

# Carnegie Mellon University

## Silicon Valley

### Toward Platformization for Connected Computing

Bob Iannucci, Ph.D.  
Distinguished Service Professor, ECE  
Director, CyLab Mobility Research Center

July 22, 2015



UNIVERSIDADE DE COIMBRA

# Agenda

- What's a Platform?
- What's an App?
- How do we build the platform for connected apps?


































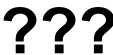
# PLATFORMS

# Wikipedia says...

*A computing platform is, in the most general sense, whatever pre-existing environment a piece of software is designed to run within, obeying its constraints, and making use of its facilities. Typical platforms include a **hardware architecture, an operating system (OS), and runtime libraries.***



# Computing Platforms: Past and Present

	Golden Age	Standard Platform	Compatibles	Value Shift to Software	Value Shift to Services
Mainframe	   CONTROL DATA 		 Magnuson Computer		  
Mini	   		 	 	
PC	   		  	 	 
"Connected"					

# Iannucci's Law

*In each generation of computing, the emergence of a standard platform transforms the industry by shifting value from hardware to software and services.*

- Hardware / software co-evolution on either side of a well-defined interface accelerates the pace of innovation
- Often (but not always), vertical gives way to horizontal
- The standard *emerges* rather than being *created*

# So, what is this Future Connected Computing Platform?

The standard platform for connected computing is a *combination* of


1. Flexible, powerful, programmable device families
2. Flexible, powerful, programmable networks

***Observation: #1 exists. #2 does not (why not?)***

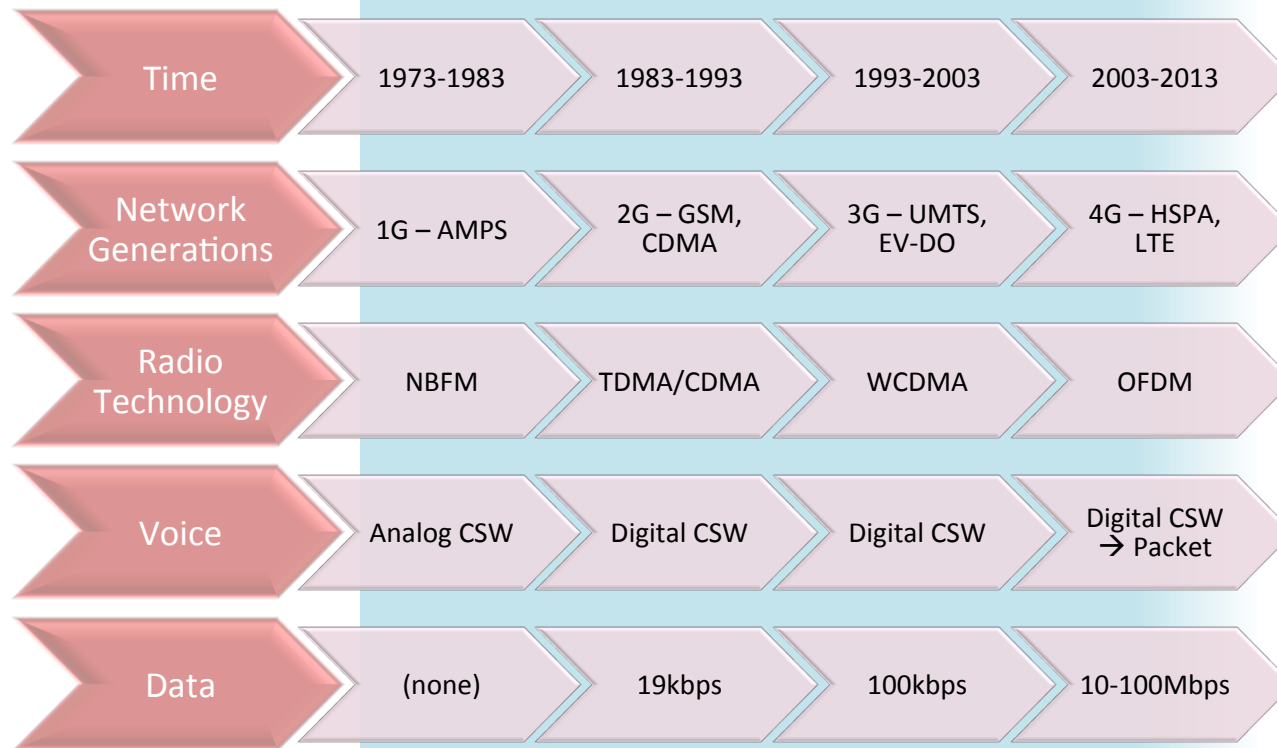
If true, the evolution of  
connected computing is stalled.

# Mobile Networks as a Computing Platform?


*Circuit switching legacy persists...*



*But then, in 2009, a funny thing happened.*



# In a network designed for CSW Voice traffic, PSW data now dominates



Source: *The State of the Internet*.  
Technical report, Akamai  
Technologies, Inc., Cambridge,  
MA, 3rd Quarter 2013.

# Past Design Assumptions Lead to Present-Day Architectural Challenges

- **Network impermeability:** apps and their cloud counterparts are separated by 150-2000 msec.
- **Hidden RAN state:** radio performance has a first-order effect on apps
- **Hidden network state:** best-efforts, no QoS tools, bufferbloat
- **Hidden app state:** app behavior has a first-order effect on RAN and network

# Won't LTE / IMT-A fix all of this?

No.

# Won't SDN / NFV fix all of this?

No.



# Summary

- Rather than providing the basis for a connected computing platform, today's networks (LAN and WAN) have worked *against* platformization
  - Computing pushed to the network endpoints
  - Significant latencies thwart cyber-physical apps
  - Closed hardware limits study, innovation

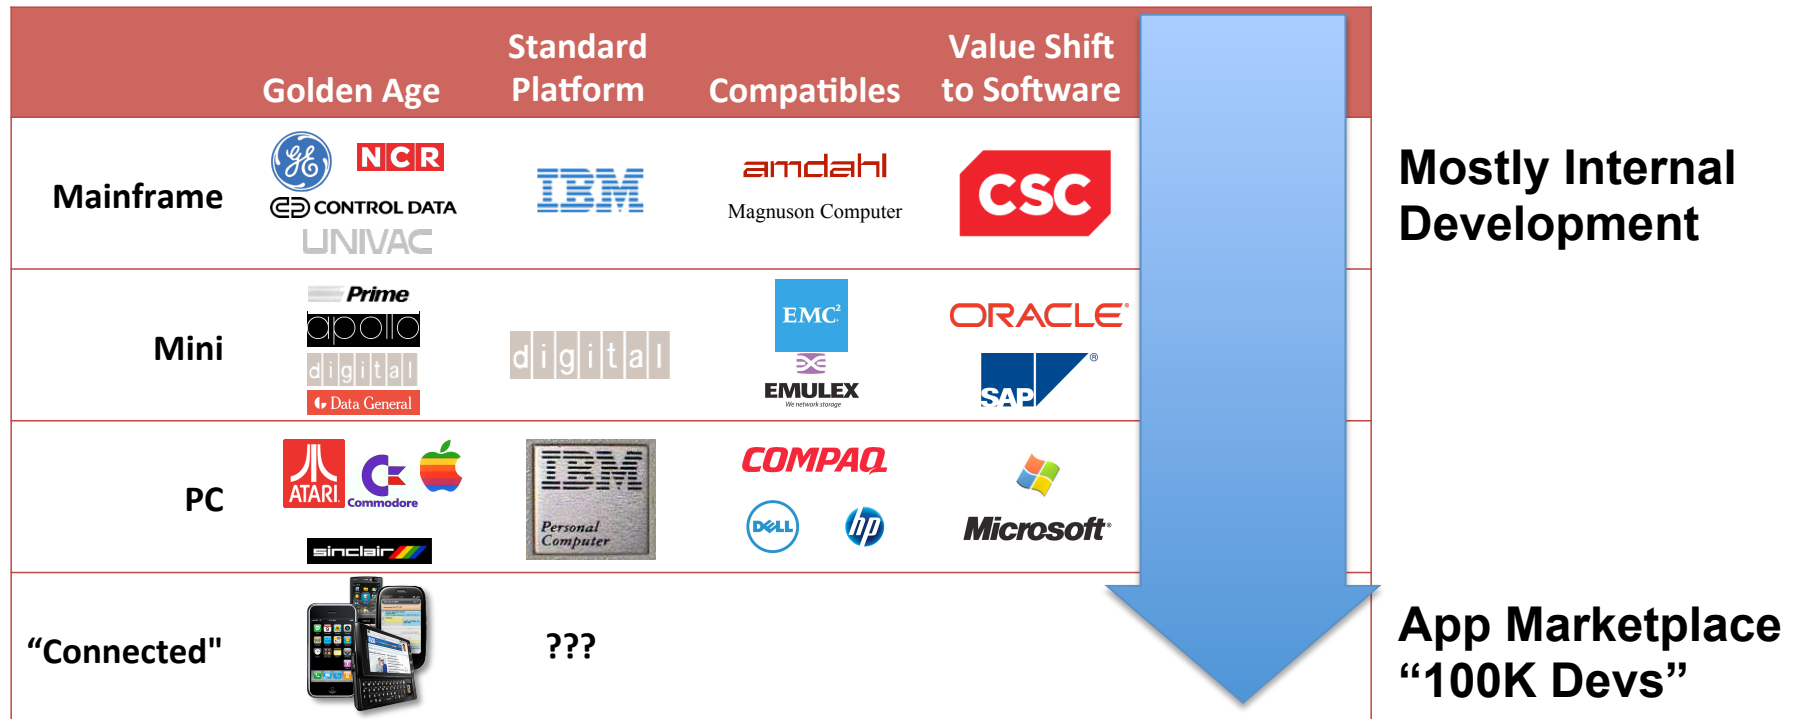
# APPS

# The Evolving Nature of Apps




























- Iannucci's Law says that value shifts to software and services in each generation of computing
- But the way that app logic is defined and the environment in which it runs has changed radically

*Let's look at some trends*





















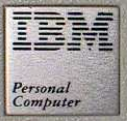






# Who Creates Apps?



# What Role does Communication Play?

	Golden Age	Standard Platform	Compatibles	Value Shift to Software	
Mainframe	   		 Magnuson Computer		Remote Job Entry
Mini	   		 	 	Modems and Dumb Terminals
PC	   		  	 	Internet and Web Apps
"Connected"		???			App Marketplace, Cloud Services

# Where Does App Logic Live?

	Golden Age	Standard Platform	Compatibles	Value Shift to Software	
Mainframe	   		 Magnuson Computer		Mainframe
Mini	   		 	 	Mini
PC	   		  	 	PC or Web server
"Connected"		???			Phone and Cloud

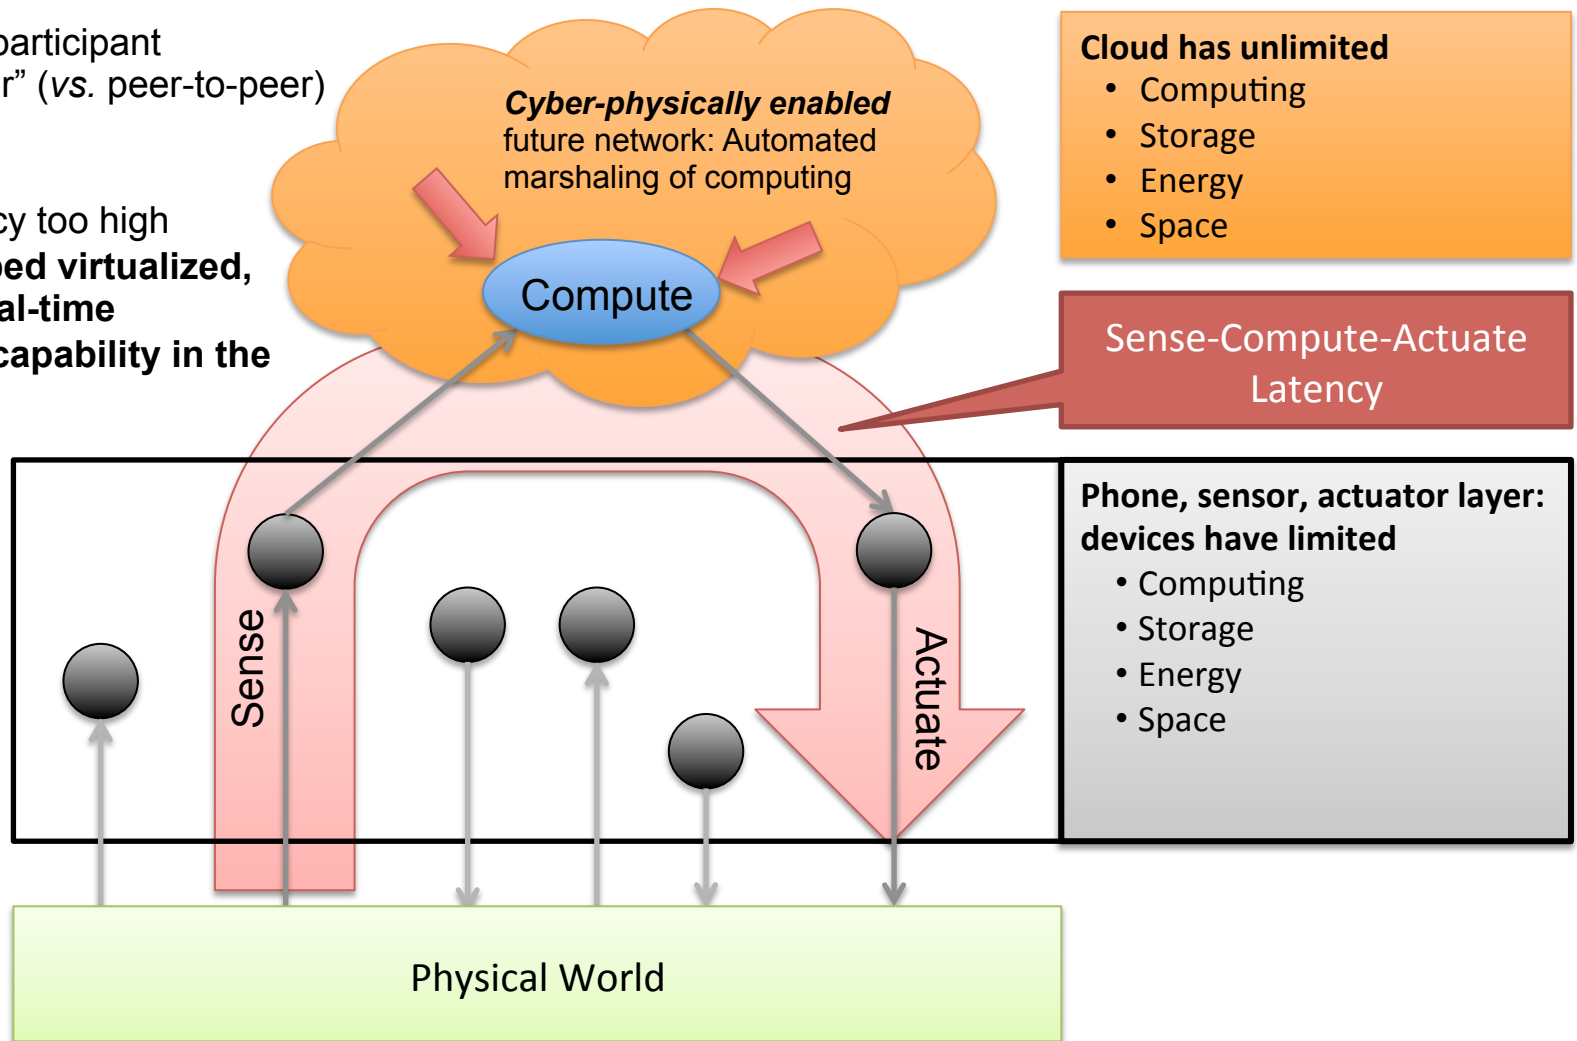
# Now, Extrapolate to the IoT

- Who creates apps?
  - Today: closed, vertical
  - Tomorrow: 100K+ developers
- Use of comms?
  - Data collection
  - Action dissemination
- Where does logic live?
  - Across a **distributed, synchronized, non-owned** pool of resources in and across a network that supports it

# The Cyber-Physically Enabled Network

Network-as-participant  
“Near to Near” (vs. peer-to-peer)

Today: latency too high  
**Future: embed virtualized,  
movable, real-time  
computing capability in the  
network**





# How Does the Network Have to Change?

- **Time-Aware:** time as a first-order concept; logical correctness includes time-accuracy; current network assumptions are leading us away from this
- **Federated:** the ability to fuse information across logical, geographical and political boundaries with the ability to manage security, privacy, precision, and provenance
- **Scalable:** the ability to scale computing capacity as IoT devices are added

# How Do We Proceed?

- **Obstacle:** existing networks are closed and are resistant to study
- **Obstacle:** network evolution is driven by a small, closed standards community
  - Contrast to evolution of the internet
- **Obstacle:** current architecture evolved from circuit switching and assumption of “wired last mile”
  - Characteristics of packet wireless are *fundamentally* different
  - Architecture did not anticipate IoT

Open instead  
of closed

Anticipate IoT  
Network  
Requirements



# TIME-AWARE

# What Does Time Have to Do with IoT?


- Device-to-device and Device-to-cloud
  - Distributed sensing: correlation of physical measurements
  - Distributed actuation: control systems
  - Deadlines have physical consequences – network an issue



Gunfire Detection



Cellular control of UAVs




Earthquake Warning

Teng, E., Falcao, J. D., Dominguez, C. R., Mokaya, F., Zhang, P., & Iannucci, B.  
*Aerial Sensing and Characterization of Three-Dimensional RF Fields.*  
In Second International Workshop on Robotic Sensor Networks. Seattle, WA. 2015.

Minson, S. E., Brooks, B. A., Glennie, C. L., Murray, J. R., Langbein, J. O.,  
Owen, S. E., Heaton, T., Iannucci, R. A. and Hauser, D. L.  
*Crowdsourced Earthquake Early Warning.* Science Advances, 1(3), 1–7.  
<http://advances.sciencemag.org/content/1/3/e1500036>

# Example: Acoustic Beamforming

- Time Difference of Arrival
- Sensor-to-sampler
  - Predictable (wires)
  - Variable (wireless)
- Distributed time sync
  - Microsecond accuracy
- Can we count on the network?



<http://www.douglas-self.com/MUSEUM/COMMS/ear/ear.htm>

Télésitemètre:  
Jean Baptiste Perrin's Acoustic Locator

# Networks and Time

- Distributing time
  - NTP, IEEE 1588 (PTP), White Rabbit, ...
- Respecting time, end-to-end
  - Deadlines
  - Latency
  - Predictability

# When Time Got Respect

- In the beginning, time was revered
- The network *was* the time authority
  - Initially Mabel (or Ernestine) at Central, then
  - Human speaking clock, then
  - Mary Moore and Jane Barbe

NIST: *“The audio portions of the WWV and WWVH broadcasts can also be heard by telephone. The time announcements are normally delayed by less than 30 ms when using land lines from within the continental United States, and the stability (delay variation) is generally  $< 1$  ms”*

Source: <http://www.nist.gov/pml/div688/grp40/ttds.cfm>

# Then Digital Happened

- Initially, time gets more respect
- A synchronous NA “T1” network
  - 1.544 Mbps ( $\sim 648$  nsec/bit)
- With a master clock
  - “Primary Reference Source”
- Dedicated digital circuits from/to anywhere; bit timing under control
  - Jitter tolerated with a modicum of protocol and buffering
  - No such thing as “wander”

Hillsboro, Missouri


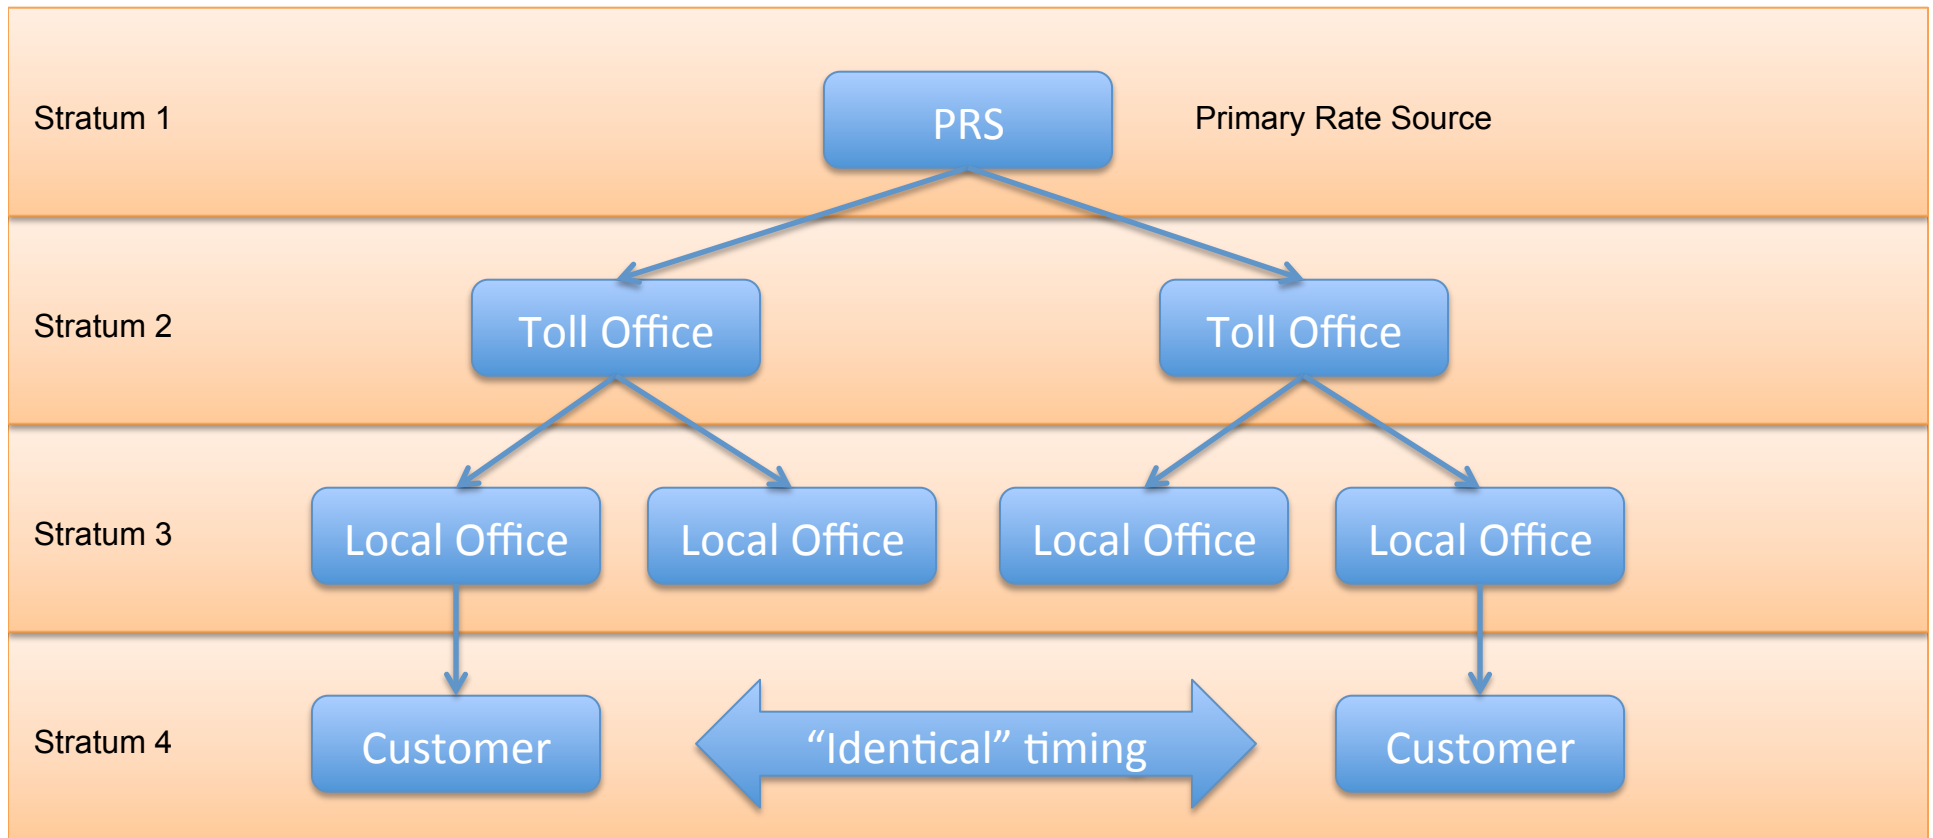


Photo credit: *Denny Reilly*



# Time as an integral element of network: Globally Synchronous



# But Then Times Changed


- AT&T Divestiture and economics drove a (de-)evolution
  - Installed plant was lots of raw copper circuits
  - Overlaid digital coding to get nailed-up, dedicated T1's: synchronous
  - Carved up and shared the T1's statically: Frame Relay
  - Supported multiple carriers / clocks: plesiochronous
  - Statistically shared but bandwidth-guaranteed: ATM (isochronous)
  - Dynamically shared with no guarantees: IP (asynchronous)
- Time did not govern the network; network could not convey time
- Buffering went through the roof; QoS went down

# IP-based Networks and QoS

- Contention-based channel access (Ethernet, 802.11)
- Best Efforts routing
  - “I’d tell a joke about UDP... but I don’t know if you’ll get it”
- Bufferbloat



*NIST: “The audio portions of the WWV and WWVH broadcasts can also be heard by telephone. The time announcements are normally delayed by less than 30 ms when using land lines from within the continental United States, and the stability (delay variation) is generally < 1 ms. **When mobile phones or voice over IP networks are used, the delays can be as large as 150 ms.**”*



There is **no substitute** for real timing guarantees in telecommunications networks

# Can Networks Support Time-Based IoT?

- 2G: things were rather orderly: GSM / TDMA (*the 217 Hz Buzz*)
  - Core network used circuit switching
- 3G: switch to CDMA
- 3.9G(LTE)/802.11: OFDM, requiring more accurate time and frequency
  - Frequency errors manifest as inter-carrier-interference (ICI)
  - Bandwidth is spent synchronizing and estimating channel
    - *e.g.*, 802.11 maxes out at 12 MHz of usable bandwidth in a 20 MHz channel
    - Can better local time reduce lost bandwidth?
- 4G and beyond: OFDMA (channel access) adds further need for accurate time
  - Statistical multiplexing – core network is IP!

# Can Devices Support Time-Based IoT?

Scheme	Time accuracy	Frequency Accuracy
HF Radio	1-10 milliseconds	down to 5 parts per $10^{12}$
Omega nav and VLF	2-10 microseconds	several parts per $10^{11}$
LORAN-C	several microseconds	1 part per $10^{11}$
Portable cesium clock	microsecond	2 parts per $10^{12}$
GPS	100-500 nanoseconds (SA)	Comparable to LORAN-C

GPS only works outdoors (we spend 80% of our time indoors)

Current GPS solutions are very power hungry

Denial of Service → Denial of Position → **Denial of Time attacks**

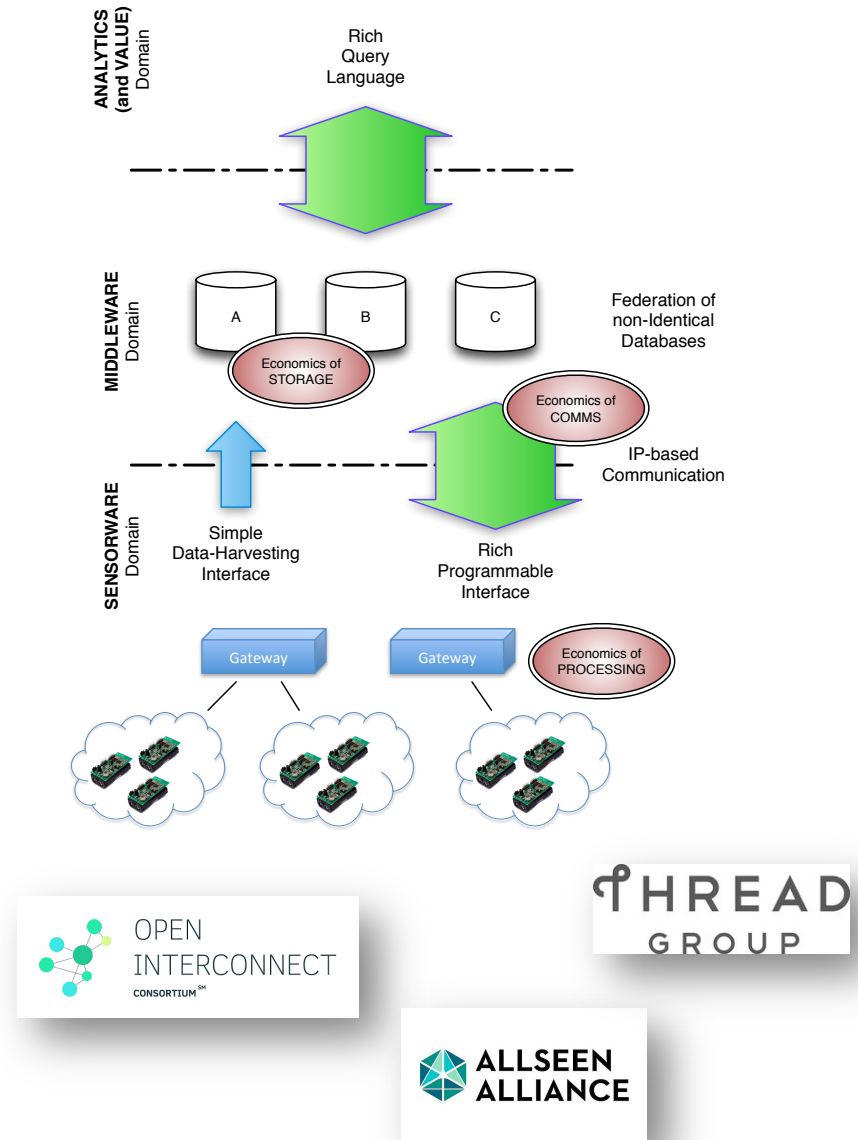
# Time Isn't Just for Communications

- For an app to be time-respecting, all of the underlying elements (computing and comms) must also be time-respecting
- As said, networks don't do this
- Modern processors are bad and getting worse
  - “*Caches are the work of the devil*” -- Burton Smith
  - OOO, speculative, predictive, interruptible, rollback, ...

# FEDERATED

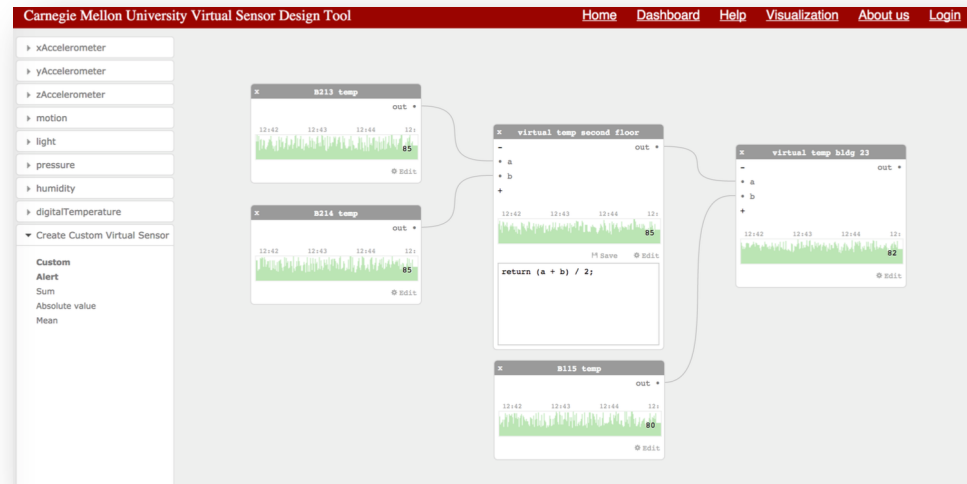
# Federation

- No one entity will build a 100 billion node network
- Nor will it be homogeneous in architecture
- And the real value will come from fusing information
- *“The great thing about standards is that everyone can have his own...”*





# A Typical Paradigm: Virtual Sensors



- Data pulled from many places
- Logic pushed into the network
- Stream-to-stream functions
  - Lazy or Eager evaluation
  - Memoization
- Software becomes part of provenance

Zhang, J., Li, Z., Sandoval, O., Xin, N., Ren, Y., Martin, R. A., Iannucci, B., Griss, M., Rosenberg, S. Cao, J. and Rowe, A.  
**Supporting Personalizable Virtual Internet of Things.**  
In 2013 IEEE 10th International Conference on Ubiquitous Intelligence & Computing (UIC 2013). Sorrento Peninsula, Italy. 2013.

# There's Just One Small Problem

- As the energy cost of computing drops, motivation for pushing logic all the way to the level of sensors will rise
- In a federated network, logic for multiple apps will come together at the sensor devices!
  - Time + Virtualization → Challenges

# SCALABLE

# Scalability

- 100 billion endpoints (or more)
- Easy to understand how comms will need to scale
- Harder to understand how network-embedded computing will need to scale
- Resource limitations will happen – how to reflect these up to the app logic?
  - We don't do this well today
- With increased sharing (federation) comes the need for mutual accountability
  - We don't do this well, either

# Clean-Sheet Design

*If we knew then what we know now...*

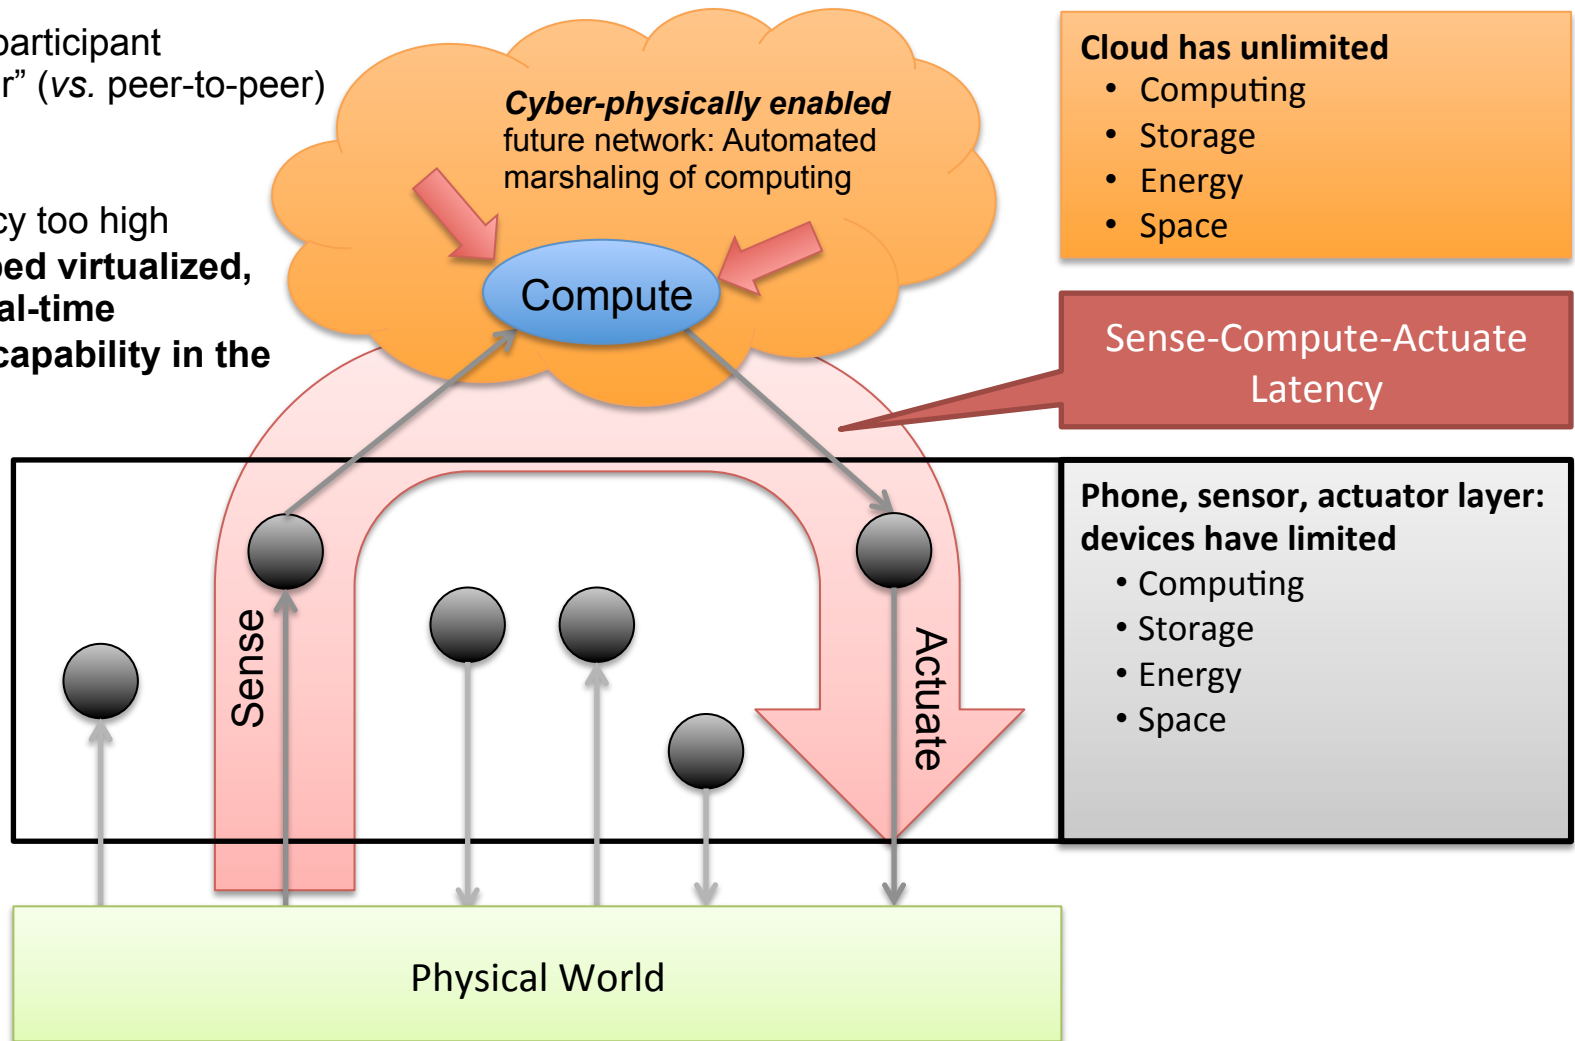
Start with the assumption of building the network side of the IoT computing platform!

- Integrated computing
- Open, compositional structure
- Application- and Service-specific resource allocation

# The Cyber-Physically Enabled Network

Network-as-participant  
“Near to Near” (vs. peer-to-peer)

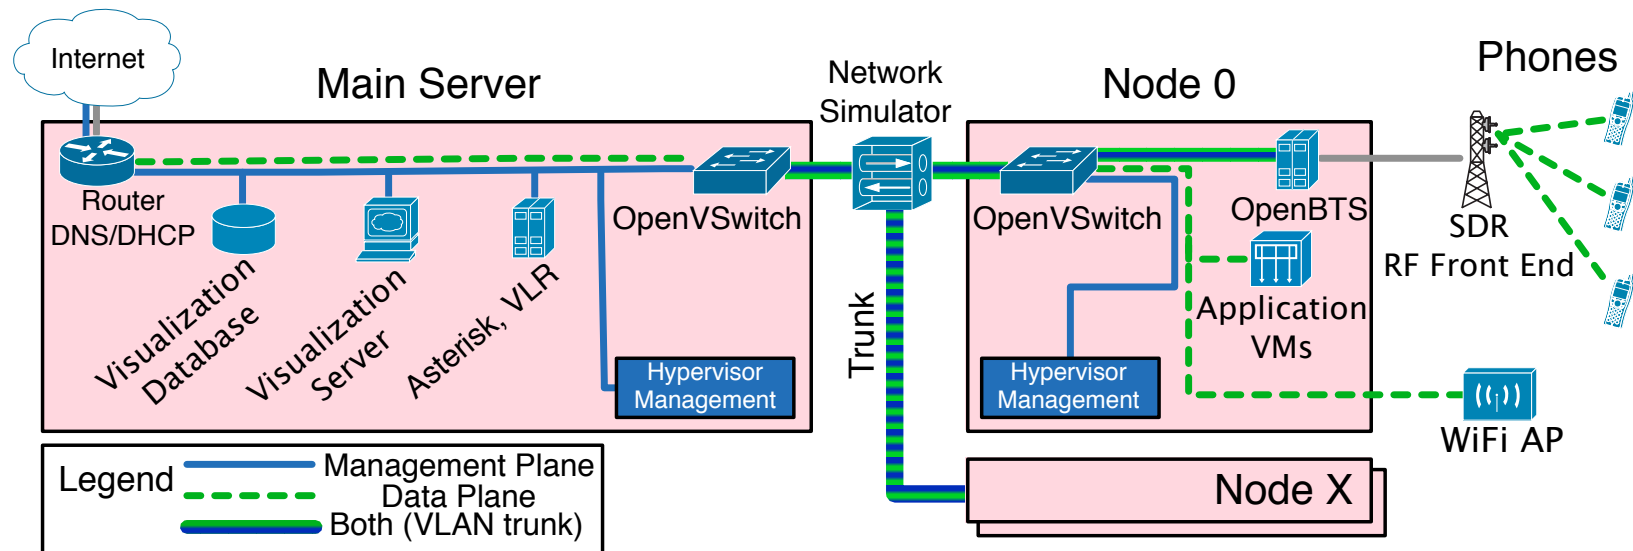
Today: latency too high  
**Future: embed virtualized,  
movable, real-time  
computing capability in the  
network**



# The CROSSMobile Project

- Shift in network organizational principles
  - Apps are treated as *part of the network*
  - Trust is dynamically established
  - App resources are dynamically negotiated
- Created with the IoT in mind
  - Embeddable, time-respecting computing that can migrate
  - Virtualized resources for network *and* computing functions
- A real testbed – for mobile platform research
  - Enough (opt-in) users to generate interesting data sets
  - Supporting research on many aspects

# Network Structure




- Voice (SIP based), SMS, data
- Multiple radio bearers with common instrumentation
- Central transaction database
- Full PSTN integration (dial-in, dial-out, CLID)

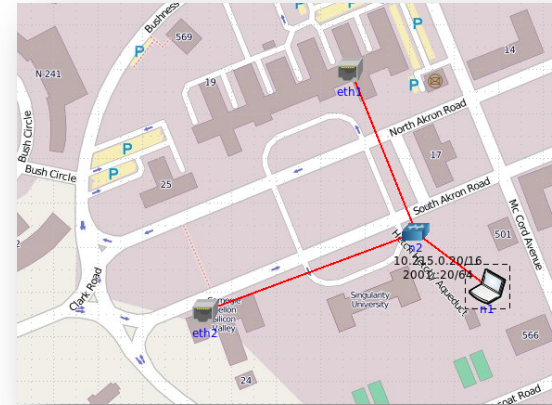
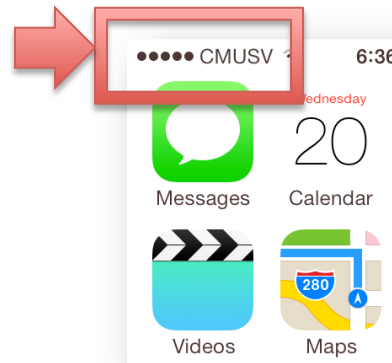



# Node Structure

- Most functions on commodity computing hardware with the exception of the RF front-end (filters, power amps, duplexer)
- Prototype: small rack
- Current: paperback book

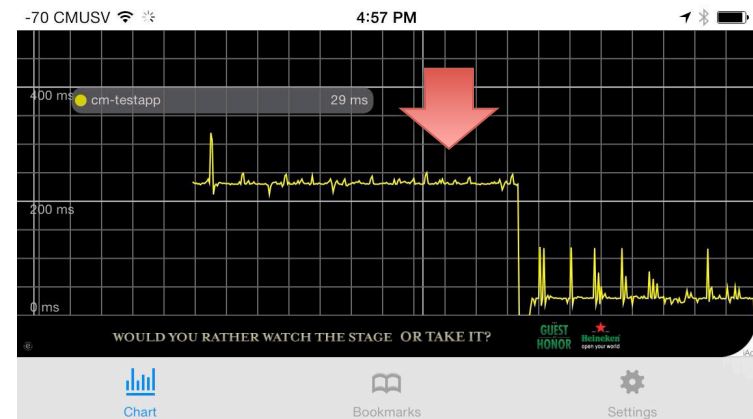



# The Testbed is Live



- CMU-SV campus, Palo Alto, and a mobile node
- Expanding to Pittsburgh campus
- FCC experimental licenses

# Testbed Supports Cloudlets / Fog Computing with Dynamic Migration



Round-trip time (msec)

- OpenStack-based VM migration today
- Moving to Docker or other lightweight containers
  - Lighter-weight than VM
  - Only move what can't be pre-cached (OS, app code, ...)

# Summary

- Platformization will re-invent connected computing
- The nature of apps is still changing
- The network is the un-sung hero
  - Time-awareness
  - Federation
  - Ability to scale
- Open testbeds in pursuit of the connected computing platform will accelerate innovation

# THANK YOU

More info: [sv.cmu.edu/bob](http://sv.cmu.edu/bob)