Intelligent Energy Aware Networks - a Content Perspective

Jaafar Elmirghani, University of Leeds, UK

j.m.h.elmirghani@leeds.ac.uk

EPSRC
Pioneering research and skills

Global Uncertainties
Security for all in a Changing World

UNIVERSITY OF LEEDS

Outline

- Introduction
- The Intelligent Energy Aware Networks (INTERNET) project
- Caching and IPTV / VoD networks
- Peer-to-peer energy efficient networks
- Distributed Energy Efficient Clouds
- Future directions

World wide ICT Carbon footprint

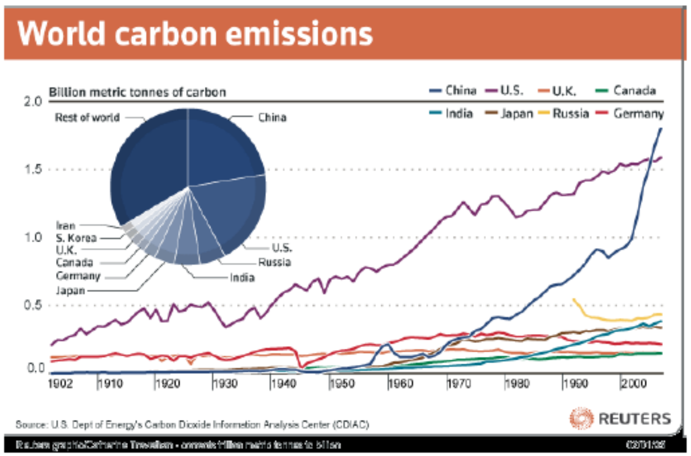

Country	Network	Energy Consumption	% of Country Total Energy Consumption
USA	Verizon 2006[1]	8.9 TWh	0.24%
Japan	NTT 2001[2]	6.6 TWh	0.7%
Italy	Telecom Italia 2005[3]	2 TWh	1%
France	France Telecom-Orange 2006[4]	2 TWh	0.4%
Spain	Telefonica 2006[5]	1.42 TWh	0.6%

S.Roy, IEEE Intelec 2008

Smart Grids

Smart Transportation

Smart Communities

Enabling a Low Carbon Economy

Smart Buildings

E-Health

360m tons CO_2

820m tons CO_2

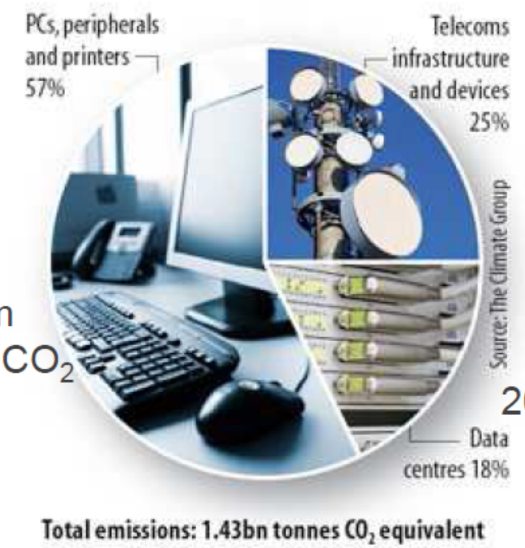

260m tons CO_2

- 2007 Worldwide ICT carbon footprint: 2% = 830 m tons CO_2
- Comparable to the global aviation industry
- Expected to grow to 4% by 2020

- Courtesy Thierry Klein, Alcatel-Lucent Bell Labs

Internet Traffic Growth Rate

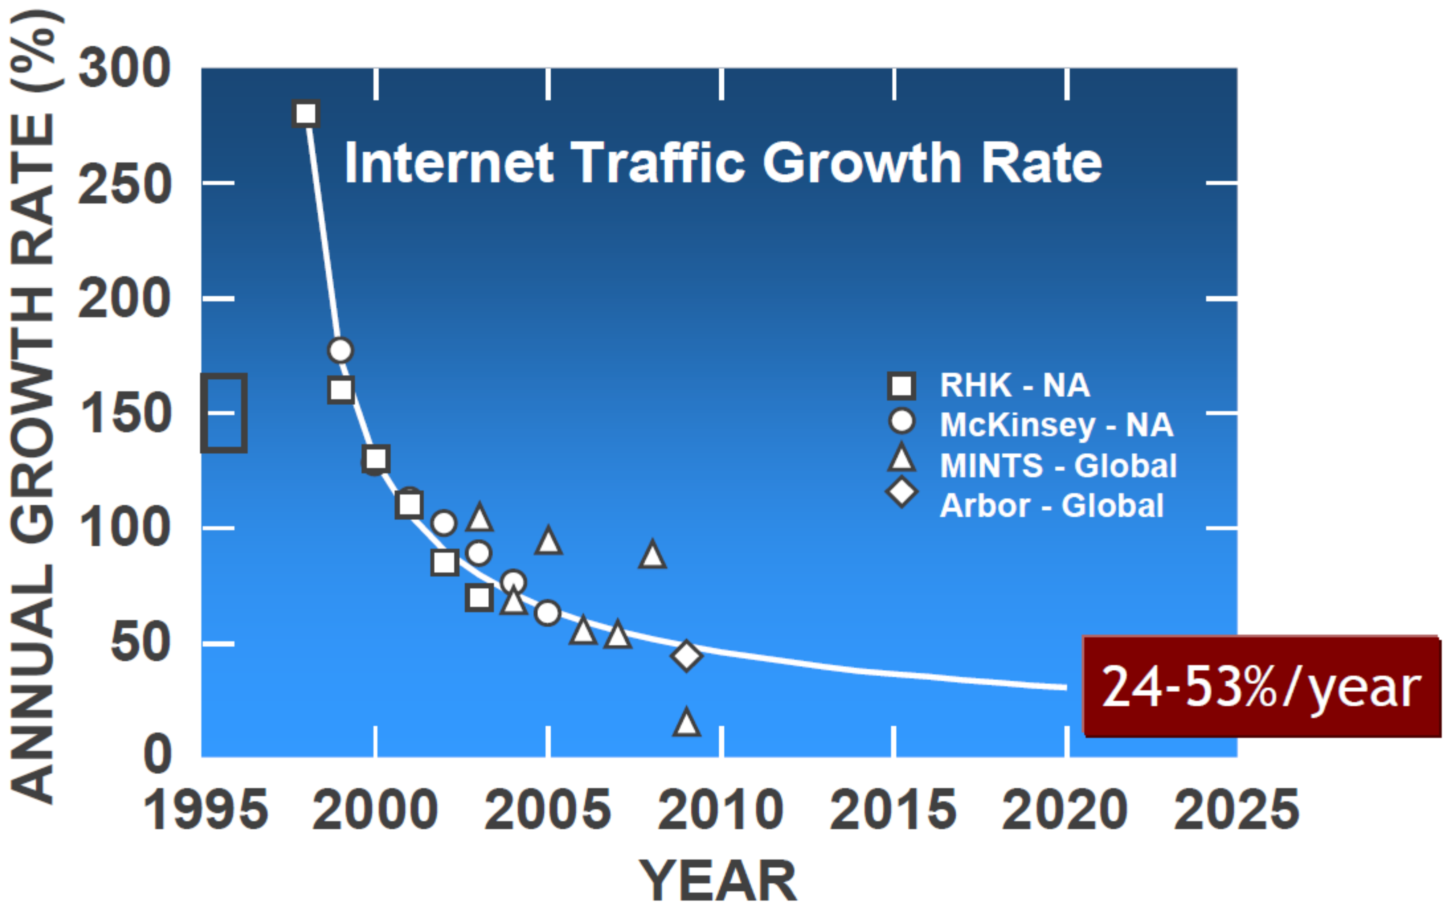

- Courtesy Thierry Klein, Alcatel-Lucent Bell Labs, Sources: RHK, 2004; McKinsey, JPMorgan, AT&T, 2001; MINTS, 2009; Arbor, 2009

Exponential traffic growth

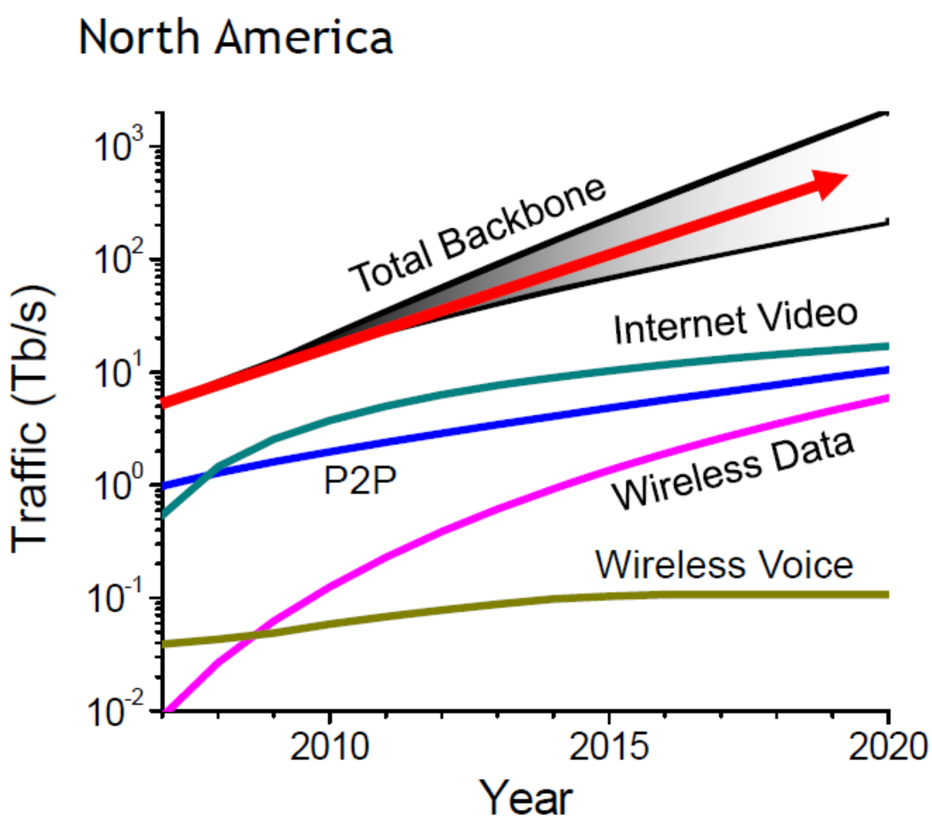

Doubling every 2 years

- 40% per year
- 30x in 10 years
- 1000x in 20 years

Mix of services is important from energy perspective:

- Mobile less efficient than fiber optics

Data from: RHK, McKinsey-JPMorgan, AT&T, MINTS, Arbor, ALU, and Bell Labs Analysis: Linear regression on log(traffic growth rate) versus log(time) with Bayesian learning to compute uncertainty

Project Goals

- The INTERNET project seeks to develop
  - New and disruptive energy efficient network architectures which are optimised for sustainable energy requirements, and are validated using national and pan-European and international models,
  - New protocols and communications techniques to support adaption within such a system, and
  - Novel hardware with low energy production and operating requirements.
- EPSRC funded, £5.9m, 2010-2015.

Collaborators

End-to-end network

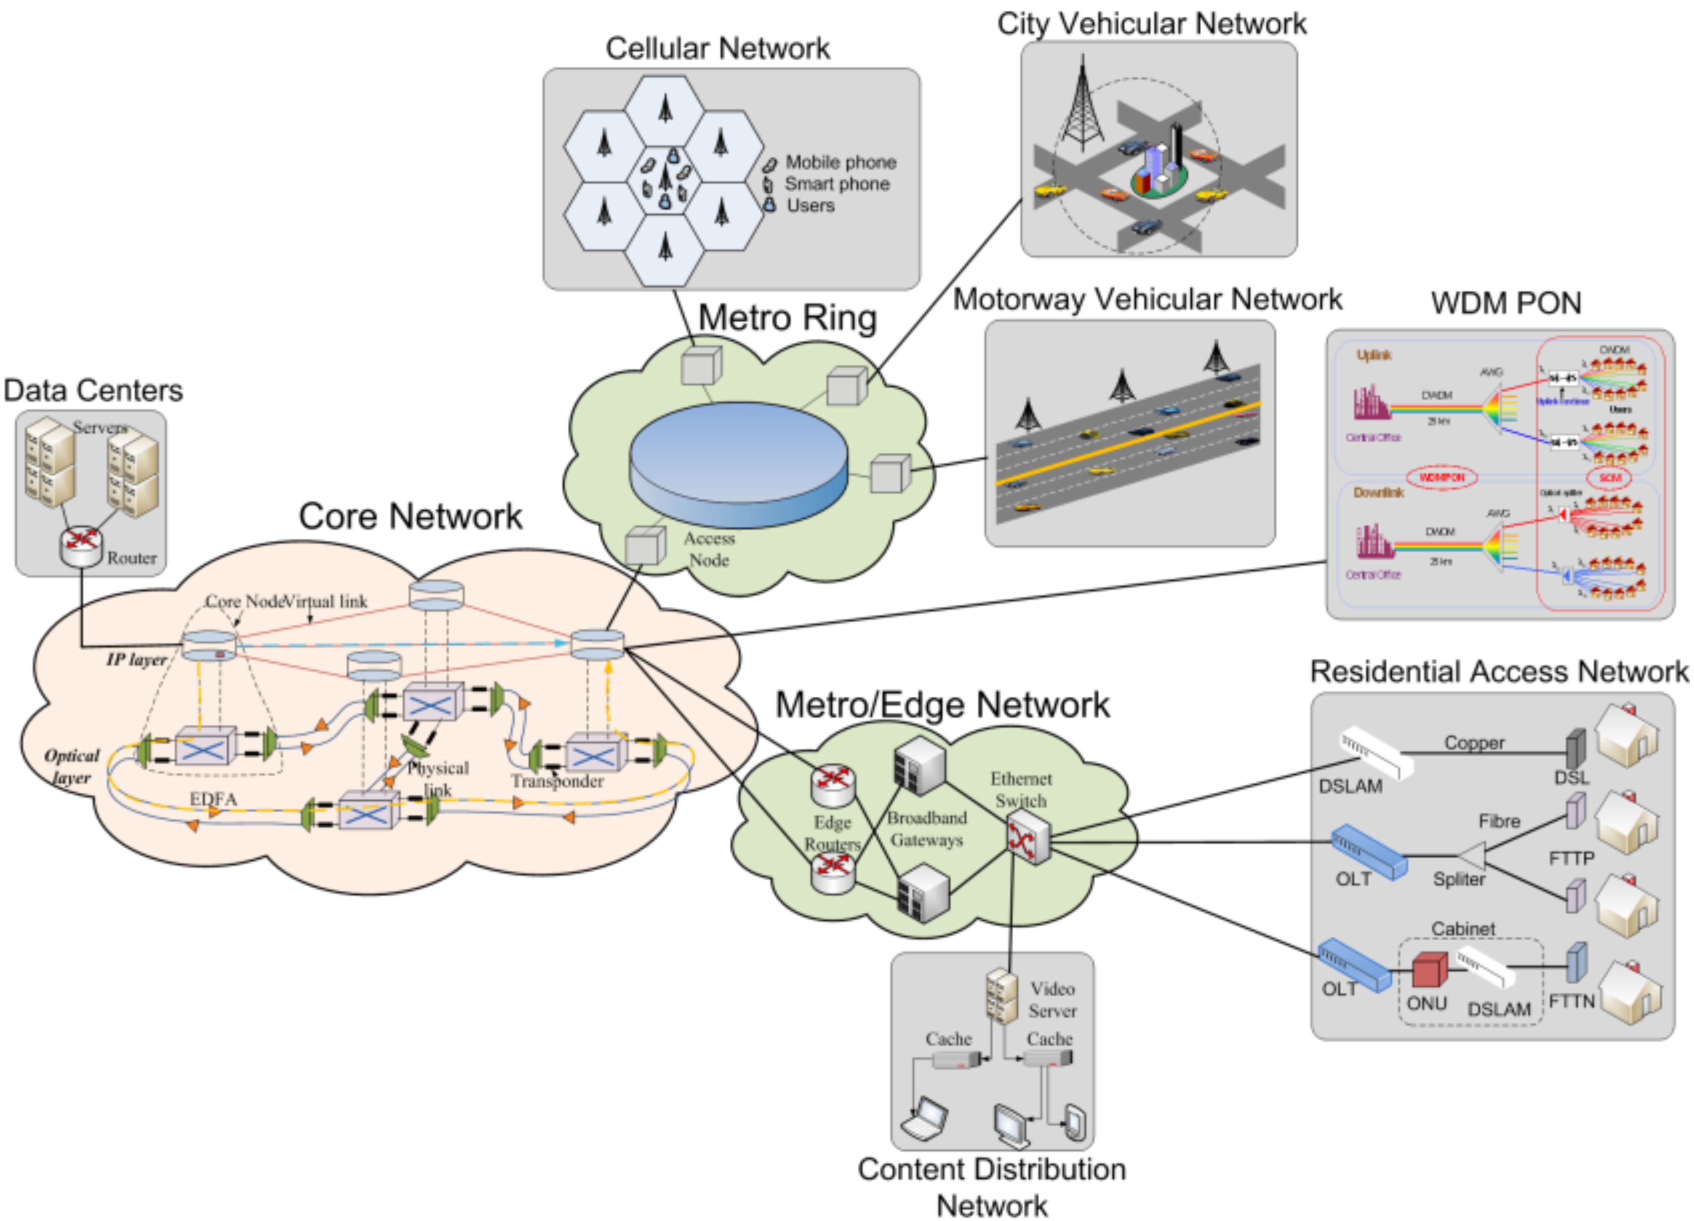

Energy Efficient Caching for IPTV On-Demand Services

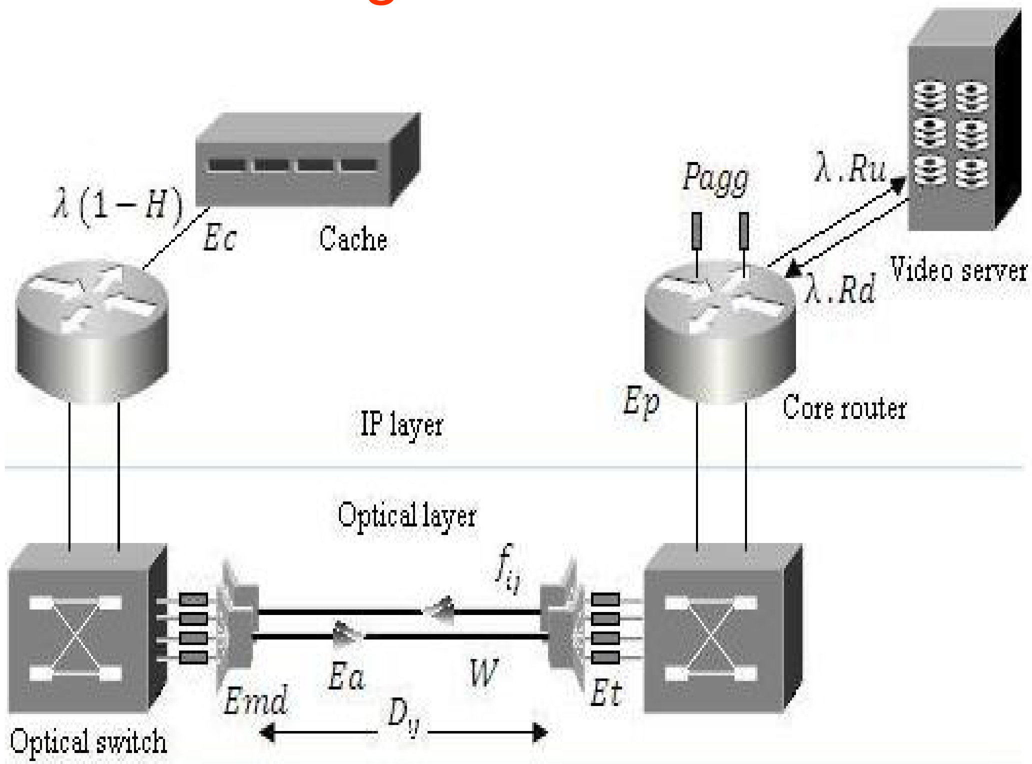

- By 2015 over 91% of the global IP traffic is projected to be a form of video (IPTV, VoD, P2P), with an annual growth in VoD traffic of 33%.
- In proxy-based architectures, proxies (or caches) are located closer to clients to cache some of the server's content.
- Our goal is to minimize the power consumption of the network by storing the optimum number of the most popular content at the nodes' caches.

Cache Size Optimization

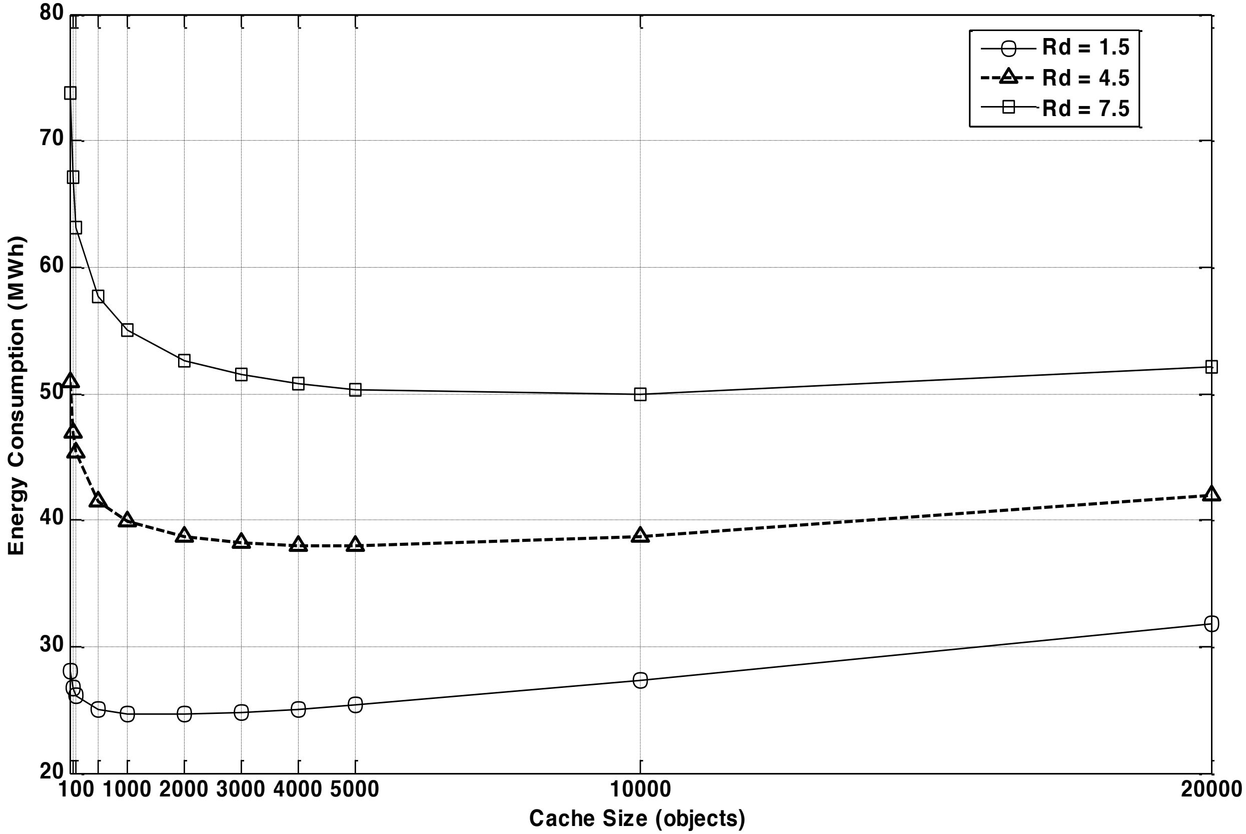

- The power consumption of the network falls with the increase in the cache size to a certain cache size after which increasing the cache size results in increasing the total energy consumption.
- In this range, the energy consumed for storage exceeds the energy consumed if some of the requests are served remotely.

Cache size optimization

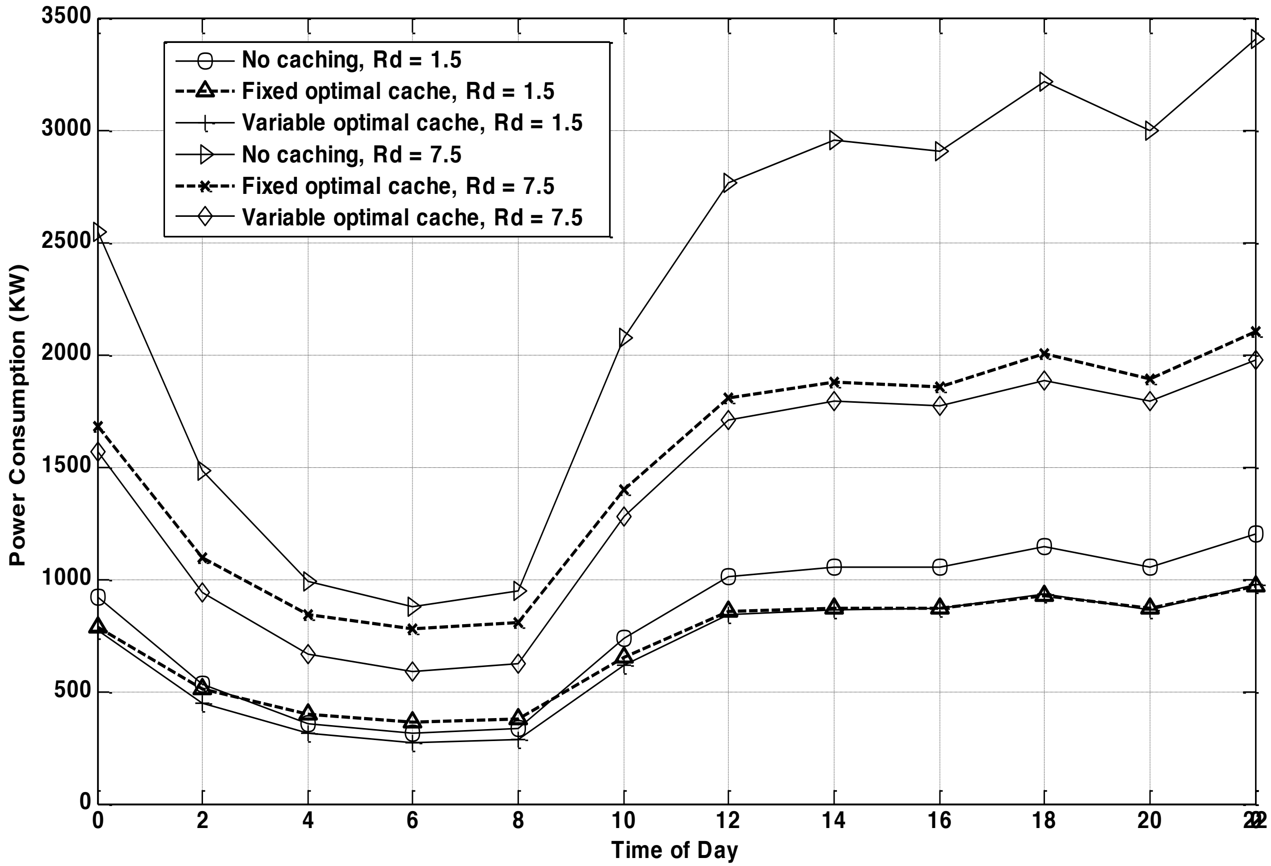

Variable size cache max network power saving 42%

- Fixed optimum cache is found considering all the nodes over the full day
- Fixed size caching reduces the network energy consumption by a maximum of 19% (average of 8%) and a maximum of 38% (average of 30%) for (Rd=1.5, Ru=0.2) and (Rd=7.5,Ru=1), respectively.

Optimum cache size at different nodes during the day (need cache size adaptation (sleep))

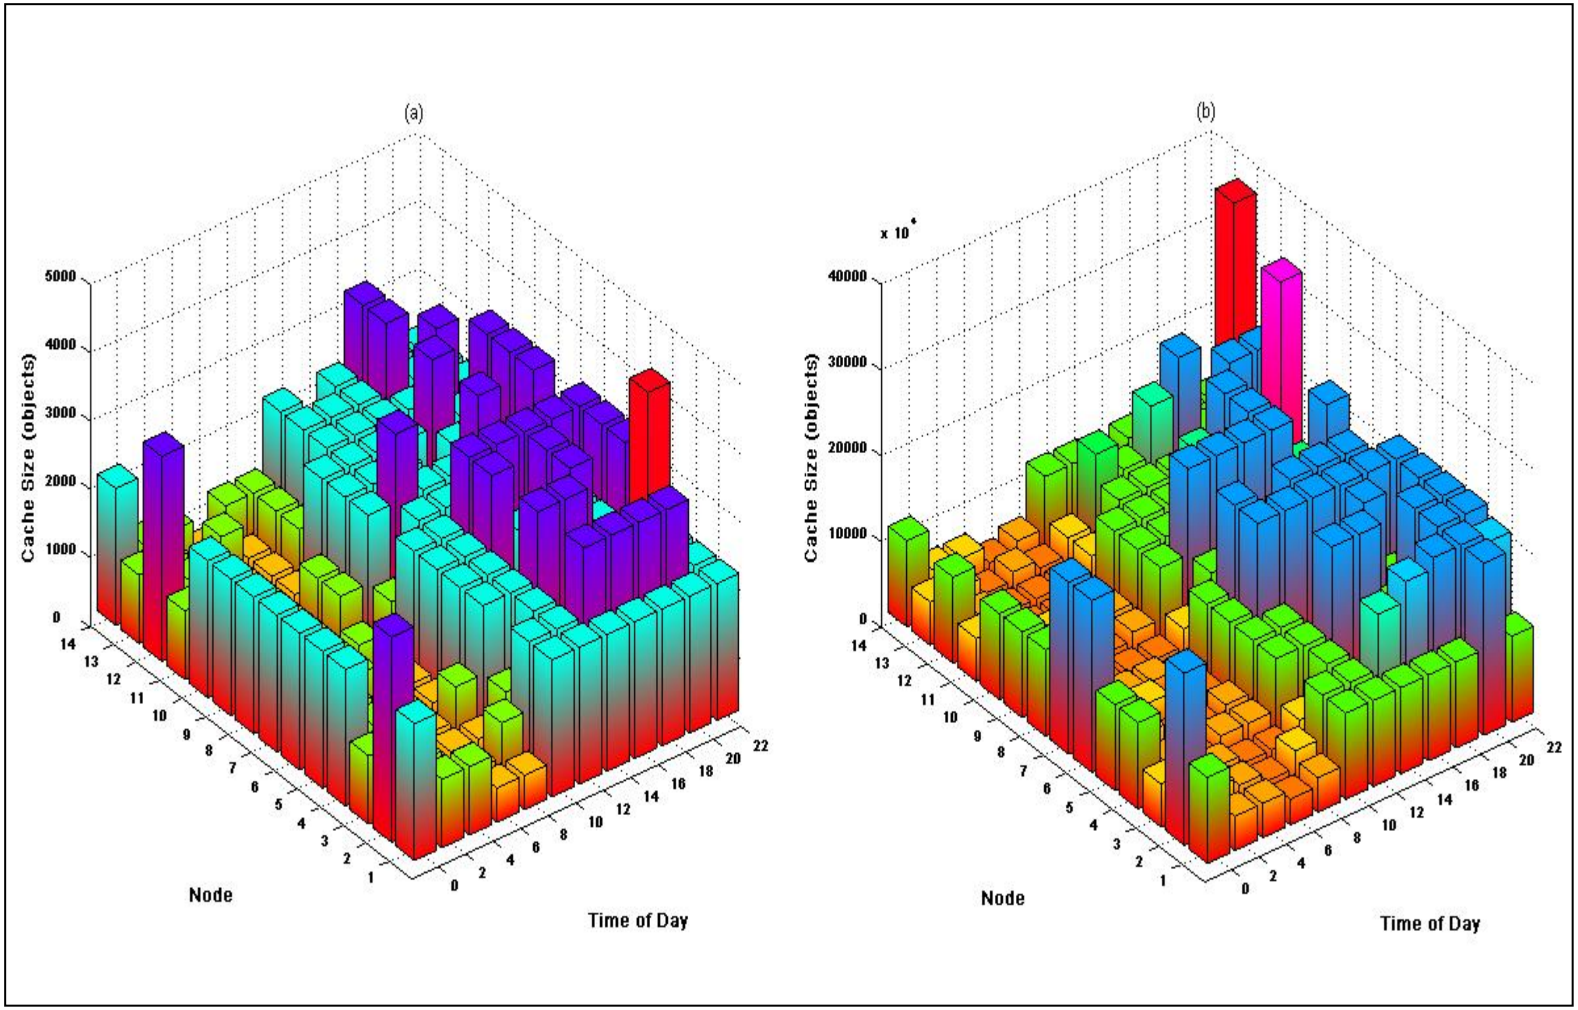

(a) Rd = 1.5

(b) Rd = 7.5

Energy-Efficient BitTorrent

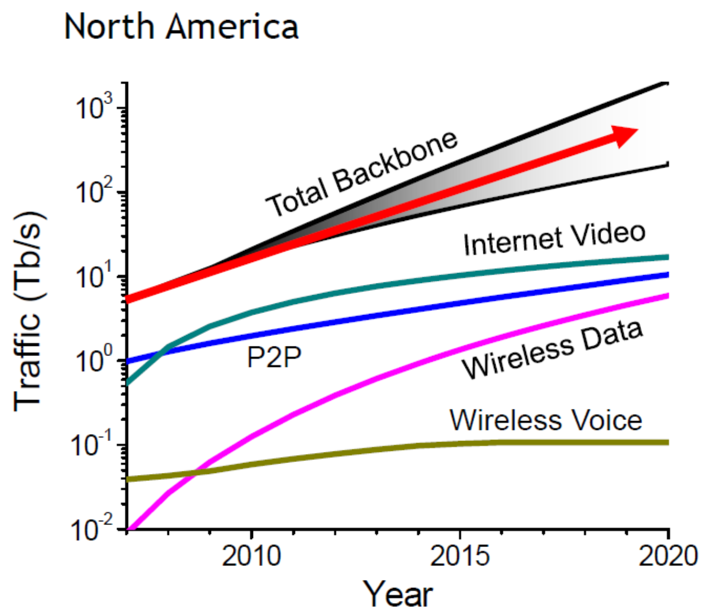

Data from: RHK, McKinsey-JPMorgan, AT&T, MINTS, Arbor, ALU, and Bell Labs Analysis: Linear regression on log(traffic growth rate) versus log(time) with Bayesian learning to compute uncertainty

Doubling every 2 years

- **40% per year**
- **30x in 10 years**
- **1000x in 20 years**

Mix of services is important from energy perspective:

- **Mobile less efficient than fiber optics**

- The two content distribution schemes, Client/Server (C/S) and Peer-to-Peer (P2P), account for a high percentage of the Internet traffic.
- We investigate the energy consumption of BitTorrent in IP over WDM networks.
- We show, by mathematical modelling (MILP) and simulation, that peers' co-location awareness, known as locality, can help reduce BitTorrent's cross traffic and consequently reduces the power consumption of BitTorrent on the network side.

Energy-Efficient BitTorrent

- The file is divided into small pieces.
- A tracker monitors the group of users currently downloading.
- Downloader groups are referred to as swarms and their members as peers. Peers are divided into seeders and leechers.
- As a leecher finishes downloading a piece, it selects a fixed number (typically 4) of interested leechers to upload the piece to, ie unchoke, (The choke algorithm).
- Tit-for-Tat (TFT) ensures fairness by not allowing peers to download more than they upload.
- We consider 160,000 groups of downloaders distributed randomly over the NSFNET network nodes.
- Each group consists of 100 members.
- File size of 3GB.
- Homogeneous system where all the peers have the same upload capacity of 1Mbps.
- Optimal Local Rarest First pieces dissemination where Leechers select the least replicated piece in the network to download first.
- BitTorrent traffic is 50% of total traffic.
- Flash crowd where the majority of leechers arrive soon after a popular content is shared.
- We compare BitTorrent to a C/S model with 5 data centers optimally located at nodes 3, 5, 8, 10 and 12 in NSFNET.
- The upload capacity and download demands are the same for BitTorrent and C/S scenarios (16Tbps).

Peer Selection
(100 Peer: 30 Seeders and 70 Leechers in Swarm 1)

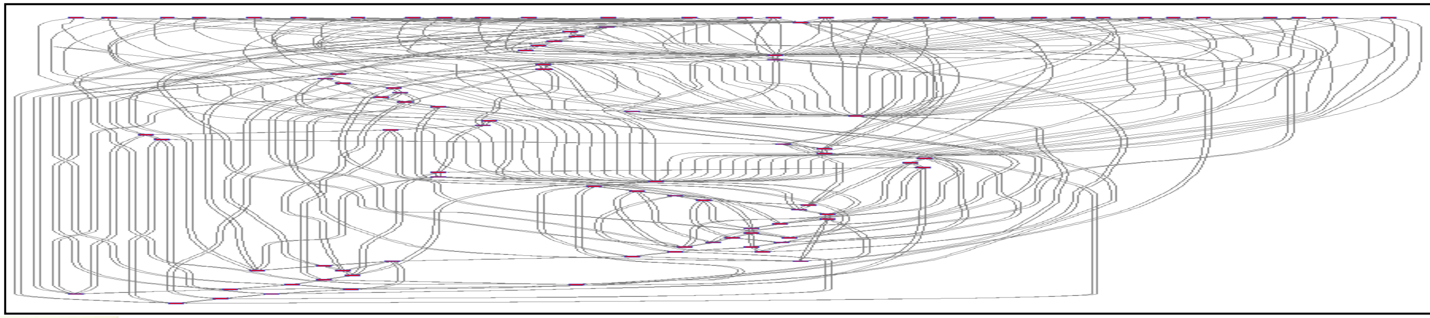

Original BitTorrent (Random Selection)

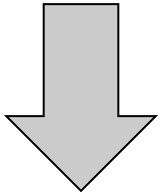

Energy Efficient BitTorrent (Optimized Selection)

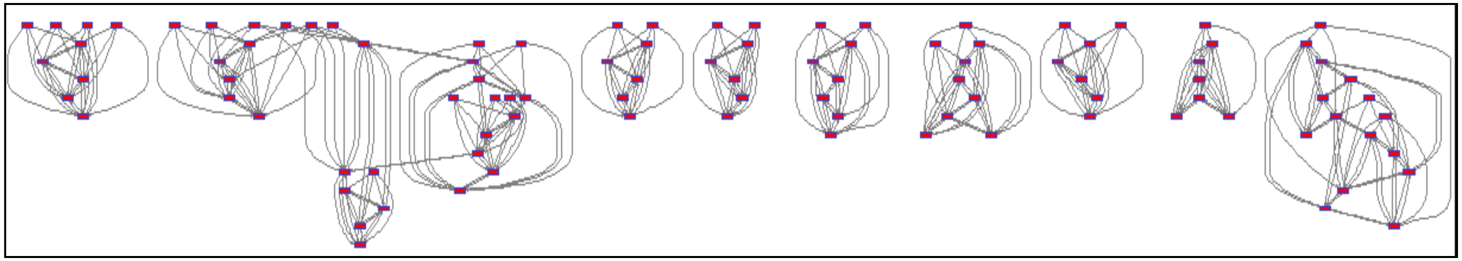

Results

Energy Consumption

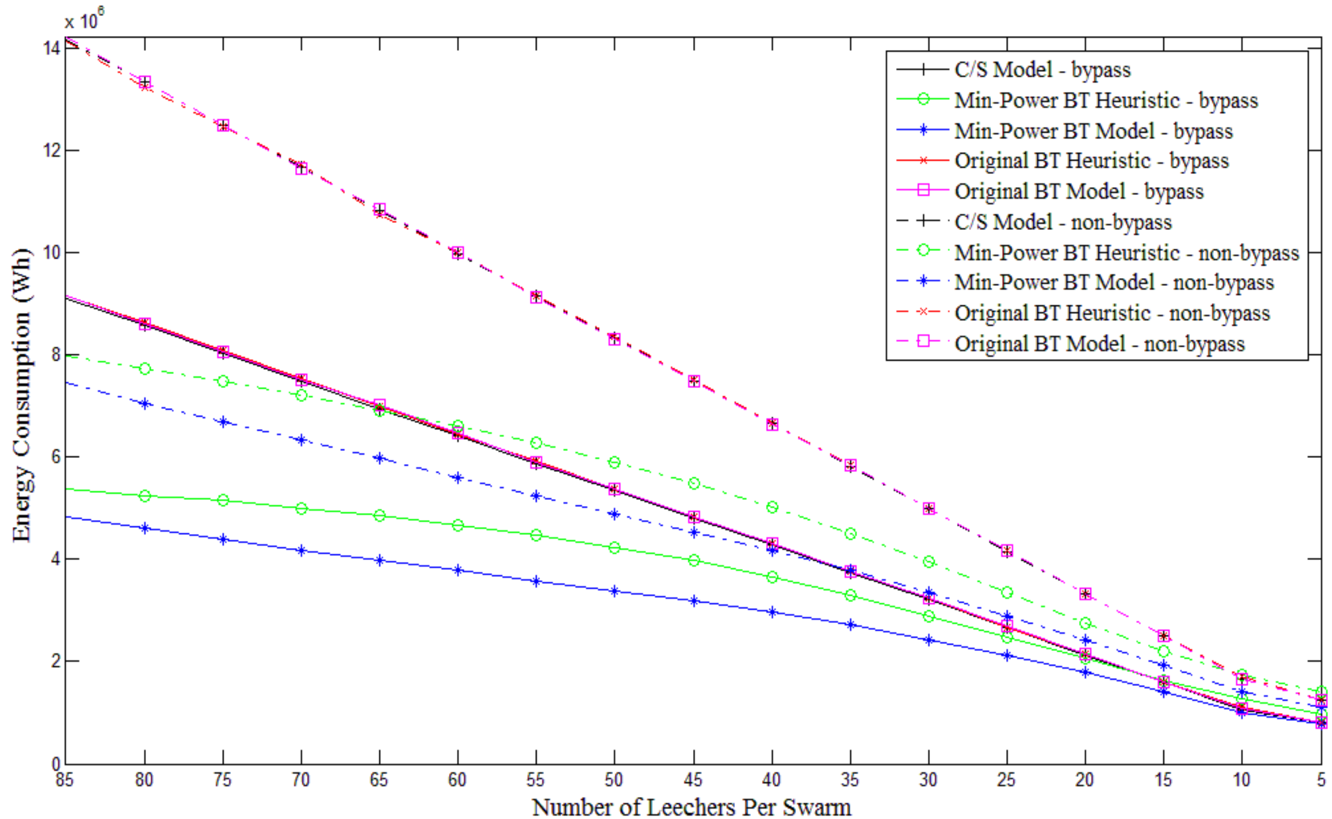

Non-bypass:
MILP average Energy Saving=36%
Heuristic average Energy Saving =25%

Bypass:
MILP average Energy Saving=30%
Heuristic average Energy Saving =15%

DEER: Distributed Energy Efficient Resources

We develop a MILP model for cloud content delivery in IP/WDM networks to answer whether centralised or distributed content delivery is the most energy efficient solution. Two kinds of decision variables are optimized for the cloud service model:

- External decision variables:
  - Number of clouds
  - Location of clouds
- Internal decision variables:
  - Number of servers
  - Number of switches
  - Number of routers
  - Storage capacity

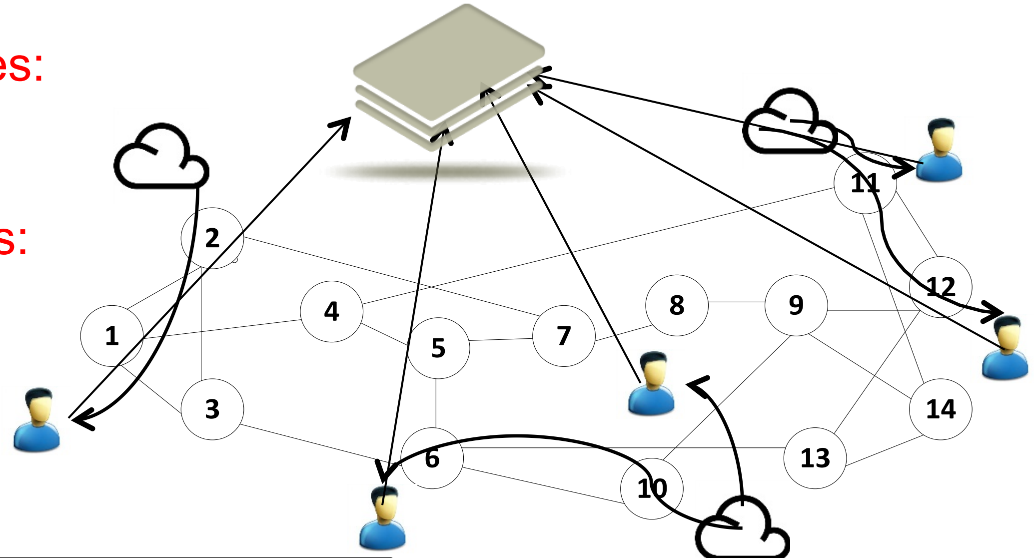

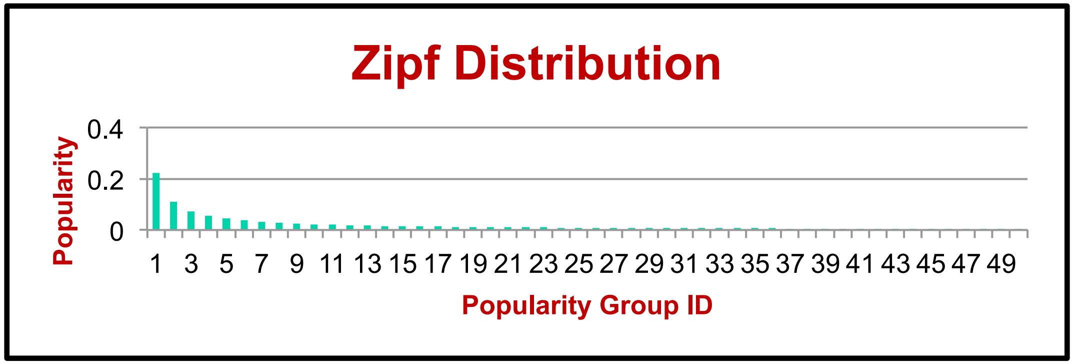

Scenarios

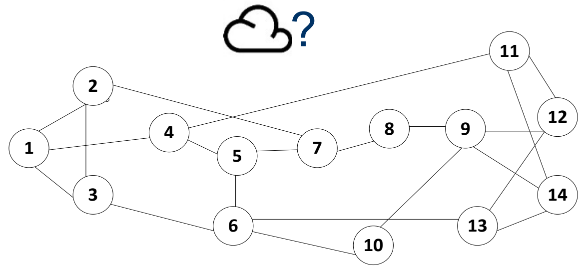

Forcing Single Cloud:
No Power Management (SNPM)
Using Power Management (SPM)

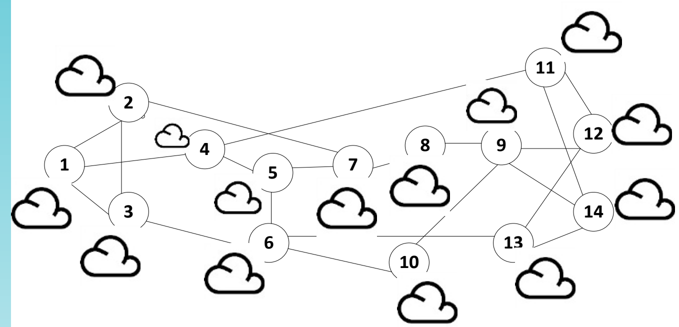

Forcing Max Number of Clouds (14):
Full Replication (MFR)
No Replication (MNR)
Popularity Based Replication (MPR)

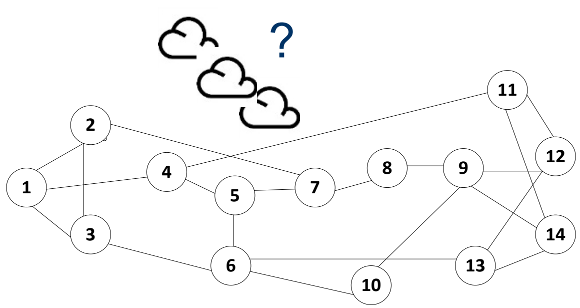

Optimal Number of Clouds:
Full Replication (OFR)
No Replication (ONR)
Popularity Based Replication (OPR)

With Power Management

Popularity Based Content Replication (OPR)
Storage=75.6*5TB

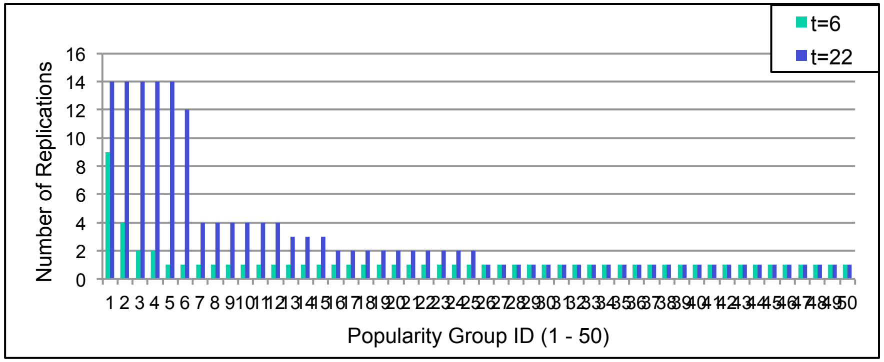

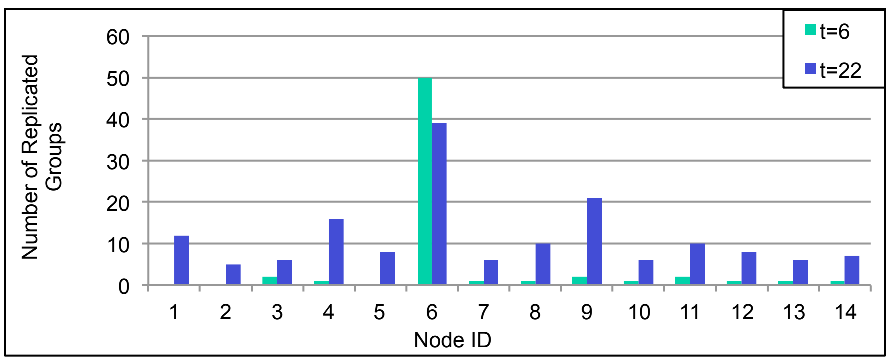

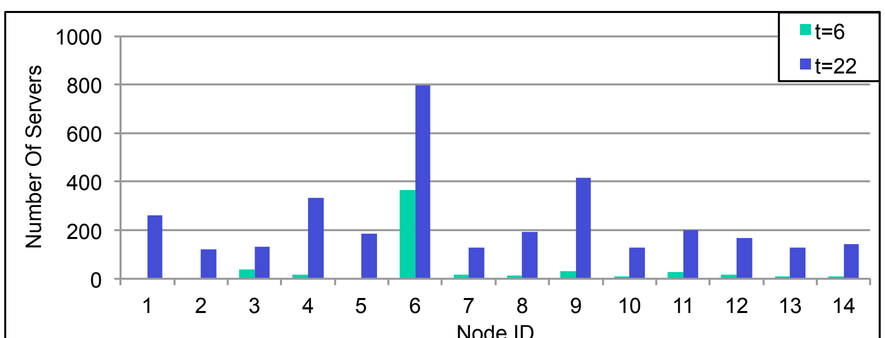

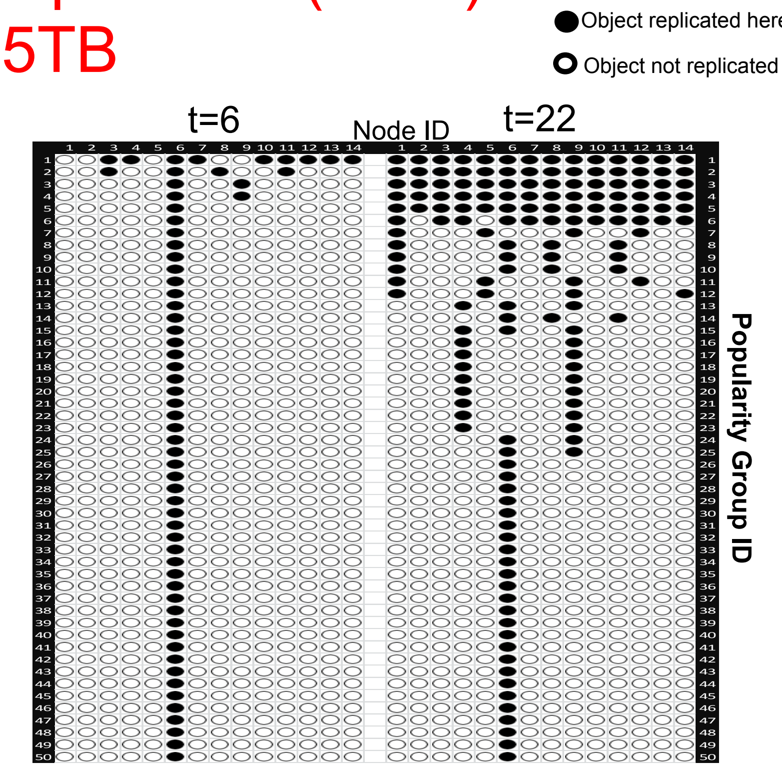

OPR Content Replication Scheme

Scenario	Total Savings	Network Saving
OPR	40%	72%
MPR	40%	72%
OFR	37.5%	56.5%
SPM	36.5%	37%
ONR	36.5%	37%
MNR	36.4%	36.5%
MFR	25.5%	99.5%

Energy Efficient Storage as a Service (StaaS)

Scenario & Assumptions

- Special case of the content delivery service where only the owner or a very limited number of authorised users have the right to access the stored content.
- All content is stored in one (or more) central locations
- StaaS should achieve trade-off between serving content owners directly from the central cloud/clouds and having clouds near to content owners.
- Upon registration for StaaS, users are granted a certain size of free storage. DropBox, for instance, grants its users 2GB.
- Different users might have different levels of utilization of their StaaS facility.
- Different users have different documents access frequency.
- High access frequency means:
    - The content owner accesses the content frequently

    and/or
    - Other authorised users become interested in the content.

Energy Efficient Storage as a Service (StaaS)

Scenario & Assumptions

- Two Average document sizes are evaluated, 45MB and 22.5MB
- Number of users evaluated are 1.2M
- Users are uniformly distributed in the network.
- Users download rate (*Drate*: in Gb/s) depends on:
  - Document access frequency (*Freq: Number of downloads per hour*)
  - Document size (*Dsize*: in Gb)

$$Drate = 2 \cdot Freq \cdot Dsize/3600$$

The factor of 2 is to take the fact that users usually re-upload their content after downloading it back to their *StaaS* drive into account.

StaaS Model Results

- **Single Cloud:** Users are served by the central cloud only.
- **Optimal Clouds:** The model selects to serve users at each node either from the central cloud or from a local cloud by migrating content from the central cloud.
- **14 Clouds:** Users at each node are served by a local cloud.

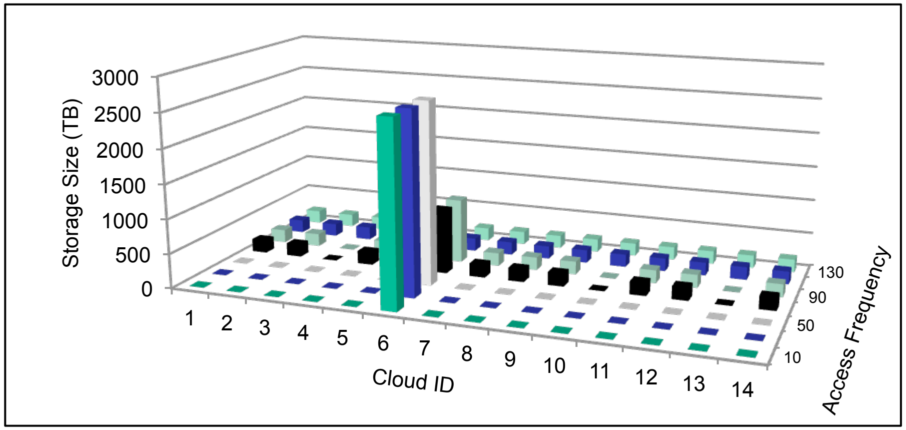

Optimal cloud scenario with the 45MB saves about 48% (averaged over the range of access frequency considered) in network power consumption compared to the single cloud scenario

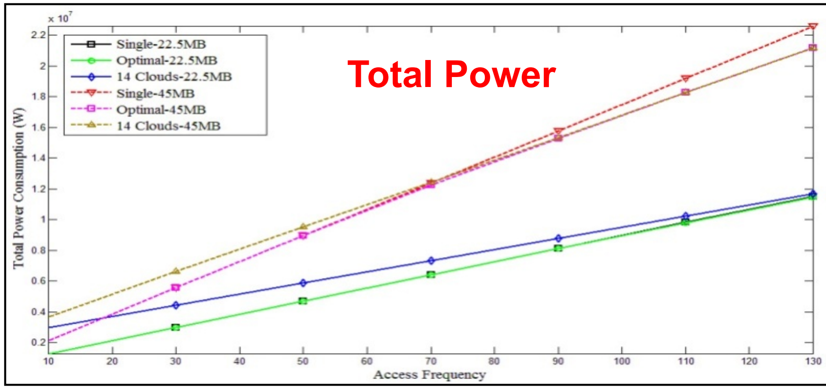

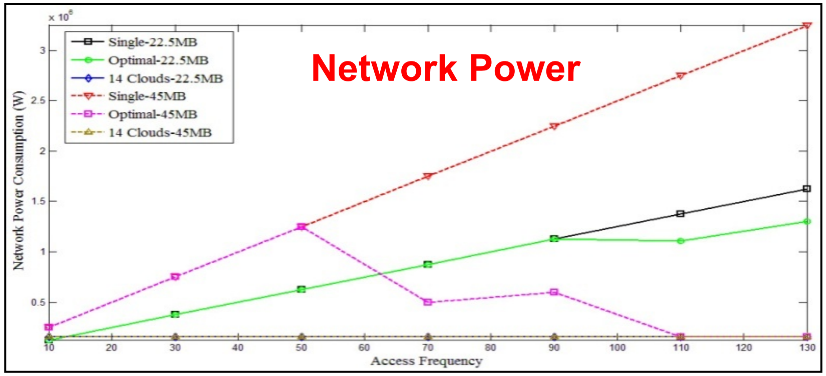

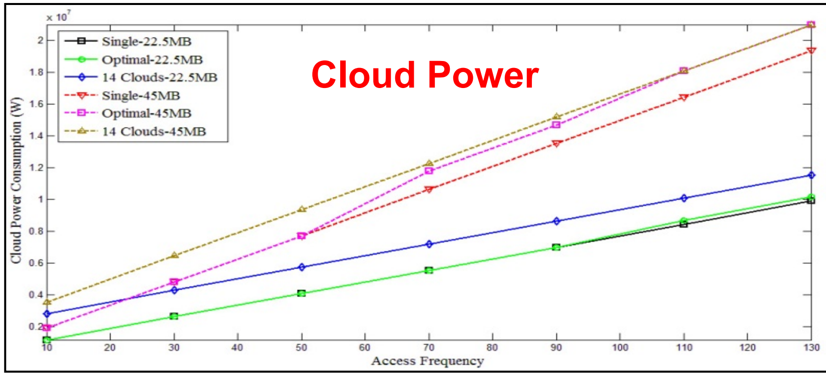

Virtual Machine (VM) Placement for Energy Efficiency

Assumptions

- Number of users fluctuates between 200k and 1200k users per day.
- Users rate 5 Mb/s,
- Users are uniformly distributed among network nodes.
- 1000 Virtual machines are evaluated due to MILP restriction on number of variables
- The problem is defined as finding the optimal location of each virtual machine

Scenarios

- **VM Migration**: Only one copy of each VM is allowed in the network
- **VM Replication**: More than one copy of each VM is allowed in the network but each copy uses full VM power
- **VM Slicing**: VMs can be divided into smaller slices to serve a smaller number of users. Sum of slices power equal VM power. We enforce a limit on the minimum size of the VM CPU utilization

Virtual Machine (VM) Placement for Energy Efficiency

We develop an MILP model to optimize cloud VM service delivery in IP/ WDM networks. Two kinds of decision variables are optimized for the cloud service model:

- External decision variables:
  - Number of clouds
  - Location of clouds
- Internal decision variables:
  - Number of servers
  - Number of switches
  - Number of routers

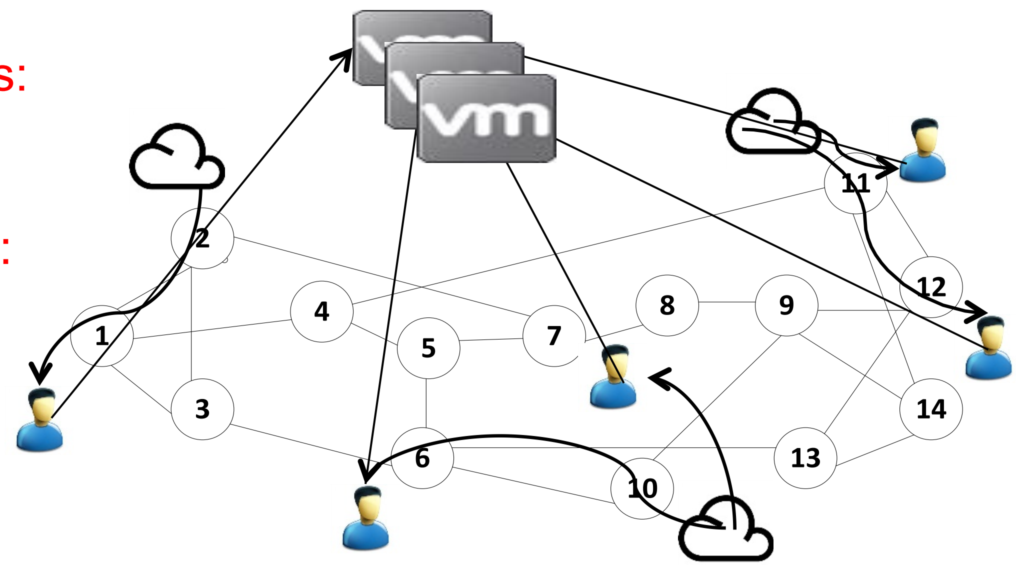

Scenario	Total Savings	Network Saving
Migrate	5%	23.5%
Replicate	6%	26%
Slice	27.5	86%

The saving are compared to single cloud at node 6

DEER-VM Heuristic

- **Migrating** VMs yields a little saving compared to single cloud solution.
- **Replicating** the full VM also yields lower saving because of the many VMs with high CPU utilization.
- **Slicing** the VMs by distributing the incoming requests among them is the most energy efficient solution.

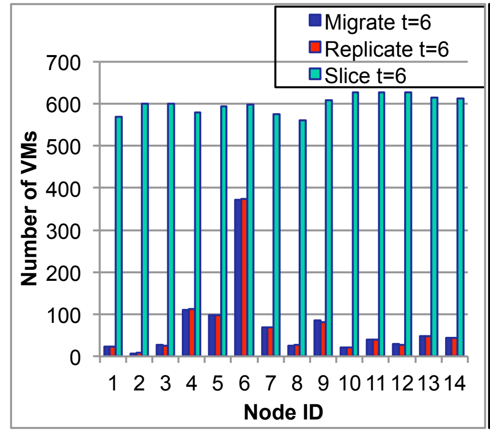

Model VMs Distribution Scheme at t=06:00

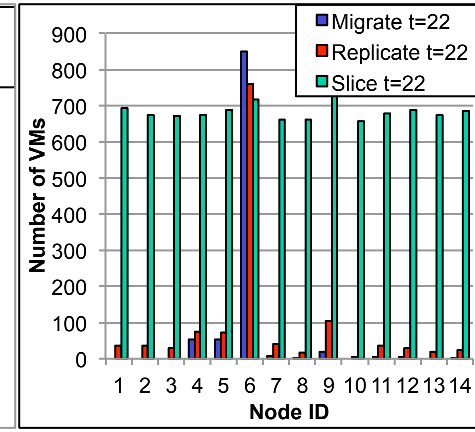

Model VMs Distribution Scheme at t=22:00

DEER-VM Heuristic

Input:	LIST= {6, 5, 4, 3, 7, 9, 13, 10, 11, 12, 14, 1, 8, 2}, $VM = \{1..NVM\}$
Output:	Optimal Placement(J'), Total Power Consumption (TPC)
	For each Virtual Machine $v \in VM$ **Do**
	For each Placement $J \subseteq$ List **Do**
	For each node $d \in N$ **Do**
	For each location candidates $\in J$ **Do**
	Add$\{cost_{sd}\} = MinHop\,(s,d)$
	$CW_{vjs} = W_v$
	End For
	Get s where: $cost_{sd} = Min\{cost_{sd}\}$
	$L_{vJsd} = D_{vd}$
	End For
	$NPC_{vJ} = MultiHopHeuristic\{N, N_m, L_{vJsd}\}$
	$CPC_{vJ} = PUE\_c \cdot (SrvPC + LANPC)$
	$TPC_{vJ} = NPC_{vJ} + CPC_{vJ}$
	End For
	$TPC_v = Min\{TPC_{vJ}\}$,
	$J' = J$
	End For
	Calculate $TPC = \sum_{v \in VM} TPC_v$

Scenario	Total Savings	Network Saving
VM-Slice-MILP	27.5%	86%
DEER-VM	21%	60%

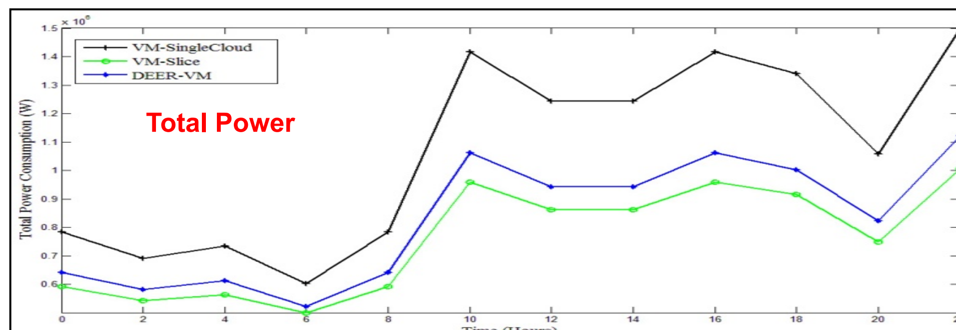

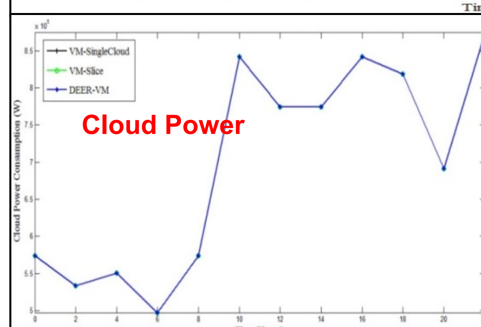

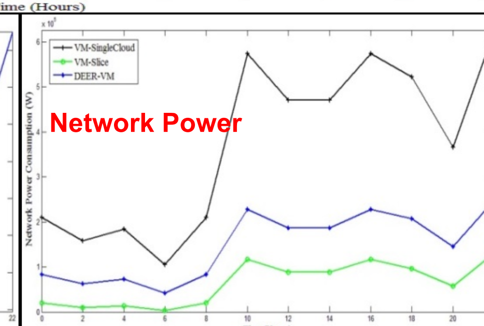

Future Directions

- Optimisation of wired wireless access architectures, metro rings - wireless mesh, PON, RoF.
- Architectures that support photonic switching instead of electronic routing.
- Auction based and self-organising dynamic architectures for energy minimisation.
- Study optimum caching location in an end-to-end network
- Develop the optimisation and simulation tools so that address energy efficiency specifically.

Related Publications

1. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., "IP Over WDM Networks Employing Renewable Energy Sources," *IEEE/OSA Journal of Lightwave Technology*, vol. 27, No. 1, pp. 3-14, 2011.
2. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., "Green IP over WDM Networks with Data Centres," *IEEE/OSA Journal of Lightwave Technology*, vol. 27, 2011.
3. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., "On the Energy Efficiency of Physical Topology Design for IP over WDM Networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 28, 2012.
4. Lawey, A., El-Gorashi, T.E.H. and Elmirghani, J.M.H., "Distributed Energy Efficient Clouds over Core Networks," *IEEE Journal of Lightwave Technology*, vol. 32, No. 7, pp. 1261 - 1281, 2014.
5. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., "Use of renewable energy in an IP over WDM network with data centres," *IET Optoelectronics*, vol. 6, No. 4, pp. 155-164, 2012.
6. Osman, N. I., El-Gorashi, T.E.H. and Elmirghani, "Caching in green IP over WDM networks," *Journal of High Speed Networks*, vol. 19, No. 1, pp. 33-53, 2013.
7. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., "Renewable Energy for Low Carbon Emission IP over WDM networks," *Proc. 15th IEEE Optical Network Design and Modelling conference* (*ONDM'11*), Bologna, Italy, 8-10 Feb 2011.

Related Publications

8. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Low Carbon Emission IP over WDM network,” *IEEE International Conference on Communications* (*ICC’11*), Koyoto, Japan, June 2011.
9. Osman, N.I., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Reduction of Energy Consumption of Video-on-Demand Services using Cache Size Optimization,” *Proc IEEE 8th International Conference on Wireless and Optical Communications Networks WOCN2011*, Paris, 24-26 May 2011.
10. Lawey, A.Q., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Impact of Peers Behaviour on the Energy Efficiency of BitTorrent over Optical Networks,” *Proc IEEE 14th International Conference on Transparent Optical Networks* (ICTON’12), 2-5 July, 2012, UK.
11. Lawey, A.Q., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Energy-Efficient Peer Selection Mechanism for BitTorrent Content Distribution,” *IEEE Global Telecom Conf* (GLOBECOM’12), Anaheim, 3-7 Dec, 2012.
12. Osman, N.I., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “The impact of content popularity distribution on energy efficient caching,” Proc IEEE 15th International Conference on Transparent Optical Networks ICTON 2013, Cartagena, Spain, June 23-27, 2013.
13. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Energy Efficiency of Optical OFDM-based Networks,” Proc. IEEE International Conference on Communications (ICC’13), Budapest, 9-13 June 2013.
14. Dong, X., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Joint Optimization of Power, Electricity Cost and Delay in IP over WDM networks,” Proc. IEEE International Conference on Communications (ICC’13), Budapest, 9-13 June 2013.
15. Lawey, A.Q., El-Gorashi, T.E.H. and Elmirghani, J.M.H., “Energy Efficient Cloud Content Delivery in Core Networks,” IEEE Global Telecommunications Conference (GLOBECOM’13), Atlanta, 9-13 Dec, 2013.