Home Broadband Access
- looking at future networking technologies

Gregor v. Bochmann

School of Information Technology and Engineering (SITE)
University of Ottawa
Canada

http://www.site.uottawa.ca/~bochmann/talks/IPS-Mome-06

L'Université canadienne
Canada's university

Invited talk at the
4th International Workshop on
Internet Performance, Simulation, Monitoring and Measurement (IPS-MoMe)

Salzburg, Austria, February 27-28, 2006

Motivation

- IPS-MoMe : Internet Performance
- What will be the Internet of the future?
  - architectures – transmission technologies – routing and bandwidth allocation – management – economic issues
- Try to capture a vision for the future
- Report on a research project on "Agile All-Photonic Networks" (AAPN)

Overview

- Community-base research planning
- Research topics for the future of Networking
  - The E-NEXT Research Network
  - An NSF workshop on the future of the Internet
  - Other research topics
- Optical networks
- Research project on "Agile All-Photonic Networks"
- Conclusions

Community-based research planning

- **Consensus building:** through mailing lists, discussions at workshops / conferences, research collaborations
- Examples:
  - The UK Grand Challenges: a perspective on long-term basic and applied research
  - NSF (USA) Workshop on Overcoming Barriers to Disruptive Innovation in Networks
  - Research program of E-NEXT (a EU - FP6 Network of Excellence)
    - "CoNEXT" conference in Toulouse, Oct. 2005 http://dmi.ensica.fr/conext/
  - Other "Grand Challenges"
    - The DARPA Grand Challenge (USA): Automated car race in the desert
    - Computer Research Association (North-America) holds workshops on "Grand Research Challenges in Computer Science and Engineering" http://www.cra.org/grand.challenges/
  - Canadian research network on Agile All-Photonic Networks (AAPN, funded by NSERC and 6 industrial partners)

UK Grand Challenges

- See http://www.ukcrc.org.uk/grand_challenges/index.cfm
- “**Definition of a Grand Challenge**
  - A grand challenge should be defined as to have international scope, so that contributions by a single nation to its achievement will raise our international profile.
  - The ambition of a grand challenge can be far greater than what can be achieved by a single research team in the span of a single research grant.
  - The grand challenge should be directed towards a revolutionary advance, rather than the evolutionary improvement of legacy products that is appropriate for industrial funding and support.
  - The topic for a grand challenge should emerge from a consensus of the general scientific community, to serve as a focus for curiosity-driven research or engineering ambition, and to support activities in which they personally wish to engage, independent of funding policy or political considerations. “ *(Note: the quotes, here and in subsequent slides, indicate that the text is copied from the source documentation)*

Grand Challenge Exercise

(from Robin Milner's talk at the IFIP World Congres 2004 in Toulouse)

WHAT IS A GRAND CHALLENGE EXERCISE?

- The community examines and adopts long-term goals . . .
- . . . from within the science, not outside it.
- Thus to develop and refine a portfolio of proposals . . .
- . . . showing the public (and funders) what we aspire to.

Ubiquitous Computing Grand Challenge

- Combination of two UK Grand Challenges: GC 2 and GC 4
  - See http://www-dse.doc.ic.ac.uk/Projects/UbiNet/GC/index.html
- **Objective:** "We propose to develop scientific theory and the design principles of *Global Ubiquitous Computing* together, in a tight experimental loop."
- **"Engineering challenges:**
  - design devices to work from solar power, are aware of their location and what other devices are nearby, and form cheap, efficient, secure, complex, changing groupings and interconnections with other devices;
  - engineer systems that are self-configuring and manage their own exceptions;
  - devise methods to filter and aggregate information so as to cope with large volumes of data, and to certify its provenience.
  - business model for ubiquitous computing, and other human-level interactions. "

Ubiquitous Computing Grand Challenge (ii)

- "**Scientific challenges**":
  - discover mathematical models for space and mobility, and develop their theories; devise mathematical tools for the analysis of dynamic networks;
  - develop model checking, as well as techniques to analyse stochastic aspects of systems, as these are pervasive in ubiquitous computing;
  - devise models of trust and its dynamics;
  - design programming languages for ubiquitous computing. "

- **A comment:** It is not clear where – in the context of ubiquitous computing – *Networking* stops and *Computing* starts. In fact, networking involves much distributed systems management (including databases); and for the Internet applications, the application layer protocols are just as important as (if not more than) the underlying networking protocols.

Overview

- Community-base research planning
- **Research topics for the future of Networking**
  - The E-NEXT Research Network
  - An NSF workshop on the future of the Internet
  - Other research topics
- Optical networks
- Research project on "Agile All-Photonic Networks"
- Conclusions

Overview of research topics

Architectural levels of Networking Technology
a narrow-waisted hourglass model:

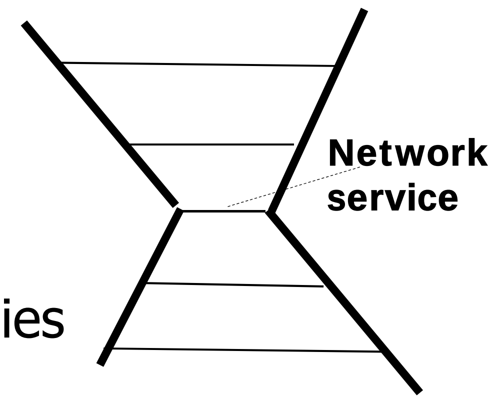

Issues

- Network layer: integration of new wireless technologies
  - cellular, LAN, PAN, ad-hoc, sensor, etc.
- Physical layer: technology push
  - Faster electronic components, e.g. 10 Gbps Ethernet
  - Fast optical switching
  - Trend: IP over Dense Wavelength Division Multiplexing (DWDM); elimination of intermediate layers of ATM, SONET; however, it may be IP over MPLS over DWDM.
- Application layer
  - Many new applications: importance of multimedia application will increase
  - New protocols for organizing applications: Web Services, Grid
  - New ways for identifying and searching services, including concern for security and trust

E-NEXT Research Network

An FP6 (EU) Network of Excellence that focuses on Internet protocols and services (see http://www.ist-e-next.net/about.php)

- 4 Work Packages:
  - Mobile and Ambient networking
    - "The main issue will come with the integration of a large set of mobile devices in a various scale network, putting a strong pressure on important functions such as routing, congestion and flow control, signalling, scalability, etc"
    - "ad-hoc networks, sensor networks, ambient networks … share in common self-organizing capabilities that require the design of mechanisms such as power control, discovery mechanisms, and auto-configuration."
  - Content Networking
    - " … to understand the various means to broadcast, distribute and manage such contents (data, and audio-visual information). Many of today's applications from online gaming to multimedia broadcasting or interactive audio-visual content distribution require a group or a distribution communication infrastructure, whether physical or virtual." - possibly using overlay networks

E-NEXT Research Network (ii)

- Work Packages (suite)
  - Self-Aware & Scalable Networking
    - "… modelling techniques and monitoring can be used together to achieve greater efficacy (of networks). Traffic engineering, for example, involves adapting the routing across network elements to the network conditions, with the joint goals of good user performance and efficient use of network resources."
    - "The name "Self-Aware Networking" derives from the consideration that the science of design, control and management of complex networked infrastructures can only be successful if the entities (nodes, terminal, software objects) composing such infrastructures are capable to have control on themselves and on the other components as well."
  - Service Aware Networking
    - " … satisfying simultaneously user and service provider interests from "just in time" elaborated communication architectures. This goal can be achieved with adaptive reconfiguration capabilities, programmable networks or, in a longer term, active networking."

NSF Workshop: Overcoming Barriers to Disruptive Innovation in Networks

Workshop organized by NSF (USA)

"Overcoming Barriers to Disruptive Innovation in Networking"
(Jan. 2005) **see** http://www.arl.wustl.edu/netv/noBarriers_final_report.pdf

Starting point: " The Internet is ossified: … Adopting a new architecture not only requires modifications to routers and host software, but given the multi-provider nature of the Internet, also requires that ISPs jointly agree on that architecture. The need for consensus is doubly damning; not only is agreement among the many providers hard to reach, it also removes any competitive advantage from architectural innovation. This discouraging combination of difficulty reaching consensus, lack of incentives for deployment, and substantial costs of upgrading the infrastructure leaves little hope for fundamental architectural change. "

NSF workshop (ii)

- Requirements for the new Internet:
  - " **Minimize trust assumptions:** the Internet originally viewed network traffic as fundamentally friendly, but should view it as adversarial;
  - **Enable user choice:** the Internet was originally developed independent of any commercial considerations, but today the network architecture must take competition and economic incentives into account;
  - **Allow for edge diversity:** the Internet originally assumed host computers were connected to the edges of the network, but host-centric assumptions are not appropriate in a world with an increasing number of sensors and mobile devices;
  - **Design for network transparency:** the Internet originally did not expose information about its internal configuration, but there is value to both users and network administrators in making the network more transparent; and
  - **Meet application requirements:** the Internet originally provided only a best-effort packet delivery service, but there is value in enhancing (adding functionality to) the network to meet application requirements. "
- Identified 7 areas of research (no time for explaining them here)

Synthesis

- Recurring themes *(in E-Next and NSF workshop)*
  - Security: "minimize trust assumptions" (NSF)
  - Integration of the Internet with mobile communication and ad hoc / sensor / ambiant networks (NSF: "edge node diversity", end host assumptions, address binding)
  - Network visibility and awareness (NSF: "network transparency"): performance monitoring, economic incentives, user-level route choice, control and management
  - Service-awareness (NSF: "meet application requirements"): service level agreements, traffic engineering
- Other: Content Networking *(use overlay networks ??)*

Note: Items in blue are the 7 research areas identified by the NSF Workshop

Meeting application requirements

- Simple Network layer service: IP

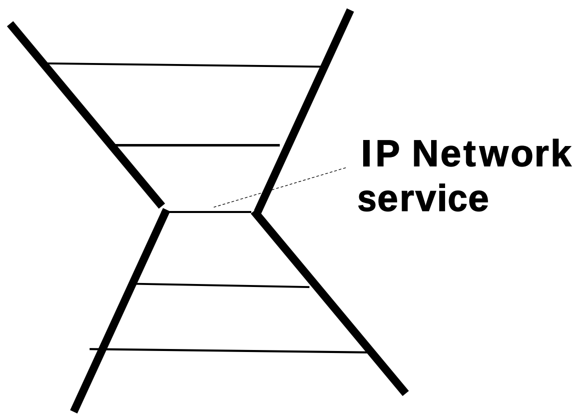

- Additional requirements:
  - "QoS control, multicast, anycast, policy-based routing, data caching …"
- Possible solutions:
  - Add more functions to IP layer
  - Use overlay networks to provide additional functions

Other topics: Overlay networks

- **Principle:** A certain number of servers connected to the Internet play the role of « virtual routers » in the overlay network.
  - Note: This is the way MBone implements multicasting over the current IP Internet service.
- The NSF workshop stresses the use of overlay networks for experimentation with new approaches
  - Could such architectures present the final solution ?
- Existing well-known applications
  - Napster and peer-to-peer media distribution applications
  - Multicasting of multimedia presentations, possibly including different quality variants
- **A Testbed:** Planetlab http://planet-lab.org/ (a global organization, but a steering committee from the USA); see also http://www.arl.wustl.edu/netv/main.html

Other topics: Lightpaths

- Experimental research networks provide high-bandwidth "lightpaths" between different sites for e-science and other applications that require guaranteed high-bandwidth connections.
  - For an overview of current applications, see http://www.internet2.edu/presentations/fall05/20050920-lambdas-sauver.htm
- *Note: The Internet is layered over the lightpaths*
- A lightpath may be realized as
  - An MPLS flow with some guaranteed bandwidth (with intermediate MLPS switches)
  - An optical wavelength (with intermediate optical switches, e.g. with OEO conversion or optically transparent ROADMS)
  - A fiber carrying several wavelength (with intermediate MEM switches)
- User-Controlled Lightpath Provisioning (UCLP) allows the e-science users to establish lightpaths dynamically through a graphic user interface or a Web Service called by an application.
  - Note: UCLP has been initiated in Canada with partial funding from Canarie (the organization responsible for the Canadian research network), see for instance http://phi.badlab.crc.ca/uclp/
- These networks make use of user-owned fibers and condominium facilities for long-haul transmission and switching

Note: Packets vs. (virtual) connections

- The old debate between packet switching and circuit switching is not dead !!
  - Distinction: In packet switching, the header of the packet/frame/cell/burst contains the destination address; in circuit switching, it contains a number (label) identifying the circuit (in TDM, this number is the timing position).
- MPLS (label switching) provides packet switching over dynamically established paths (virtual connections)
- Optical lightpaths are connection-oriented. It is expected that existing ROADM (Reconfigurable optical add/drop multiplexers) technology will be widely deployed within a few years; see for instance

  http://lw.pennnet.com/Articles/Article_Display.cfm?Section=ARTCL&ARTICLE_ID=203231&VERSION_NUM=1
- An optical lightpath at a given wavelength is very large, typically 10 Gbps.
  - Sub-multiplexing of a lightpath in the time domain is proposed by many research projects; usually by sharing the bandwidth between several virtual connections.

Overview

- Community-base research planning
- Research topics for the Future of Networking
  - The E-NEXT Research Network
  - An NSF workshop on the future of the Internet
  - Other research topics
- **Optical networks**
- Research project on “Agile All-Photonic Networks“
- Conclusions

Optical networks

- Currently deployed:
  - optical transmission with DWDM
  - Some optical switching
    - Note: most "optical switches" convert the optical signal into the electrical domain and perform the switching in the electrical domain (oeo conversion).
- Expected to be deployed:
  - Transparent optical switching at speeds of milliseconds
    - ROADMs: switching individual wavelengths
    - MEM switches: switching all light in a fiber
- Being developed: optical switching in microseconds

Burst switching

- **Question:** Can one do packet switching in the optical domain (without oeo conversion)?
  - At a switching speed of 1 μs, one could switch bursts of 10 μs length (at 10 Gbps, containing 100 Kbits; typically many IP packets)
  - For each burst, control information is sent through a separate control channel.
- **Note:** Burst switching allows to share the large optical bandwidth among several virtual connections.
- The term "burst switching" originally meant "no reservation":
  - in case of conflict for an output port, one of the incoming bursts would be dropped or could be buffered.
  - Should one introduce optical buffers in the form of delay lines?
- **Note:** There is **no** consensus that the future evolution of the Internet will go towards burst switching. See for instance http://www.nren.nasa.gov/workshop8/ppt/Level3_ONT2_7_v1.ppt and http://www.nren.nasa.gov/workshop8/

AAPN

An NSERC Research Network

The Agile All-Photonic Network : An Architectural Overview

Project leader: David Plant, McGill University

Theme 1: Network architectures
Gregor v. Bochmann, University of Ottawa

Theme 2: Device technologies for transmission and switching

AAPN Professors (Theme 1 in red)

- McGill: Lawrence Chen, Mark Coats, Andrew Kirk, Lorne Mason, David Plant (Theme #2 Lead), and Richard Vickers
- U. of Ottawa: Xiaoyi Bao, Gregor Bochmann (Theme #1 Lead), Trevor Hall, and Oliver Yang
- U. of Toronto: Stewart Aitchison and Ted Sargent
- McMaster: Wei-Ping Huang
- Queens: John Cartledge (Theme #3 Lead)

The AAPN research network

- Our vision: Connectivity "at the end of the street" to a dynamically reconfigurable photonic network that supports high bandwidth telecommunication services.
- Technical approach:
  - Simplified network architecture (overlaid stars)
  - Specific version of burst switching
    - Fixed burst size, coordinated switching at core node for all input ports (this requires precise synchronization between edge nodes and the core)
      - See for instance http://beethoven.site.uottawa.ca/dsrg/PublicDocuments/Publications/Hall05a.pdf
    - Burst switching with reservation per flow (virtual connection), either fixed or dynamically varying
      - See for instance http://beethoven.site.uottawa.ca/dsrg/PublicDocuments/Publications/Agus05a.pdf

Starting Assumptions

- Avoid difficult technologies such as
  - Wavelength conversion
  - Optical memory
  - Optical packet header recognition and replacement
- Current state of the art for data rates, channel spacing, and optical bandwidth
- Simplified topology based on overlaid stars
- Edge based control in small/medium size edge nodes

Starting Assumptions (ii)

- No distinction between long-haul and metro networks
- Fast optical space switching (<1 μsec)
- Slotted Time Division Multiplexing (TDM) or slotted burst switching
- Need for fast compensation of transmission impairments (<1 μsec)

Agile All-Photonic Network

Edge node with slotted transmission
(e.g. 10 Gb/s capacity per wavelength)

Fast photonic core switch
(one space switch per wavelength)

Opto-electronic interface

- *Provisions sub-multiples of a wavelength*
- *Large number of edge nodes*

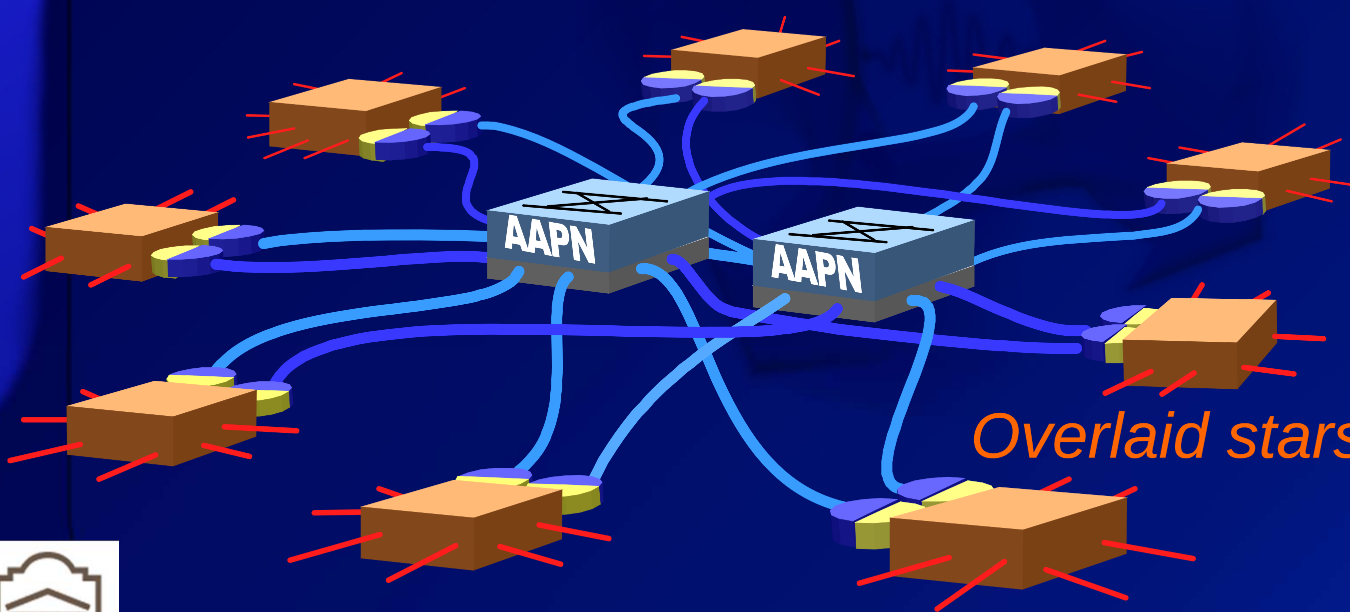

Overlaid stars architecture

Architecture (ii)

- A wavelength stack of bufferless transparent photonic switches is placed at the core nodes
  - a set of space switches, one switch for each wavelength
- Port sharing is required to allow a core node to support large numbers of edge nodes
  - A *selector* may therefore be used between edge and core nodes

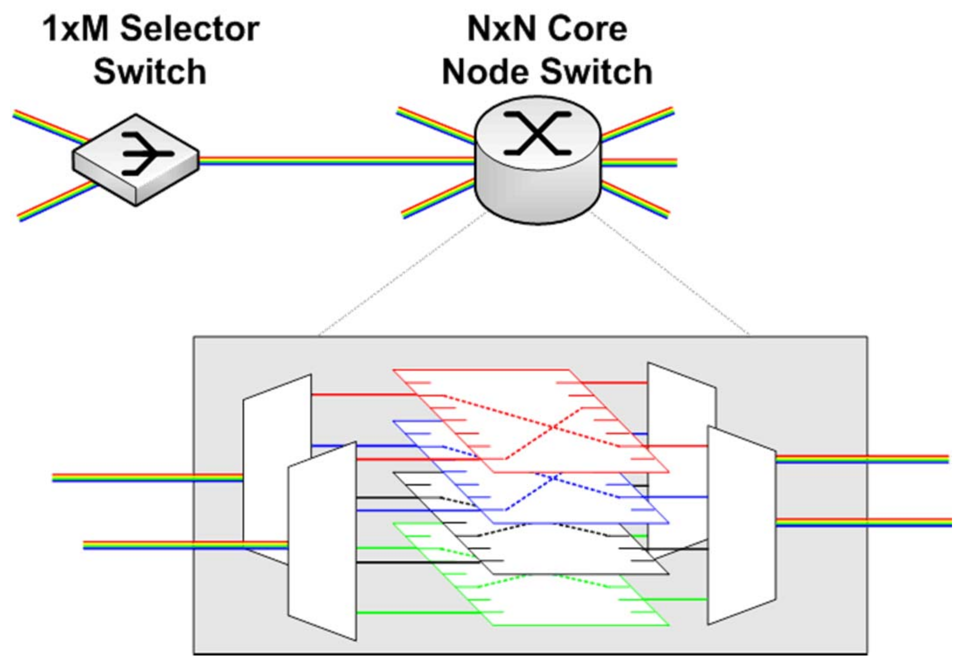

Bandwidth allocation schemes

For flows between edge nodes

- Optical wavelength: Whole wavelength (for large bandwidth flows) – like the PetaWeb explored by Nortel Networks
- Optical circuit: One or several time slots within each TDM frame
- Burst switching: individual bursts (with or without reservation)

- Coordination by controller at core node
  - Signaling protocol between edge and core node

(suitable for metro and long-haul networks)

Integration with above (MPLS and IP)

- MPLS flows passing through the AAPN
- With N edge nodes, there are N x N links in the AAPN (scalability problem for IP routing protocol)
- "Virtual router" star architecture
- OSPF sub-areas
- How to find optimal inter-area route

  (work sponsored
  by Telus)

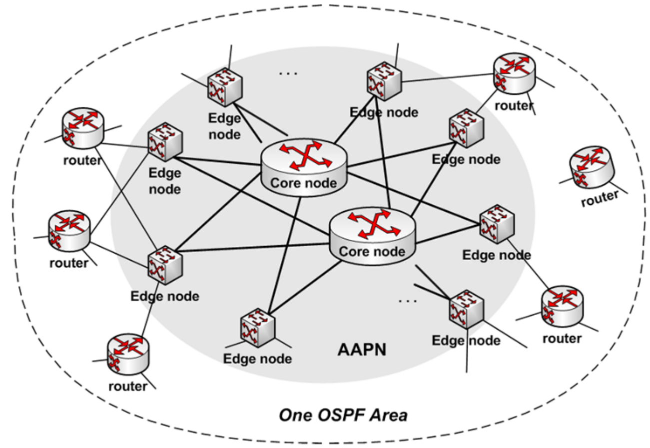

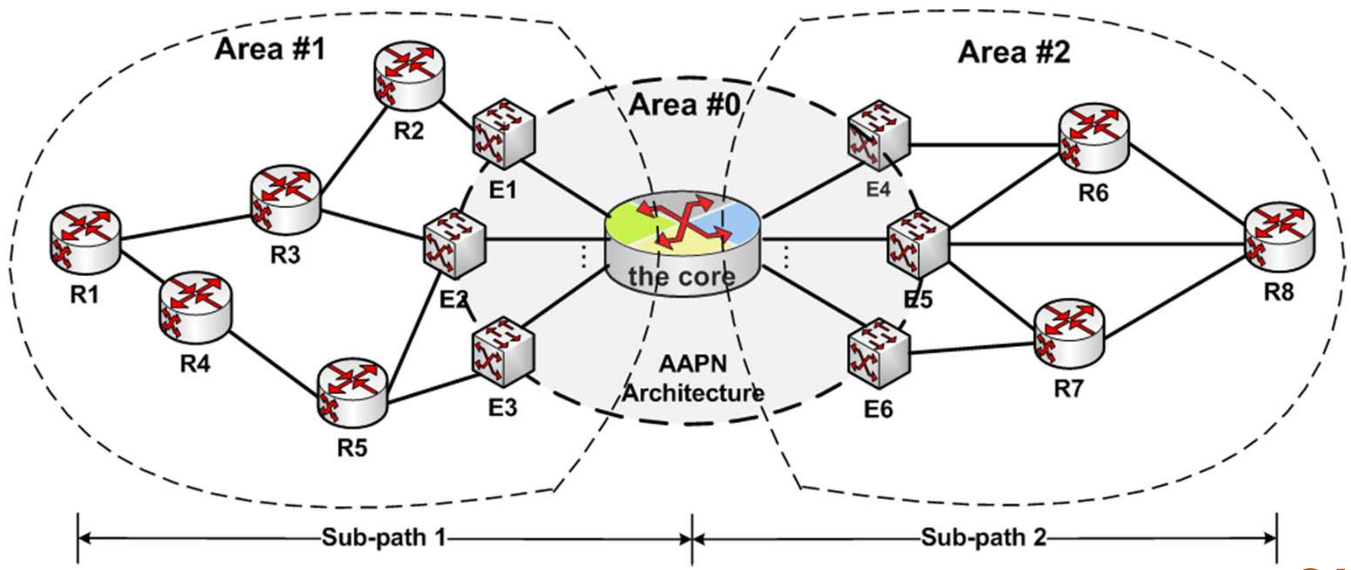

Deployment aspects - Questions

- Long-haul or Metro ?
  - connectivity "at the end of the street"; to a server farm
  - AANP as a backbone network ?
- High capacity (many wavelengths) or low capacity (single or few wavelengths) ?
- Multiple core nodes ?
  - For reliability
  - For load sharing
- Transmission infrastructure ?
  - Using dedicated fibers
  - Using wavelength channels provided by ROADM network

Conclusions

- The simple IP layer service is an advantage, although
  - there is danger of “ossification”
  - there are richer application requirements (e.g. QoS, multicasting) and a need for user-level choice
- New networking technologies introduce many changes
  - Wireless communication - ambient / sensor / ad hoc networks
  - Optical switching
- Other important issues
  - Trust and security
  - Scalable management
  - Economic issues: tariff structure, ownership