CTVR / the telecommunications research centre

TRINITY COLLEGE DUBLIN

Tackling the capacity crunch

An attempt to rationalisation and mitigation

May 14th 2015, Royal Society: Communications networks beyond the capacity crunch

Marco Ruffini, CONNECT reasearch centre, Trinity College Dublin, Ireland
marco.ruffini@tcd.ie

DISCUS

Acknowledgments

- SFI-funded CTVR telecommunications research centre and EU-funded DISCUS FP7 project.
- David B. Payne (Trinity College Dublin)
- Emanuele Di Pascale (past PhD Student at Trinity College)

CTVR | the telecommunications research centre

DISCUS

sfi Science Foundation Ireland

TRINITY COLLEGE DUBLIN

SEVENTH FRAMEWORK PROGRAMME

Summary

- Past and current data on capacity crunch
- Capacity crunch rationalisation: where will happen and when will happen?
  - Access network analysis
  - Backhaul network analysis
- Mitigation possibilities:
  - Architectural innovations
  - Bypass backhaul through edge caching

THE CAPACITY CRUNCH...

THE CAPACITY CRUNCH… RELOADED

Business Week, October 8th 2000:

At the Speed of Light

…Internet traffic is doubling every three months, and optical technology is the only practical way to carry it all..

http://www.businessweek.com/printer/articles/78818-at-the-speed-of-light?type=old_article

Reed Hundt, Federal Communications Commission chairman 1993-1997. Book written in 2000 titled:

You Say You Want a Revolution

…[i]n 1999, data traffic was doubling every 90 days …

http://www.amazon.com/You-Say-Want-Revolution-Information/dp/0300181930

Keith Mitchell, executive chairman of London Internet Exchange.

March 2000

…[LINX] traffic doubles every hundred days or so…

- LYNX own data contradicted this: Mar'99-Mar'00

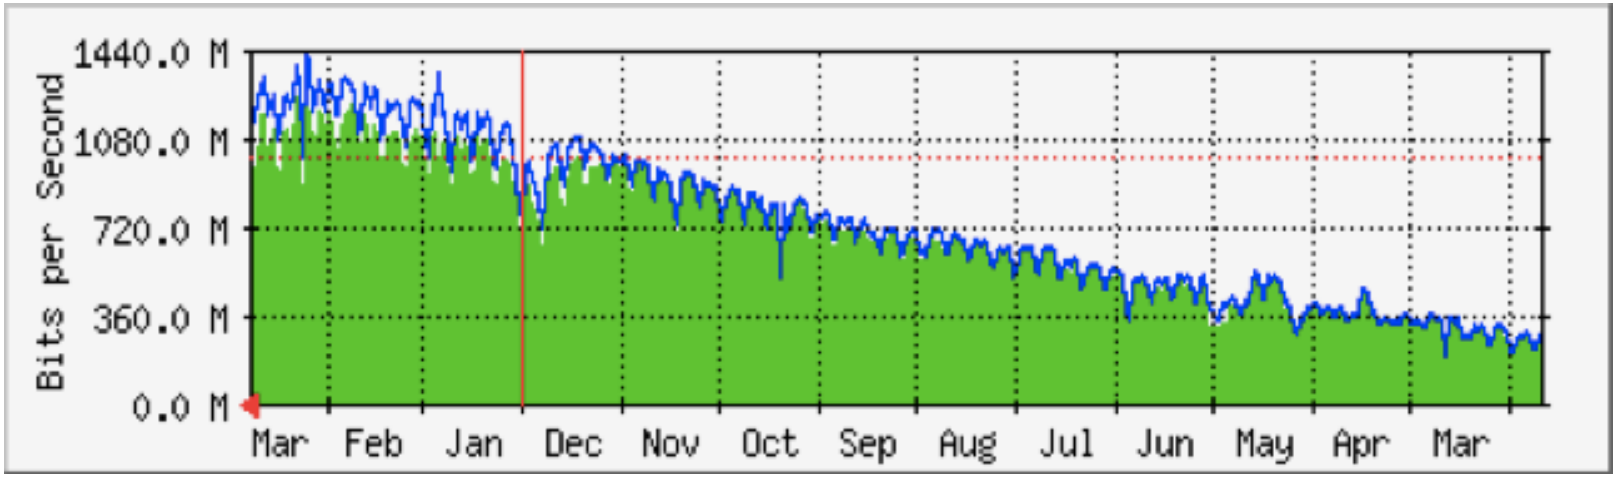

- For more information: Andrew M. Odlyzko "Internet traffic Growth: Sources and implications"

http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.81.5738&rep=rep1&type=pdf

What's different from today?

- At that time traffic was growing fast, with correct estimations of about 100% yearly growth.

- Today (based on CISCO's VNI) 2013-2018
  - West Europe: Internet yearly growth rate of 20%
  - USA: Internet yearly growth rate of 23%
  - Highest is middle east and Africa with 41%, but overall traffic still 5 times less than USA

- So measured growth (except in BT ;-) is indeed moderate..

What happens after 2018?

TRINITY COLLEGE DUBLIN

- Forecast based on past extrapolation don't really work past 3 years

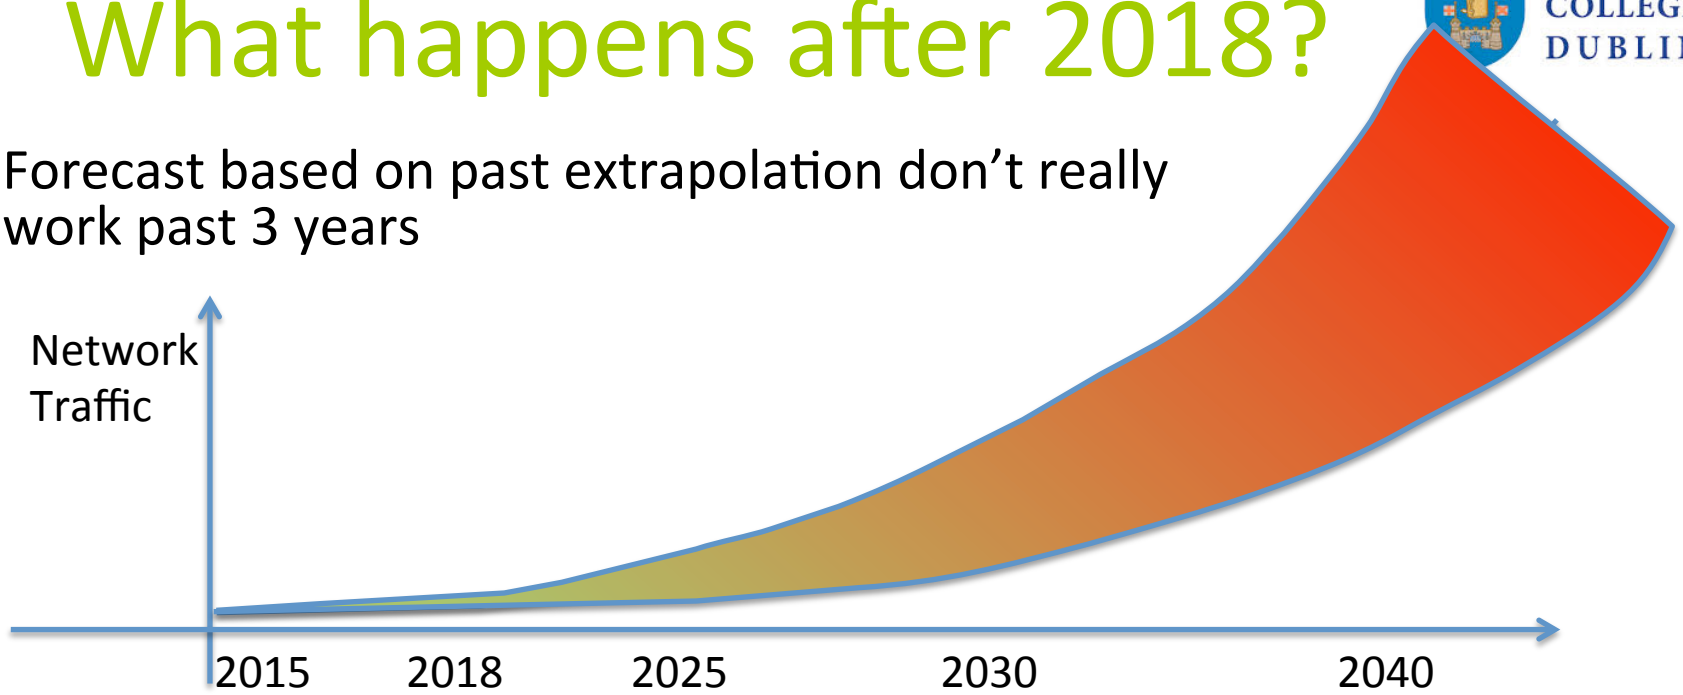

- So we can start fantasizing of completely new applications which will require by 2025 much larger capacity, we'll have multiple 8K pixels screens, IoT will require unimaginable amount of communication between machines…

- However reality might proceed much slower
  - Example: HD TVs have been around for over 15 years, but very few channels are today in 1080p.. And that's still highly compressed too..

Capacity crunch rationalisation

- We developed a tool that considers statistical user behaviour and potential applications used.
  - Define a number of available applications (general categories)
  - Define a type of user (e.g., moderate residential, high residential, small business)
  - Build scenarios linking user types with applications and daily usage (statistical distributions for: number of application sessions, session duration, session start time,…)

Tool interface

- For each application and user type, identify user behaviour through beta distributions:
  - number of daily instances,
  - session duration,
  - session start time
  - Time interval between sessions

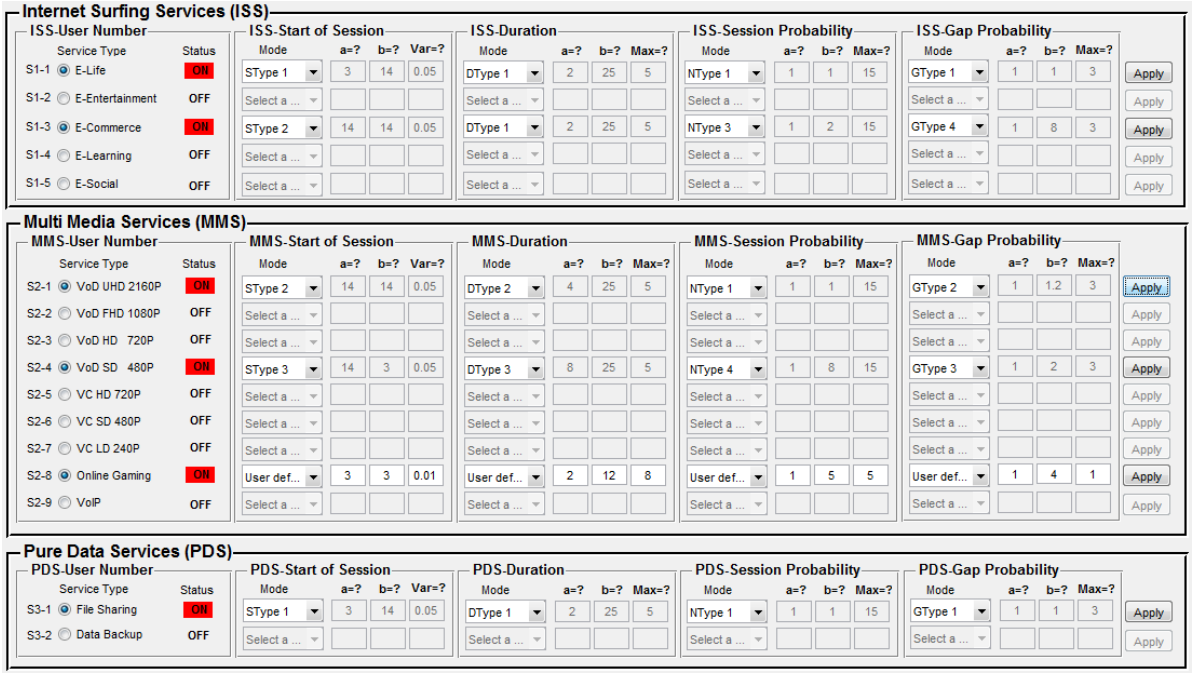

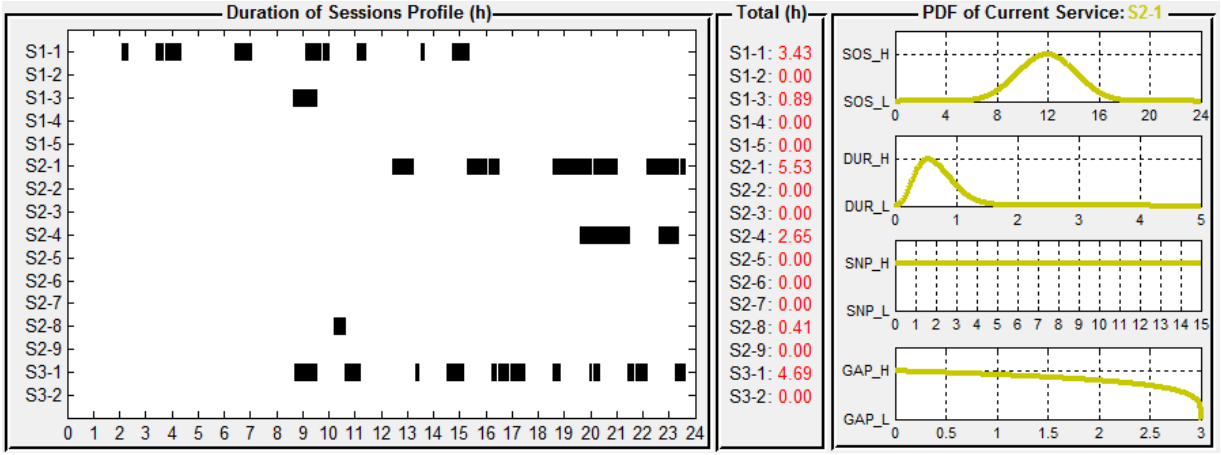

- Example of Beta distribution for one application (right hand side)
- Example of daily instances for all active applications (left hand side)

Tool operations

TRINITY COLLEGE DUBLIN

- For core traffic:
  - Give number of nodes and population per node
  - Give number and "location" of Data Centres and Internet Exchanges
  - Calculate demand matrix
  - Add separately inter-DC traffic and leased lines

Obtain a set of matrices giving node to node demand every 5-minute interval for one day

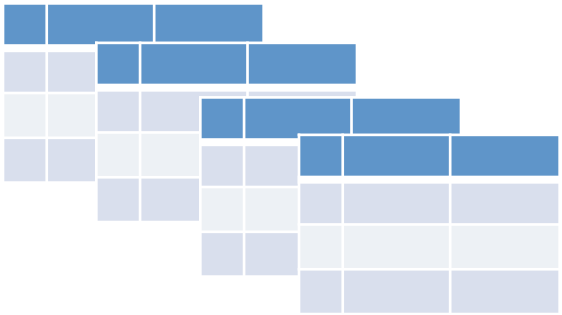

- For access traffic:
  - examine traffic at aggregated PON level (using 512-way split)

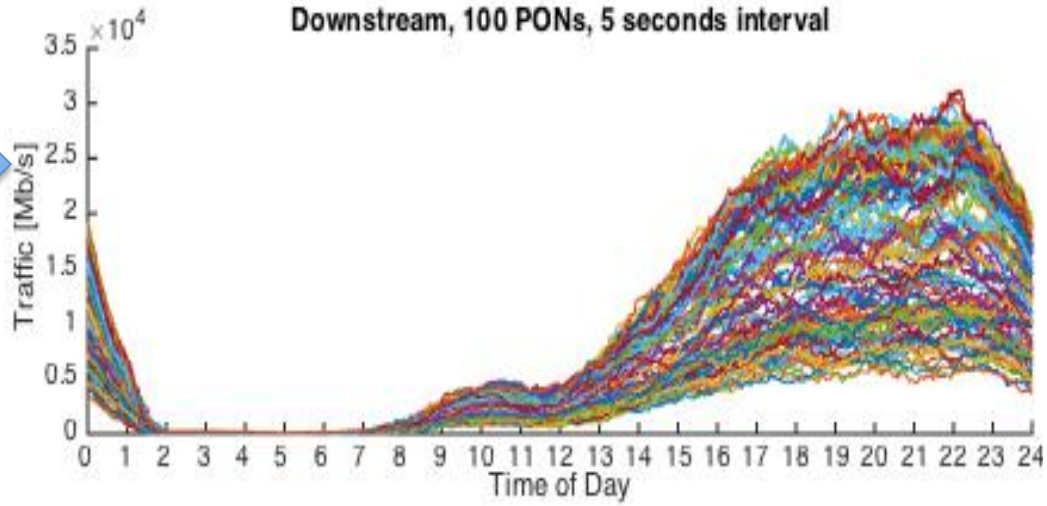

Access network requirements

- The tool becomes particularly interesting for evaluating access networks requirements and compare different technologies: FTTCab, Vectored-FTTCab, GPON, NG-PON2, LR-PON, G.FAST, XG.FAST…
  - Differences are considered in terms of peak rate, split factor and aggregate rate.

- A modification was carried out so that the tool includes streaming type of applications (relatively constant rate) and file download type (variable rate with bursting to max available capacity)

Disclaimer

- **The tools does not predict the future!**
- It cannot predict new applications…
- … but it allows us to test **what if** scenarios

- What happens if we reduce compression ratio of videos, e.g. 4K, 3D high fidelity at over 100Mb/s
- What if we include e-health applications
- What if we have multiple parallel video streams..

- In addition it operates over statistical variables to include randomness of user behaviour.

Test scenarios

- We assumed three scenarios:
  - Scenario 1: (Cisco ~2018) sustained busy peak 1.3Mb/s, daily usage of 4.5 GB
    + download file size distribution 100-1000MB
  - Scenario 2: sustained busy peak 7.5Mb/s, daily usage of 38 GB
    + download file size distribution 500-5000MB
  - Scenario 3: sustained busy peak 58Mb/s, daily usage of 275 GB
    + download file size distribution 3500-35000MB

- When will they occur?

Growth Rate	Scenario 2 timeframe	Scenario 3 timeframe
20%	2028	2039
40%	2023	2029
60%	2021	2026

Technologies under scrutiny

Technology	Peak Rate distribution	Aggregated capacity	Split ratio
FTTCab	Uniform:20-70Mb/s	1 – 10 Gb/s	512
Vectored-FTTCab	Uniform: 100-150Mb/s	1 – 10 Gb/s	128
GPON	2.5Gb/s	2.5 Gb/s	32-64
LR-PON	10 Gb/s	10 Gb/s per wavelength	512
NG-PON2	10 Gb/s	10 Gb/s per wavelength	128
G.FAST	Uniform: 300-700 Mb/s	GPON and NG-PON2 backhaul	64-128
XG.FAST	Uniform: 3-10Gb/s	10 Gb/s per wavelength (NG-PON2)	128

Is there a real capacity crunch?

What some operators think:

"Today's GPON will allow satisfactory capacity for the foreseeable future…

… after that (unforeseeable future) there's already a standard ready to deliver over an order of magnitude more capacity…"

What our result show:

- Scenario 1 – short term:
  - All technologies generally work, FTTCab needs a 10Gb/s backhaul for 500 users, and in average has download times over an order of magnitude slower than PONs
- Scenario 2 – medium to medium-long term
  - FTTCab not effective unless vectored, and two orders of magnitude slower than 10G PONs.
  - G.FAST Ok (it is assumed G.FAST is backhauled by GPON)
  - All other technologies OK, LR-PON requires second wavelength.
- Scenario 3 – long to very long term
  - FTTCab not effective, G.FAST works but slower compared to PONs
  - LR-PON requiring 8+ wavelengths
  - NG-PON2 requiring 2+ wavelengths
  - XG.FAST works well (it is assumed XG.FAST is backhauled by NG-PON2)

Growth Rate	Scenario 2 timeframe	Scenario 3 timeframe
20%	2028	2039
40%	2023	2029
60%	2021	2026

Summary of results

Technology	Short term	Medium to medium-long term		Long to very long term
FTTCab				
Vectored-FTTCab				
GPON		64-split	32-split	
LR-PON		Requires second wavelength		Using 8+ wavelengths
NG-PON2		Considering 1 wavelength		2+ wavelengths needed
G.FAST				NG-PON2 backhaul
XG.FAST				

Legend

- Unsuitable
- OK but sensibly slower
- Suitable
- Much spare capacity

- FTTCab has limited peak rate which eventually will also limit sustained rate.
- G.FAST OK up to medium to medium-long term
- GPON OK up to medium to medium-long term
- NG-PON not using full capacity even in very long term
- **Suggests that split ratios higher than 128 would be required for a more efficient aggregation.**

- WDM-PON (1G) not considered but would prove over-dimensioned and yet slower than PON systems

<u>But in general access technology good enough to avoid capacity crunch in (fixed) access</u>

But the network is end-to-end…
… so what about the backhaul+core?

TRINITY COLLEGE DUBLIN

- Optical access creates an unprecedented change in the network: same technology in access and core!

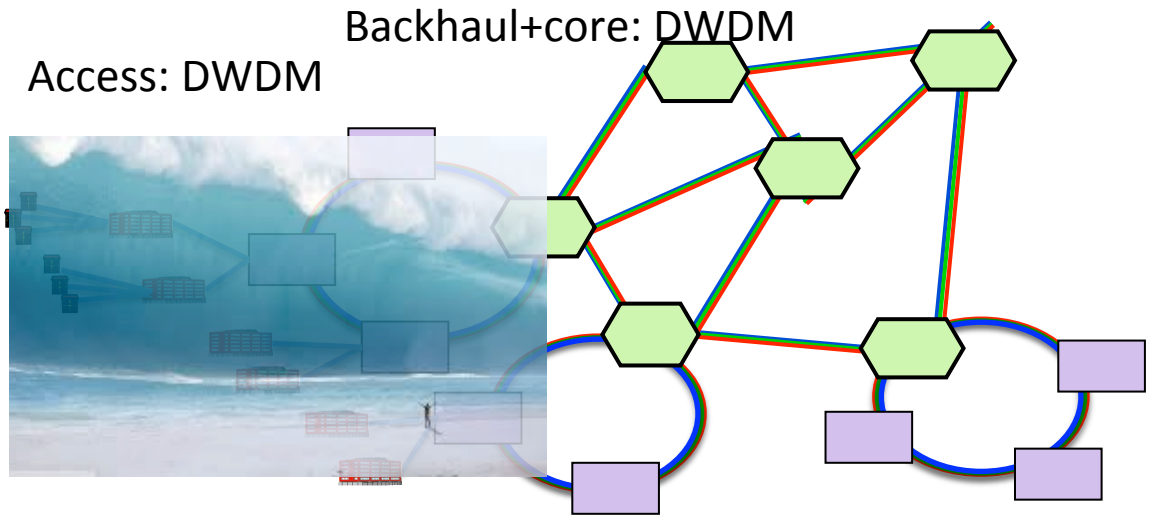

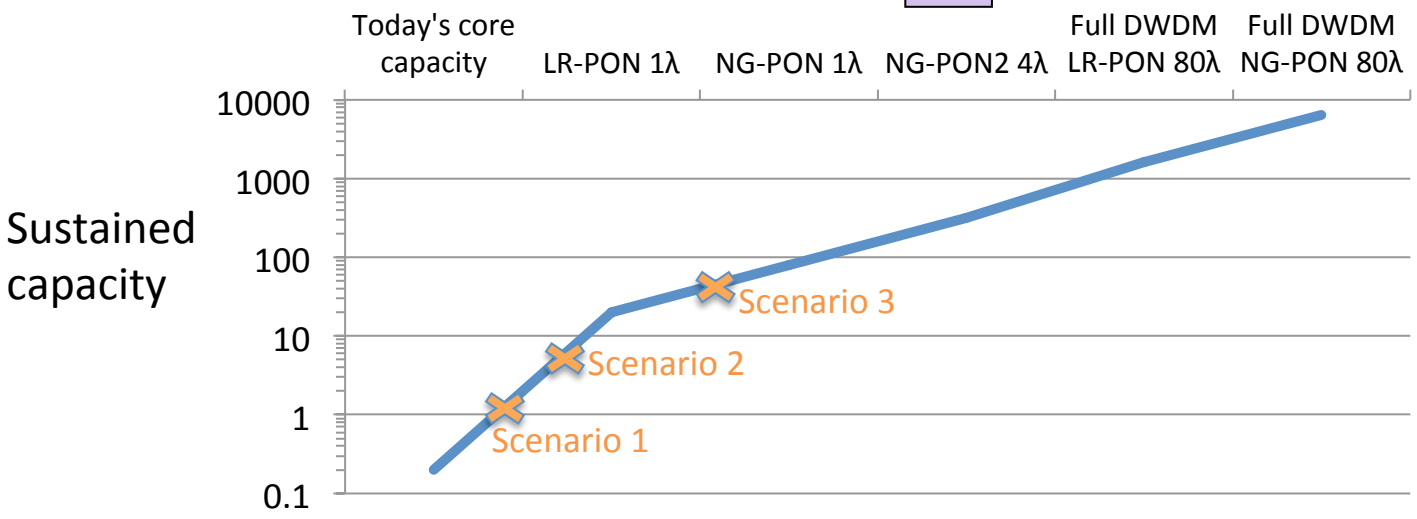

Core capacity crunch mitigation

- How can the core sustain 2,3,4 orders of magnitude increase in access capacity?

Maybe it can't*, but mitigation is possible:

1. Access will always need to operate considerable capacity aggregation (statistical multiplexing)
2. Architectural innovations
3. Bypass core through edge caching

*It might be achievable technically, but not sustainable economically

Architectural innovations

TRINITY COLLEGE DUBLIN

- In DISCUS we are studying effectiveness of Long-Reach access and flat optical core
  - Electronics do represent the bottleneck in networks, so reduce OEO conversions through:
    - Long-Reach access (consolidates local exchanges by a factor of ~50)
    - Flat optical core (reduces number of electronic hops)

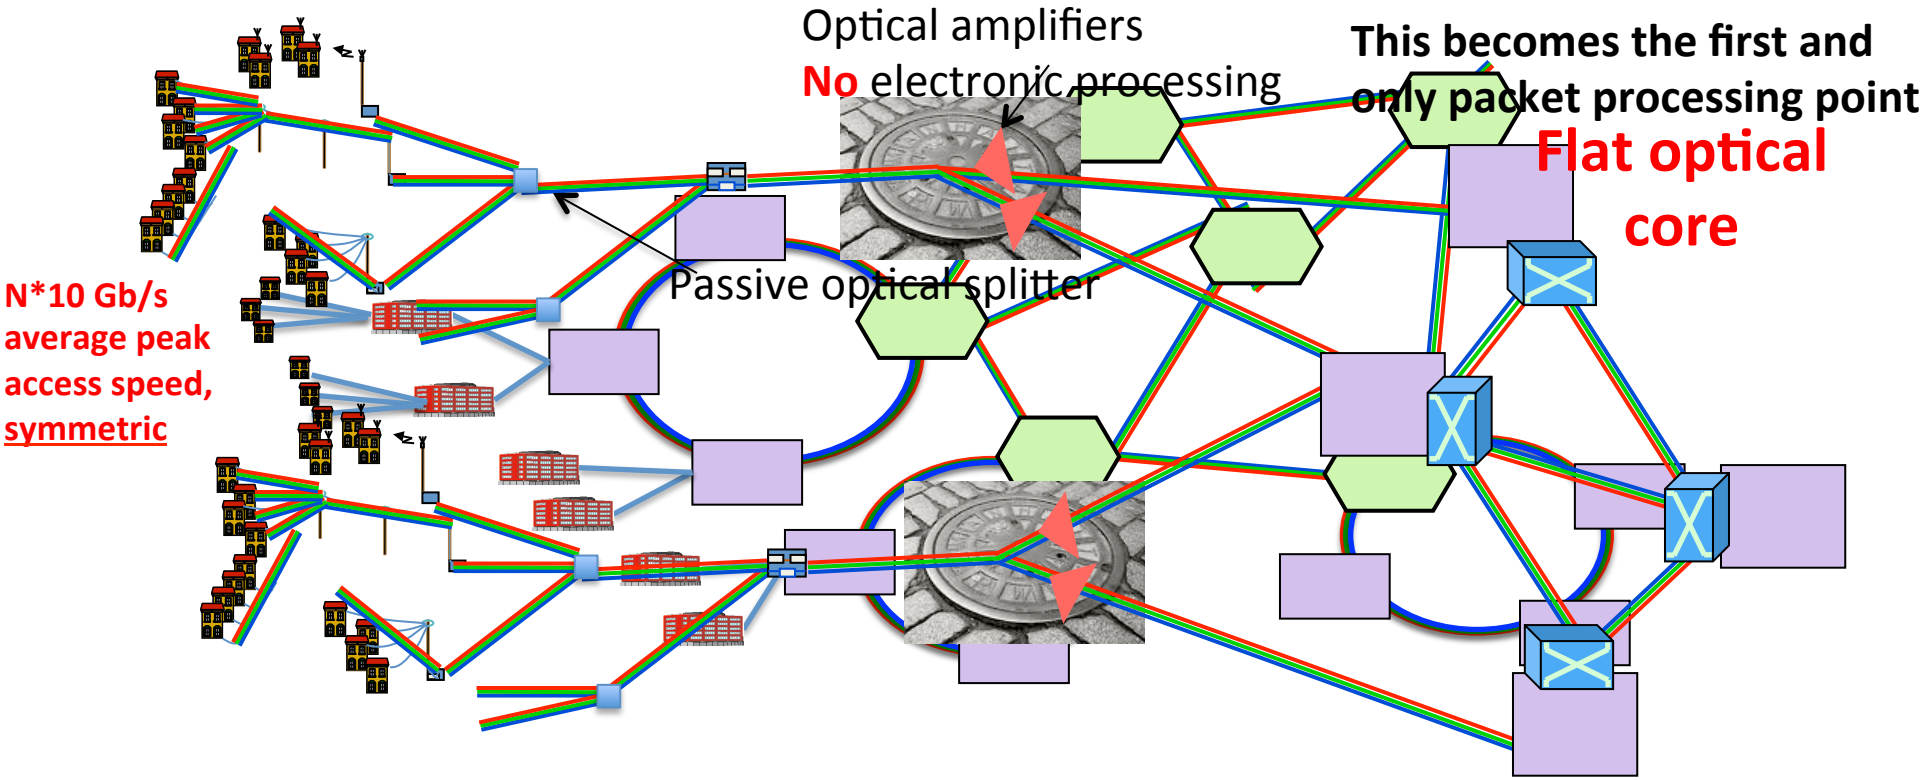

Metro-Core node

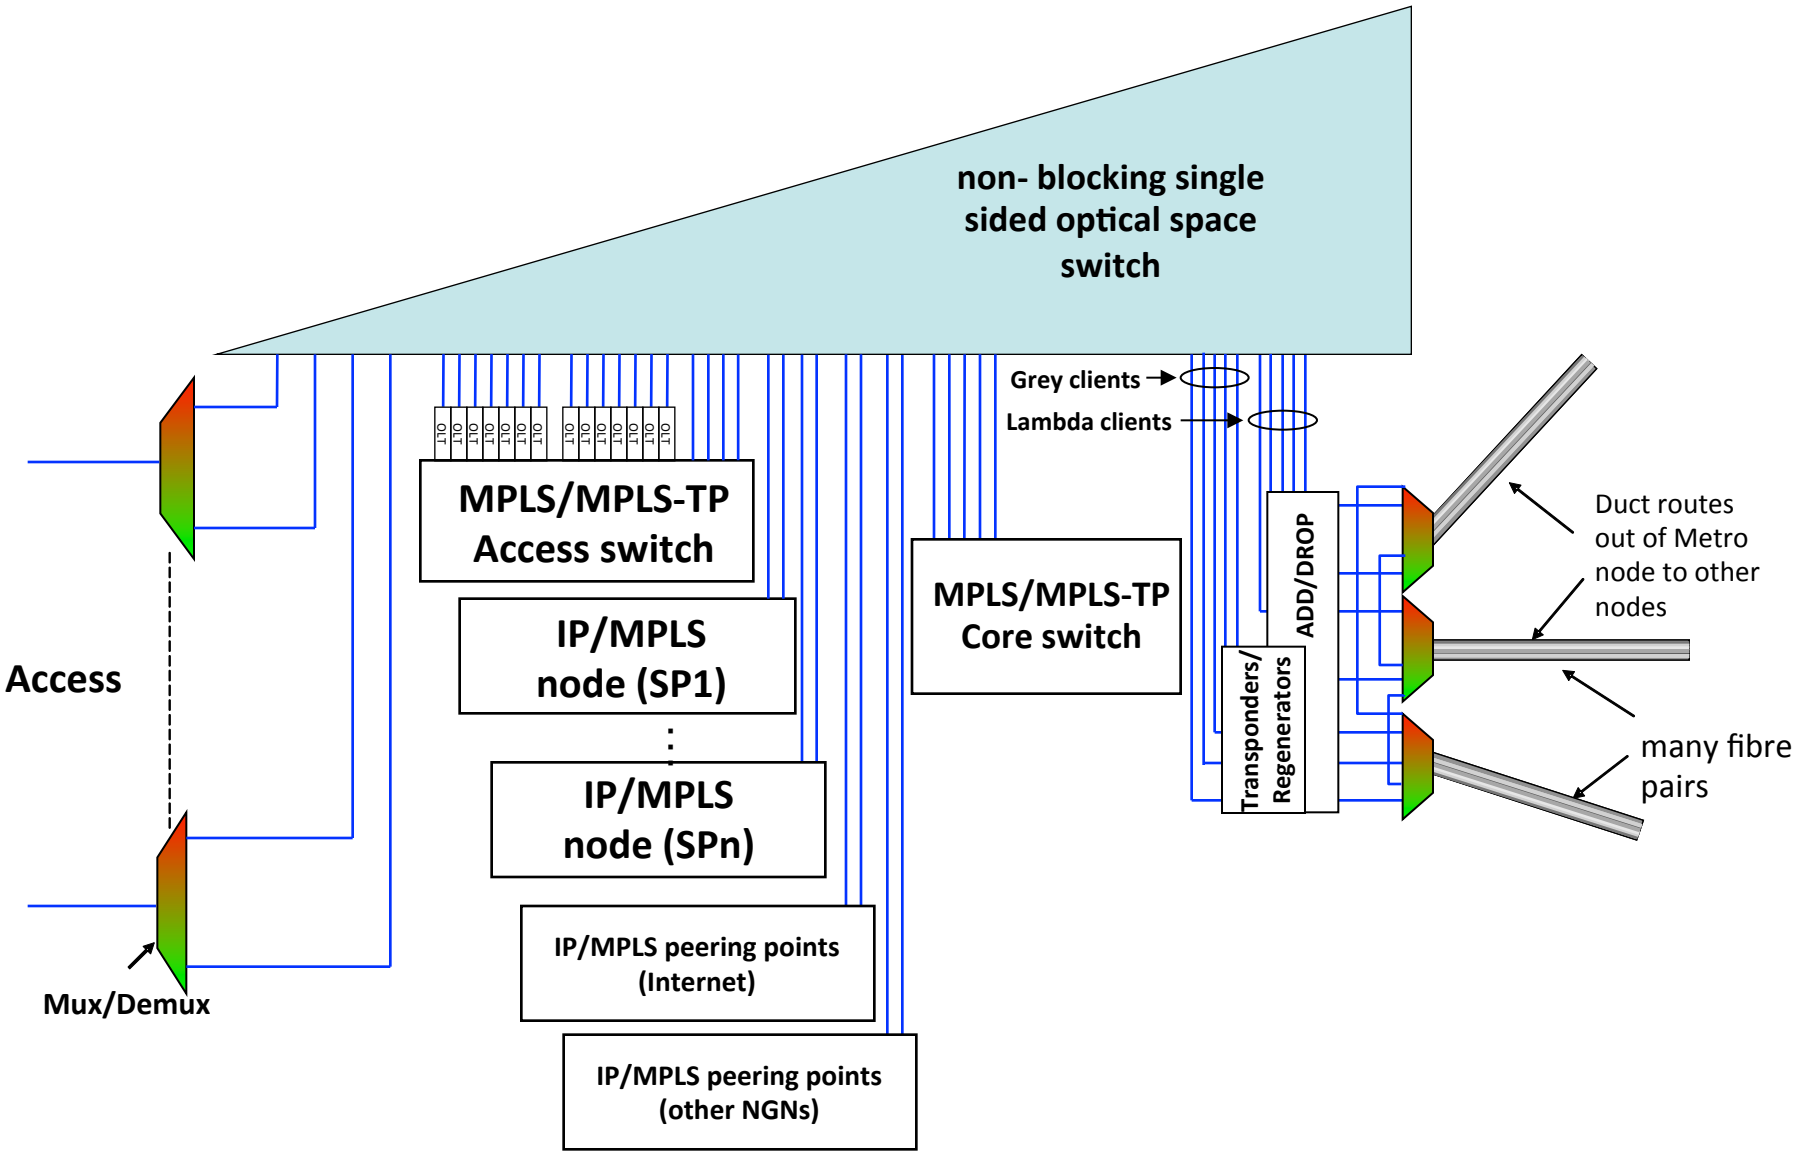

Flat vs. hierarchical core comparison

- Model for comparison of flat (optical island) vs. 2-hierarchy core network
  - Only considers number/cost of core ports
  - Assumes 300K users per node
- Current limitations:
  - Only one rate in the network and user uniformly distributed
  - Cost only based on port numbers
- Model based on real population data coming soon

Port speed	Cost (£)
10G port	4300
40G port	12500
100G port	25000
400G port	72000

Cost of core ports per customer

Minimum cost curves:
Hierarchical core ports from 100 to 400G
Flat core from 10 to 400G

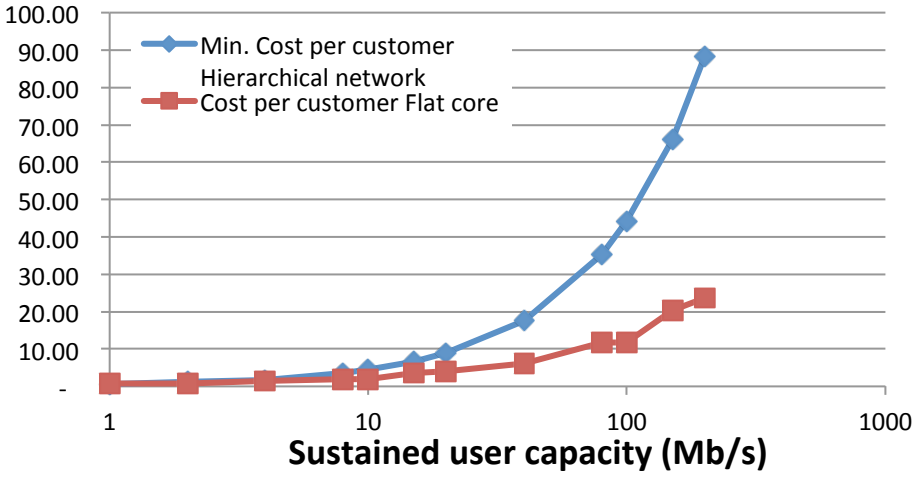

Both hierarchical and flat core ports at 100G

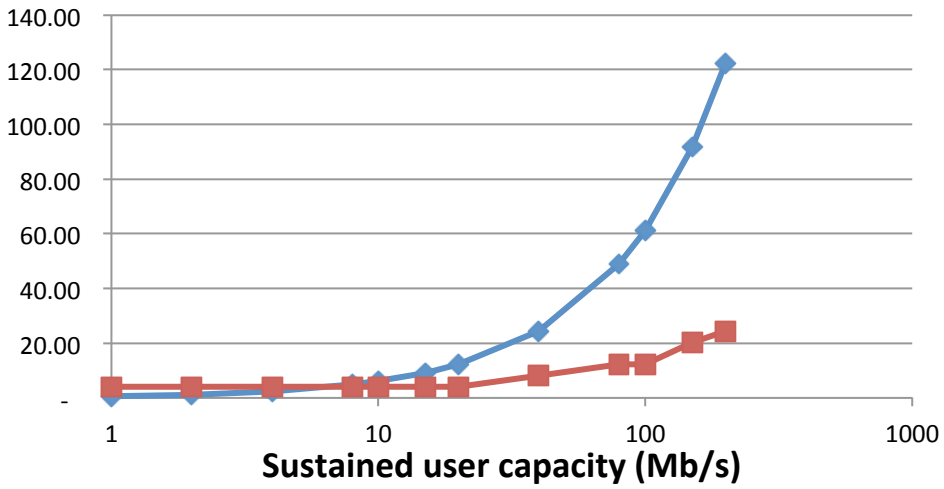

Bypass core when possible

TRINITY COLLEGE DUBLIN

- We studied the possibility of **network-managed** peer-to-peer (e.g., managed by a network operator):
  - Data stored in user set top box and used to distribute content
- Application for PON networks:
  - Trade-off large available upstream capacity to reduce use of core capacity
- Focus on Video on demand and time-shifted TV
- Simulation results based on 7 million users, using content catalog of 60K elements, on LR-PON network.
  - Realistic Zipf-Mandelbrot distribution for content popularity

Core traffic reduction

- Peer2peer can reduce core traffic by up to 99%, compared to Centralised Server (CS) solutions

E. Di Pascale, D. B. Payne, L. Wosinska and M. Ruffini. Locality-Aware P2P Multimedia Delivery over Next-Generation Optical Networks. JOCN, vol. 6, no. 9, September 2014.

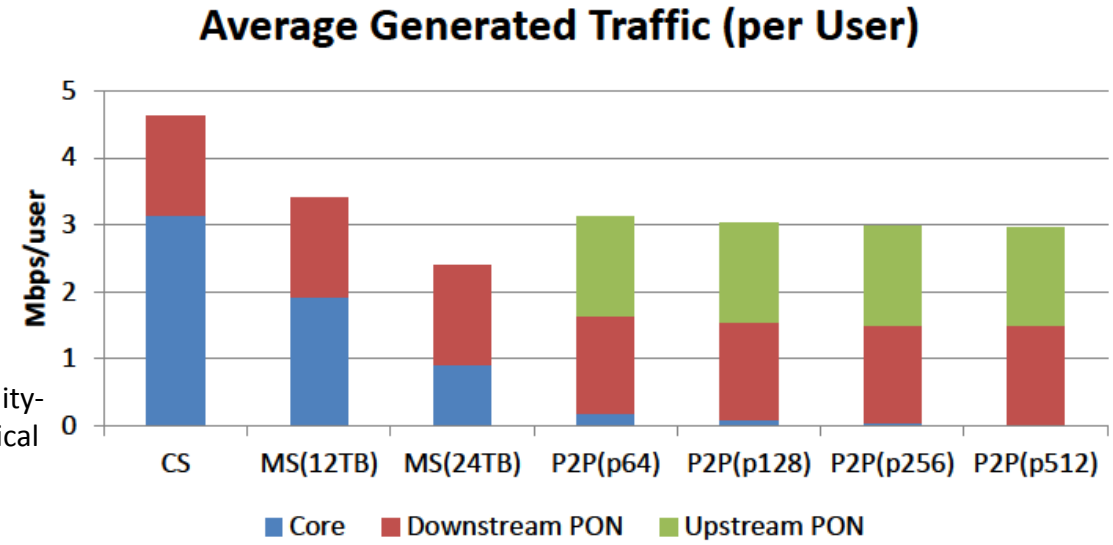

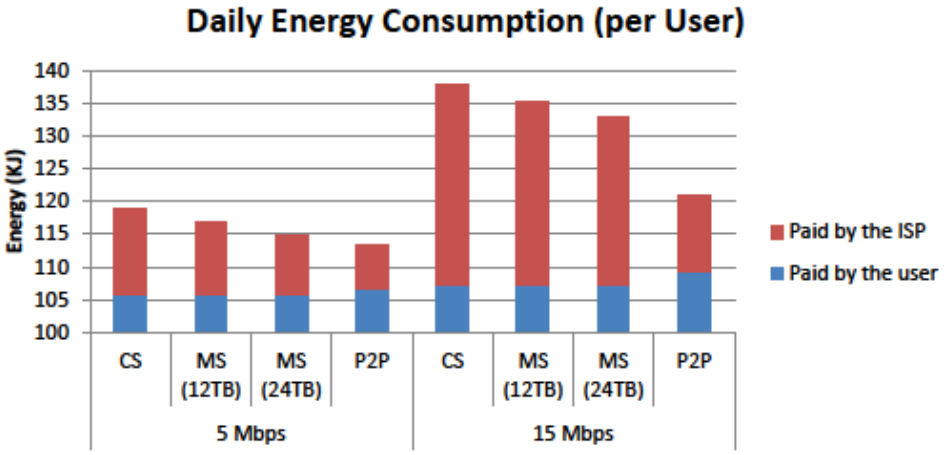

- Energy consumption lower if you assume the set top box would have storage in any case
- But more importantly the user pays the electricity bill

- Simple caching optimisation algorithms can reduce used storage capacity by 3-4 times (looking at about 1-2GB)

E. Di Pascale and M. Ruffini. Cache Storage Optimization for Locality-Aware Peer-to-Peer Multimedia Distribution. To appear in ICC 2015

Further advantages

- Simple caching optimisation algorithms can reduce used storage capacity by 3-4 times (looking at about 1-2GB)

E. Di Pascale and M. Ruffini. Cache Storage Optimization for Locality-Aware Peer-to-Peer Multimedia Distribution. To appear in ICC 2015

- In principle could be used for more general data storage: distributed data centre

- Its highly distributed nature makes it inherently resilient (more work should be done on this topic)

Conclusions

- Currently standardised FTTH access technology more then enough for covering foreseeable future requirements
- Capacity crunch might occur in the core (from an economic viability perspective) following a steep capacity increase in the access
- Architectural network innovation can help with the economic capacity crunch (we looking at LR-PON and Flat core to reduce cost of access and core network -- preliminary analysis, more results coming soon)
- Bypass of core already happening, but worth further investigation: edge caching seems very effective in reducing core traffic (especially on architectures with high degree of node consolidation)

Future work

More work required on network consolidation, over three dimensions:

- Consolidation in the space domain (less nodes):
  - Use fibre low loss and optical amplifiers in the access to reduce number of nodes (current network architecture inherited from the **copper age**)
- Consolidation in the services domain (one network for all services):
  - Run diverse services on same (FTTH) infrastructure (residential broadband, small-cell backhauling, business services)
  - No need for a separate network, but a more intelligent one (capacity virtualisation)
  - Reduce cost of building and maintaining/operating different networks
- Consolidation in the ownership domain (one physical infrastructure for all operators):
  - Share cost of physical infrastructure, lower barriers for new entrants, facilitate competition
  - More work on virtualisation of access network to create mutually independent virtual slices
- SDN seems the right framework for achieving all this