Research on energy-efficient networking in the TREND Network of Excellence

Fabio Neri (Politecnico di Torino, Italy)
neri@polito.it

iMinds 2010: Innovation, Incubation, Inspiration
16 December 2010, Ghent, Belgium

Towards Real Energy-efficient Network Design

Grant agreement n. 257740
THEME [ICT-2009.1.1]
[The Network of the Future]

The aim of the TREND Network of Excellence (started September 2010, with duration 3 years) is to foster the **integration** of the EU research community in *green networking* with a long term perspective to consolidate the European leadership in the field

http://www.fp7-trend.eu

SEVENTH FRAMEWORK PROGRAMME

Power consumption of ICT

- ICT is today responsible for 2% to 10% of the worldwide power consumption
- Some studies predict that 50% of the world electricity will be used by ICT in a few years: the new generations of Internet applications may require amounts of electricity that cannot be generated nor transported to major metropolitan areas
- This may end up limiting the penetration of broadband networking

- Energy-efficient networking has recently become a very fashionable research topic
- Too many buzzwords and too few data
- Need to understand which are the real possible gains, if any, of power awareness and power efficiency in networking

Why a Network of Excellence?

- The European technical community can claim both a strong background in networking and a high level of environmental consciousness
- There exists a significant market potential for green networking equipment
- In the above framework, a NoE is the perfect instrument for promoting awareness and contributing to build a sound knowledge base in the scientific community
- TREND is a small NoE, with strong industrial participation and commitment, and a significant commitment on (joint) research activities

Questions addressed by TREND

- *What is the real power consumption of ICT?*
- *In which sections of network infrastructures (home, access, core, data center) can energy-conscious approaches be more effective?*
- *Can optical technologies help in significantly reducing power consumption?*
- *Is the current Internet architecture, and TCP/IP in particular, energy friendly?*
- *What are the best engineering criteria and principles to actively support energy efficiency along the sequence of network design, planning, and operation?*
- *What kind of mutually beneficial incentives can be proposed to network operators, service providers, and users, in order to maximize energy efficiency?*

TREND topics

Energy efficiency in network infrastructures:

- Estimation of power consumption in ICT and collection of data
- Assessment of the fundamental energy-saving potential of network technologies and protocols
- Guidelines for policies and incentives to stimulate energy-efficiency in networks
- Redesign the home equipment for energy efficient communications
- Organizing the flying bits: saving energy on wireless access
- Power on/off strategies for energy saving and transparent connectivity
- Energy-efficient networking equipment
- Energy-efficient network design and control
- Energy-efficient service provisioning and content distribution
- Energy-efficient protection schemes
- Tools for power management and power management protocols

TREND partners

(1) POLITECNICO DI TORINO
(2) ALCATEL - LUCENT BELL LABS FRANCE
(3) HUAWEI TECHNOLOGIES DUESSELDORF GmbH
(4) TELEFONICA INVESTIGACION Y DESARROLLO SA
(5) FRANCE TELECOM SA
(6) FASTWEB SPA
(7) UNIVERSIDAD CARLOS III DE MADRID
(8) INTERDISCIPLINARY INSTITUTE FOR BROADBAND TECHNOLOGY
(9) TECHNISCHE UNIVERSITAT BERLIN
(10) ECOLE POLYTECHNIQUE FEDERALE DE LAUSANNE
(11) CONSORZIO NAZIONALE INTERUNIVERSITARIO PER LE TELECOMUNICAZIONI
(12) PANEPISTIMIO THESSALIAS (UNIVERSITY OF THESSALY)

TREND organization

- 12 partners (2 manufacturers – 3 telecom operators – 7 university groups) + 7 Collaborating Institutions (participating to research and integration activities with no budget allocation), all with significant previous experience in the technical topics
- duration: 3 years
- estimated effort: 446 person-months
- project budget: 4.4 MEuro (34% for research, 35% for integration and structuring, 21% for spreading of excellence, 8% for project management)
- EC contribution: 3.0 MEuro

TREND organization

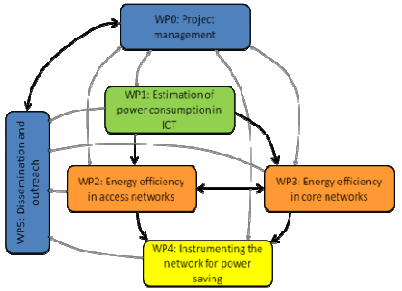

- 6 workpackages:
  - WP1: Assessment of power consumption in ICT
  - WP2: Energy efficiency in access and home networks
  - WP3: Energy efficiency in core networks
  - WP4: Instrumenting the network for power saving
  - WP5: Dissemination and outreach
  - WP6: Project organization and management
- Integration enablers:
  - Integrated Research Actions and Joint Experimental Activities
  - mobility and joint publications
  - joint education and dissemination

Part 2: An interesting question

« Is circuit switching more energy-friendly than packet switching?»

Packet versus circuit switching

- Is circuit switching to be preferred to packet switching in terms of energy efficiency?

- Interesting questions/research issues arise in the domains of:
  - Transmission
  - Switching
  - Networking

Packet vs circuit: Transmission (I)

- Packet switching requires more information overheads (packet headers) hence more power
- Continuous-mode versus burst-mode operation of transmission lines:
  - Current physical layers are mainly continuous-mode: line always active; idle bits/cells transmitted between useful information
  - Even Ethernet has evolved to continuous-mode
  - Burst-mode operation: no power between packets; new synchronization and receiver adaptation at every packet

Packet vs circuit: Transmission (II)

- Burst-mode operation can save significant power with typical link loads. But burst-mode operation:
  - Is more complex and power costly
  - Has more synchronization issues
- Hybrid continuous-/burst-mode operation is an interesting option
  - Similar to powering off idle continuous-mode lines
  - Fast resynchronization of waking-up lines is necessary

- Is burst-mode operation applicable to circuit switching? Idle detection techniques (which are application dependent) should be devised …

Packet vs circuit: Transmission (III)

- Digital operation: take decisions on every bit at the receiver (*Is this signal a '0' or a '1'?*)
- Is digital operation on every link along the source-destination path necessary?
- Analog operation (*amplify and forward the received signal waveforms*) is less power-hungry
- What about network management and security?

Packet vs circuit: Switching (I)

Packet switching not easy to define:

"*distributed, progressive, header-based resource allocation, with storing and discarding of data units*"

Current circuit switching devices consume less power than packet switching devices

Packet vs circuit: Switching (II)

- Packet switching was not conceived for energy efficiency
- Issues related to power consumption
  - Large buffers are needed to solve contentions in a distributed fashion
  - Header processing and routing/forwarding operations are needlessly repeated on all packets belonging to the same flow
  - The switching fabrics must be reconfigured at rates that depend on line rates and on packet sizes, not on Quality of Service requirements

Packet vs circuit: Switching (III)

- TCP/IP was not conceived for energy efficiency
- Issues related to power consumption
  - The stateless IP paradigm asks for independency between datagrams belonging to the same flow, hence causes needlessly repeated switching operations
  - TCP generates acknowledgement packets
- But TCP/IP is very difficult to modify …

Packet vs circuit: Switching (IV)

- Other interesting questions for pkt switching:
  - Packet sizes:
    - Larger packets require smaller switching rates (less power) but larger buffers (more power)
  - Synchronous vs. asynchronous switch operation:
    - Segmentation/reassembly of variable-size packets into fixed-size switched data units has power costs
    - Independent clock domains for different subsystems in switching devices require more adaptation buffers but less energy costs for clock distribution
- Packet switching easily permits dynamic voltage (hence power?) scaling

Packet vs circuit: Networking (I)

- Centralized resource allocation schemes and traffic schedulers avoid distributed contention resolution:
  - Less need for buffering in switches
  - Less intelligence in switches
  - More signaling traffic
  - Does it scale?
- Fast circuit switching (e.g., one circuit per web page) appears as a good solution
- Will the Internet become the control plane of an analog circuit-switched WDM network?

Packet vs circuit: Networking (II)

- Switching vs multiplexing: given the bandwidth abundance of optical fibers, and the flexibility permitted by wavelength agility, contention resolution and traffic grooming can be avoided in most switching devices
- For example: two packets/flows competing for the same link are immediately forwarded on different (wavelength) circuits, and not buffered and transmitted at different times
- Extreme case: one end-to-end lightpath (i.e., one optical circuit) for every information flow ➔ no packet switching

Packet vs circuit: Networking (III)

- Buffering vs discarding packets: retransmissions have large power costs, and require extra buffering at edge nodes
- Should we look for lossless packet transport schemes (e.g., flow controlled FiberChannel)?

- Shared media are not energy-efficient:
  - useless load on receivers
  - coordination overheads (e.g., packet collisions, or token circulation)

Packet versus circuit switching: Summary

- Interesting questions/research issues arise in transmission, switching and networking domains
- Do we still have to design networks for optimal bandwidth utilization, or rather minimize the amount of switching/power?
- Are we stepping back from packet to (some) circuit switching, and from digital to (some) analog operation?

Part 3: A specific research activity …

- Optical fibers offer bandwidth abundance
- Packet switching permits to efficiently use bandwidth by efficiently grooming information
- Can we trade bandwidth efficiency for energy?
- How much multiplexing and how much switching?

A. Bianco, E. Bonetto, D. Cuda, G.A. Gavilanes Castillo, M. Mellia, F. Neri, "Energy-Efficient Design of Wavelength-Routing Networks," 35th European Conference on Optical Communication (ECOC 2010), Torino, Italy, Sept. 2009

WDM networks

- Wavelength Division Multiplexing (WDM) networks transparently carry several information-bearing signals in optical fibers
- Wavelength Routing (WR) networks (also called Wavelength Switched Optical Networks - WSON)
  - Are the most used paradigm in WDM networks
  - Lightpaths (optical circuits) are routed over WDM fiber links thanks to (transparent, but slow) optical cross-connects, which operate in the optical domain
  - Optical cross-connects are less energy hungry than packet switches and routers
  - A sequence of one or more lightpaths is traversed from source to destination; electronic switching/grooming is performed at nodes terminating lightpaths

WR network design

- The design of Wavelength Routed networks is usually separated into two sub-problems:
  - Logical Topology Design (LTD)
  - Routing and Wavelength Assignment (RWA)

Given a traffic matrix, find a set of lightpaths (optical circuits) satisfying the traffic requests and optimizing a given target function (regardless of the physical topology)

PA-Logical Topology Design

- LTD: given the traffic matrix, find the best (in terms of cost/performance) lightpath set
- Power-Aware LTD (PA-LTD) target:
  - find the best trade-off between the utilization of electronic (for switching) and optical resources (for transmission)
  - Minimizing the total power consumption

Power due to optical transmission

Power due to electronic switching

- Power consumption contributions (for a given bitrate Rb)
  - P^{SW}: electronic switches for grooming, add and drop traffic
  - P^{TX} : optical transceivers for electro-optical conversion and transmission

Logical topologies

- When $P^{TX} \approx P^{SW}$, PA-LTD (OPEX based) and traditional (CAPEX based) LTD find similar solutions, since they simply minimize the number of transceivers
- For $P^{TX} < P^{SW}$, the best topologies have more transmission and less switching, i.e., a larger number of lightpaths and of optical transceivers
- PA-LTD looks for the best compromise between the number of transceivers and the amount of traffic switched electronically (or between the amount of multiplexing and the amount of switching)

Our PA-LTD

Each node is equipped with an electronic switch and a set of optical transceivers

- Transceivers (lightpaths) have a fixed capacity (10 Gbps) and consume a fixed nominal power (P^{TX} = 8W)
- Traffic is switched electronically at the edges of a lightpath, and the power cost of switching depends linearly on the amount of traffic

The parameter ν^O ("efficiency of optics") equals the ratio between optical transmission power and electronic switching power (at the same bitrate)

$$P^{SW} = \nu^O P^{TX}$$

Our PA-LTD

- PA-LTD was formulated as a Mixed Integer Linear Programming (MILP) problem
- We solved it using the CPLEX optimization environment
- We compared PA-LTD against an LTD model with CAPEX minimization as a target (CA-LTD)
- The impact of PA-LTD over CAPEX and OPEX was analyzed
- We characterized the logical topologies obtained for different power consumption scenarios of electronic switching vs. optical transmission

Transmission vs. switching

Depending on the power consumption of switching with respect to transmission, the logical topology can be partially connected to minimize the number of transceivers (equivalent to cost-optimized) …

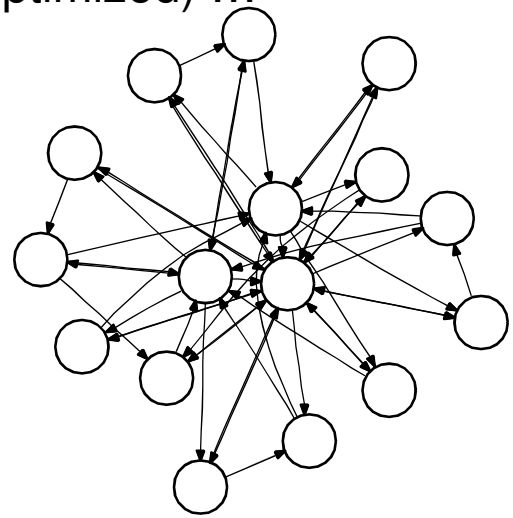

Partial mesh:
forward traffic
on most of the nodes

But for less power-efficient switches, more links become convenient (cost- and power- optimized solutions differ)

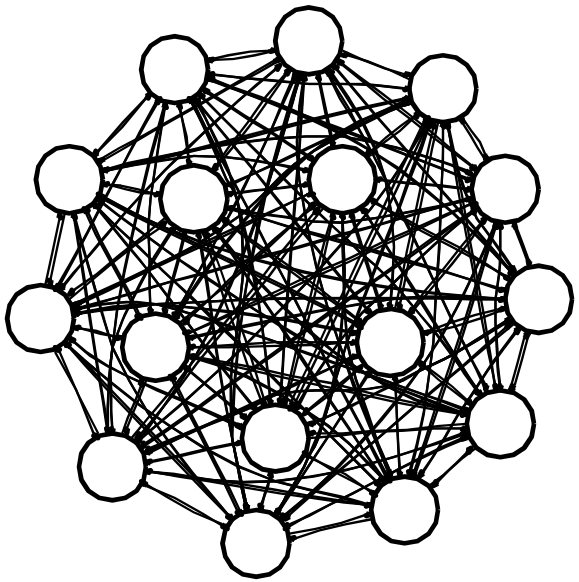

Full mesh:
no ftraffic
forwarding

Test scenarios

- We considered uniform traffic matrices with different traffic loads: $\lambda_{0.1}$, $\lambda_{0.2}$, and $\lambda_{0.5}$, corresponding to 1, 2 and 5 Gb/s per node
- The optics efficiency (ν^O) can change as new technologies become available

+Laser
+Modulator
+Driver
+Control
+Receiver
+CDR

Generic optical transceiver ~8Watt @ 10 Gb/s

Commercial tunable optical transceiver

- **CAPEX**
  - In this first study, only the cost of all transceivers is considered
  - One transceiver at 10Gbps costs ~100 $
- **OPEX**
  - Energy consumption is considered as the main contribution
  - Energy cost is set at 0.2 $/kWh

In a low-traffic scenario

- All the traffic from every node can be transmitted by a single lightpath …

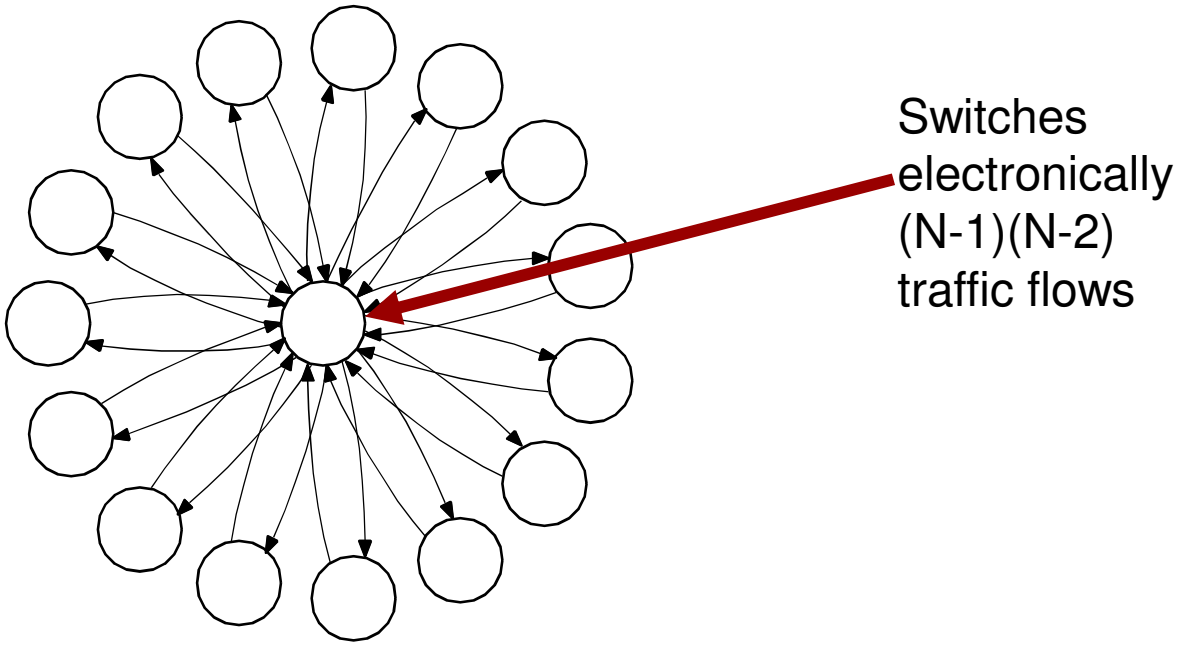

... and electronic switching is concentrated in a single node

For a network with 16 nodes

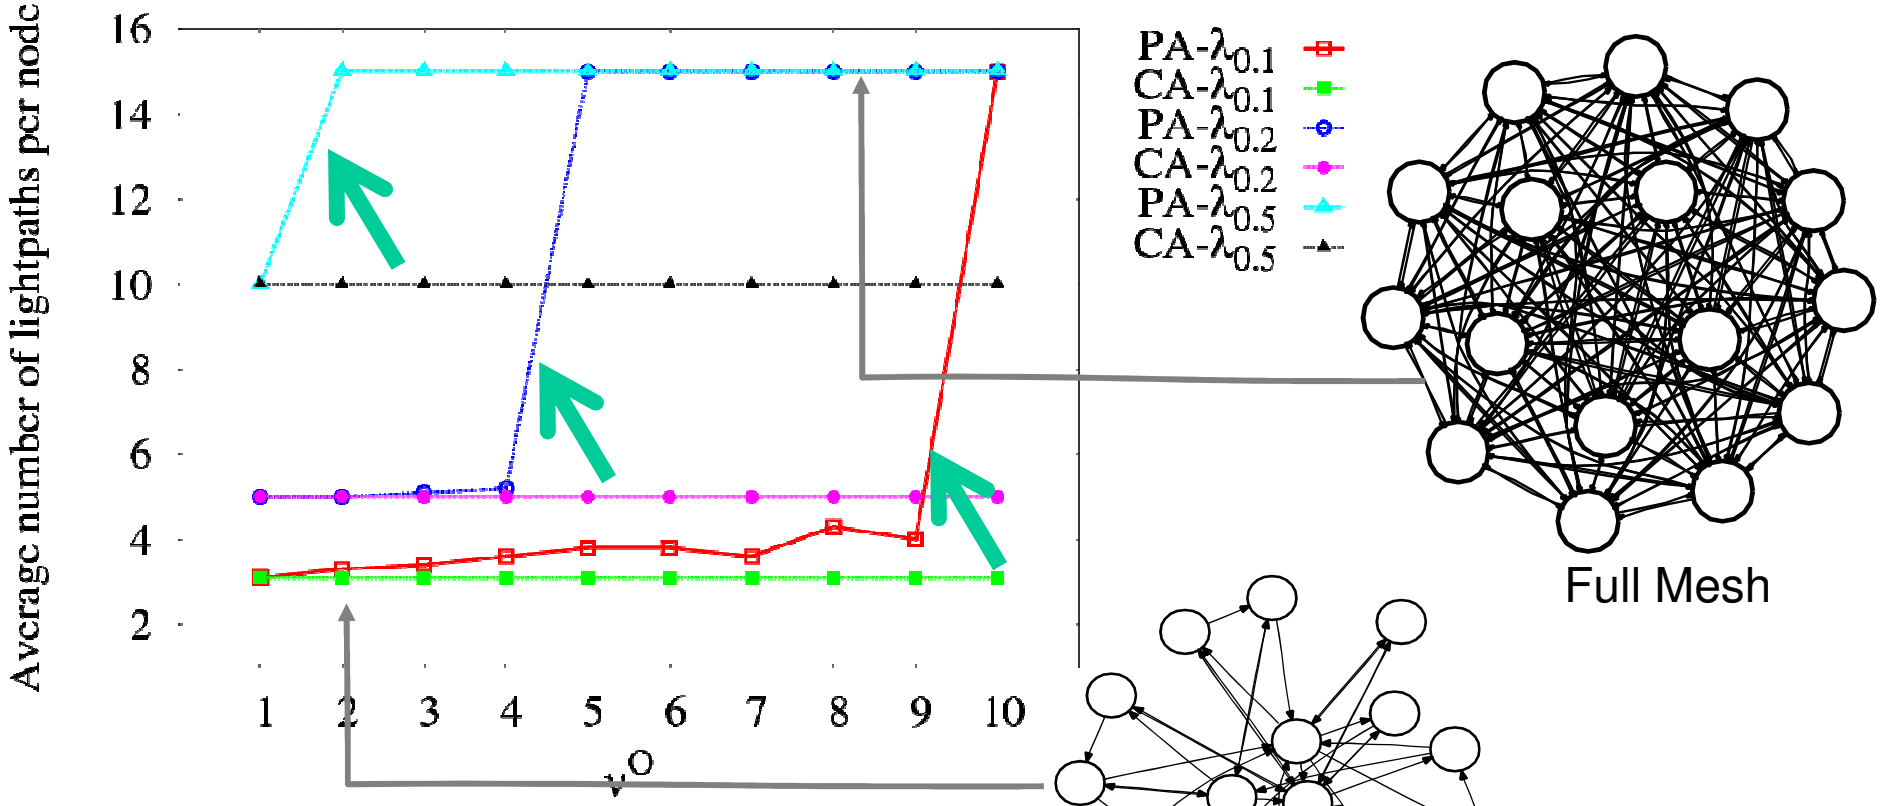

Following the traffic requirements, the logical topologies become more connected

Power advantage of PA-LTD

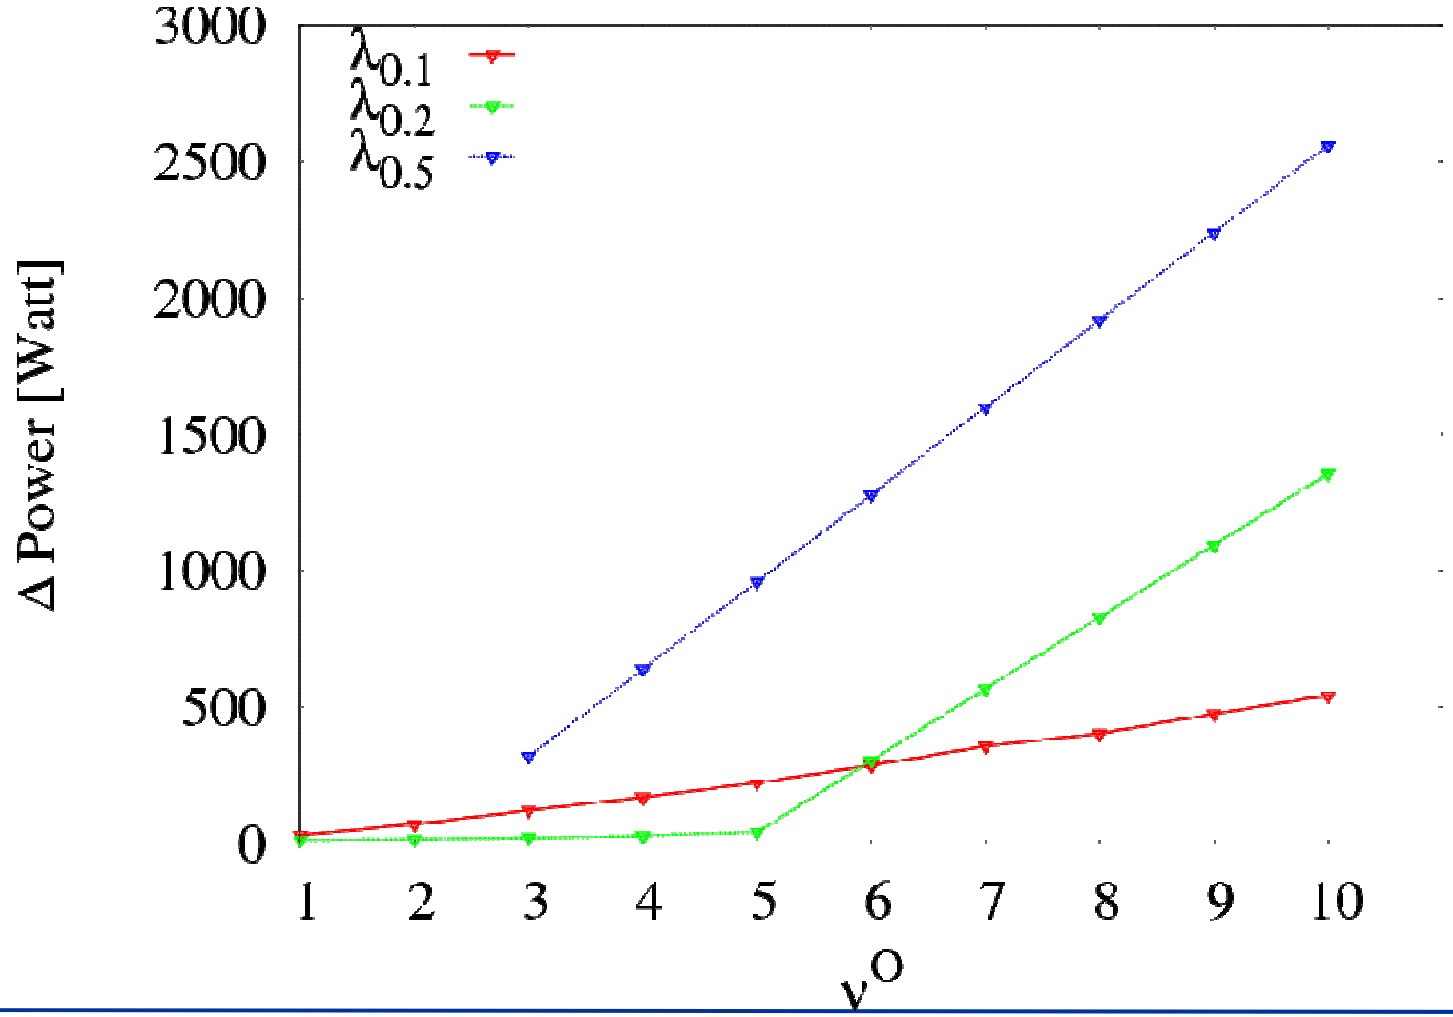

Economical analysis

- For $P^{TX} < P^{SW}$, lower power consumption means higher costs (due to the larger number of transceivers), but also money savings by energy-efficiency
- PA-LTD finds results having usually larger CAPEXs than CA-LTD results
- It is possible to recover the difference due to a smaller OPEX
- Break-Even Point (BEP):
  - Time required to start having an economical advantage for a PA-LTD solution respect to the CA-LTD one

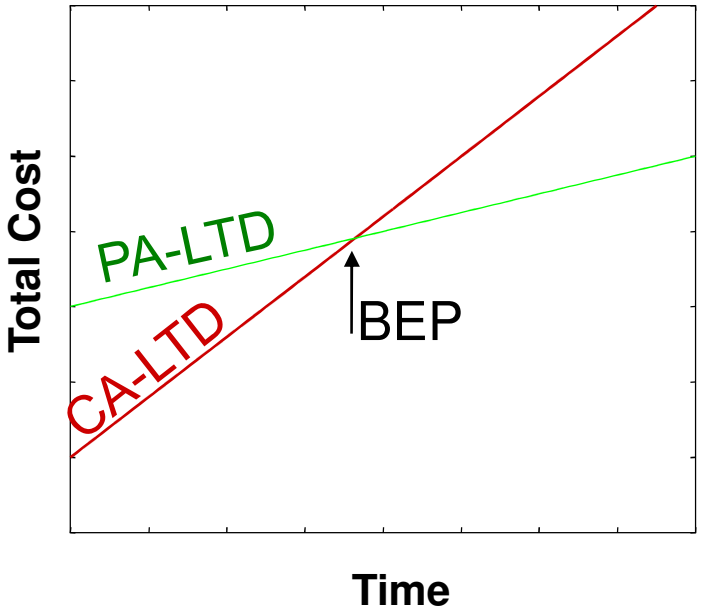

Break Even Point

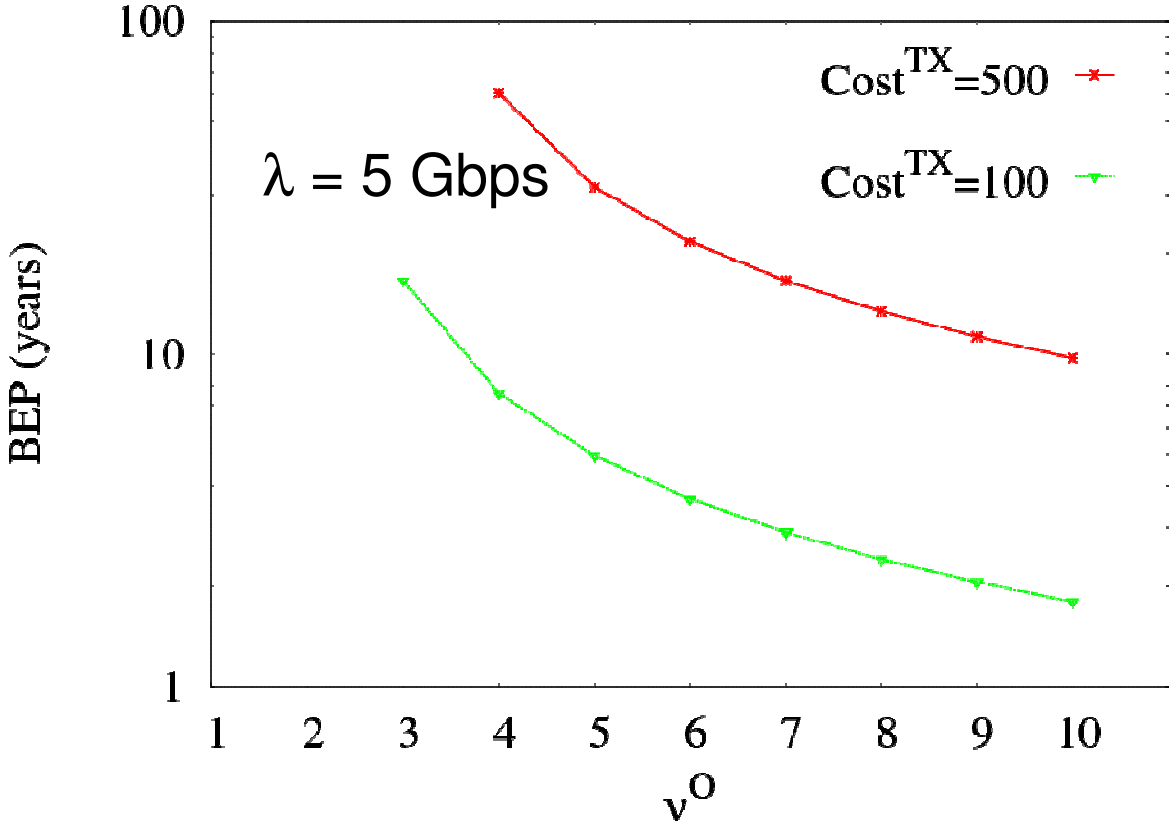

As transceiver's cost reduces, the BEP time becomes shorter

Conclusions

- Power saving can be a driver for a deeper penetration of optical technologies in networking beyond point-to-point transmission
- Results suggest that a “more multiplexing, less switching” (or a “more circuit switching, less packet switching”) evolution of network architectures may be convenient
- Is circuit switching coming back in the network core?

Towards Real Energy-efficient Network Design

Grant agreement n. 257740
THEME [ICT-2009.1.1]
[The Network of the Future]

More information coming soon at:

http://www.fp7-trend.eu

Contact:

fabio.neri@polito.it

SEVENTH FRAMEWORK PROGRAMME