Dynamic Resource Management in Core Networks

Jaafar M. H. Elmirghani, University of Leeds

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

EPSRC

Pioneering research and skills

Global Uncertainties

Security for all in a Changing World

UNIVERSITY OF LEEDS

Overview

- Dynamic adaptation to available renewable energy
- Optimum placement of data centres and content
- Dynamic content caching
- Dynamic peer-to-peer content distribution
- Future directions

End-to-end network

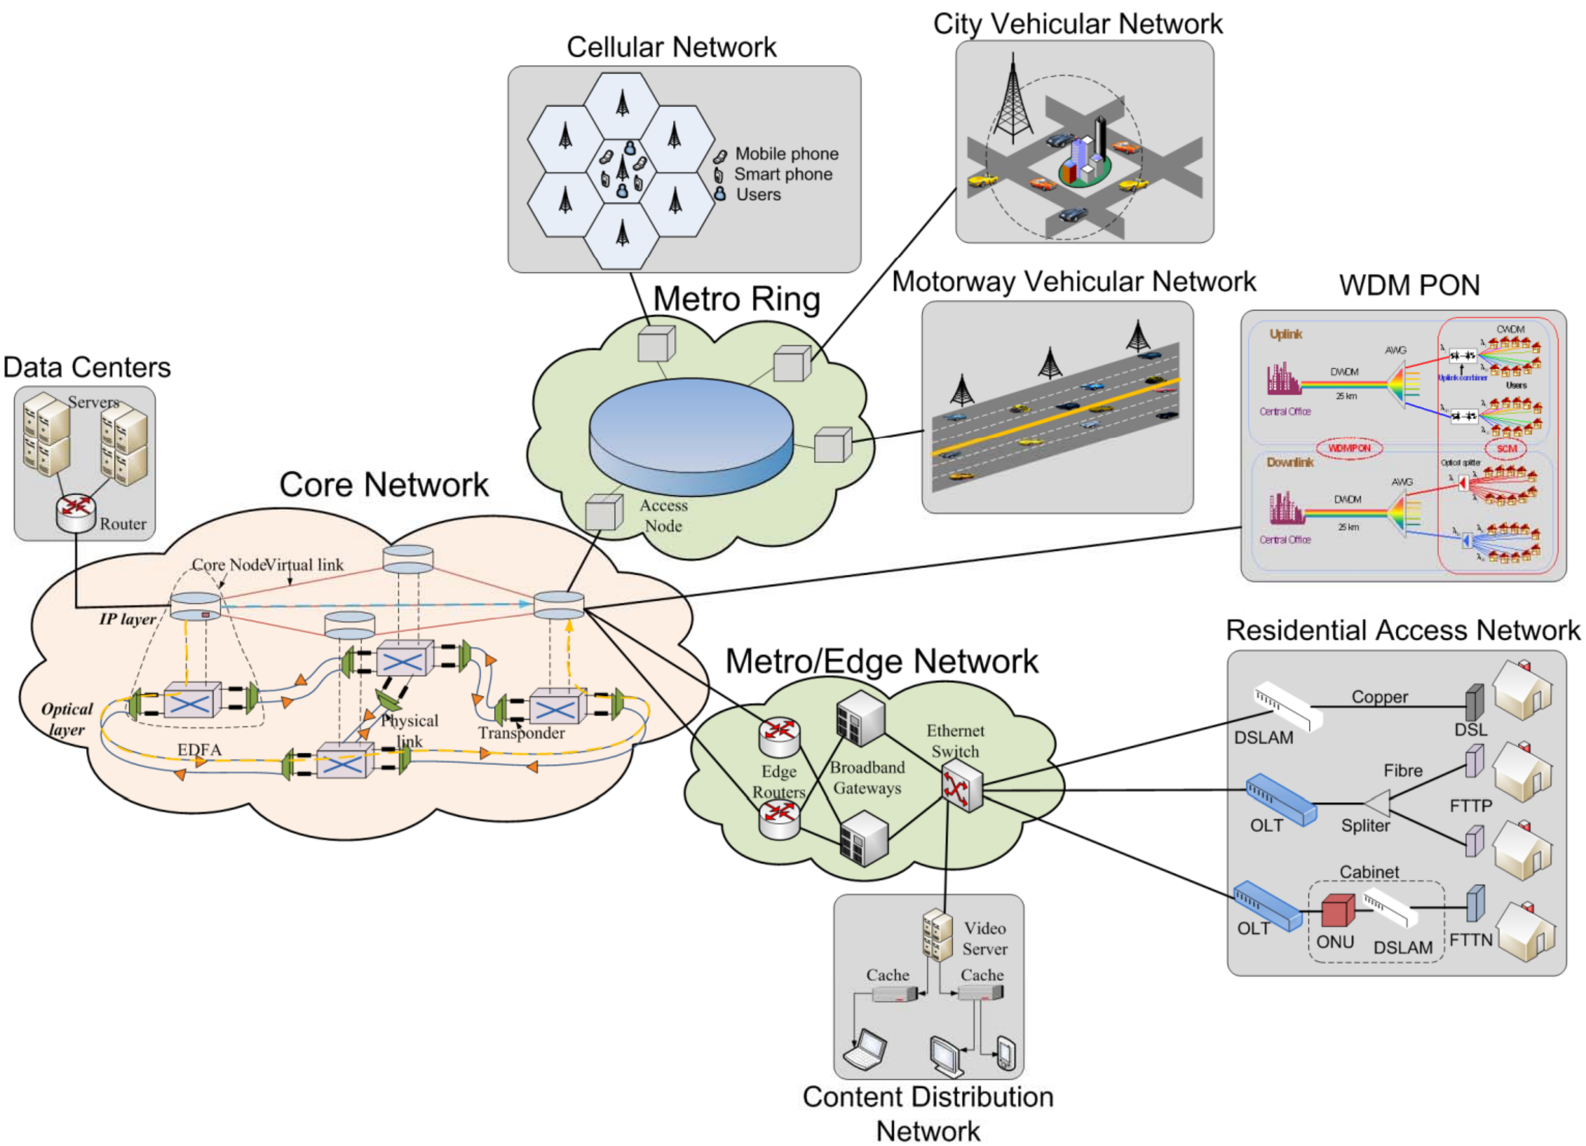

“Hybrid-power” IP over WDM network architecture

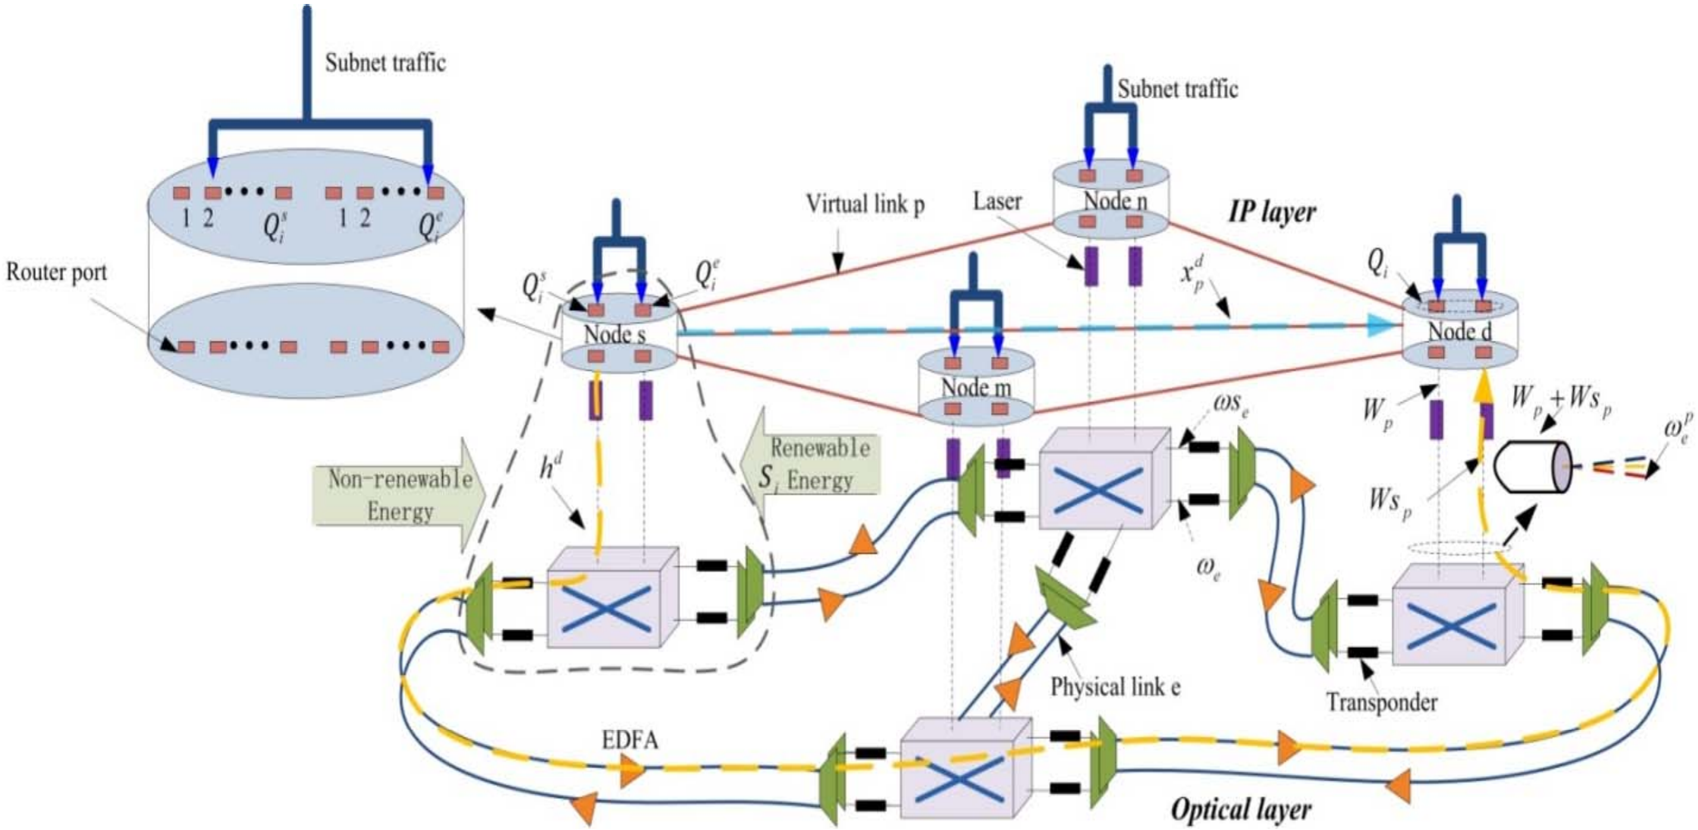

The Scenario

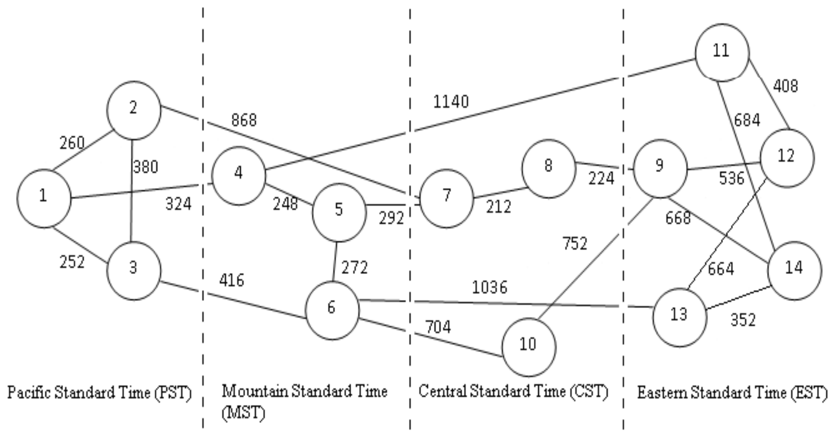

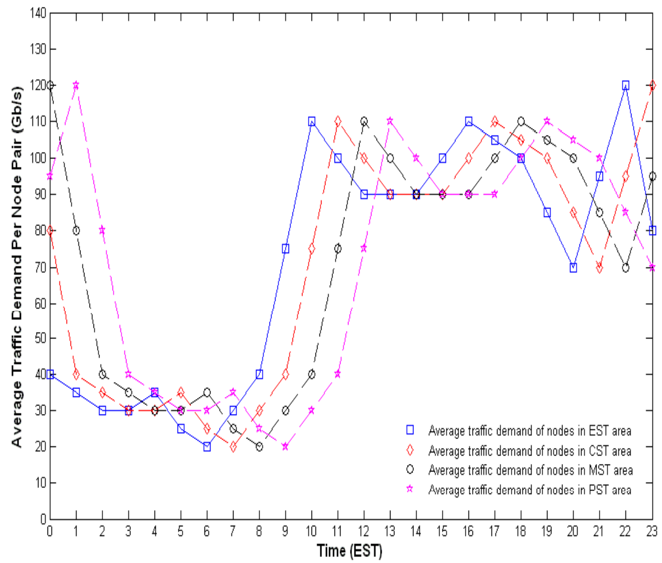

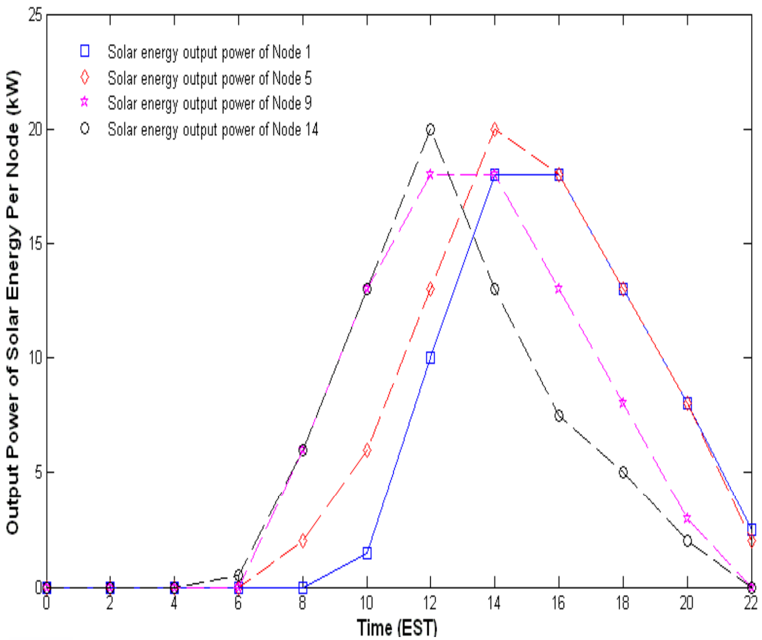

Distance between two neighbouring EDFAs	80 (km)
Capacity of each wavelength (B)	40 (Gb/s)
Power consumption of a router port (PR)	1000 (W)
Power consumption of a transponder (PT)	73 (W)
Power consumption of an EDFA (PE)	8 (W)
Power consumption of an optical switch (PO)	85 (W)
Power consumption of a multiplexer/demultiplexer (PMD)	16 (W)

Energy saving under ALR with the REO-hop heuristic

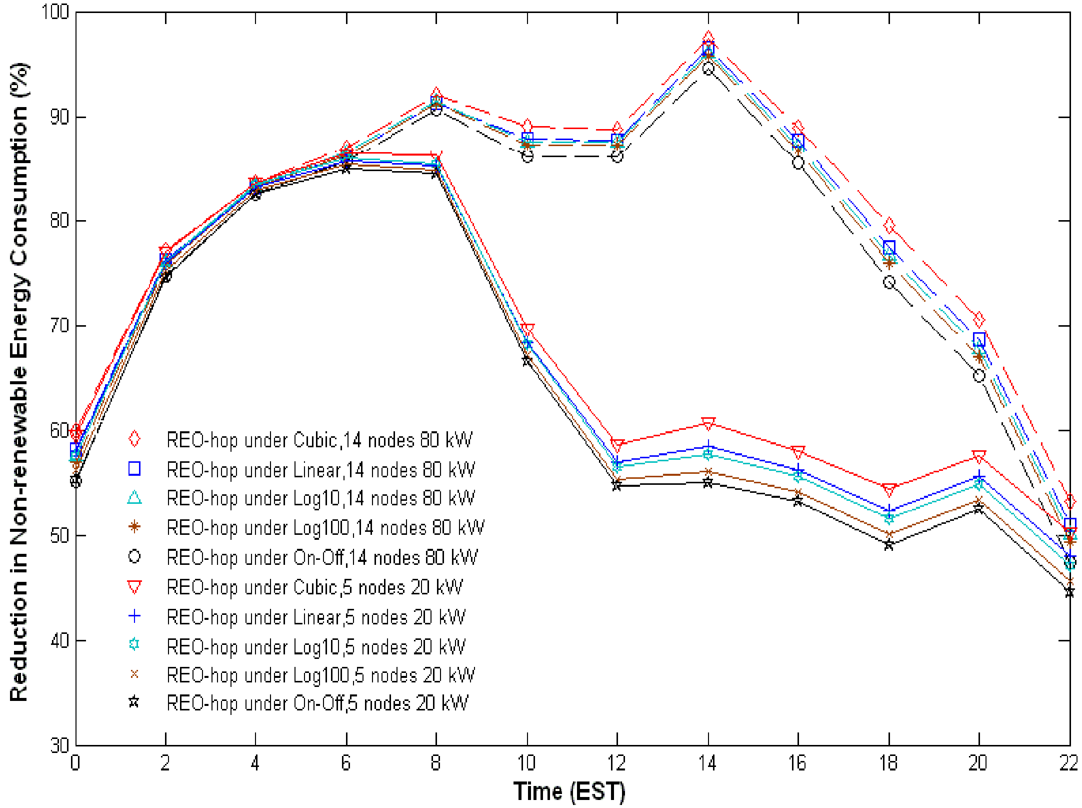

- With only 20 kW renewable in 5 nodes the energy saving compared to the non-bypass case without solar energy is approximately 85% (maximum) and 65% (average).
- Note that the 85% and 65% savings are almost real energy savings since the renewable energy is low here and has limited effect.
- When all nodes use 80 kW renewable energy, the energy saving is approximately 97% (maximum) and 78% (average).

Network design with data centres, energy-efficiency

- Three problems are investigated:
  - **Firstly**, the optimization, Linear Programming (LP), of the data centres locations to minimize the Power consumption.
    - Investigate the IP over WDM routing approach (bypass and non-bypass), the regularity of the network topology and the number of data centres in the network.
  - **Secondly**, we investigate the energy savings introduced by implementing a data replication scheme in the IP over WDM network with data centres, where frequently accessed data objects are replicated over multiple data centres according to their popularity.
  - **Thirdly**, we investigate introducing renewable energy sources (wind and solar energy) to the IP over WDM network with data centres.
    - We evaluate the merits of transporting bits to where renewable energy is (wind farms), instead of transporting renewable energy to where data centres are.
    - We consider the impact of the electrical power transmission losses, network topology, routing, traffic.

Data centres in an IP over WDM network

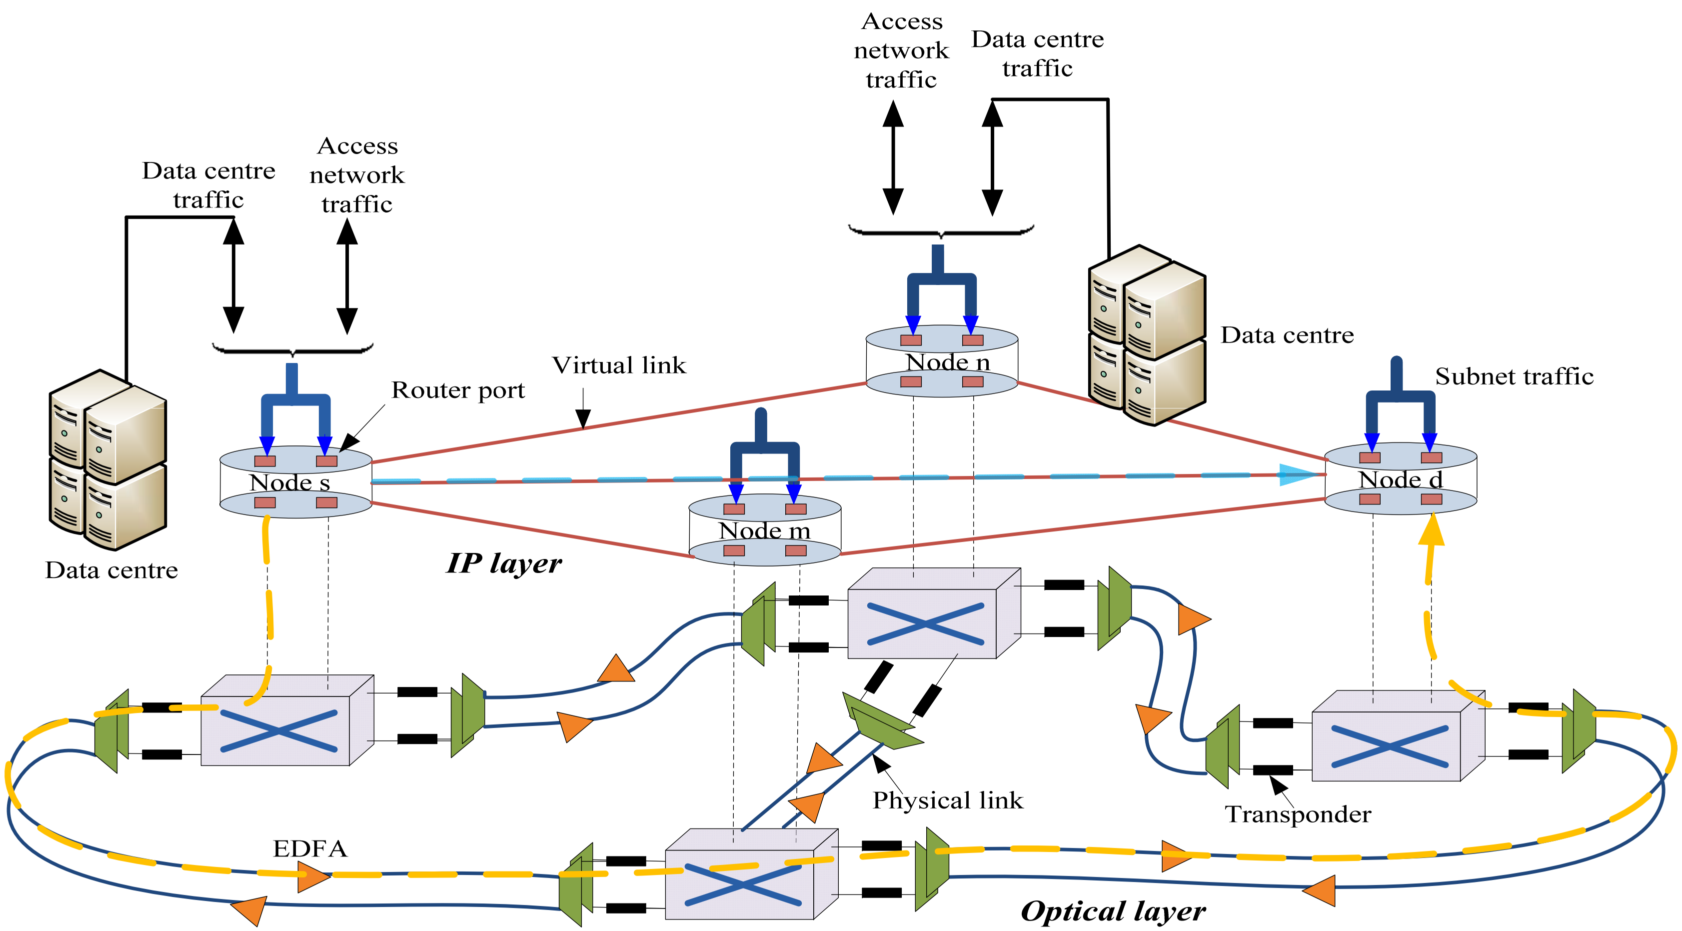

Summary of power savings as a result of data centre location optimisation

Topology	Data centre traffic only	Data centre traffic and regular traffic
Irregular topology under the non-bypass heuristic	37%	11%
Irregular topology under the multi-hop bypass heuristic	17%	6.3%
NSFNET topology with a single data centre under the non-bypass heuristic	26.6%	12.7%
NSFNET topology with a single data centre under the multi-hop bypass heuristic	8.6%	4.6%
NSFNET topology with 5 data centres under the non-bypass heuristic	11.4%	4.4%
NSFNET topology with 5 data centres under the multi-hop bypass heuristic	6.5%	1.7%

Data replication in IP over WDM networks with data centres

- Large operators have multiple data centres.
- Content (that has different popularity) can be replicated to reduce delay and power consumption.
- A MILP model is developed to optimize the selection of data centres to replicate data objects under the lightpath bypass approach.
- A Zipf distribution is assumed for content popularity.
- With 5 data objects, the popularities are: 43.7%, 21.8%, 14.5%, 10.9% and 9%.

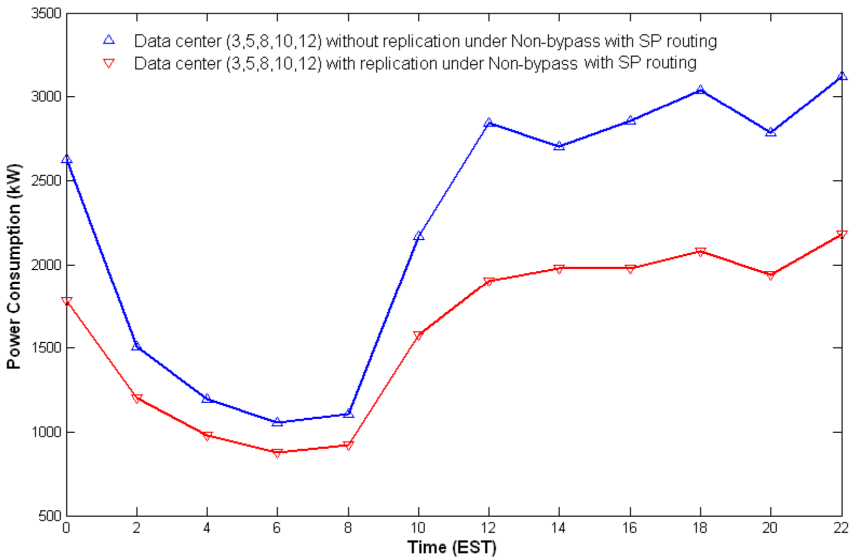

DC & regular traffic
Non-bypass:

LP optimal DC nodes = (5,6,8,10,13)

LP determines where each object is replicated

Power saving=28%

Renewable energy in IP over WDM networks with data centres

- We compare moving bits to where renewable energy is (wind farms) to transporting renewable energy to data centres.
- We have selected only 3 wind farms based on their location and maximum output power to power the data centres in the network: 1) WF1: Cedar Creek Wind Farm, 2) WF2: Capricorn Ridge Wind Farm, 3) WF3: Twin Groves Wind Farm in blue. The maximum output power of the three wind farms is 300MW, 700 MW and 400 MW, respectively.
- We assume the transmission power loss is 15% per 1000km [25] and the percentage of the power of wind farms allocated to data centres is assumed to be 0.3%.

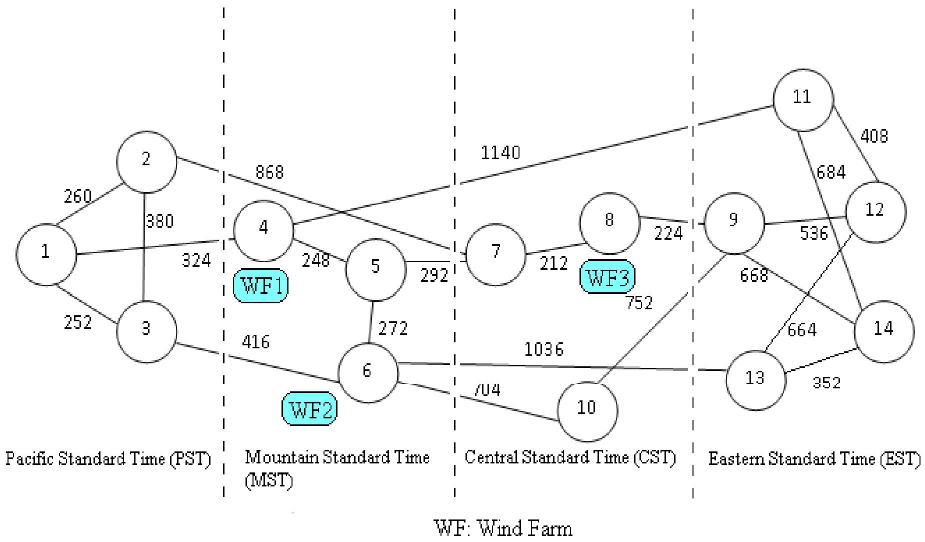

Data centre, computing, cooling and lighting power usage

- The cooling & lighting power consumption of a typical data centre is 150-200W/ft^2. Assuming a 3500ft^2 data centre, the total power consumed in a typical data centre for cooling is 700kW and the computing power consumption in a data centre is assumed to be 300kW which is typical for this data centre size.

- The power allocated by a wind farm to a data centre is known and is assumed here to be 1.4MW. This corresponds to a power usage efficiency (PUE) of 2 which is typical for a data centre.

- The renewable energy available to a data centre is a function of the transmission losses and these are location dependent. Furthermore the network topology, traffic, components' power consumption also play an important role in determining the optimum data centre location.

- Therefore the LP model here takes into account the previous trade-offs as well as the trade-offs introduced by the losses associated with the transmission of renewable energy to the data centre locations.

Renewable energy in the IP over WDM network with data centres

LP, Simulation and Results

- We run the LP model with five data centres (*Ndc*=5) under the previous assumptions.

- The optimal locations of data centres obtained from the LP model are as follows (4, 5, 6, 7, 8) where data centres 4 and 5 are powered by WF1, data centre 6 and 7 are powered by WF2, and data centre 8 is powered by WF3.

- The LP model results are such that all the data centres are located in the centre of the network.

- It can be observed that the optimum data centres locations are next to or near wind farms.

Energy efficient caching for IPTV on-demand services

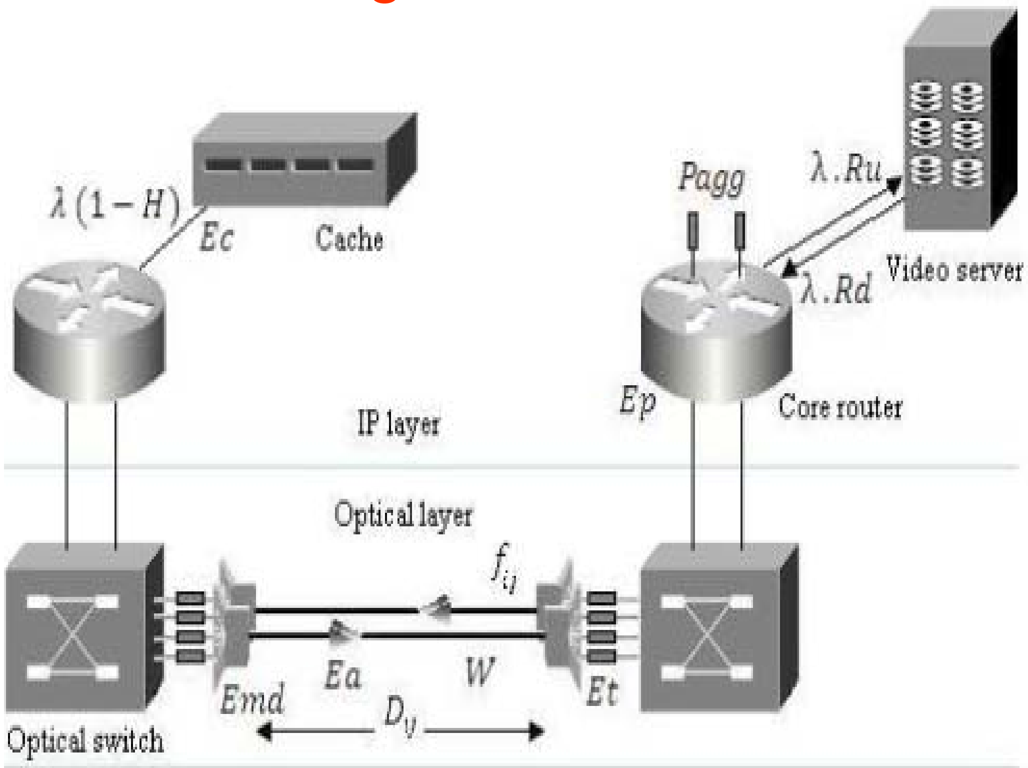

- By 2014 over 91% of the global IP traffic is projected to be a form of video (IPTV, VoD, P2P), with an annual growth in VoD traffic of 33%.
- In proxy-based architectures, proxies (or caches) are located closer to clients to cache some of the server's content.
- Our goal is to minimize the power consumption of the network by storing the optimum number of the most popular content at the nodes' caches.

Cache size optimization

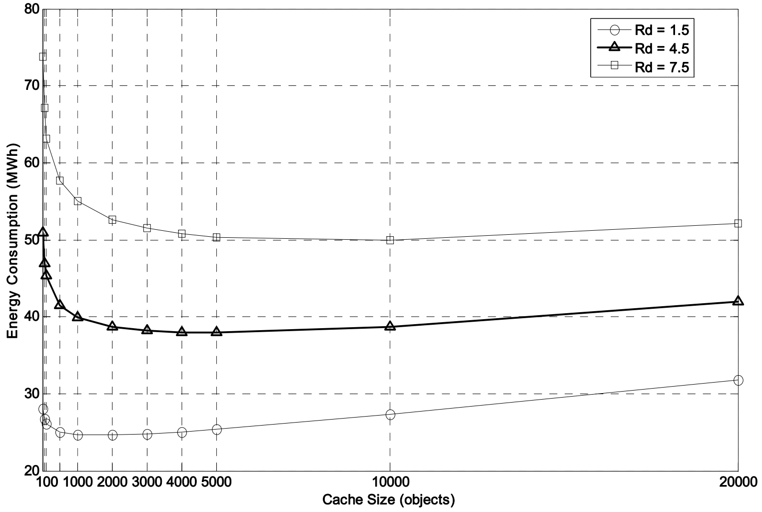

Variable size cache max network power saving 42%

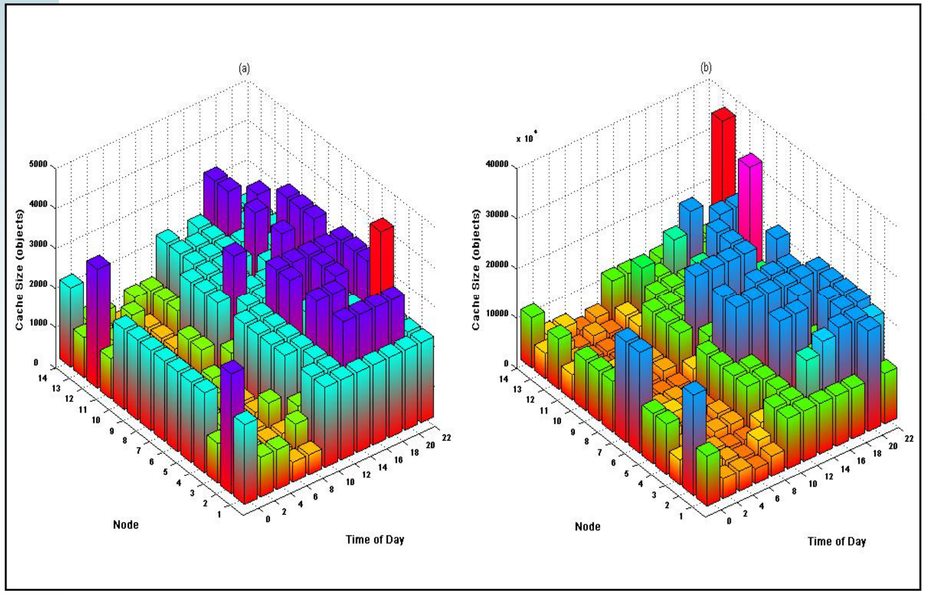

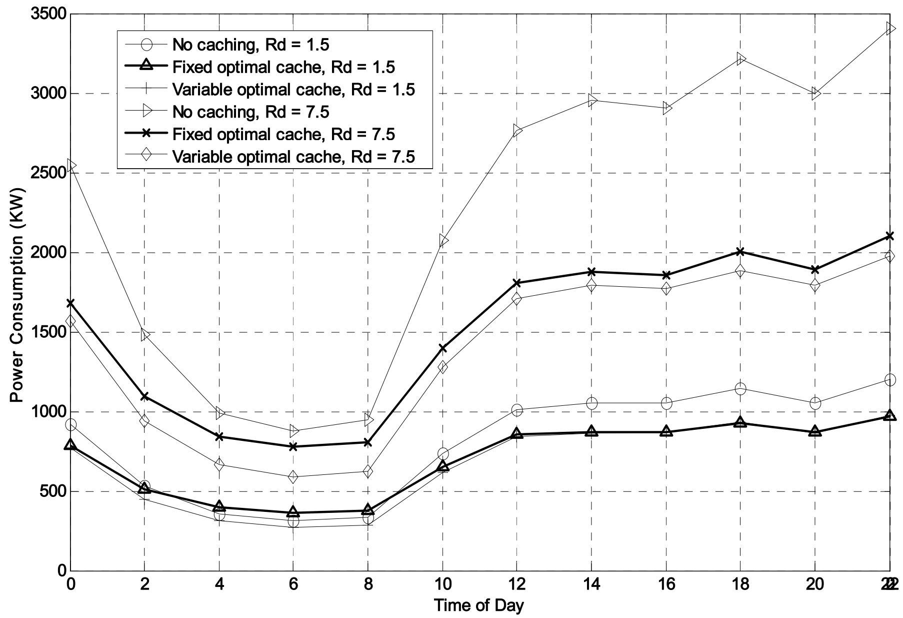

Energy Efficient BitTorrent over IP over WDM Networks

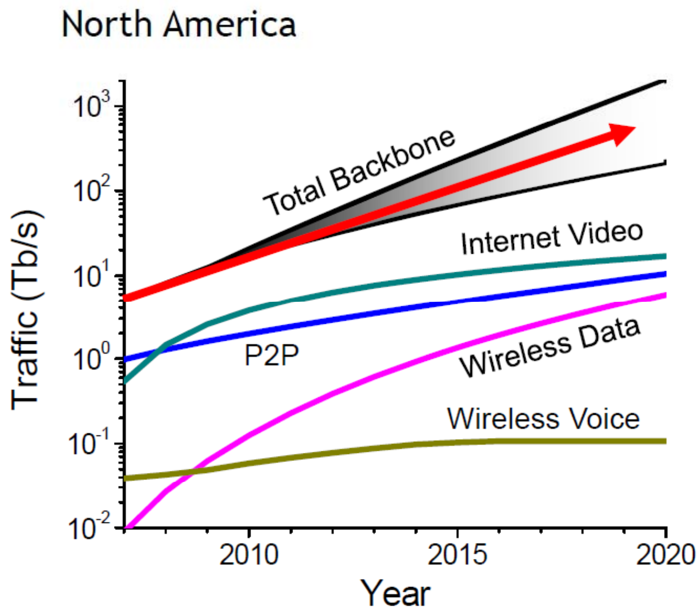

Doubling every 2 years

- 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

- 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 over IP over WDM Networks

- 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.

Energy Efficient BitTorrent over IP over WDM Networks

- 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).

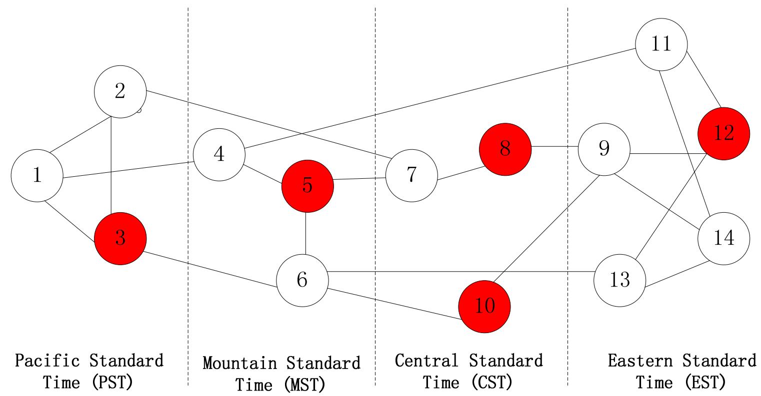

NSFNET Network

Peer Selection

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

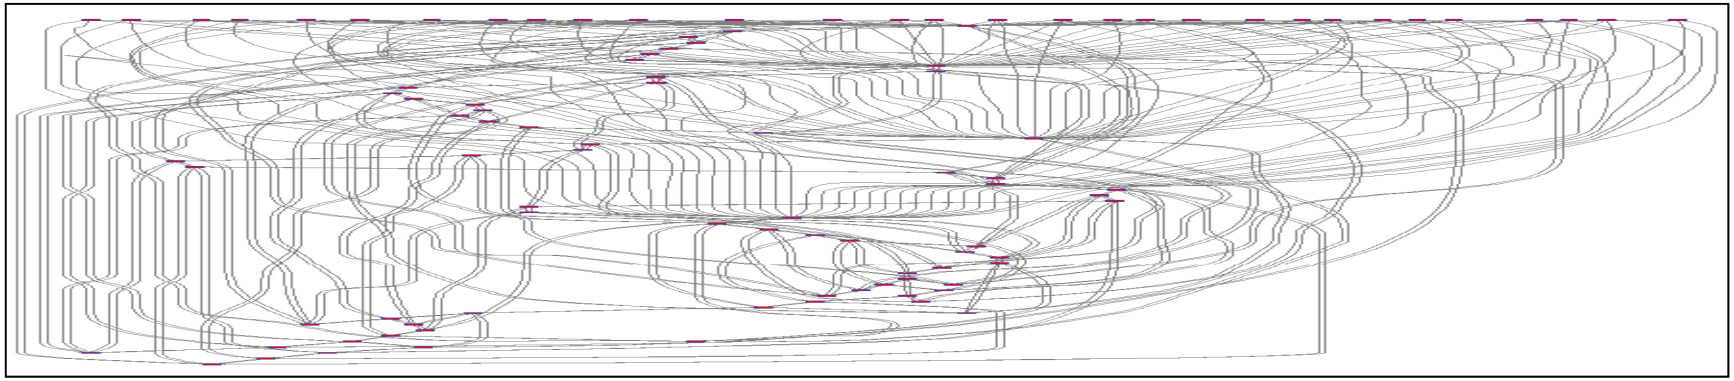

Original BitTorrent (Random Selection) β=0

Energy Efficient BitTorrent (Optimized Selection) β=1

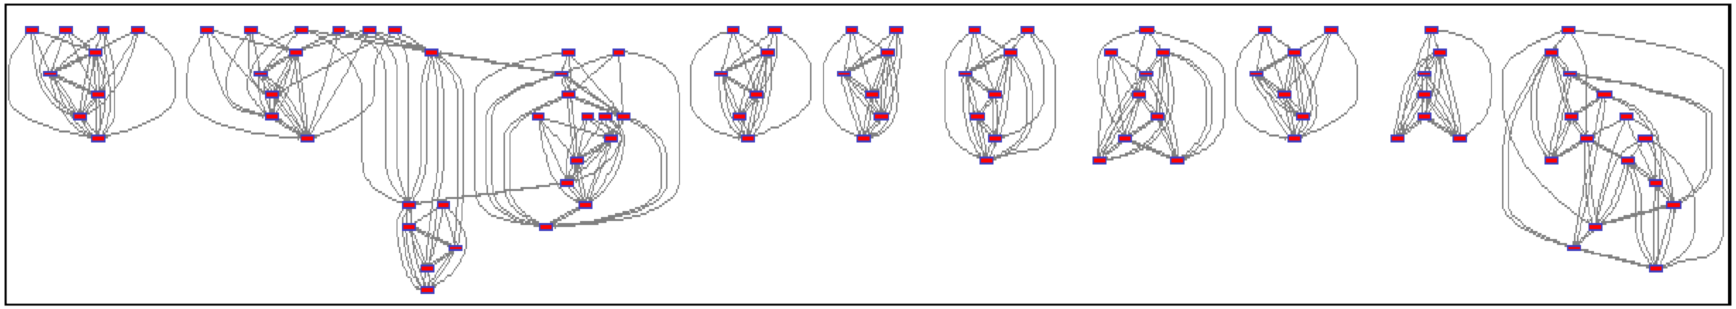

Average Download Rate

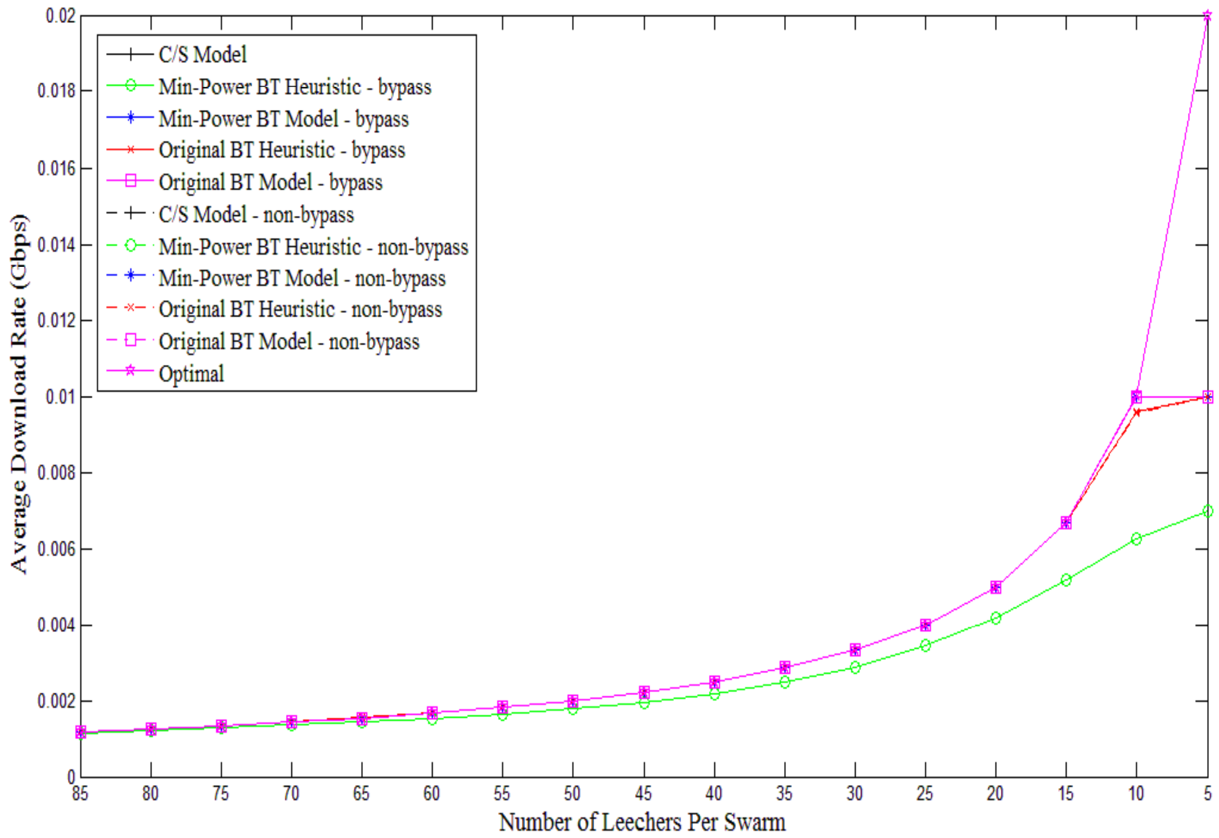

- All models reach optimal performance
- Energy-efficient heuristic reduce performance by 13%

Power Consumption

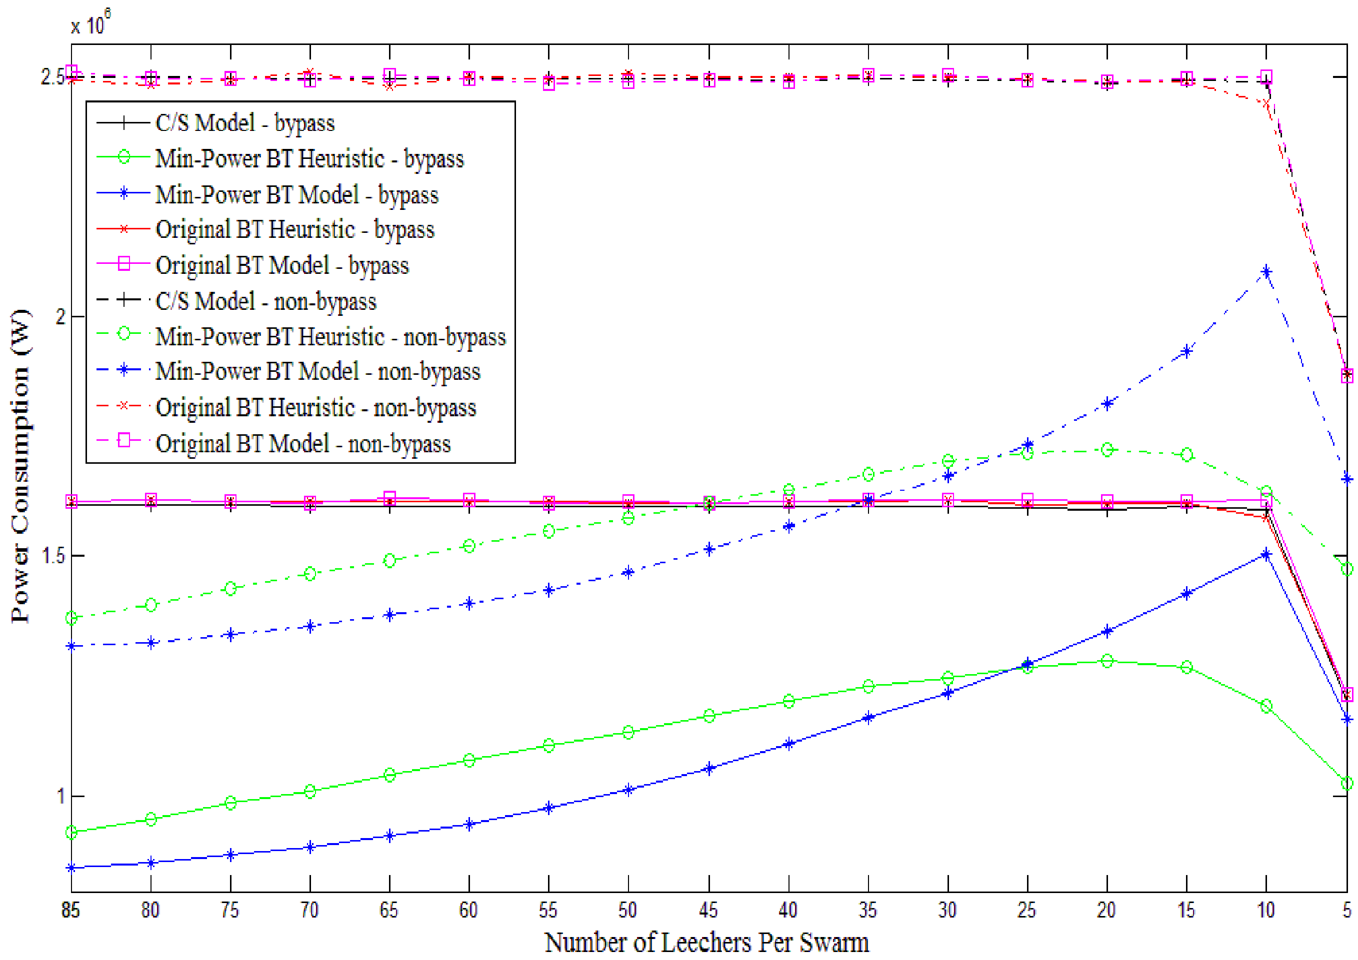

Non-bypass:
Model Power Saving=36%
Heuristic Power Saving =36%

Bypass:
Model Power Saving=30%
Heuristic Power Saving =28%

Energy Consumption

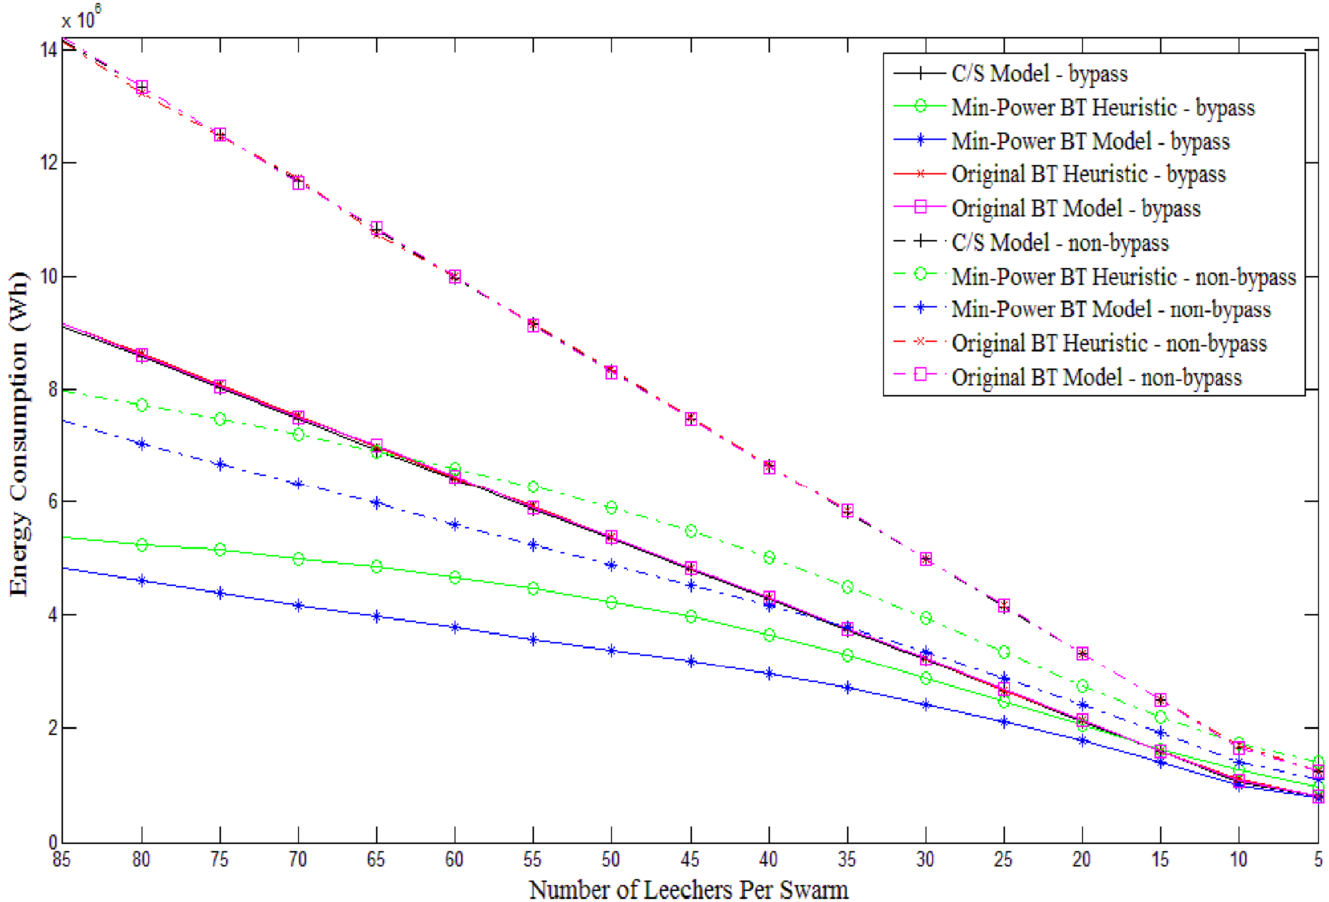

Non-bypass:

MILP average Energy Saving=36%

Heuristic average Energy Saving =25%

Bypass:

MILP average Energy Saving=30%

Heuristic average Energy Saving =15%

Future directions

- Dynamic load migration to match energy supply and demand; availability of renewable energy
- Traffic shaping to enhance dynamic resource adaptation and energy saving
- Hybrid P2P and C/S content distribution networks
- Cross layer resource adaptation (physical layer impairments, adaptive and mixed line rates; application awareness of physical and network layers)
- Dynamic resource adaptation in clean slate architectures (eg. time switched, subcarrier switched, time-subcarrier switching)