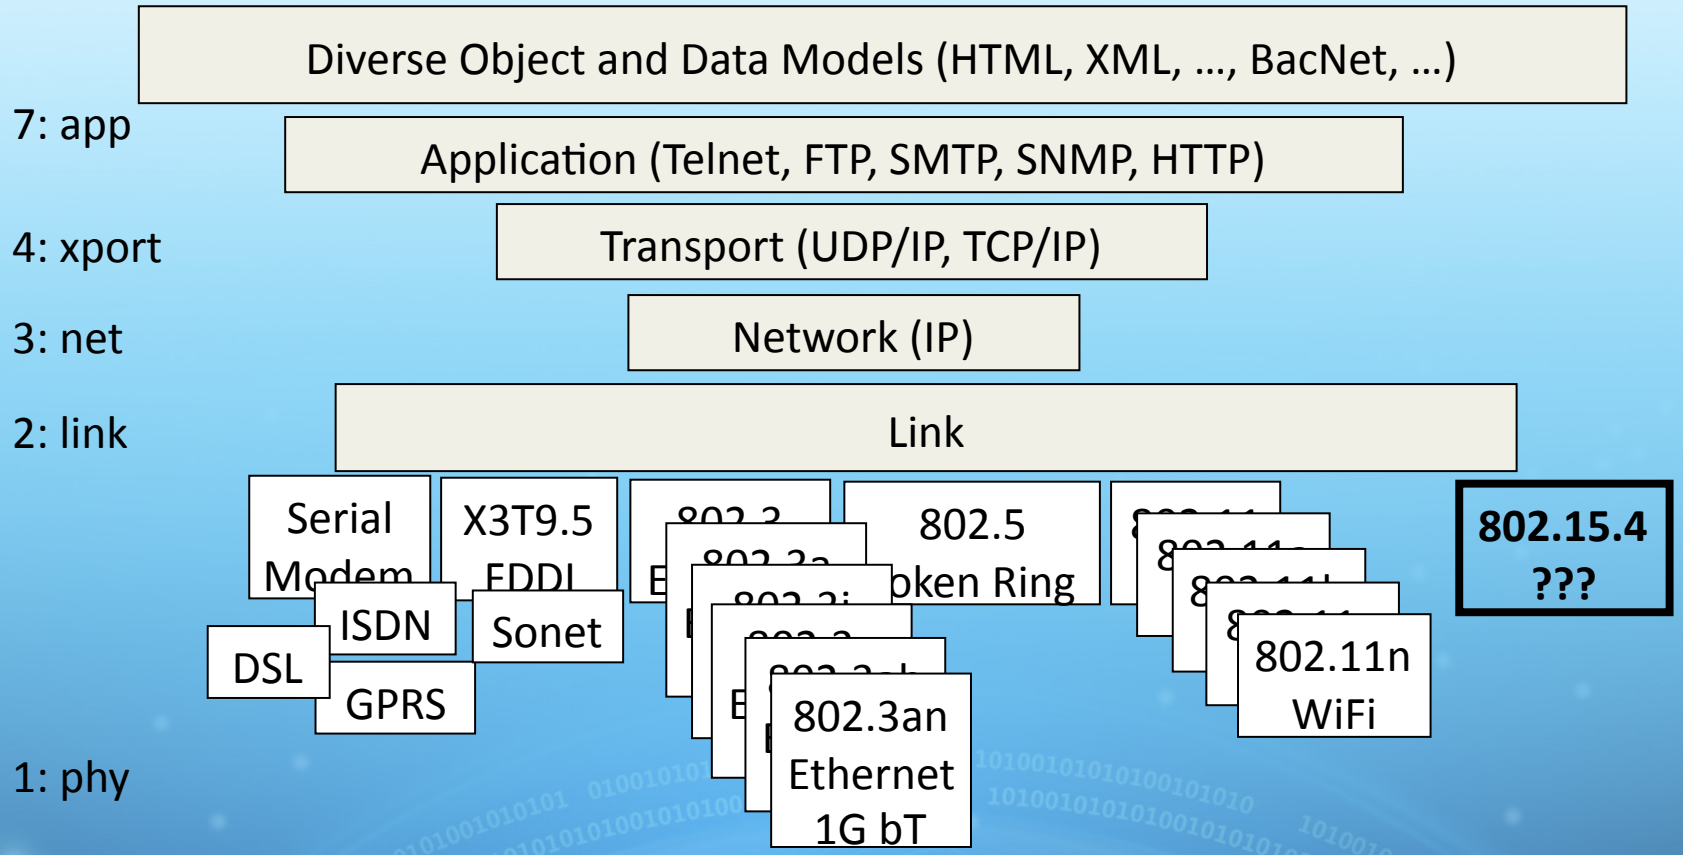

Leading Internet Research Perspective - a decade ago

- "Resource constraints may cause us to **give up the layered architecture**."
- "Sheer numbers of devices, and their unattended deployment, will preclude reliance on broadcast communication or the configuration currently needed to deploy and operate networked devices."
- "There are significant robustness and scalability advantages to designing applications using localized algorithms."
- "Unlike traditional networks, a sensor node may not need an identity (e.g. address)."
- "It is reasonable to assume that sensor networks can be tailored to the application at hand."

We were wrong…

Microsoft Research

The Mote/TinyOS revolution…

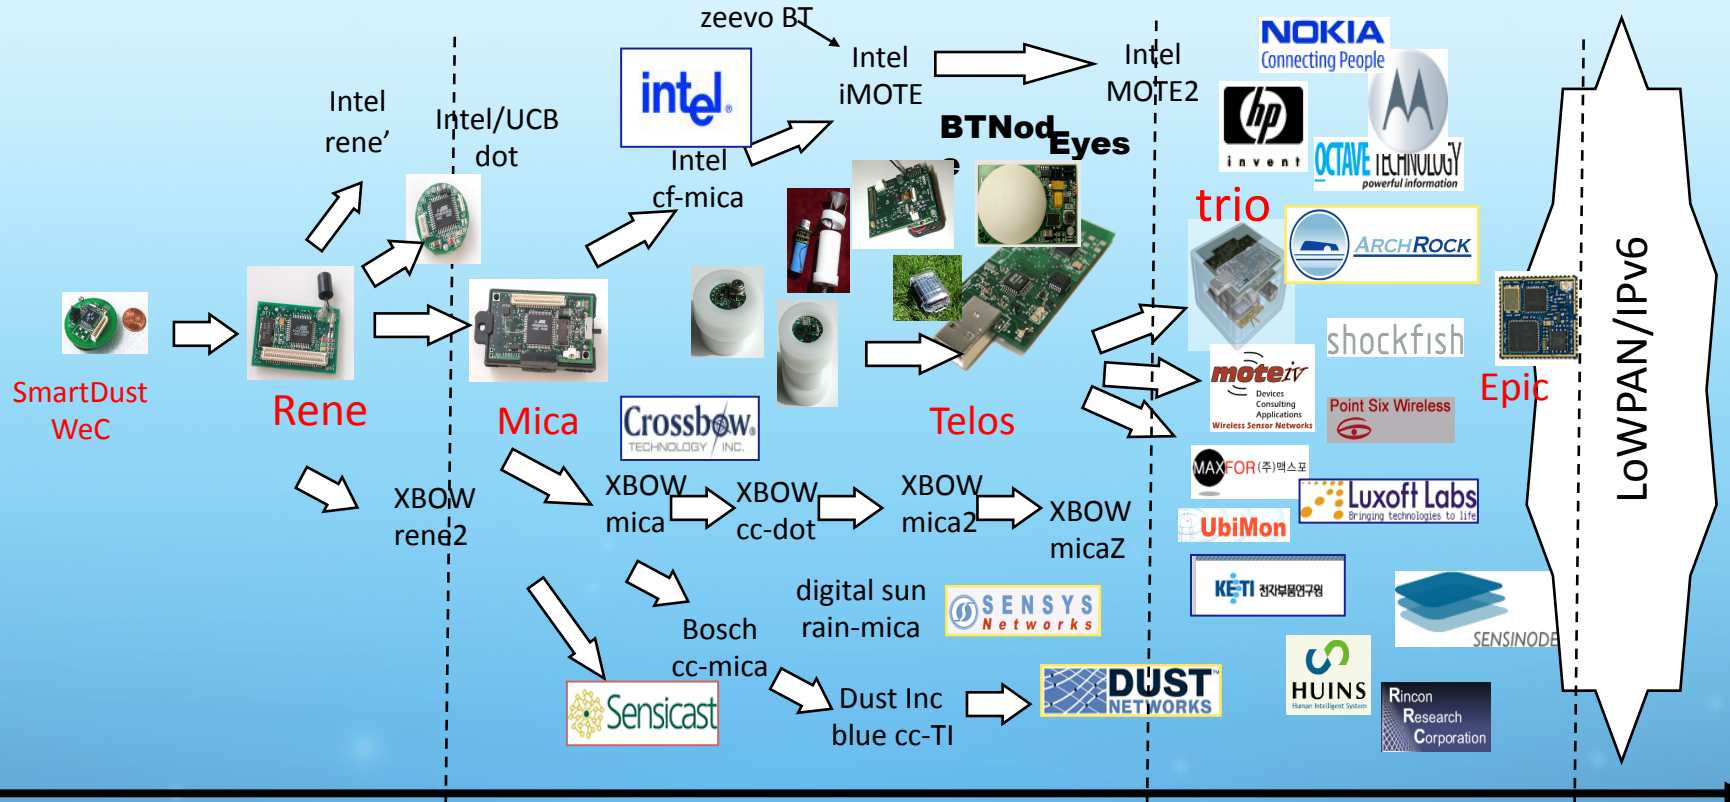

Internet of Every Thing – Realized 2008

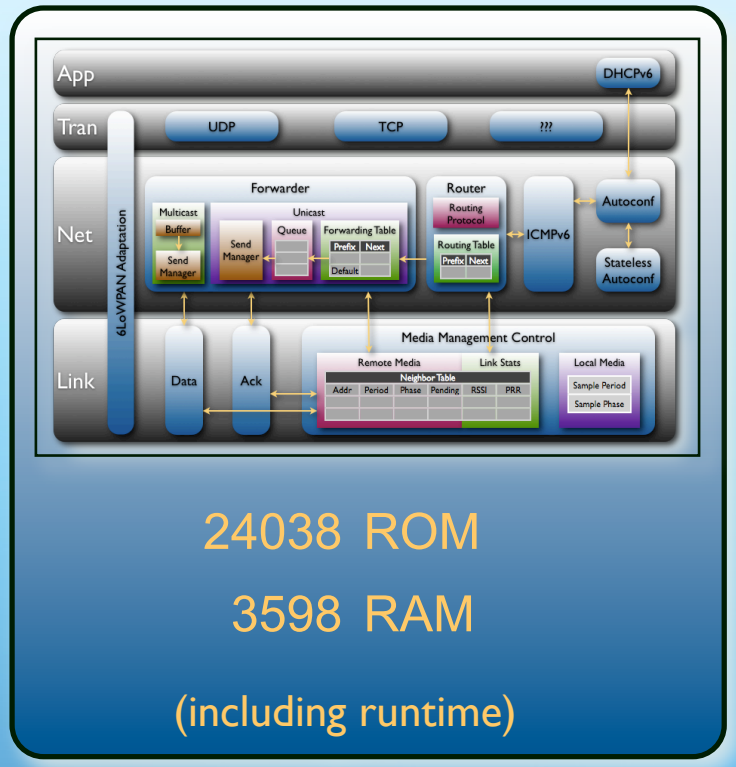

24038 ROM

3598 RAM

(including runtime)

* Production implementation on TI msp430/cc2420

- Footprint, power, packet size, & bandwidth

	ROM	RAM
CC2420 Driver	3149	272
802.15.4 Encryption	1194	101
Media Access Control	330	9
Media Management Control	1348	20
6LoWPAN + IPv6	2550	0
Checksums	134	0
SLAAC	216	32
DHCPv6 Client	212	3
DHCPv6 Proxy	104	2
ICMPv6	522	0
Unicast Forwarder	1158	451
Multicast Forwarder	352	4
Message Buffers	0	2048
Router	2050	106
UDP	450	6
TCP	1674	50

Microsoft Research

Internet of Every Thing – standardized 2010

ROLL
Internet-Draft
Intended status: Standards Track
Expires: April 4, 2011

T. Winter, Ed.

P. Thubert, Ed.
Cisco Systems
A. Brandt
Sigma Designs
T. Clausen
LIX, Ecole Polytechnique
J. Hui
Arch Rock Corporation
R. Kelsey
Ember Corporation
P. Levis
Stanford University
K. Pister
Dust Networks
R. Struik

JP. Vasseur
Cisco Systems

October 1, 2010

RPL: IPv6 Routing Protocol for Low power and Lossy Networks
draft-ietf-roll-rpl-12

Abstract

Low power and Lossy Networks (LLNs) are a class of network in which both the routers and their interconnect are constrained. LLN routers

2008-02-15 charter

Routing Over Low power and Lossy networks (roll)

Charter

Current Status: Active Working Group

Chair(s):
JP Vasseur <jpv@cisco.com>
David Culler <culler@eecs.berkeley.edu>

IPSO Alliance
Enabling the
INTERNET OF THINGS

ZigBee Alliance

ZigBee Smart Energy Version 2.0 Documents

ZigBee Smart Energy version 2.0 will be IP-based and offer a variety of new features.

ISA arch

Networking in the Physical World

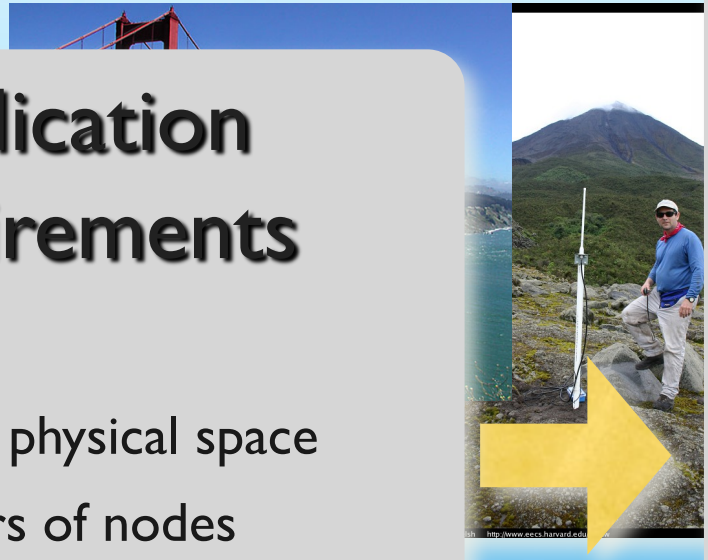

Application Requirements

- Embedding in physical space
- Large numbers of nodes
- Low total cost of ownership

Networking Challenges

- Variable communication
  - Unknown obstacles
  - Variable density and loss
- Constrained resources
  - Limited routing state
  - Limited throughput
  - Limited buffering
- Low power wireless
  - Multihop
  - Low SNR
  - Small MTU

Confluence Research, Industry, and Standards

Research Community

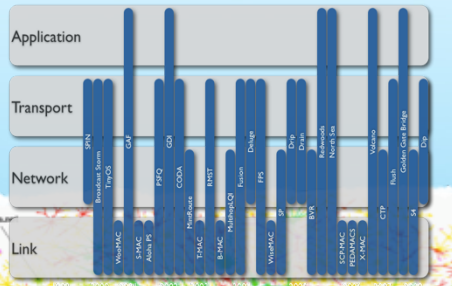

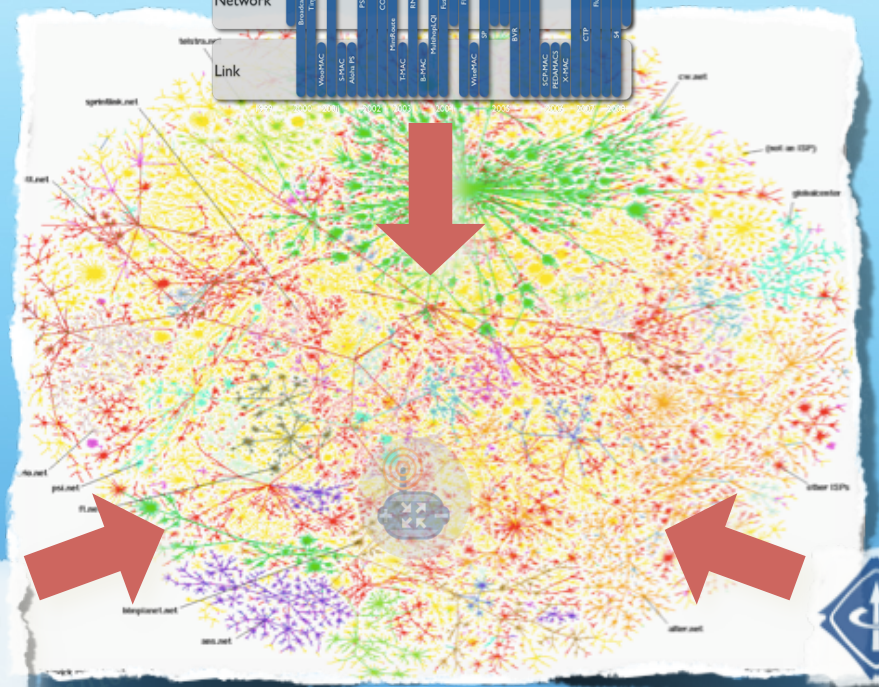

IETF

IEEE

Microsoft Research

Key Research Developments

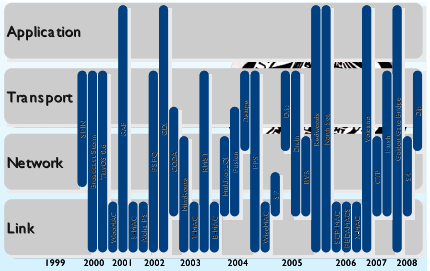

- Event-Driven Component-Base Operating System
  - Framework for building System & Network abstractions
  - Low-Power Protocols
  - Hardware and Application Specific
- Idle listening
  - All the energy is consumed by listening for a packet to receive

  => Turn radio on only when there is something to hear
- Reliable routing on Low-Power & Lossy Links
  - Power, Range, Obstructions => multi-hop
  - Always at edge of SNR => loss is common

  => monitoring, retransmission, and local rerouting
- Trickle – don't flood (tx rate < 1/density, and < info change)
  - Connectivity is determined by physical points of interest, not network designer.
  - never naively respond to a broadcast
  - re-broadcast very very politely

Microsoft Research

TinyOS – Framework for Innovation

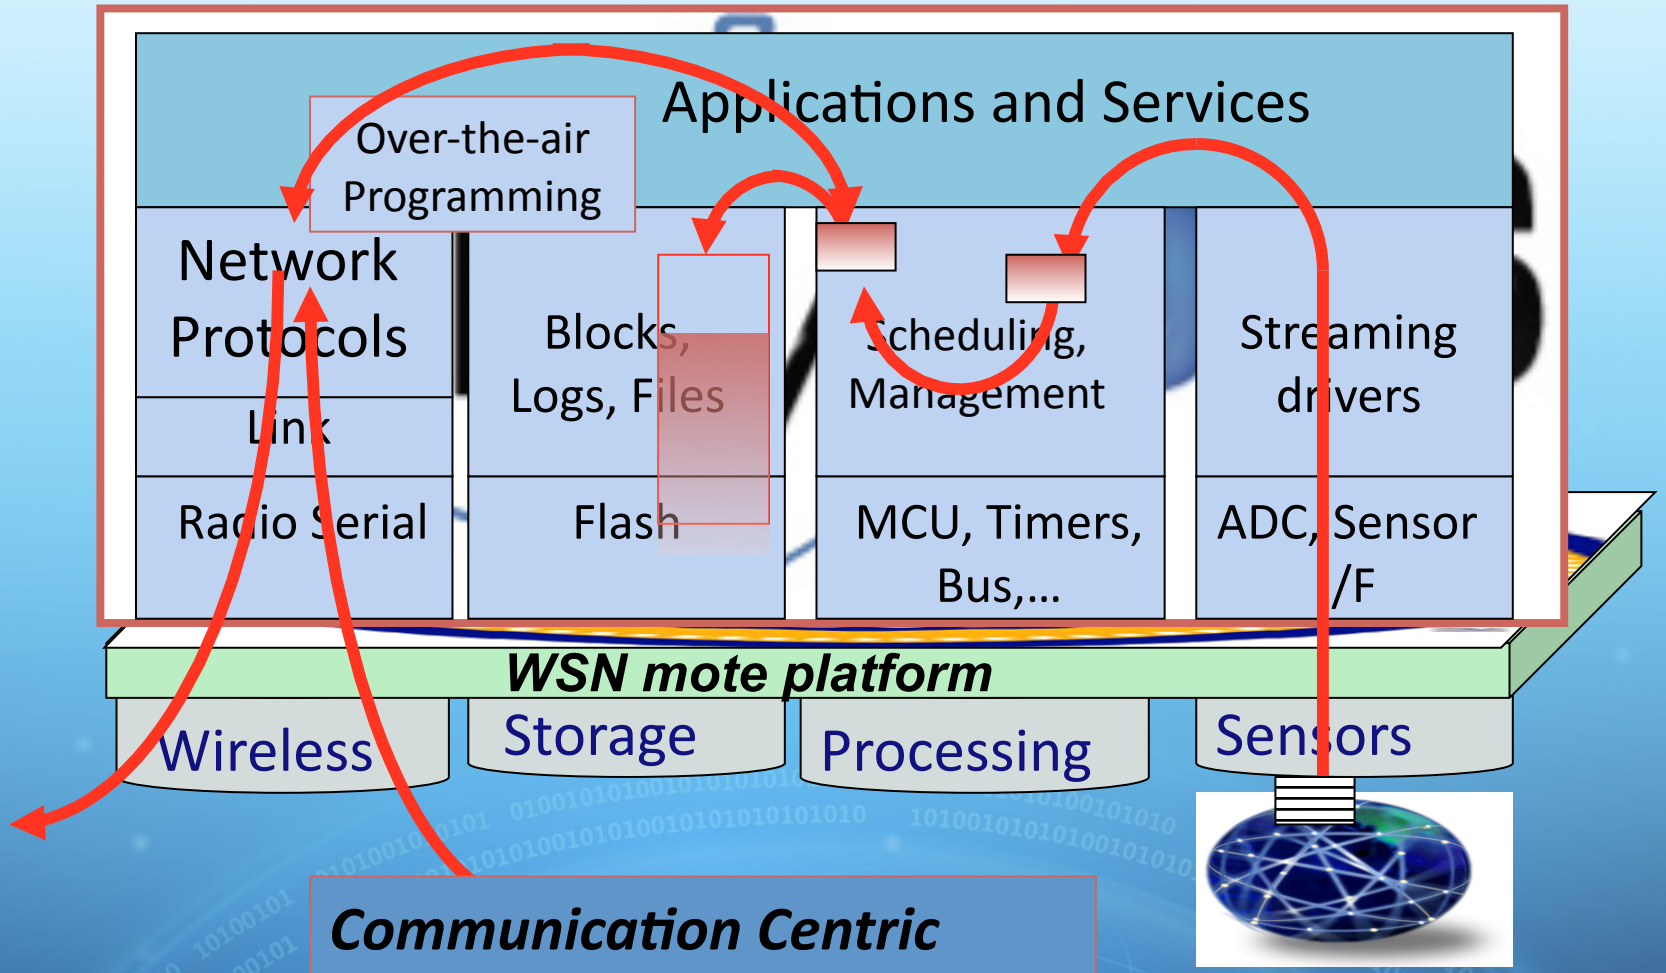

A worldwide community

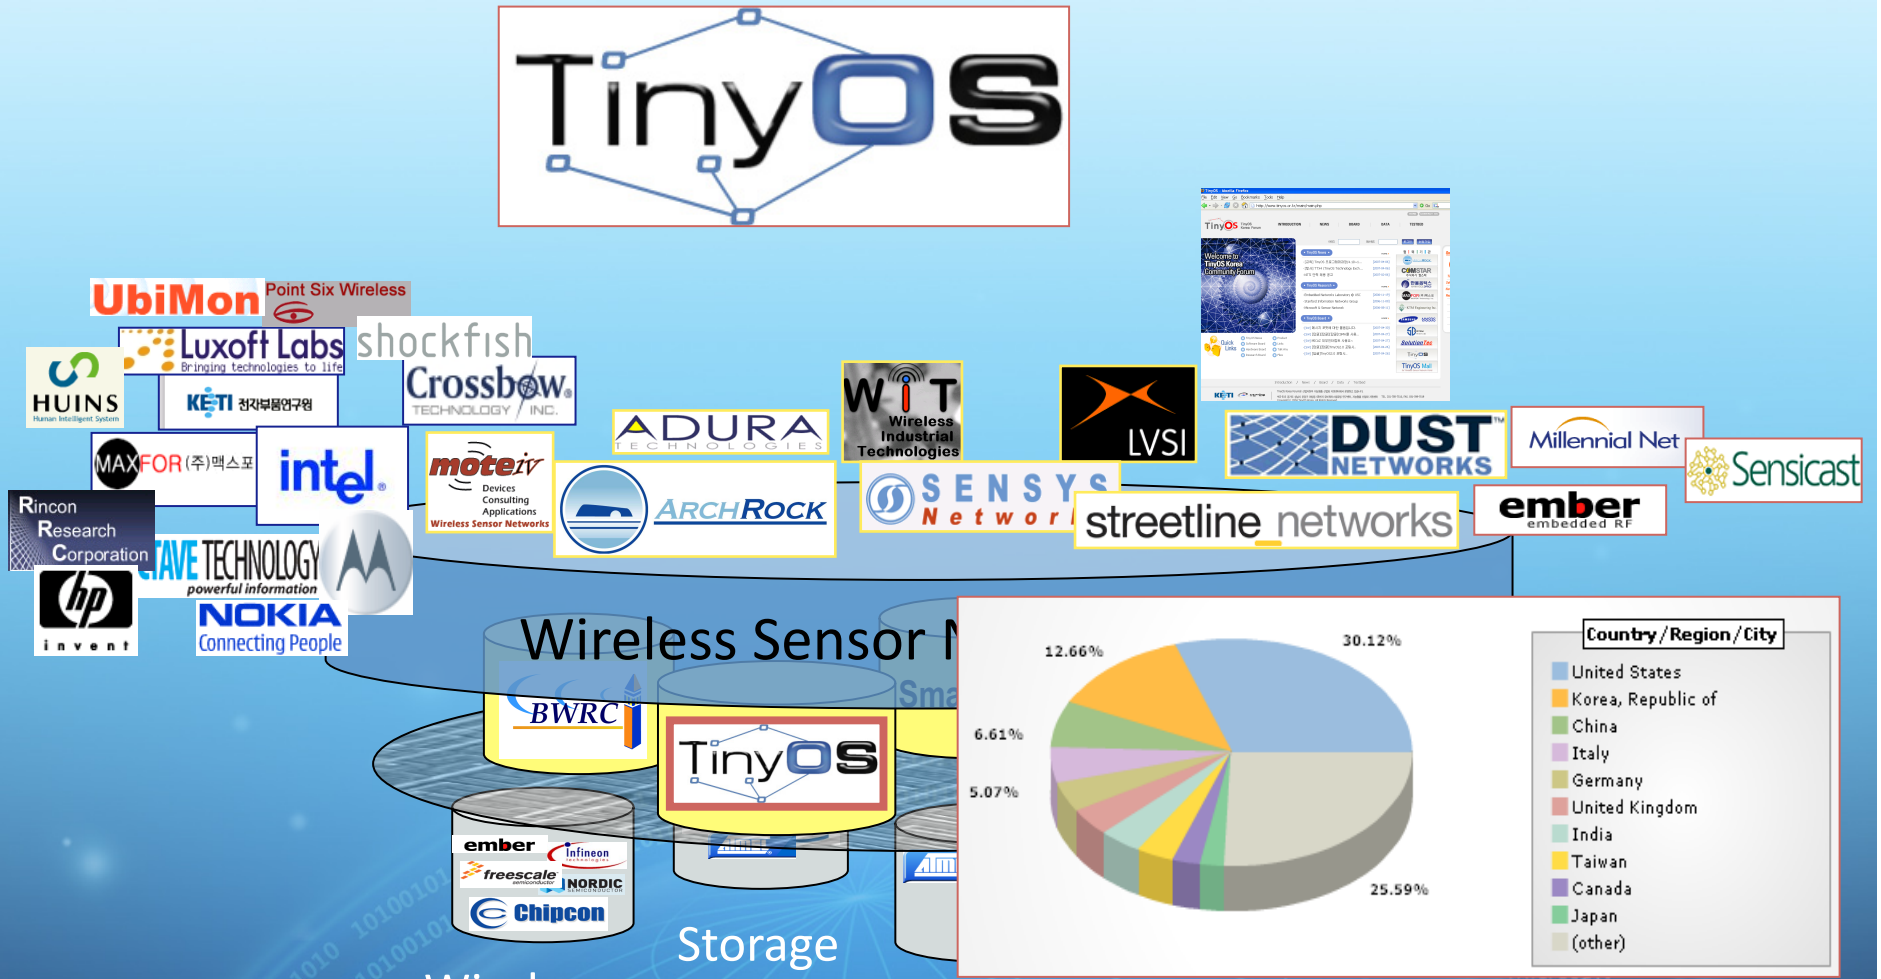

Low Power Networking in the Real World

Applications

Network

system

architecture

Technology

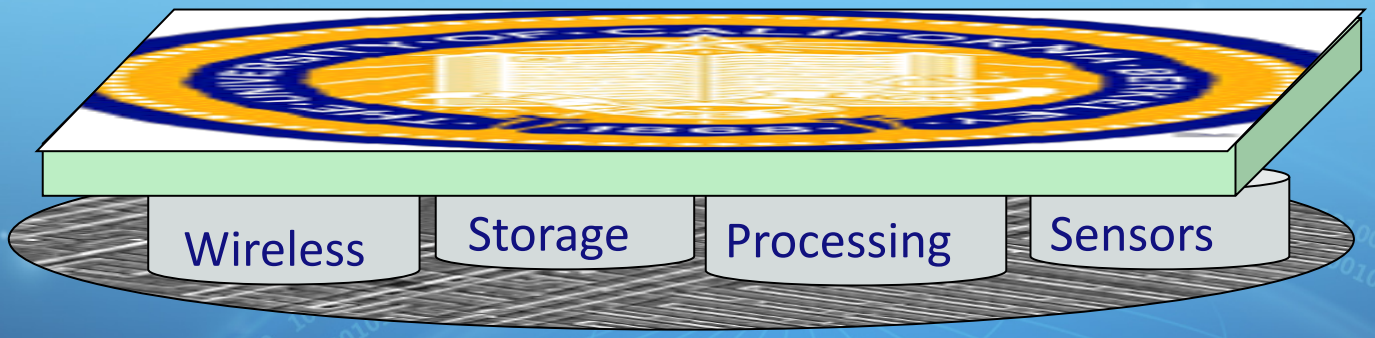

Microsoft Research

A Low-Power Standard Link

	802.15.4	802.15.1	802.15.3	802.11	802.3
Class	WPAN	WPAN	WPAN	WLAN	LAN
Lifetime (days)	100-1000+	1-7	Powered	0.1-5	Powered
Net Size	65535	7	243	30	1024
BW (kbps)	20-250	720	11,000+	11,000+	100,000+
Range (m)	1-75+	1-10+	10	1-100	185 (wired)
Goals	Low Power, Large Scale, Low Cost	Cable Replacement	Cable Replacement	Throughput	Throughput

- Low Transmit power, Low Signal-to-noise Ratio (SNR), modest BW, Little Frames

Microsoft Research

Decade of Networking (sans Architecture)

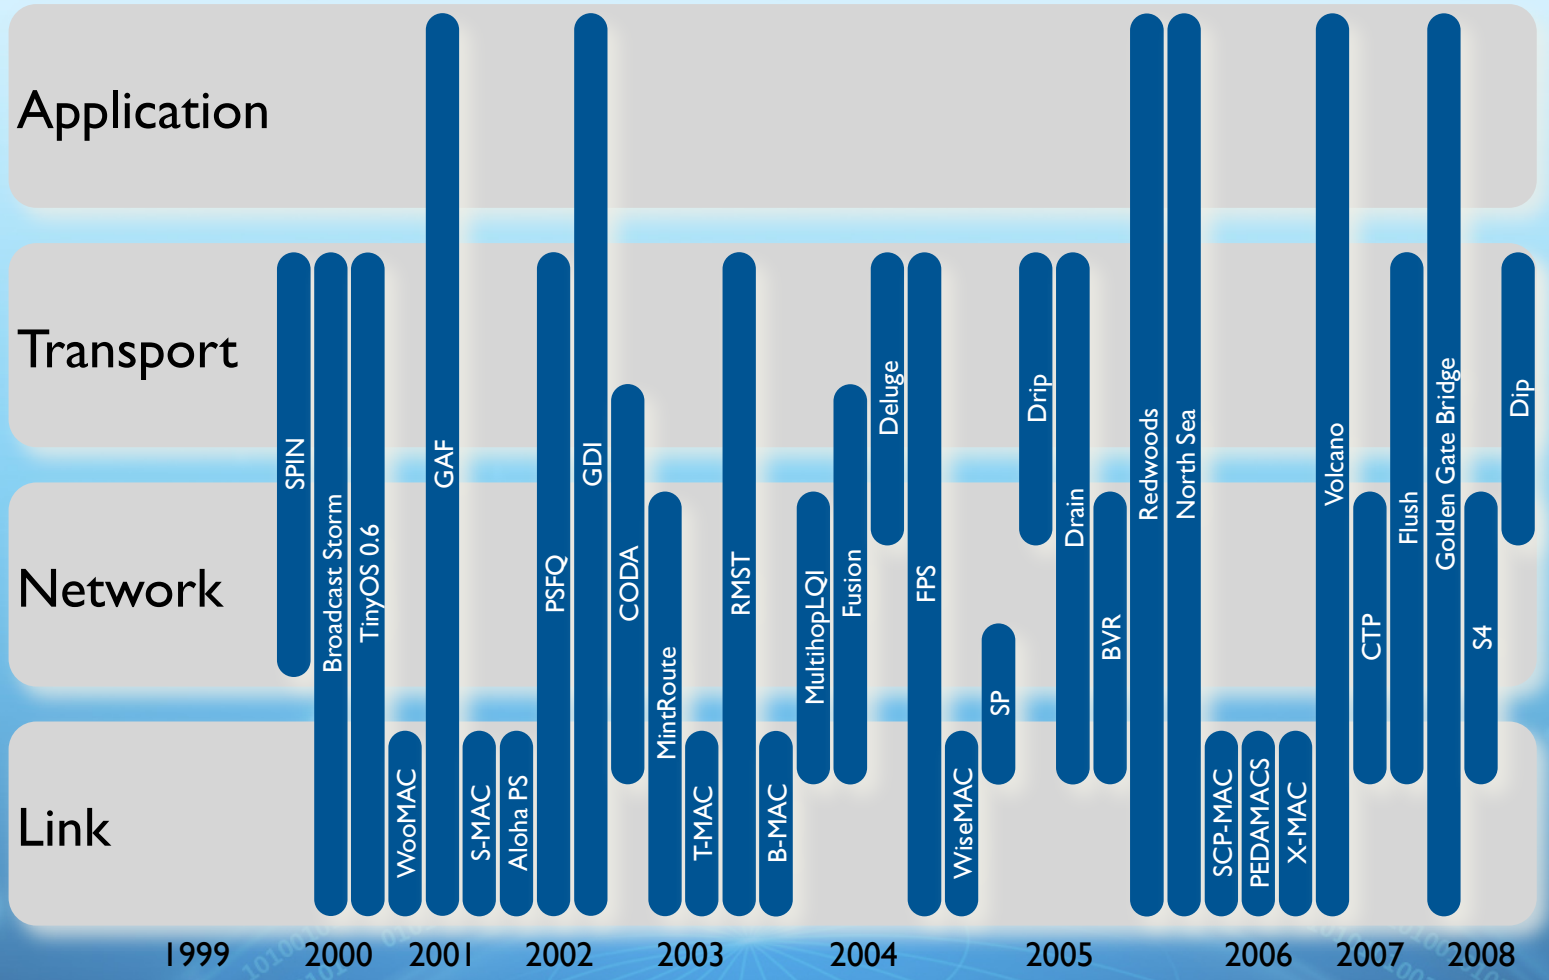

The "Idle Listening" Problem

- The power consumption of "short range" (i.e., low-power) wireless communications is roughly the same when
  - transmitting,
  - receiving,
  - or simply ON, "listening" for potential reception.
  - IEEE 802.15.4, Zwave, Bluetooth, …, WiFi
- Radio must be ON (listening) in order receive anything.
  - Transmission is rare
  - Listening happens all the time

⇒ Energy consumption dominated by *idle listening*

⇒ *Do Nothing Well*

Microsoft Research

Communication Power – Passive Vigilance

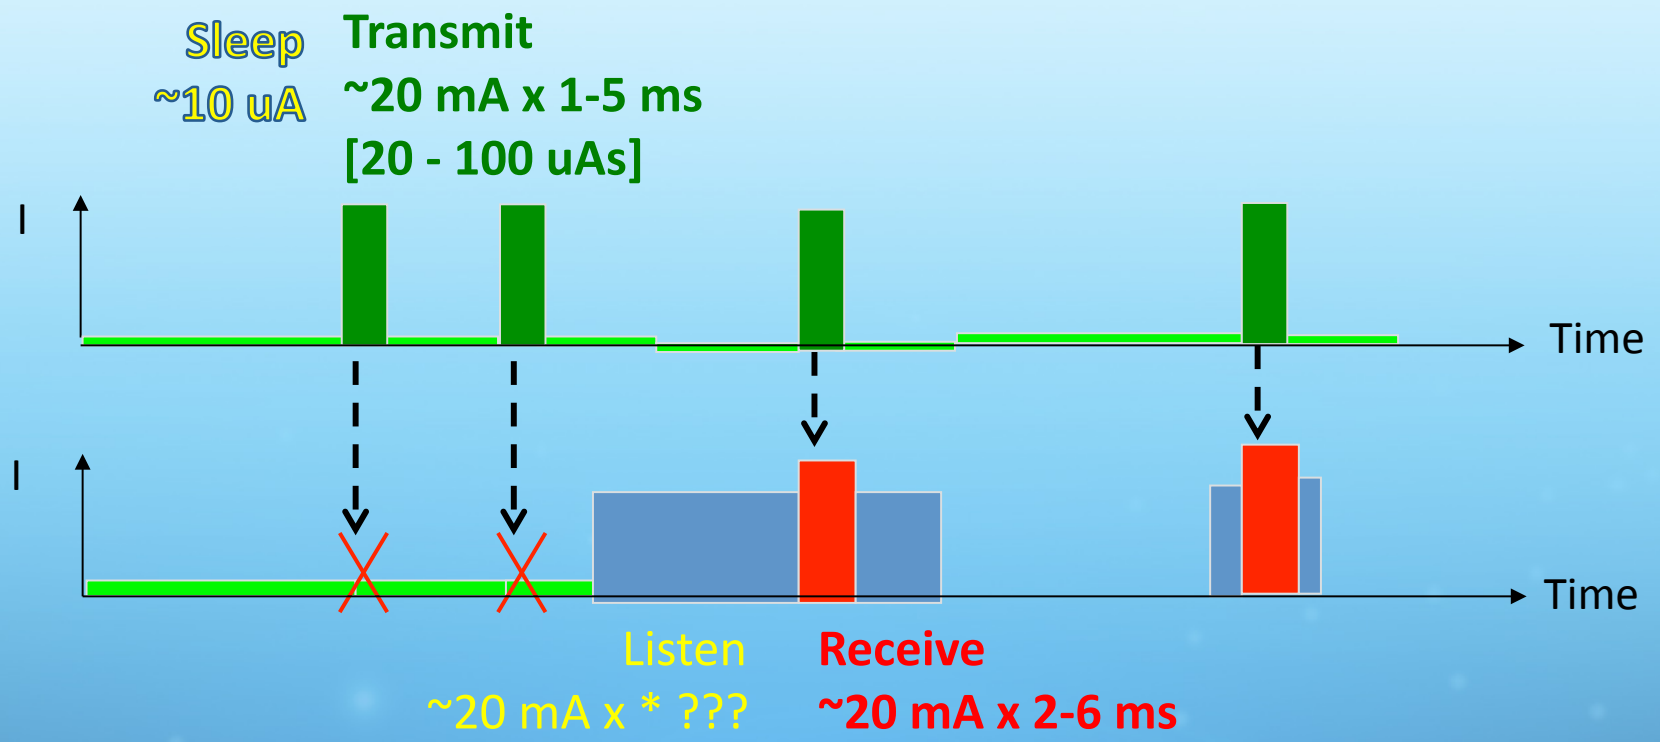

- Listen just when there is something to hear …

3 Basic Solution Techniques

- Scheduled Listening
  - Arrange a schedule of communication Time Slots
  - Maintain coordinated clocks and schedule
  - Listen during specific “slots”
  - Many variants:
    - Aloha, Token-Ring, TDMA, Beacons, Bluetooth piconets, …
    - S-MAC, T-MAC, PEDAMACS, TSMP, FPS, …
- Sampled Listening
  - Listen for very short intervals to detect eminent transmissions
  - On detection, listen actively to receive
  - DARPA packet radio, LPL, BMAC, XMAC, …
  - Maintain “always on” illusion, Robust
- Listen after send (with powered infrastructure)
  - After transmit to a receptive device, listen for a short time
  - Many variants: 802.11 AMAT, Key fobs, remote modems, …
- Many hybrids possible

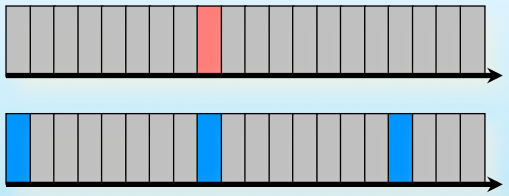

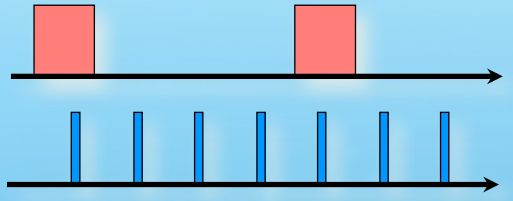

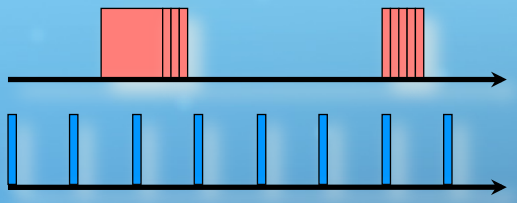

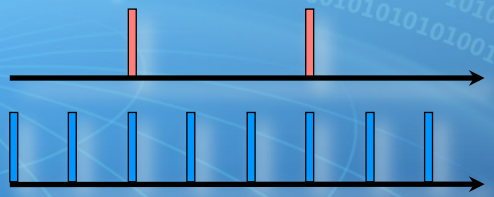

Microsoft Research

Scheduled Listen

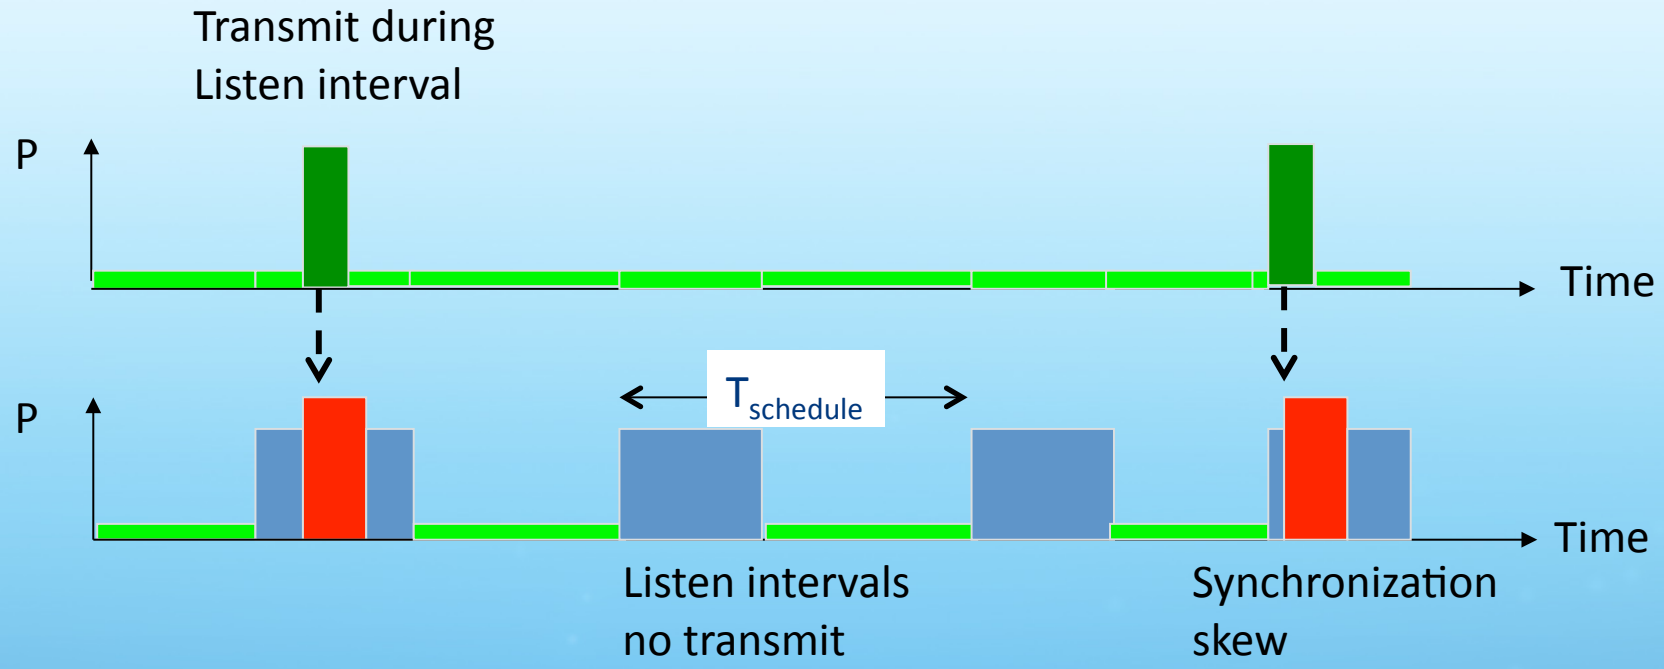

- Communication to maintain synch. sets lower bound.
- Full power listen to discover and join schedule. (!!!)
- Listen in slot, even when there is nothing to hear
- Partitions limited bandwidth, adds to large latency

Sampled Listen

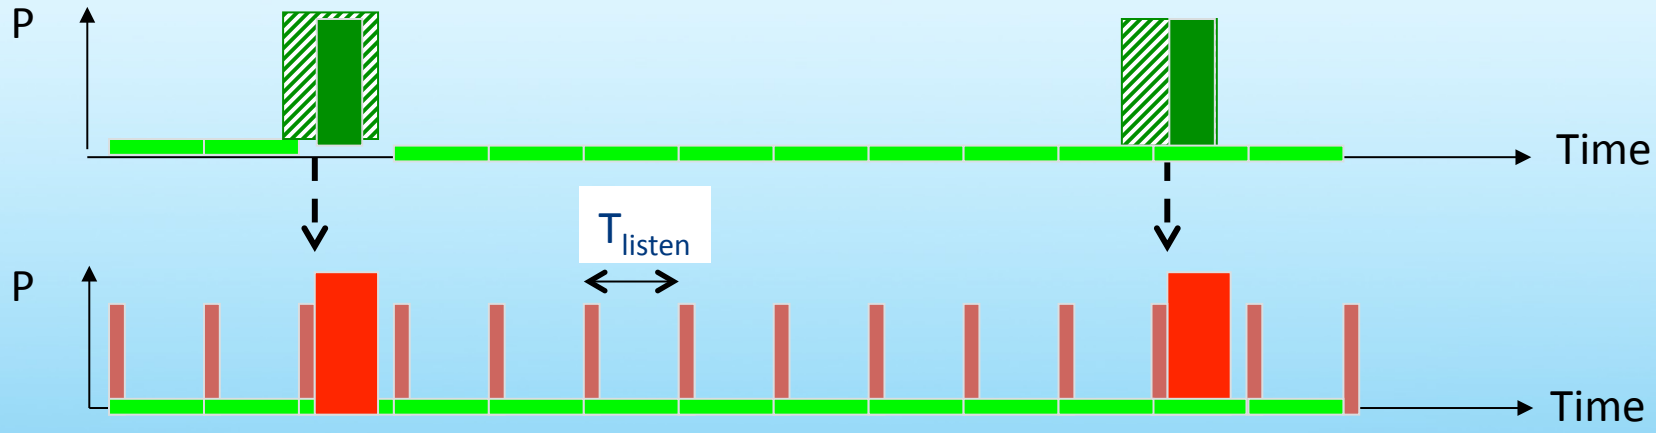

- Frequent, very short "Listen Samples"
- Increase (infrequent) Transmit to cover sampling
- No synchronization required.
- Rapid join and at low power.
- Optimizations built on top
  - Scheduled sleep, wake times, time synch,

=> Always-On illusion at very low power

Microsoft Research

3 Basic Solution Techniques

- Scheduled Listening
  - Arrange a schedule of communication Time Slots
  - Maintain coordinated clocks and schedule
  - Listen during specific "slots"
  - Many variants:
    - Aloha, Token-Ring, TDMA, Beacons, Bluetooth piconets, …
    - S-MAC, T-MAC, PEDAMACS, TSMP, FPS, …

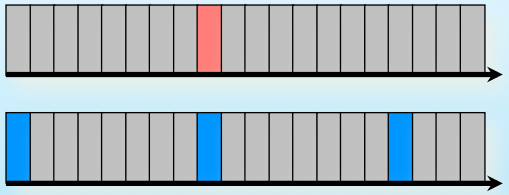

- Sampled Listening
  - Listen for very short intervals to detect eminent transmissions
  - On detection, listen actively to receive
  - DARPA packet radio, LPL, BMAC, XMAC,
  - Maintain "always on" illusion, Robust

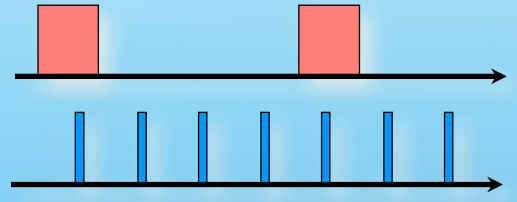

- Listen after send (with powered infrastructure)
  - After transmit to a receptive device, listen for a short time
  - Many variants: 802.11 AMAT, Key fobs, remote modems, …
- Many hybrids possible

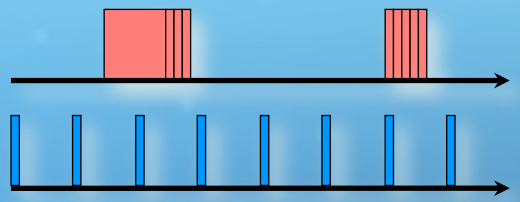

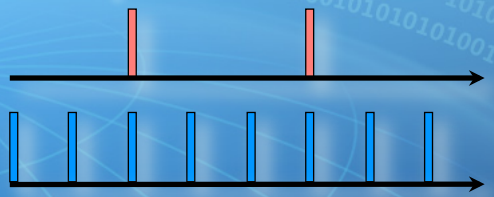

Microsoft Research

Routing in Low Power Wireless Networks

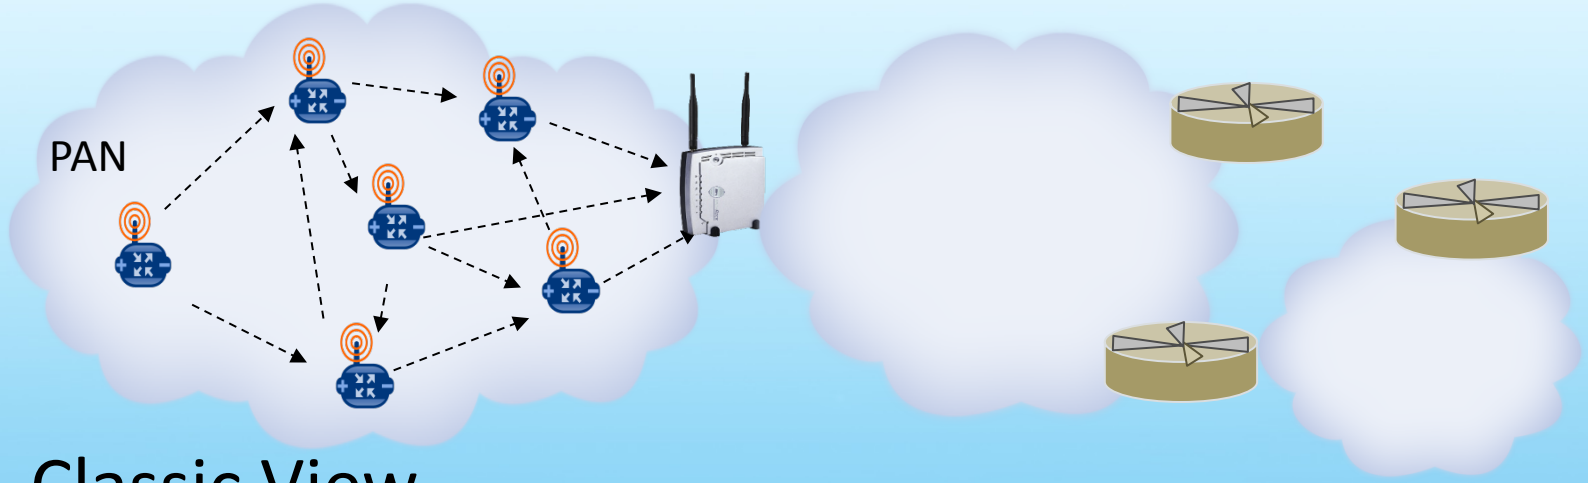

- Classic View
  - Network = Graph of routers and links
    - Like a street map
  - Routing is a (distributed) algorithm for finding good paths in this (slowly changing) graph
  - Realized (hop by hop) by tables and addressing
- But, … there is no graph
  - Discover it by attempting to communicate
  - Changes due to environment

Microsoft Research

Self-Organized Routing - nutshell

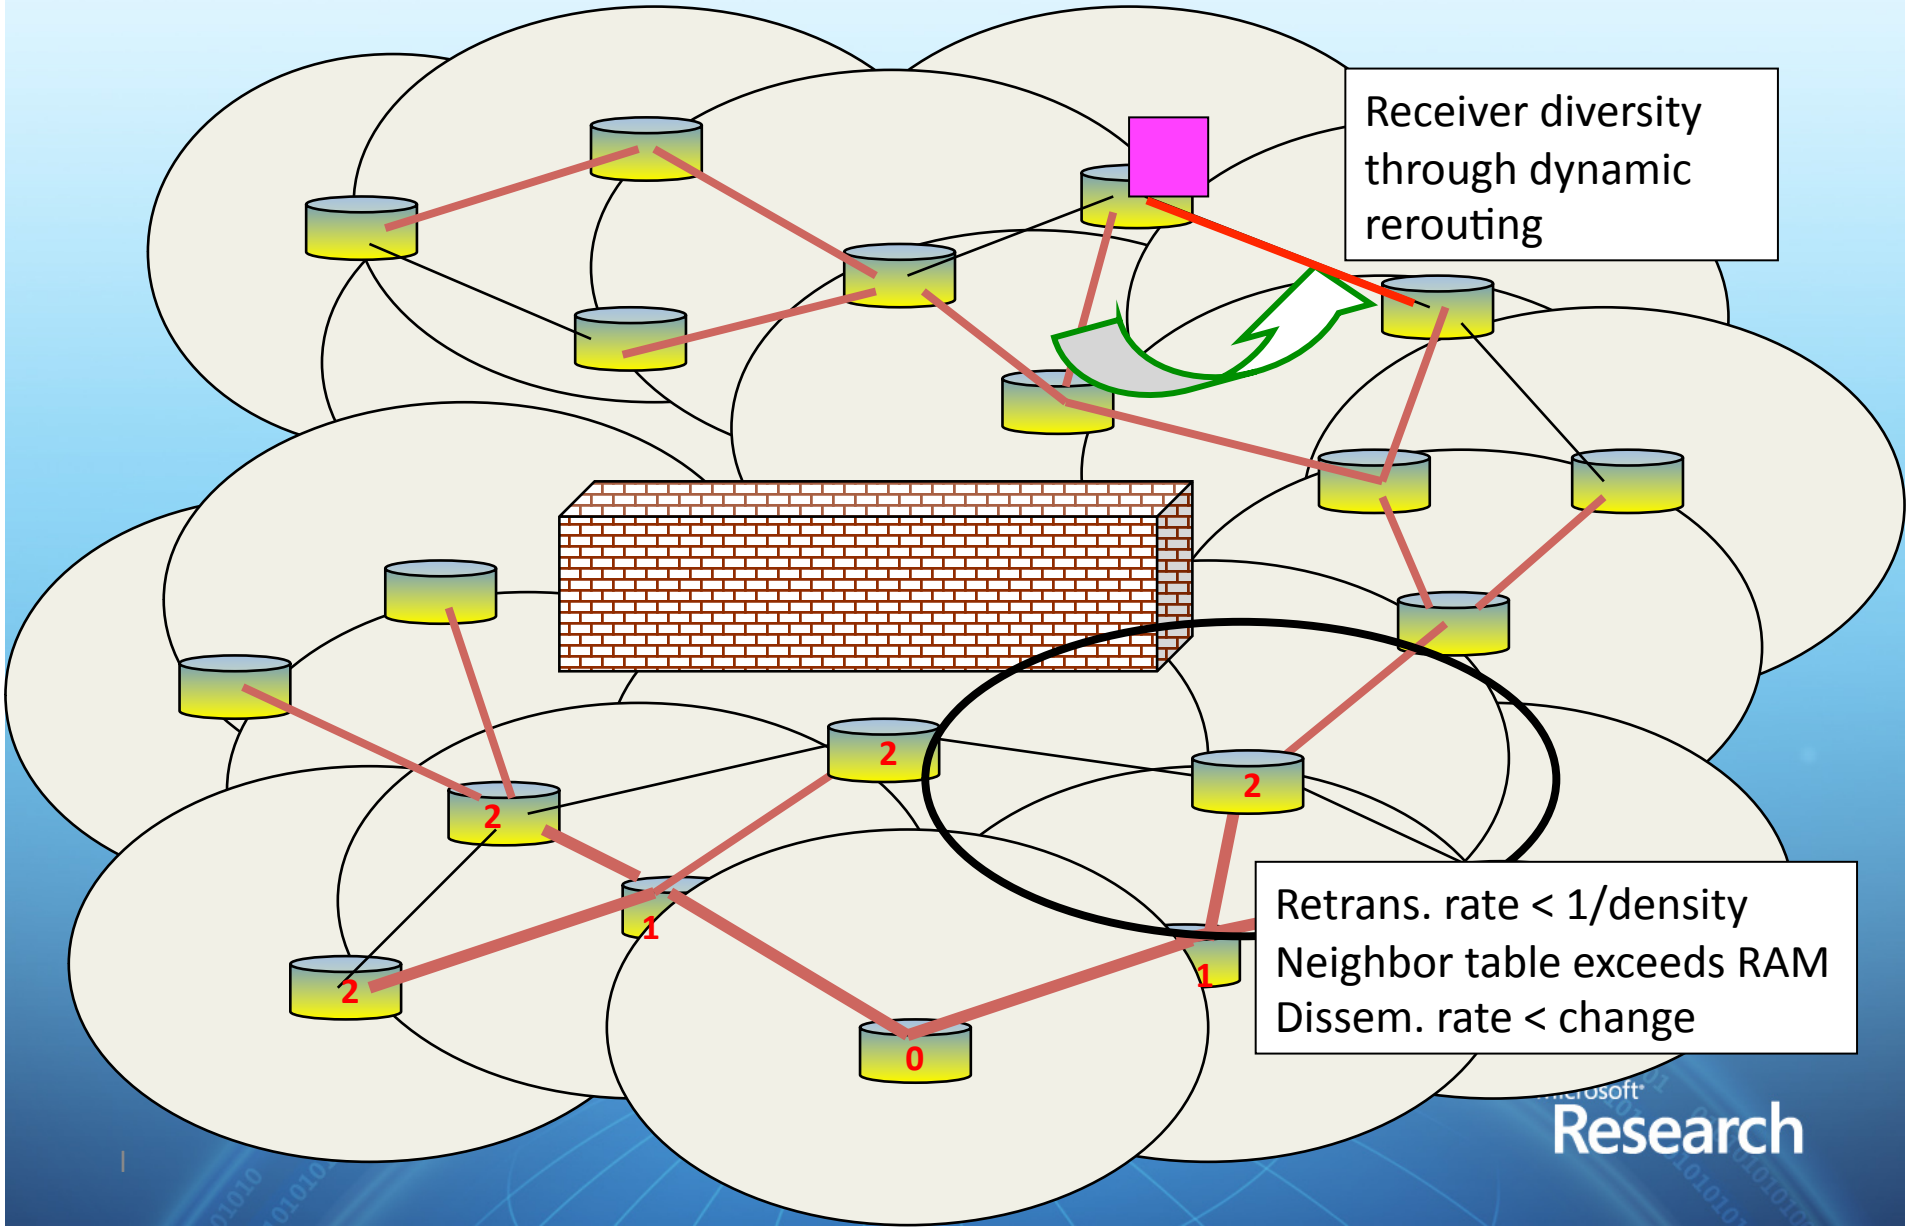

Key IPv6 Contributions

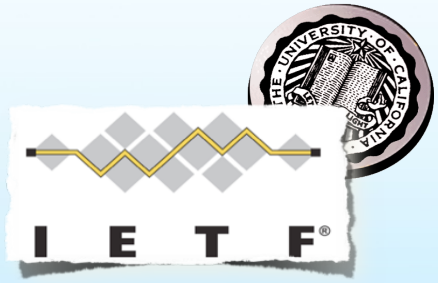

- Large simple address
  - Network ID + Interface ID
  - Plenty of addresses, easy to allocate and manage
- Autoconfiguration and Management
  - ICMPv6
- Integrated bootstrap and discovery
  - Neighbors, routers, DHCP
- Protocol options framework
  - Plan for extensibility
- Simplify for speed
  - MTU discovery with min
- 6-to-4 translation for compatibility

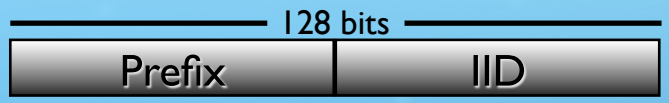

Microsoft Research

6LoWPAN – IPv6 over 802.15.4

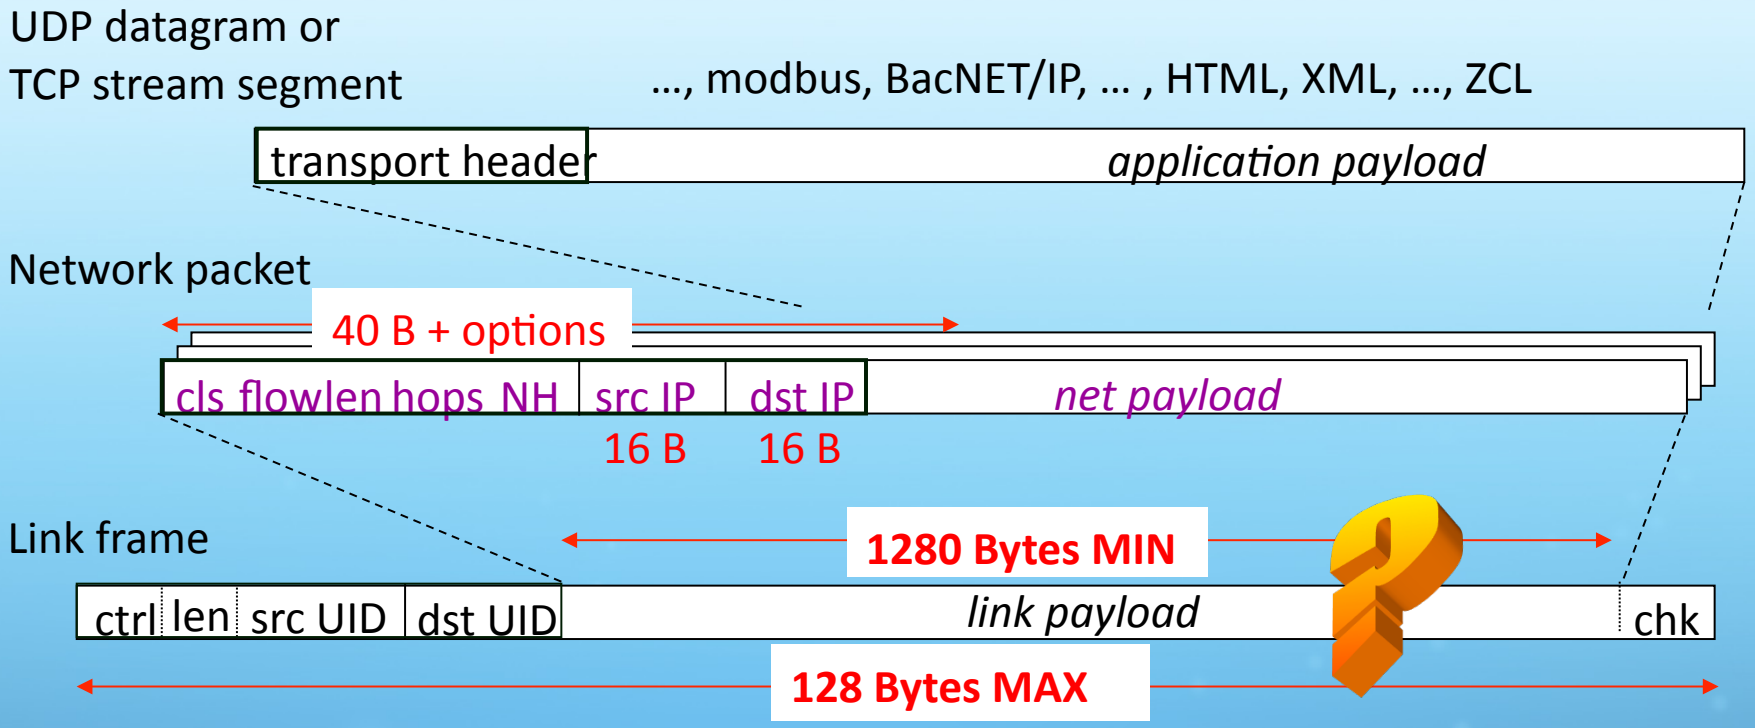

- Large IP Address & Header => 16 bit short address / 64 bit EUID
- Minimum Transfer Unit => Fragmentation
- Short range & Embedded => Multiple Hops

Microsoft Research

6LoWPAN adaptation layer

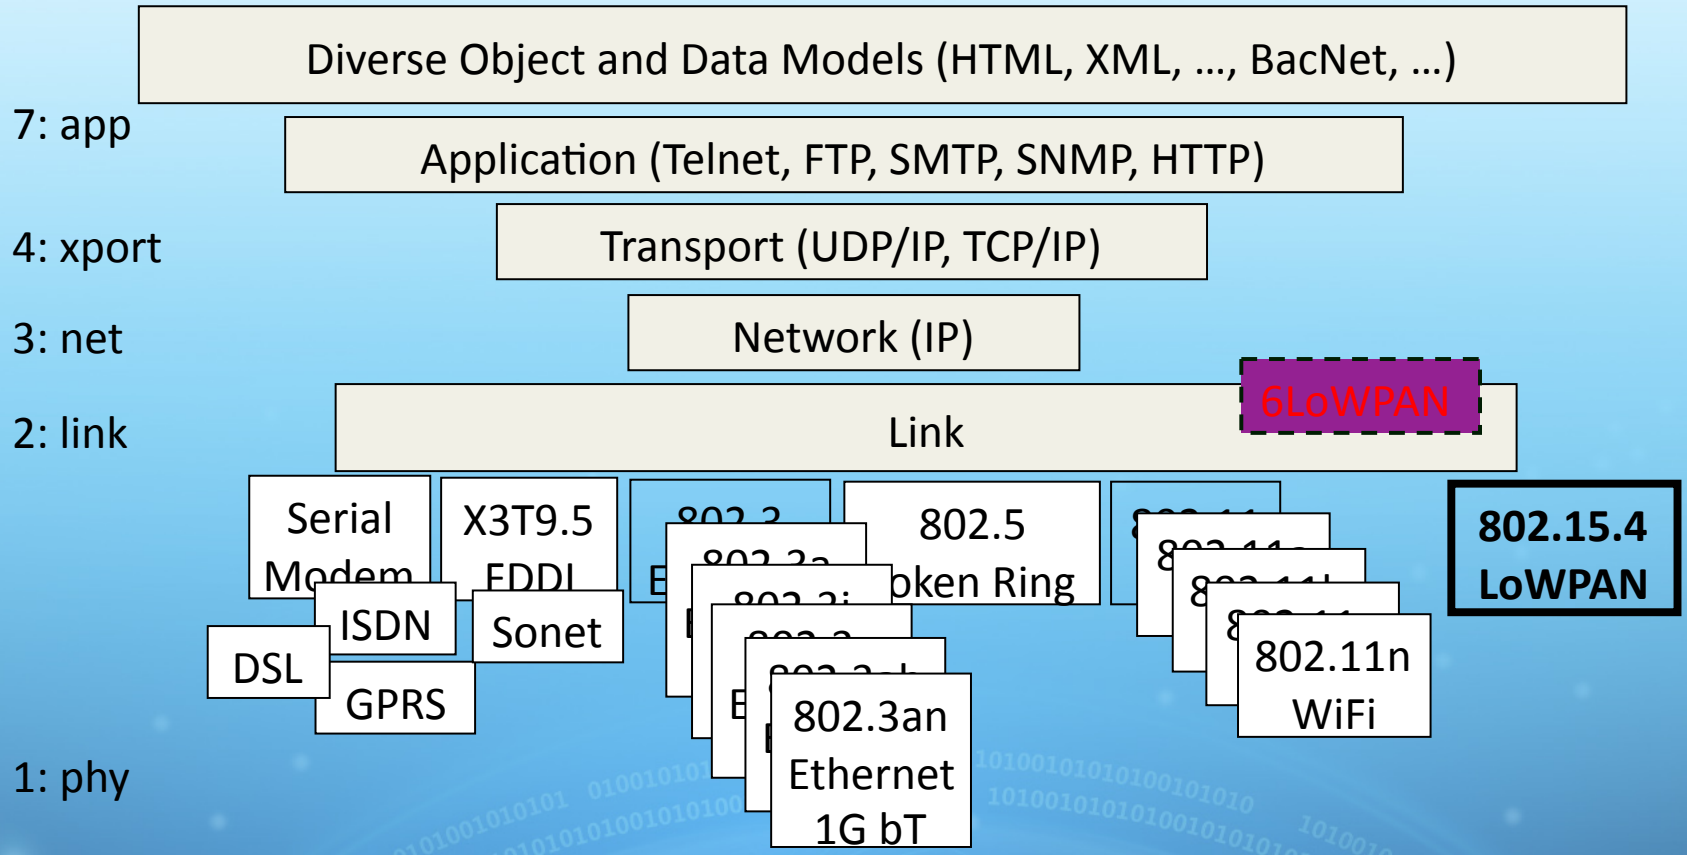

6LoWPAN – IP Header Optimization

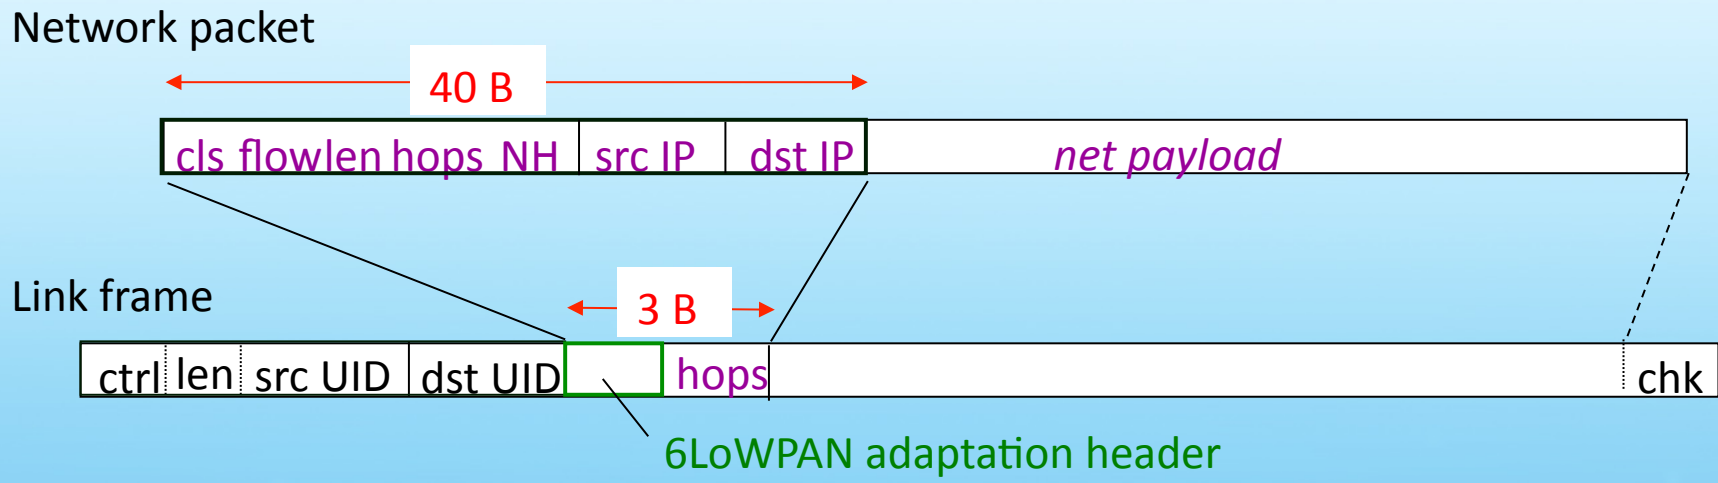

- Eliminate all fields in the IPv6 header that can be derived from the 802.15.4 header in the common case
  - Source address : derived from link address
  - Destination address : derived from link address
  - Length : derived from link frame length
  - Traffic Class & Flow Label : zero
  - Next header : UDP, TCP, or ICMP
- Additional IPv6 options follow as options

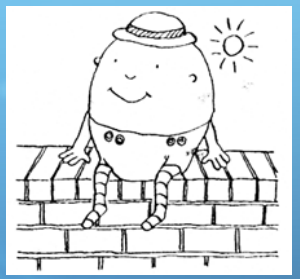

Microsoft Research

IP Routing Everywhere

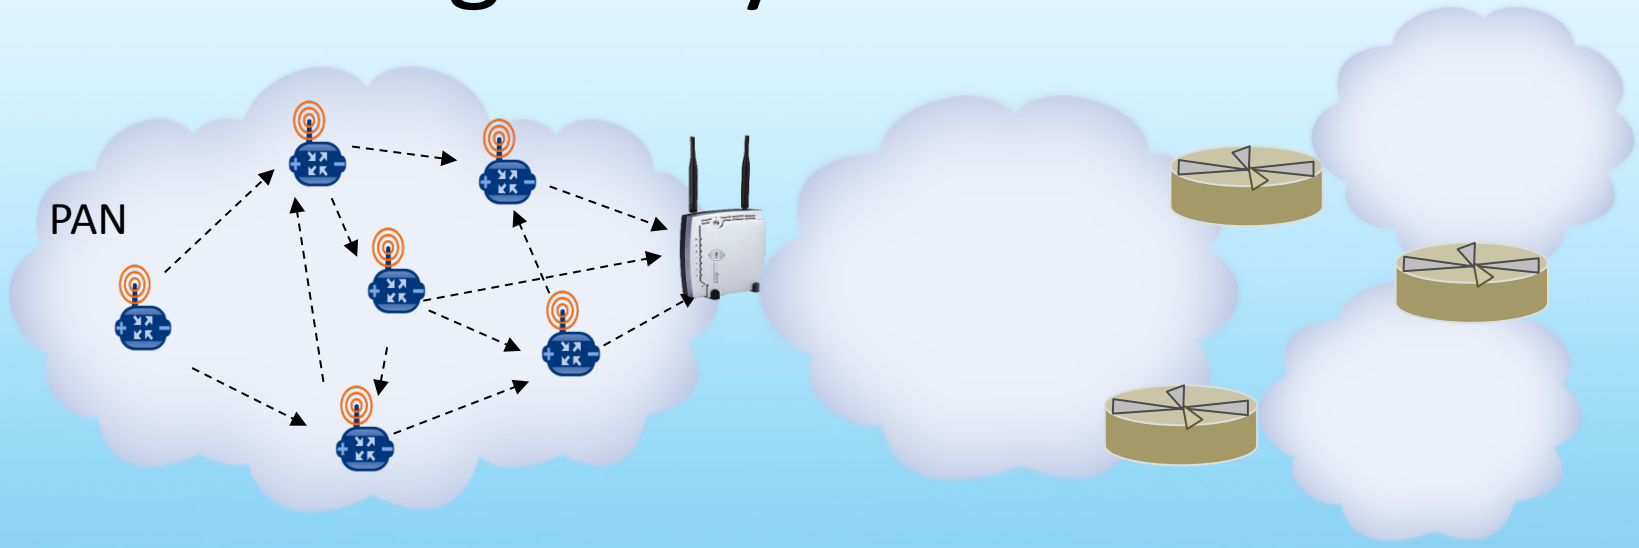

- Conventional IP link is a full broadcast domain
  - Routing connects links (i.e, networks)
- Many IP links have evolved from a broadcast domain to a “mesh” with emulated broadcast
  - ethernet => switched ethernet
  - 802.11 => 802.11s
- Utilize high bandwidth on powered links to maintain the illusion of a broadcast domain
- 802.15.4 networks are limited in bandwidth and power so the emulation is quite visible.

Microsoft Research

Embedded IPv6 in Concept

Structured Decomposition

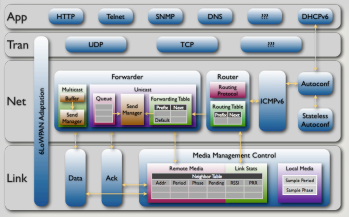

Retain strict modularity
Some key cross-layer visibility

IP Link ⇒ Always On

Retain illusion even when always off

IP Link ⇒ "Reliable"

Retain best-effort reliability over unreliable links

IP Link ⇒ Broadcast Domain

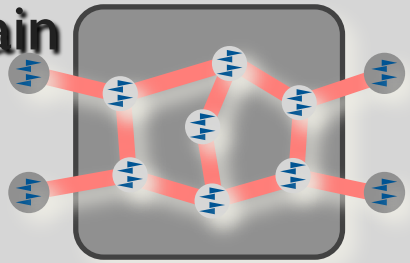

IPv6 can support a semi-broadcast link with few changes

Microsoft Research

Example: Autoconfiguration

Configuring Large Numbers of Interfaces

RFC 4861 – Neighbor Discovery
RFC 4862 – Stateless Addr Autoconf
RFC 3315 – DHCPv6

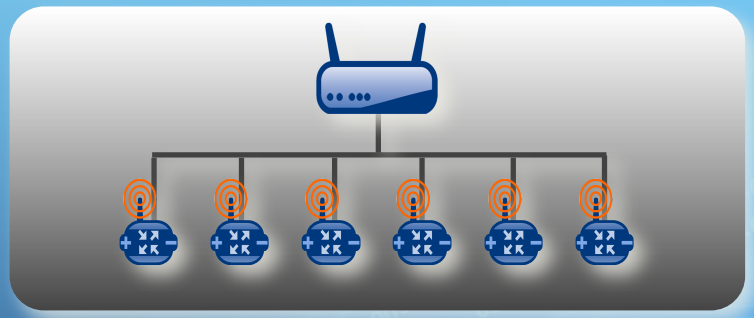

	Existing Options		New Options
ICMPv6 Hdr	Router Adv	Prefix Info	MHop Info
	Cur Hop Limit	Prefix Length	Network ID
	Managed Addr Config	Autonomous Config	Sequence Number
	Other Config	Valid Lifetime	Router Hops
	Router Lifetime	Preferred Lifetime	Flags
	Reachable Time	Prefix	

Microsoft Research

Example: Routing with IPv6 options

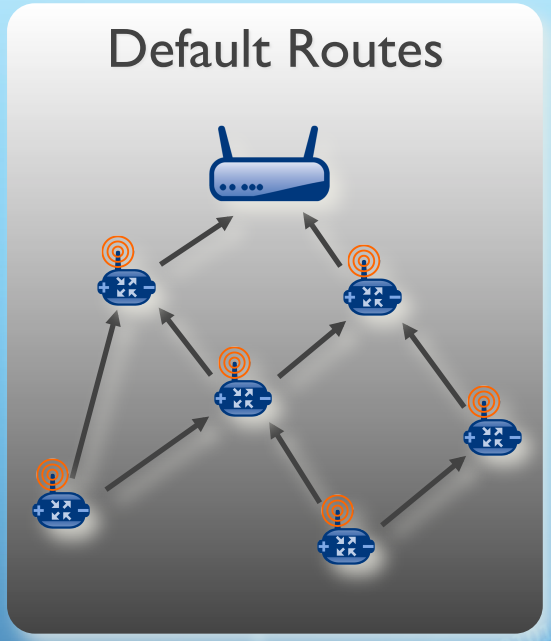

Discovering Links

ICMPv6 Hdr	**Router Adv**	MHop Info

Building a Connectivity Graph

Low Routing Cost

High Routing Cost

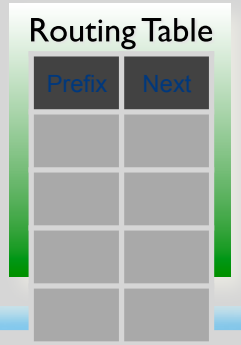

Selecting the Next Hop

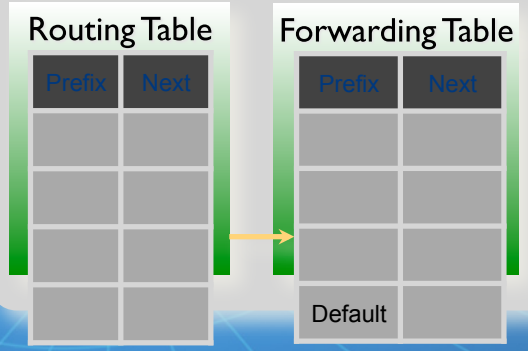

- Default route
- Hop-by-hop retry
- Reroute on loss

Research

Complete Embedded IPv6 Stack

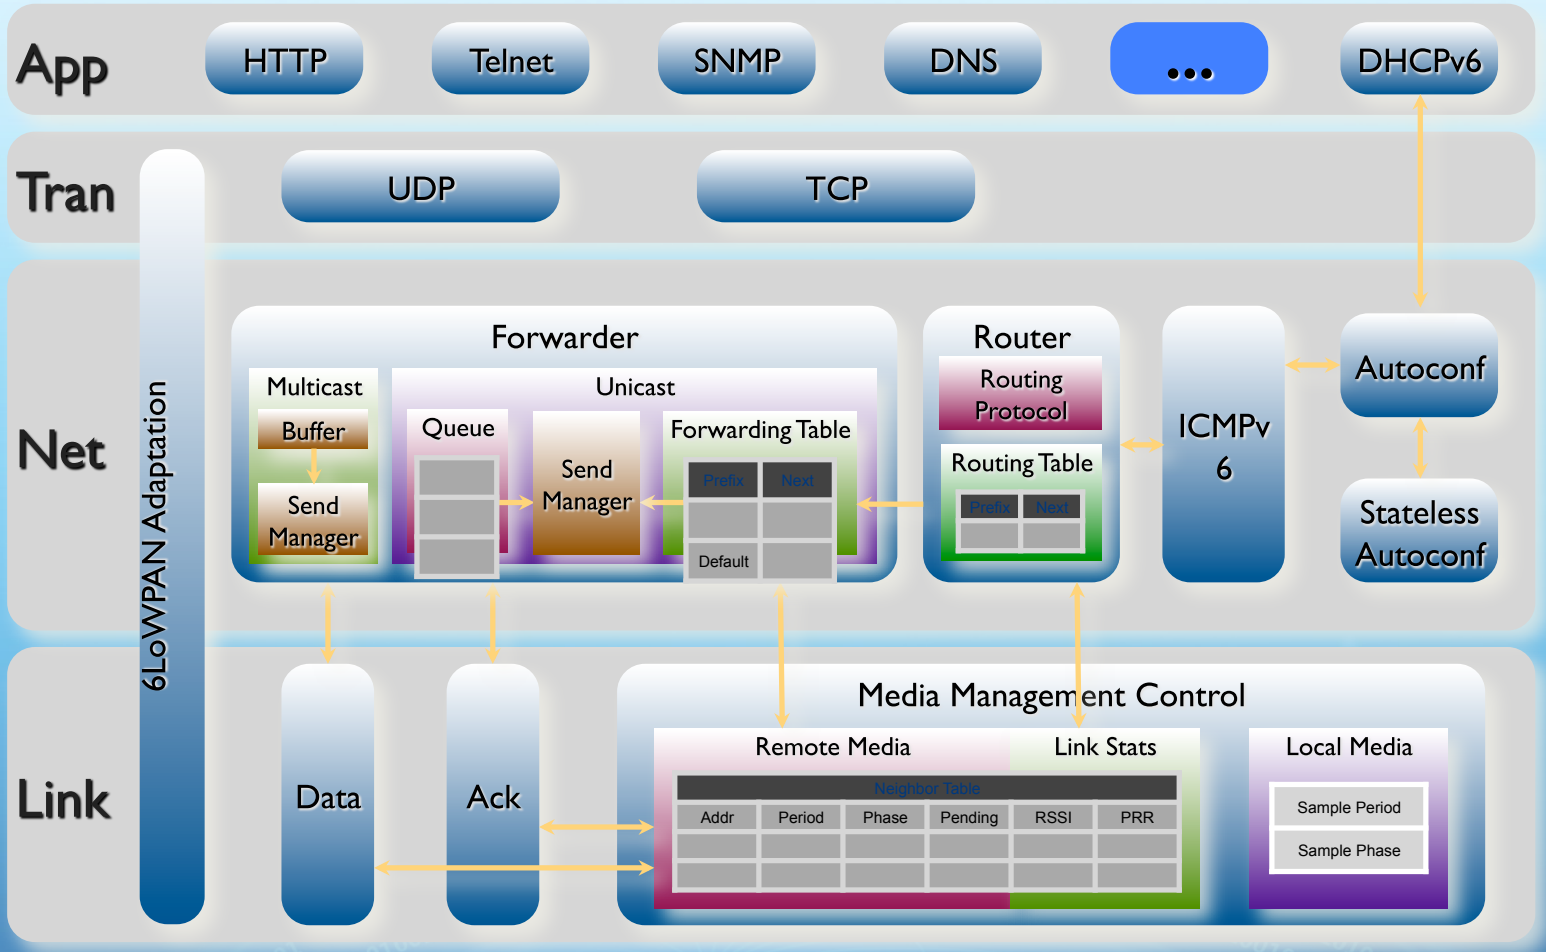

Microsoft Research

Adding up the pieces

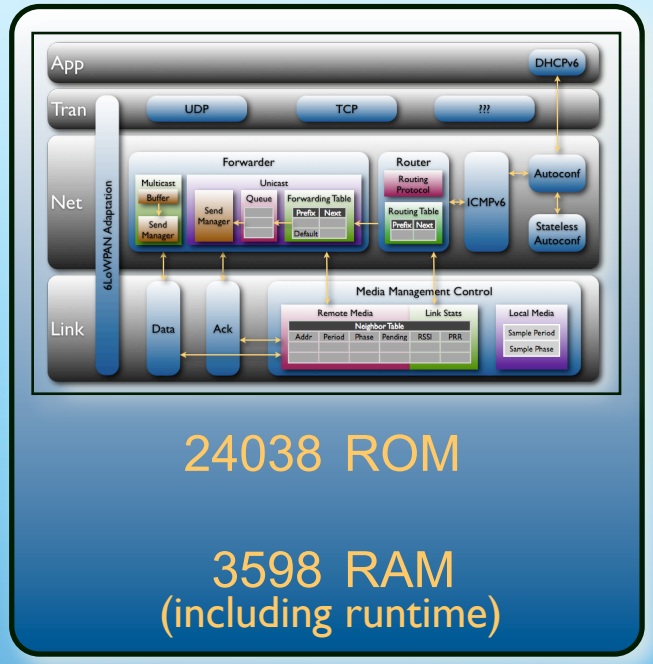

24038 ROM

3598 RAM
(including runtime)

* Production implementation on TI msp430/cc2420

	ROM	RAM
CC2420 Driver	3149	272
802.15.4 Encryption	1194	101
Media Access Control	330	9
Media Management Control	1348	20
6LoWPAN + IPv6	2550	0
Checksums	134	0
SLAAC	216	32
DHCPv6 Client	212	3
DHCPv6 Proxy	104	2
ICMPv6	522	0
Unicast Forwarder	1158	451
Multicast Forwarder	352	4
Message Buffers	0	2048
Router	2050	106
UDP	450	6
TCP	1674	50

Microsoft Research

and Power and reliability …

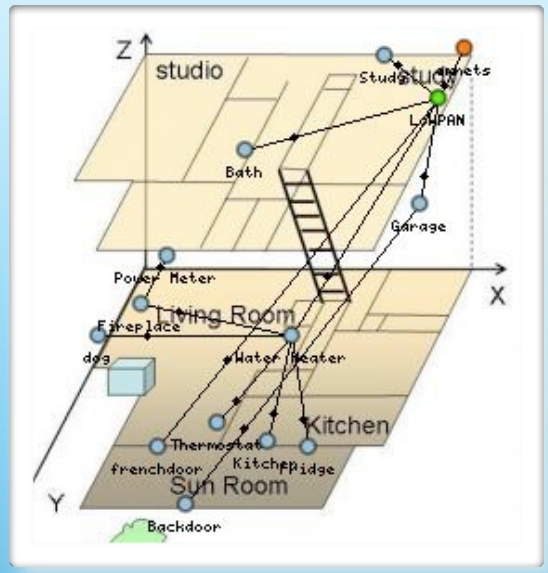

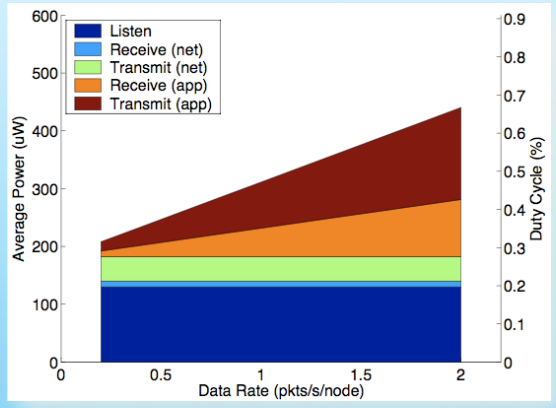

Data Rate Sensitivity
(Router)

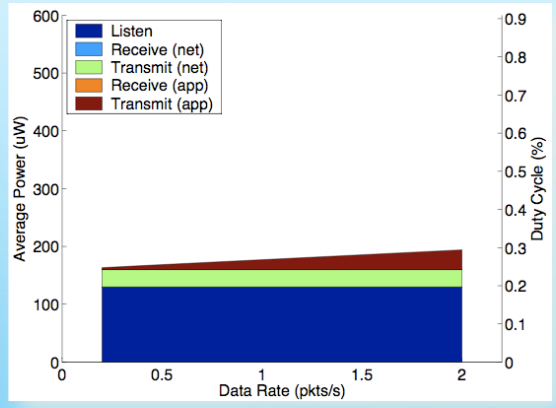

Data Rate Sensitivity
(Edge)

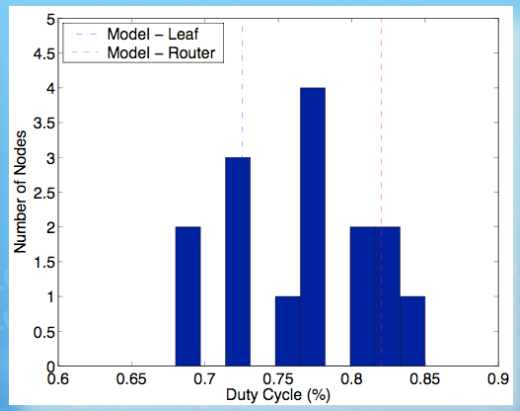

Deployment Duty Cycle

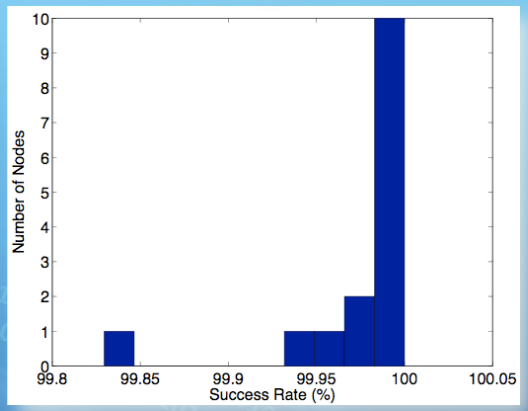

Deployment Reliability

The "Killer App" for Wireless Sensor Networks

- Energy and the environmental impact of extraction, use, and disposal
- THE problem of the Industrial Age
- We need to find Information Age solutions to THE Industrial Age Problem

- it starts with the Faustian bargain of oblivious consumption

Microsoft Research

The Grid: Marvel of Industrial Age Design

- Deliver high quality low-cost power
- To millions of customers over thousands of miles
- Synchronized to <<16 ms cycle (60 Hz)
- With no orders, no forecasts, no plans
- No inventory anywhere in the supply chain

- To enable rapid economic & industrial growth through oblivious consumption

A New Reality …

1. Energy becoming increasingly dear
   - increased cost of acquisition
   - inclusion of environmental costs
2. Improvements in energy efficiency cause high dynamic variability in the load
   - high peak-to-ave ratio, bursty
3. Limitations of existing grid present transmission and distribution bottlenecks
4. Incorporation of renewable resources reduces control over supply
   - most are non-dispatchable (solar, wind)

Microsoft Research

Aware Co-operative Grid

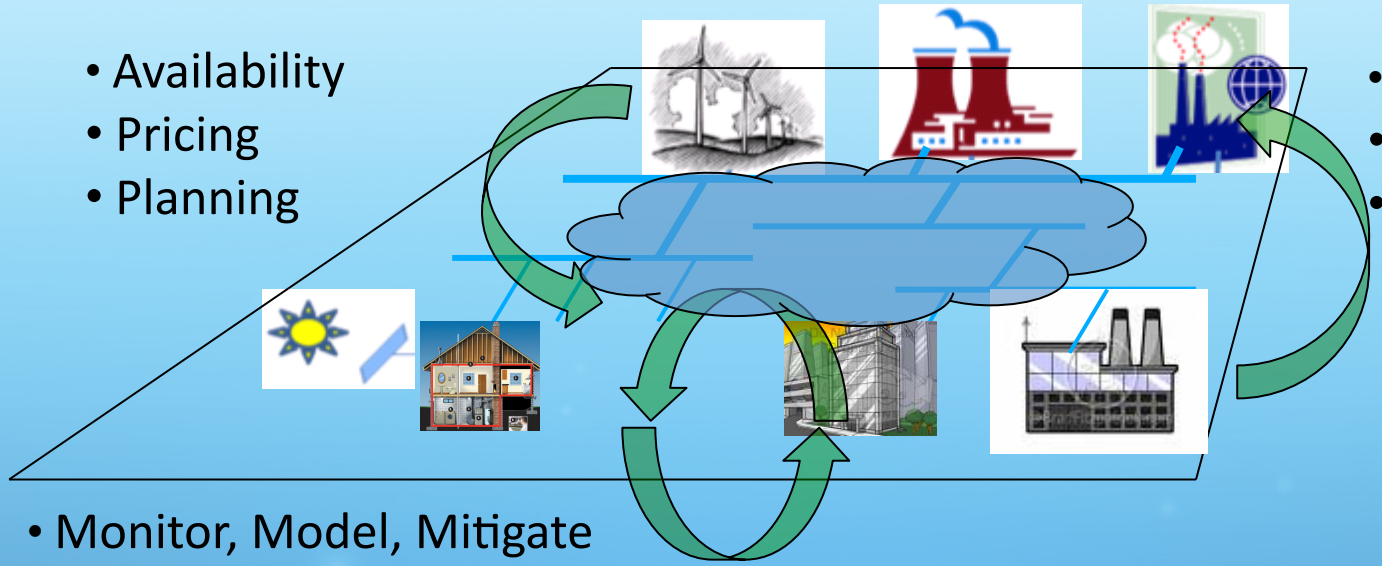

- Availability
- Pricing
- Planning

- Forecasting
- Tracking
- Market

- Monitor, Model, Mitigate
  - Deep instrumentation
  - Waste elimination
  - Efficient Operation
- Shifting, Scheduling, Adaptation

Microsoft Research

The “Macroscope”

- Observe complex interactions over time and space

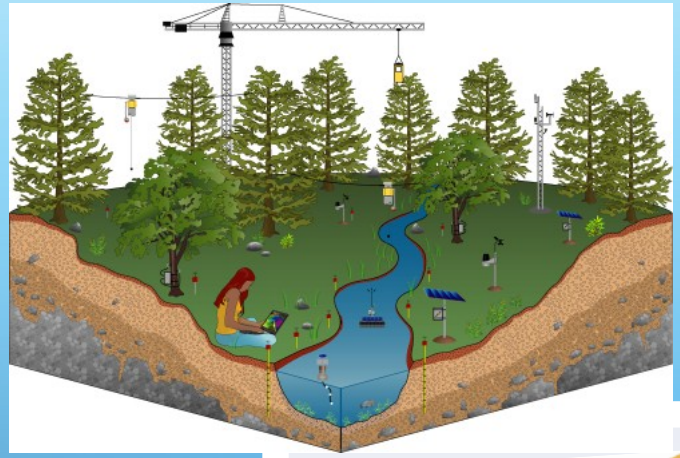

Slope (Spatial Analyst)

Aspect (Spatial Analyst)

Daily Average Temperature (Geostatistical Analyst)

Elevation (Calculated from Contour Map)

Aerial Photograph (10.16cm/pixels)

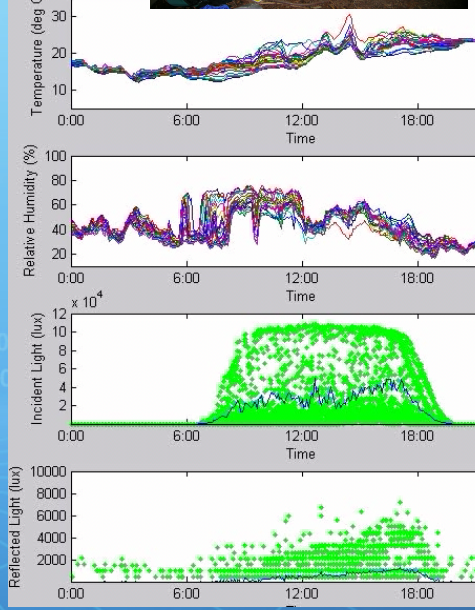

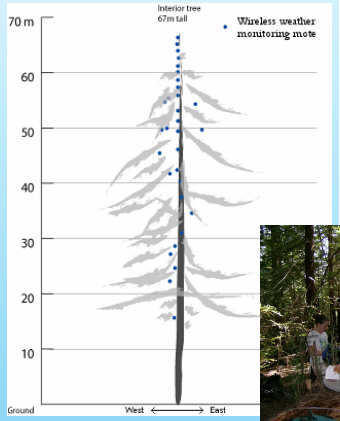

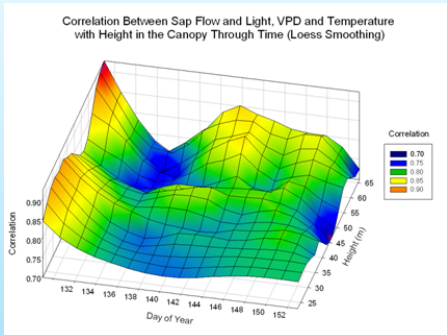

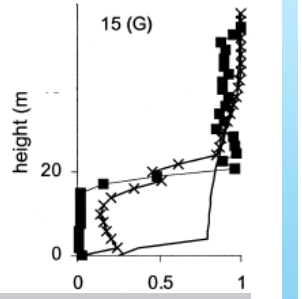

Where to Start?

- Buildings
  - 72% of electrical consumption (US),
  - 40-50% of total consumption,
  - 42% of GHG footprint
  - US commercial building consumption doubled 1980-2000, 1.5x more by 2025 [NREL]
- Where Coal is used
- Prime target of opportunity for renewable supplies

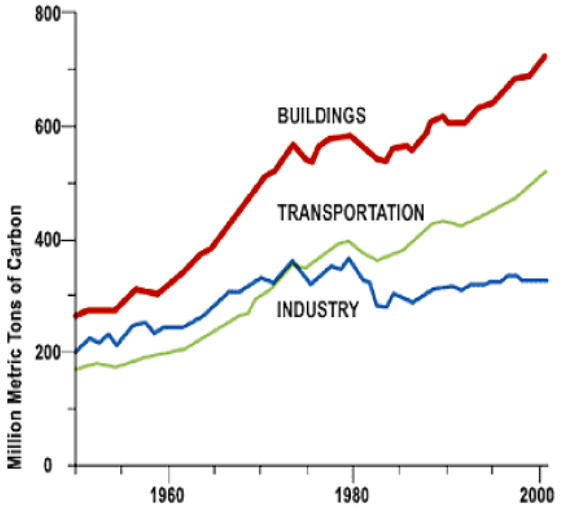

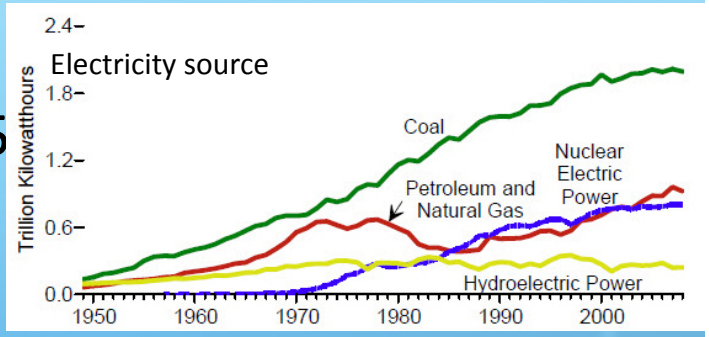

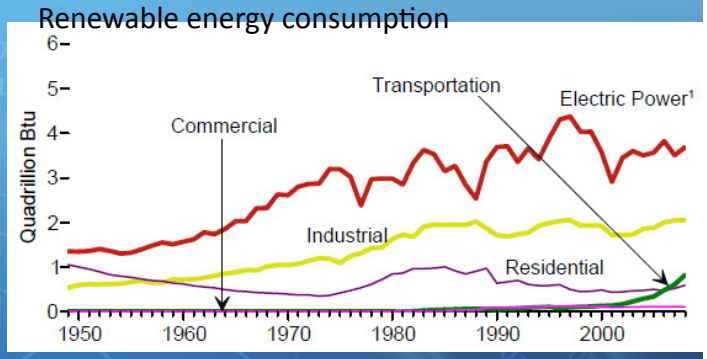

Our Buildings

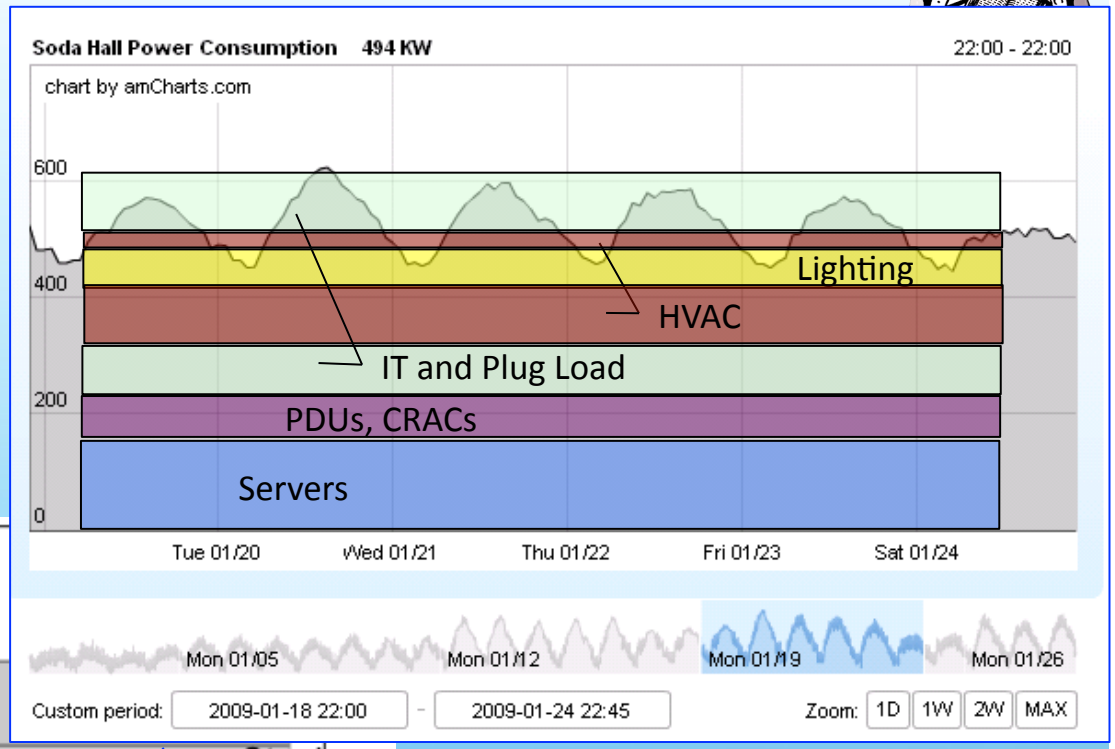

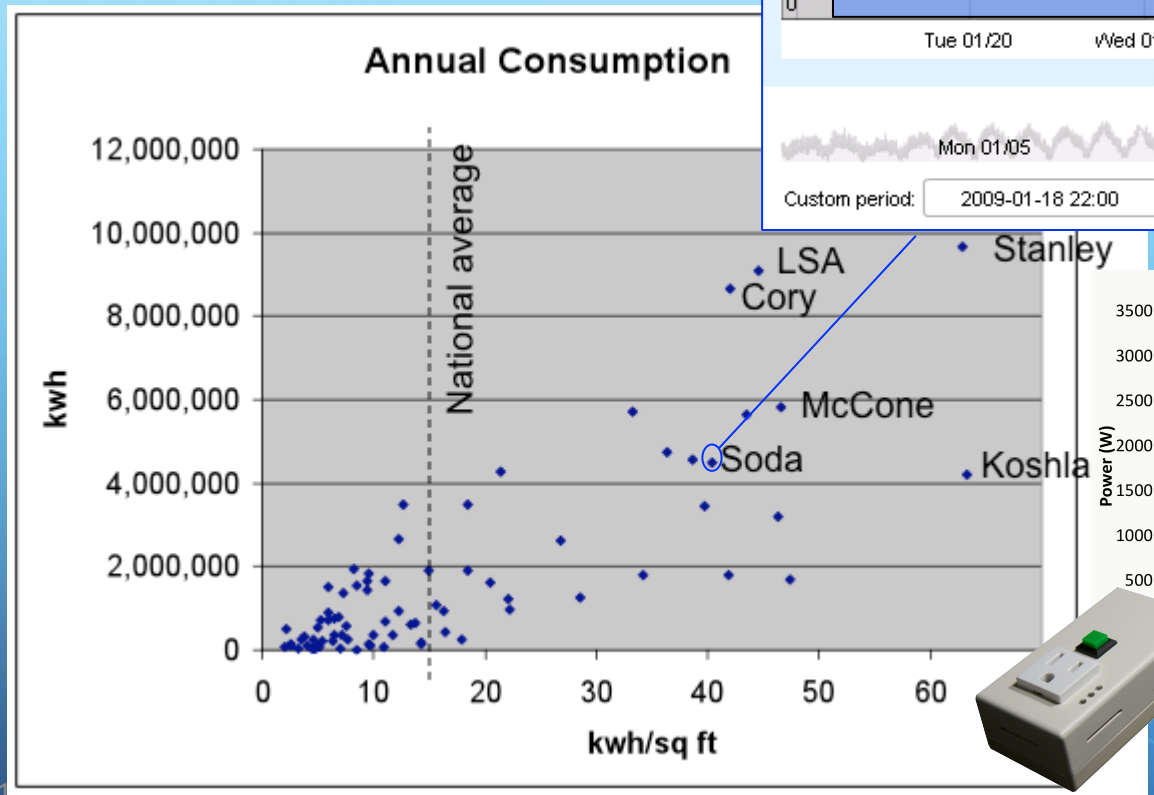

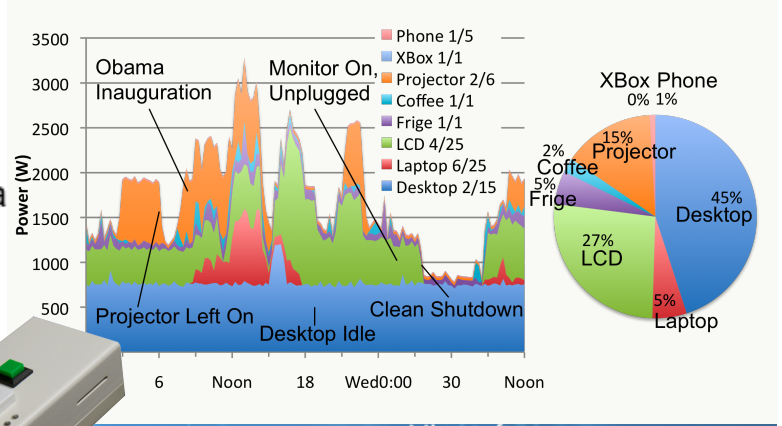

Microsoft Research

Energy Transparent Building

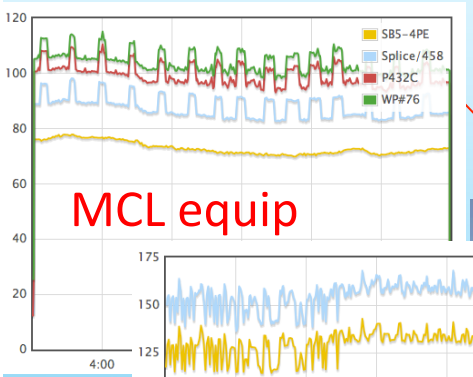

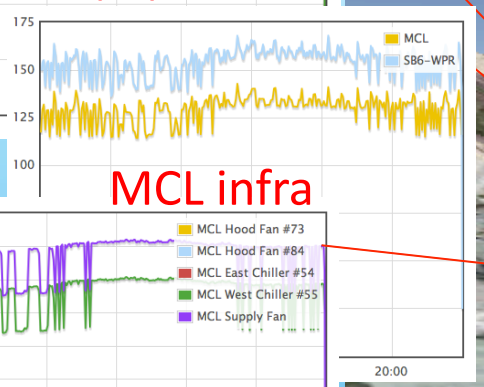

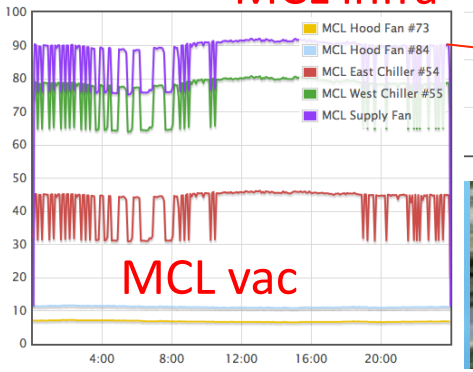

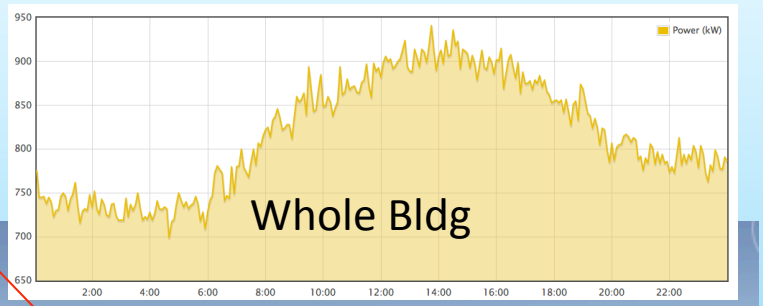

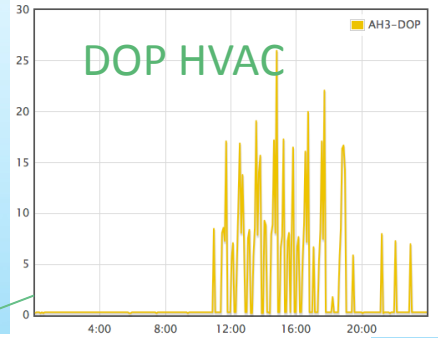

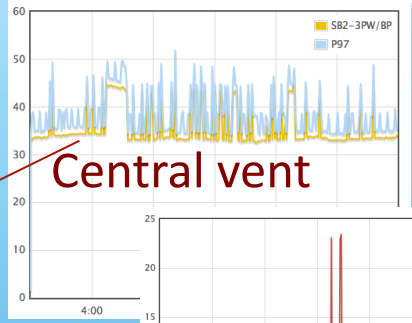

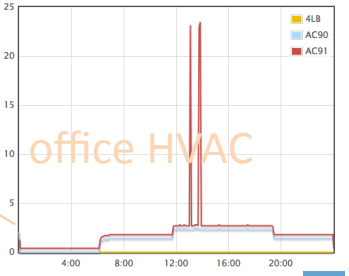

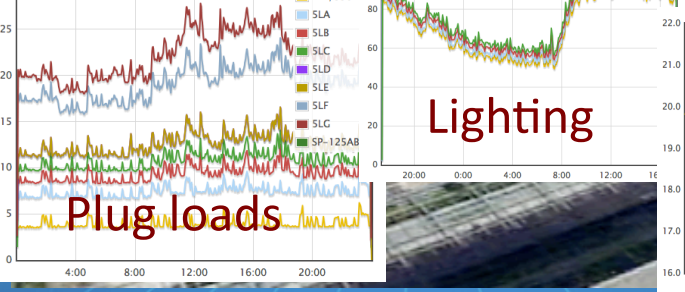

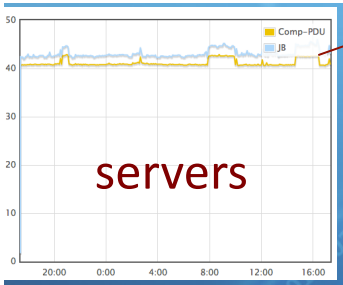

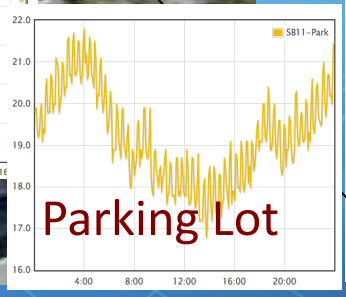

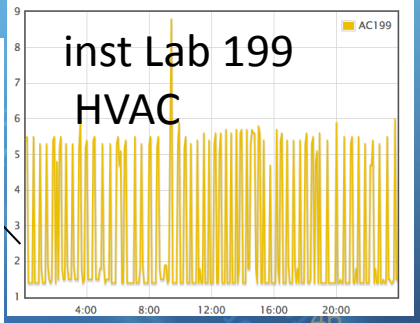

IP everywhere / Real World Web

Microsoft Research

Cyber / Physical Buildings

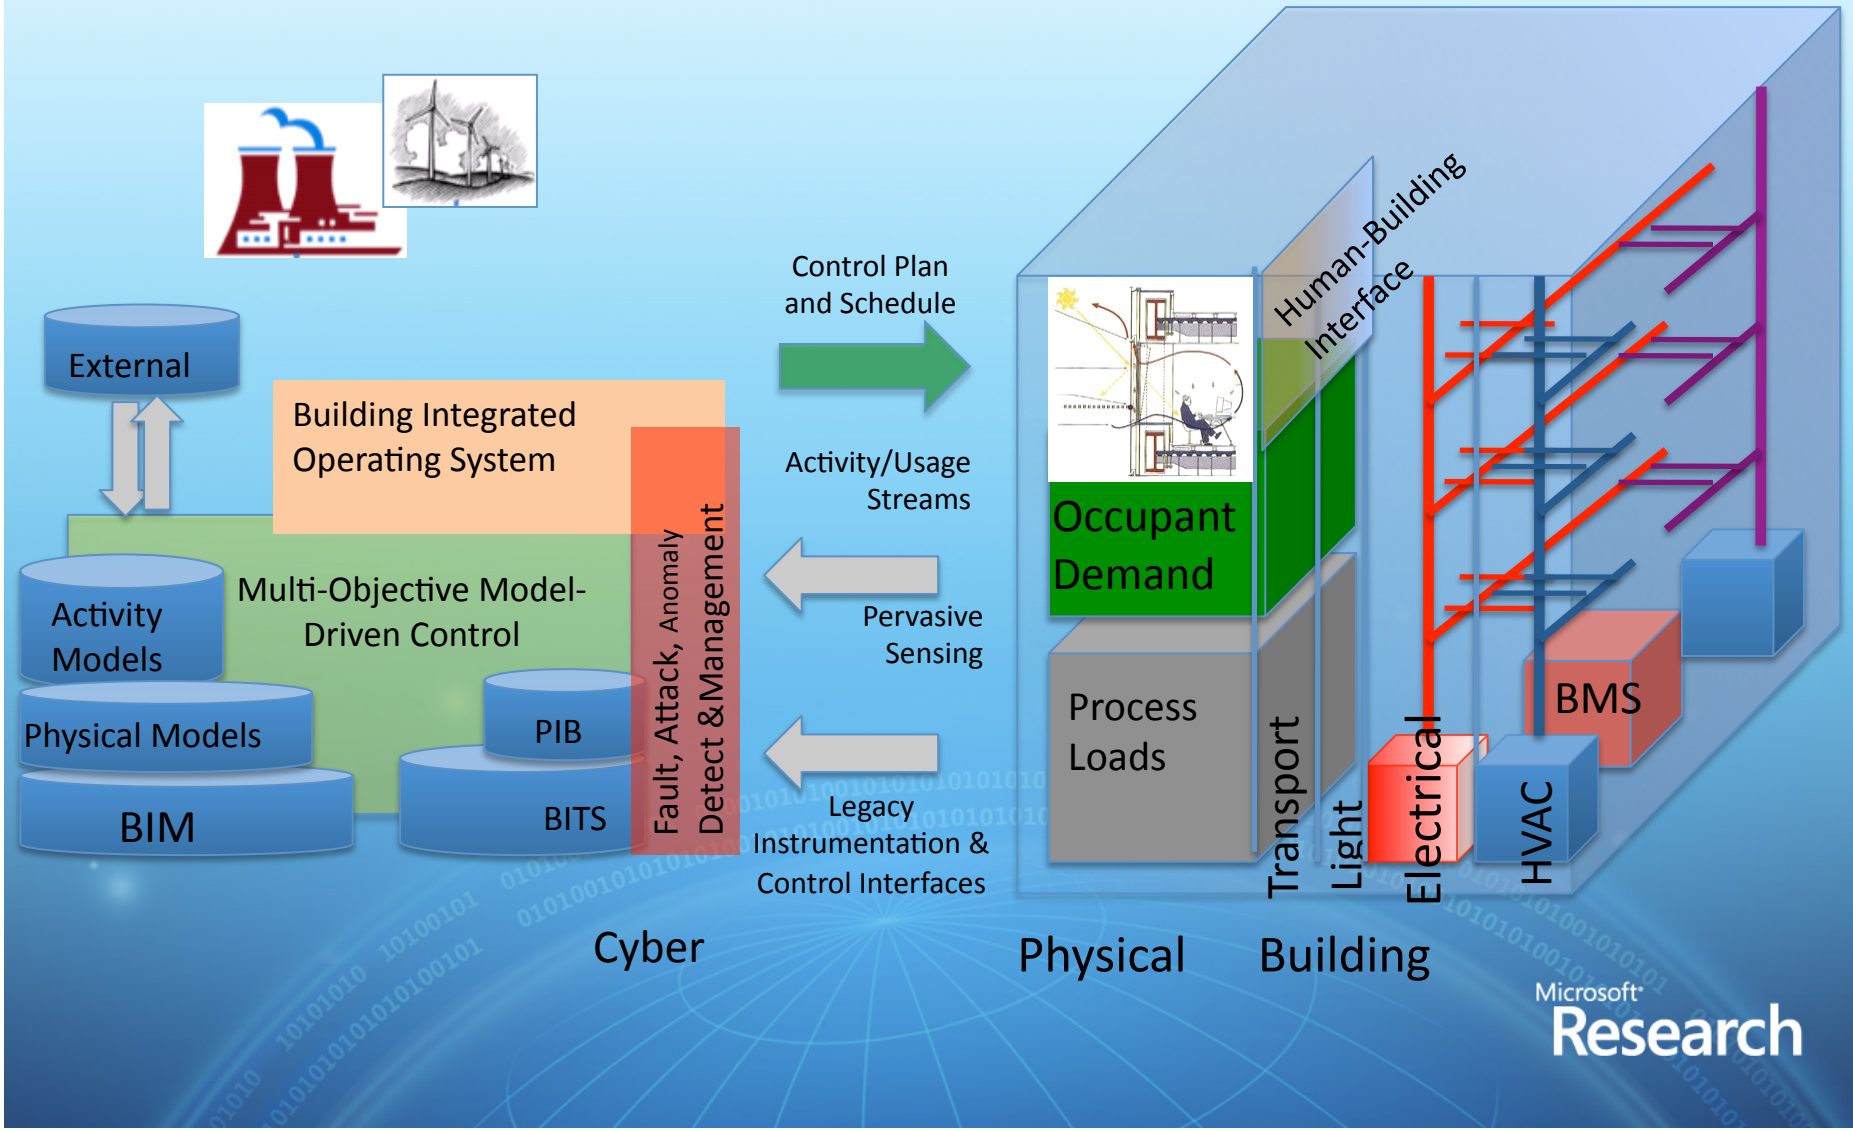

Microsoft Research

Conclusions

- The Internet is Every Thing is Here
- 15 years of deep innovation and research
  - Critical WSN breakthroughs
  - Key IPv6 developments
- Worldwide community of students, faculty, and industry
- Engagement of International organizations
- Fundamentally a new Scientific Instrument
- Focus it on the World's most important problem
  - Energy, productivity, and the environment

Microsoft Research

Thank You

Microsoft Research