

IK1350 Protocols in Computer Networks/ Protokoll i datornätverk Spring 2008, Period 3 Module 14: Future and Summary

Lecture notes of G. Q. Maguire Jr.



KTH Information and
Communication Technology

© 2008 G.Q.Maguire Jr. .

All rights reserved. No part of this course may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission of the author.

Last modified: 2008.02.07:09:57

Outline

- Third generation of networking
- QoS
- Interface trends
- IP SANs (Storage Area Networks): iSCSI, ...
- A glimpse into the future.

Third generation of networking

Van Jaconson describes the three generations as:

- "Generation 1: the phone system - focus on the wires.
- Generation 2: the Internet - focus on the machines connected to the wires.
- Generation 3? dissemination - focus on the data flowing between the machines connected to the wires."

-- slide 2: "A Brief History of Networking" of
Van Jacobson, "If a Clean Slate is the solution what was the problem?",
Stanford Clean Slate Seminar, February 27, 2006

<http://cleanslate.stanford.edu/seminars/jacobson.pdf>

Dissemination not conversation

On slide 17 of the same talk, Van Jacobson states:

- "The raison d'être of today's networking, both circuit switched and TCP/IP, is to allow two entities to have a conversation.
- The overwhelming use (>99% according to most measurements) of today's networks is for an entity to acquire or distribute named chunks of data (like web pages or email messages).

Acquiring named chunks of data is not a conversation, its a *dissemination* (the computer equivalent of “Does anybody have the time?”)”

-- slide 17 of
Van Jacobson, "If a Clean Slate is the solution what was the problem?",
Stanford Clean Slate Seminar, February 27, 2006

<http://cleanslate.stanford.edu/seminars/jacobson.pdf>

Quality of Service (QoS)

QoS refers to statistical performance guarantees that a network can make regarding packet loss, delay, throughput, and jitter.

Best effort delivery means no QoS guarantee.

QoS is thought to be more and more important these days.

Many proposals, implementations and studies.

Does Internet need QoS? How can IP network provide it?

Service Differentiation

Integrated Services (InteServ):

- RSVP: connection request
- All nodes IntServ-capable
- Scalability
- Complicated network management

Differentiated Service (DiffServ): end of one-size-fits-all

- Classes of Service
- QoS based Routing
- Classes of Service at Gigabit rates
- New Pricing and Billing Policies
- New Resource Allocation Methods

See: [148]

Constraint-based Routing

QoS routing: selects network routes with sufficient resources for the requested QoS parameters

- to satisfy the QoS requirements for every admitted connection;
- to achieve network efficiency in resource utilization.

Policy-based Routing: e.g. Virtual Private Networks (VPN)

How can we combine this with IP mobility?

Performance

Routers:

1/2 to 1 Million packets per second (pps) for every gigabit per second of aggregate bandwidth

more than 250,000 routes

PC interfaces

Standard I/O ports of PCs:

- Accelerated Graphics Port (AGP)
- PCI
 - Version 2.1 PCI bus - 64 bit, 66MHz, can burst to 528 Mbps
 - PCI-X 2.0: “High Performance, Backward Compatible PCI for the Future” [149]
 - PCI-X 533, offering up to 4.3 gigabytes per second of bandwidth
- Universal Serial Bus (USB)
 - USB: 12Mbps - with plug and play
 - USB 2.0 [150]
- Apple Computers' Firewire™ ⇒ IEEE 1394
 - supporting more than 400 Mbps
 - P1394B (Gigabit 1394) defined in IEEE Std 1394b-2002
- 10/100/1000 Ethernet

Fibre Channel

From the X2T11 standards activity

Topologies: Point-to-Point, Fabric, and Arbitrated loop

Addresses: Loops, LANs, and worldwide addresses

Fibre Channel Profiles

Fibre Channel products


- Disk drives
- Network interfaces

IP Storage Area Networks (SANs)

Using IP in conjunction with storage:

- Fibre Channel Over IP (FCIP)


JBOD == Just a Bunch of Disks



Note that this approach simply interconnects the two Fibre Channel switches. The connection between the two switches is TCP and it simply encapsulates a FCIP header and a Fibre Channel Frame.

- Internet Fibre Channel Protocol (iFCP)


JBOD == Just a Bunch of Disks



Note that this approach interconnects Fibre Channel devices. The connection between the two switches is TCP and it encapsulates a iFCP header and a Fibre Channel Frame; note that iFCP devices can simply be attached to the internet or an intranet. This means that there has to be a mapping between Fibre Channel addresses and IP addresses.

- Internet SCSI (iSCSI)

JBOD == Just a Bunch of Disks



Once a SCSI initiator has logged-in to a SCSI target, it can simply issue SCSI commands, just as if the device were on a local SCSI chain!

For more information see [151]

Clustering

Myricom, Inc. <http://www.myri.com/>

American National Standards Institute (ANSI) Standard -- ANSI/VITA 26-1998

- Started by
 - Prof. Charles L. Seitz - Caltech, now President and CEO
 - Dr. Robert Felderman - Director of Software Development
 - Mr. Glenn Brown - Engineer and programmer

Clusters used to form high performance servers, using commodity networks and hosts. For performance numbers see: <http://www.myri.com/myrinet/overview/>

“Beowulf-class” machines

Using large numbers of commodity machines to make high performance computational systems by interconnecting them with a network.

- LANL's Loki <http://loki-www.lanl.gov>
- LANL's Avalon <http://cnls.lanl.gov/avalon/>
- JPL's Hyglac <http://hpc.jpl.nasa.gov/PS/HYGLAC/hyglac.html>
- INRIA's PopC (Pile of PCs)
- ...

Very high-speed Backbone Network Service (vBNS)

vBNS project (<http://www.vbns.net/>) created to provide a backbone for the US high-performance computing users and their SuperComputer Centers.

- mostly OC12C, but now adding OC48C links (2.4Gbps)
- connections to all NAPs
- provide for multimedia services (provides multicast)
- participate in developing advanced routing technologies
- supports IPv4 and IPv6

vBNS backbone network (<http://www.stanford.edu/group/itss-cns/i2/vbns.html>)

There is some interest in using vBNS for inter-gigapop interconnections.

Internet2

<http://www.internet2.org/>

- World class **research**
Driven by computational physics, biology, chemistry, ... and scientific visualization, virtual “experiments”, and remote control of real experiments.
- Networking R&D - focused on exploiting the capabilities of broadband networks media integration, interactivity, real time collaboration, ...
- Improve **production** Internet services and applications for **all** members of the academic community, both nationally and internationally.

Purpose: support national research objectives, distance education, lifelong learning, and related efforts.

<http://www.hpcc.gov/white-house/internet/background.html>

Gigapops

Who will be operating them?

Where will they be?

How many will there be?

What is the aggregate throughput that they will require?

What is the maximum per port throughput?

How many ports will they need to support?

Will they support "mixing"? (mixing is used to defeat traffic analysis)

Whose hardware and software will they use? What is the required functionality?

Speed through Silicon

FPGAs used in many routers - for flexibility and to allow near hardware speed implementations of protocols.

ASICs: Vertex Networks, Inc. , MMC Networks, Inc, Galileo Technology, TI, ...

Future networks

Terabit per second == 10^{12}

Readily achievable via combining multiple Gigabit per second streams using Wavelength Division Multiplexing (WDM).

Petabit per second == 10^{15}

Differentiated Services: Classes of Service, Multimedia

Constraint-based Routing (QoS Routing)

Ad Hoc Networking

Auto-configuration (Plug and Play Internet)

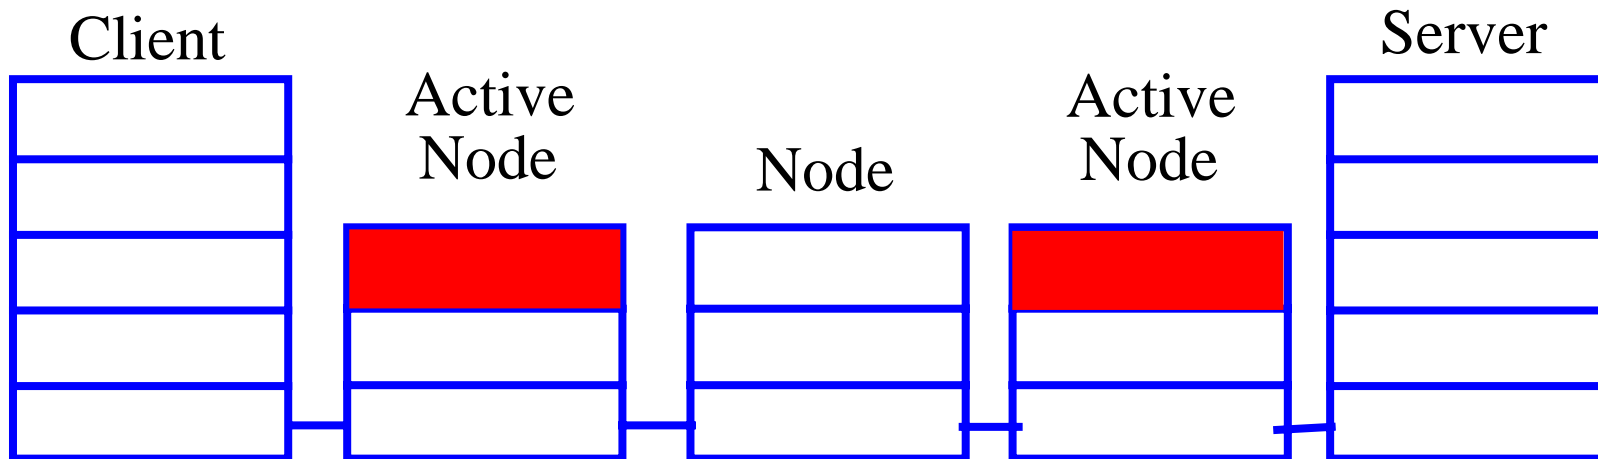
Active Networking

Smart Networking

Knowledge-based Networking

Active Networks


- Network nodes can perform customized computations on the messages flowing through them.
- Can change, modify the contents of the messages.
- Potentially Mobility Enabling Routing using active network concept



Smart Networks with Sensors


(Ren and Maguire, KTH)

- Context/situation-aware Systems
- Smart services: active + user-awareness networking
- Knowledge-based Networking

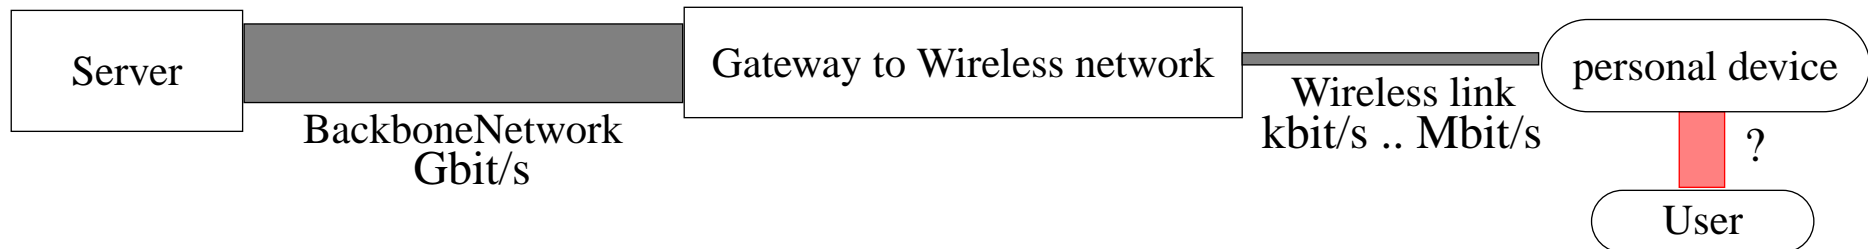


Bottlenecks

- Server and Network Bandwidth and **latency**



- User Bandwidth and **latency**



- Power and Energy ⇒ need a computational theory of $O(\text{energy})$
- **Imagination!**

Near Future systems

Personal Portal




Figure 138: Vision-2, 2000 - high level of integration

Evolution of new varieties of networks

Already we have: **WANs** (Wide Area), **MANs** (Metropolitan Area), **LANs** (Local Area Networks)

VANs **Vehicle Area Networks**

Very local networks

DANs **Desk Area Networks**

The computer/printer/telephone/... will all be part of a very local area network on your desk.

- ◆ wireless links ⇒ No longer will you have to plug your printer into your computer (PDA/...) into your computer
- ◆ active badges ⇒ No longer will you have to sign in/out of areas, write down peoples names at meetings, ... the system can provide this data based on the active badges

Olivetti and Xerox are exploring “Teleporting” your windows environment to the workstation nearest you, on command, if there are multiple choices probe each one (currently a “beep” is emitted to tell the user which).

BANs **Body Area Networks**

Users will be carrying multiple devices which wish to communicate:

- ◆ thus there will be a need for a network between these devices which you carry around; and
- ◆ personal devices will wish to interact with fixed devices (such as Bankomat machines, vehicle control systems, diagnostic consoles (for a “mechanic” or repairman), ...) and other peripherals.

Situational awareness and Adaptability




Figure 139: Where am I? What am I? Who am I?

Where am I going? When will I be there? What should I become? Who should I become?

- Location dependent services
- Predicting location to reduce latency, reduce power, hide position, ...
- Adapting the radio to the available mode(s), purposely changing mode, ...
- Reconfigure the electronics to adapt, for upgrades, for fault tolerance, ...; Reconfiguration vs. powering up and down fixed modules (what are the “right” modules, what is the “right” means of interconnect, what is the “right” packaging/connectors/..., needed speed of adaptation)
- “right” level of independence; spectrum from Highly Independent \Rightarrow Very Dumb

Location Dependent service(s)

How do I know where I am?

- Outdoors: GPS or from the network operators knowledge [resolution: 100m to sub-centimeter]
- Indoor: IR and RF beacons, triangulation, knowing what you can **see** or **hear**

What can I do with this knowledge?

KTH students built a JAVA Applet which gets data from GPS unit and dynamically displays a list of the information available - as a function of where you are:

- ◆ if near bus, subway, train stop - you get transit information - potentially with real-time schedule - since the system knows current location of vehicles
- ◆ list of restaurants, shops, etc. where you are and in the direction you are headed
 - ◆ the scope is based on your **velocity vector** - so if you move quickly it reduces detail, but increases the scope
- ◆ map information with updated position

How do I know who I'm with or what I'm near?

- Olivetti, Xerox, and MIT - using IR emitters as "ID" tags
 - ◆ Olivetti put them on people, equipment, ...
 - ◆ Xerox put them on electronic notepads, rooms, ...
 - ◆ MIT Media Lab is putting them on people + lots of inanimate objects (clock, fish tank, ...)

Human centered

- Computer - human interaction is currently focused on the computer (computer-centric)
 - ◆ Currently computers know little about their environment
 - ◆ **Where** are we?
 - ◆ **Who** is using me?
 - ◆ Is the user **still** there?
- Evolving Environment awareness
 - ◆ Give computers senses via sensors
 - ◆ **Environment**
 - ◆ User **identity** and **presence**
- Badge as a smart card replacement
 - ◆ biometric signature of the person currently using the badge
 - ◆ the badge ensures that only you can use it
- You wear your own personal user interface
 - ◆ interface can be consistent across all appliances
 - ◆ not because each appliance supports the interface, but because the user's own interface provides consistency
- Make the **human** the focus of the computer's interaction (\Rightarrow human-centric)

Requirements

- Systems with which humans wish to interact:
 - ◆ traditional computers, desktop workspaces, domestic appliances, building and automotive systems, doors, elevators (lifts), environmental control, seats and mirrors, etc.
- Systems to provide sensor data:
 - ◆ location, orientation, light, heat, humidity, temperature, gas analysis, biomedical, ...
- Systems to correlate the sensor information and provide it in a useful way to the computer systems:
 - ◆ Spatial and temporal sensor fusion,
 - ◆ 3D and 4D databases,
 - ◆ Machine Learning, and
 - ◆ Prediction (based on pattern extraction)
- Agents and actuators to provide intelligent control of the environment
- wireless/wired/mobile communications **infrastructures** to link it all together
 - ◆ must assure privacy and security

Dumb Badge, Smart Badge, and Intelligent Badge

- Dumb Badge just emits its ID periodically
- Smart Badge - [an IP device] Location and Context Aware (i.e., a sensor platform)
- Intelligent Badge - add local processing for local interaction by the user


Acknowledgment:

All of the badge work is done in cooperation with:

- Dr. Mark T. Smith - Hewlett-Packard Research Laboratories, Palo Alto, California, USA
- Dr. H. W. Peter Beadle
 - ◆ Formerly: University of Wollongong, Wollongong, Australia
 - ◆ Currently: Director, Motorola Australian Research Centre, Botany, NSW, Australia

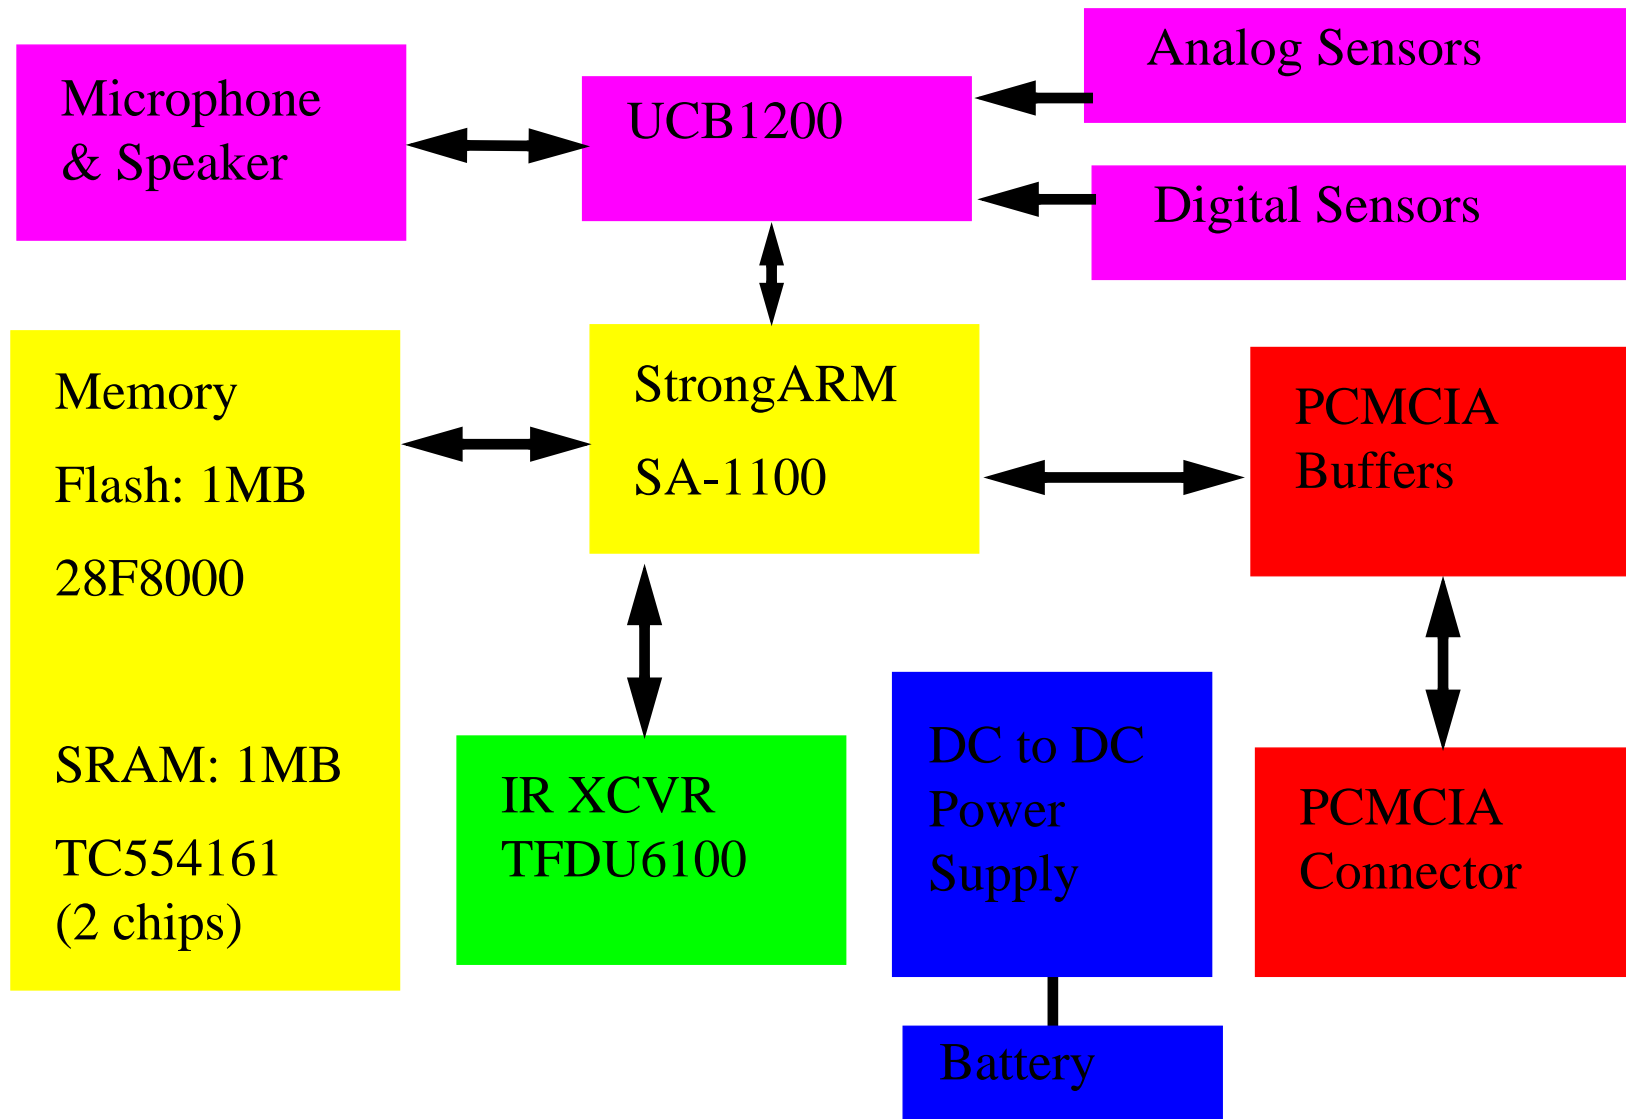
Badge Communications Model

Badges are IP devices (or should be), they communicate via network attached access points.




⇒ Banks as intermediaries (**if** they have **any** future role)

Smart Badge 3





Smart Badge Sensors



Details of the 3rd version:

<http://www.it.kth.se/edu/gru/Fingerinfo/telesys.finger/Mobile.VT98/badge3.html>

Badge 3




A view of the packaged badge

As shown by HP at Comdex'98, November 16-20, 1998



MEDIA

High integration (goal of MEDIA project)





Partners:


- Kungl Tekniska Högskolan (KTH/ELE/ESDlab and KTH/IT/CCSlab)
- Tampere University of Technology (TUT)
- GMD FOKUS (GMD)
- Technische Universität Braunschweig (UBR)
- Interuniversity Microelectronics Centre (IMEC)
- Ericsson Radio Systems AB (ERA)

See <http://www.ele.kth.se/ESD/MEDIA> for more information

Split the functions between access point and access point server



Future home/office/... network accesspoints



Personal Computing and Communication (PCC)

Upper limit of bandwidth: saturate the senses: sight, sound, touch, smell, taste
⇒ ~1 Gbit/sec/user

Current workstations shipping with 1 Gbit/sec interfaces for LAN!

Telepresence for work is the long-term “killer” application

-- Gordon Bell and James N. Gray¹

1. “The Revolution Yet to Happen” in Beyond Calculation: The Next Fifty Years of Computing, Eds. Denning and Metcalfe, Copernicus, 1997.

Uploading ourselves to the net

In Bob Metcalfe's speech at MIT: <http://web.mit.edu/alum/president/speech.html>

One of great insights of this talk is that the internet is the way to **immortality**¹:

Now, for the next 50 years, the web will drive electronic commerce into the information age, ubiquitous computers will disappear into the woodwork, and we'll start uploading ourselves into the Internet to become at last immortal.

-- Robert M. Metcalfe
June 26, 1997

1. Robert M. Metcalfe, "Internet Futures", MIT Enterprise Forum, June 26, 1997.

Future Systems




Figure 140: Vision-3, 2005-2015 - very high level of integration

Bionic Technologies, Inc.'s Intracortical Electrode Array

Acute microelectrode assembly (10x10 array, 100 active electrodes) \$1,250.00




Figure 141: 10 x 10 silicon electrode array (each electrode: 1.5mm long, 0.08mm wide at base, 0.001mm tip), Built at the Univ. of Utah, by Richard A. Normann, et al.; from Scientific American, March 1994, pg. 108.

Non-metallic bi-directional neural interfaces

Neurochip: Neuron silicon circuits <http://mnphys.biochem.mpg.de/> :

(a) Silicon-Neuron Junction (input to the nerve)

(b) Neuron transistor (output from the nerve)




Figure 142: (a) Capacitive coupling of data into nerve and
(b) using the charge in the nerve to control a transistor's gate for getting data out of the nerve

(a) Peter Fromherz and Alfred Stett, “Silicon-Neuron Junction: Capacitive Stimulation of an Individual Neuron on a Silicon Chip” *Phys.Rev.Lett.* 75 (1995) 1670-1673

(b) P.Fromherz, A.Offenhäusser, T.Vetter, J.Weis, “A Neuron-Silicon Junction: A Retzius-Cell of the Leech on an Insulated-Gate Field-Effect Transistor” *Science* 252 (1991) 1290-1293.

What is *your* time line?

- What is going to be your planning horizon?
- What will be the depreciation time for your equipment/software/infrastructure/... ?
- How fast:
 - ◆ can you change?
 - ◆ should you change?
 - ◆ will you change?

Summary

- Telecom operators are **reinventing themselves and their infrastructures**
- Things to watch IPv6, IPsec, Mobile-IP, DHCP, the new domain name registries, appliances, ...
- Low cost access points which exploit existing or easily installed infrastructure are key to **creating a ubiquitous mobile infrastructure with effectively infinite bandwidth.**
- Smart Badge is a vehicle for exploring our ideas:
 - Exploits hardware and software complexity by hiding it.
 - Explores allowing devices and services to use each other in an extemporaneous way.
 - Enables a large number of location and environment aware applications, most of which are service consuming.
 - Keep your eyes open for the increasing numbers of sensors which will be on the network.
 - Service is where the money is!
- Personal Communication and Computation in the early 21st century: **“Just Wear IT!”**
- Coming in 20-30 years: “Just implant IT!”
- Remember: The internet will be what **you** make it.

Further Reading

- [148] Kalevi Kilkki, *Differentiated Services for the Internet*, Macmillan Technical Publishing, 384 pages, June 1999, ISBN: 1578701325.
- [149] PCI-SIG, *PCI-X 2.0: High Performance, Backward Compatible PCI for the Future*, May 19, 2005 http://www.pcisig.com/specifications/pcix_20
- [150] USB.org, *Universal Serial Bus Revision 2.0 specification*, May 19, 2005
http://www.usb.org/developers/docs/usb_20_02212005.zip
- [151] Tom Clark, *IP SANs: A Guide to iSCSI, iFCP, and TCIP Protocols for Storage Area Networks*, Addison-Wesley, 288 pages, 2002, ISBN: 0-201-75277-8

General: <http://www.ietf.org/>

Thanks

Best wishes on your written assignments (or projects).