Intelligent energy aware networks

Y. Audzevich[1], A. Moore[1], A. Rice[1], R. Sohan[1], S. Timotheou[1], J. Crowcroft[1], A. Wonfor[2], H. Wang[2], R. Penty[2], I. White[2], X. Dong[3], T. El-Gorashi[3] and J. Elmirghani[3]

Computer Laboratory, University of Cambridge, UK[1]
Department of Engineering, University of Cambridge, UK[2]
School of Electronic and Electrical Engineering, University of Leeds, UK[3]

Summary

Today the energy consumption of Information and Communication Technology (ICT) industry is a significant contributor to the total energy demand in many developed countries. Recent studies show that the ICT industry is responsible for about 2% of the global emission of CO_2 and this percentage is predicted to increase as the Internet expands in bandwidth and reach. In this chapter we highlight different approaches for energy efficiency in communication networks. Firstly we review the techniques proposed to reduce the energy consumption of communication networks at the equipment and network levels. Secondly we investigate the use of renewable energy to reduce the CO_2 emission of IP over WDM networks. Issues including how to use renewable energy (solar in this work) more effectively, how to reduce the non-renewable energy consumption of transponders (the second most energy consuming device in a node), how to select the location of nodes using renewable energy, and load dependent energy consumption are considered. Thirdly we discuss workload migration using virtualization technologies in data centers as an approach of energy consumption minimization. Finally we consider some of the photonic systems advances which have the potential to reduce significantly the energy consumption within Ethernet switches and IP routers in the datacenter, showing how integrated photonic switch fabrics are starting to have the performance required for energy efficient high switching applications.

I. Introduction

Energy efficient processes are increasingly key priorities for ICT companies with attention being paid to both ecological and economic drivers. Although in some cases the use of ICT can be beneficial to the environment (for example by reducing journeys and introducing more efficient business processes), countries are becoming increasingly aware of the very large growth in energy consumption of telecommunications companies. In particular, the predicted future growth in the number of connected devices and the Internet bandwidth of an order of magnitude or two, is not practical if it leads to a corresponding growth in energy consumption. Regulations may therefore come soon, particularly if governments mandate increasing moves towards carbon neutrality, as has already begun to occur in the public sector in British Columbia and New Zealand. Indeed the UK has highlighted this as a priority [1]. This represents a significant departure from accepted practices where ICT services are provided to meet the growing demand, with little regard for the energy consequences of relative location of supply and demand. It also departs from existing attempts to constrain energy demand as it has been shown that increased equipment power efficiency leads to more consumption, the Khazzoom-Brookes postulate [2].

Internet power usage has continued to increase over the past decade due to (i) an ever increasing number of connected devices (ii) higher per device power consumption and (iii) growing device usage (per day) [3]. Drivers for future bandwidth expansion and power demand include streaming and video on demand, cloud computing and proposals to locate user computing and data within the network. New services such as IPTV will increase capacity demand in the backbone, meaning that the power consumption in the backbone will become more important relative to the access network. Internet traffic growth of between 40% and 300% is expected in the next 3 years [4], fuelled mainly by these broadband and video applications. In addition to the ecological impact, the economic impact is compounded by the increasing cost of electricity which is expected to rise by 200% over the next 7 to 8 years [4].

A report issued by the Ministry of Internal Affairs and Communications in Japan concluded that ICT equipment (routers, servers, PCs and network systems) consumed 4% of the total electricity generated in Japan in 2006, a figure of 45 TWh. Over the past five years the figure has grown by more than 20% [5]. Similar trends are observed in Europe. In Italy Telecom Italia is the second largest consumer of electricity (2TWh/year) after railways, consuming about 1% of total country demand growing by 8% and 12% with respect to 2005 and 2004 respectively [4], [6]. BT used 0.7% of the total UK energy in 2007 making it the largest single energy consumer in the nation [7]. In 2007, about 10% of the total UK power consumption was estimated to be related to ICT equipment [8] and ICT has been observed to produce more greenhouse gases than aviation [9]. In the US, ICT accounts for 3% of the total countrywide electrical energy consumption [10, 11].

The power consumption of the telecommunication equipment has rapidly grown over recent years. The typical router capacity increased from 100 Gb/s in 2000 to 10 Tb/s in 2008 with its power consumption increasing from 1.7 kW to 50 kW [12]. Furthermore, the rate of increase in router and telecom equipment power density is starting to accelerate and air cooling solutions are coming to the end of their capabilities with calls for the consideration of expensive liquid cooling solutions. ICT and telecommunications equipment continue to consume significant power even in the idle (not sleep) state. The power wastage is significant given that typical router utilisation figures are 20% - 50% in the Internet Service Provider (ISP) backbone, 8% - 25% in enterprise networks and less than 1% in LANs. In ICT equipment about 20% of the lifetime energy consumption is used in manufacturing and 80% during the use of the product in its life time [13]. Therefore it is important to reduce the network power consumption. It is worth observing that current estimates are that 37% of energy in ICT is due to transport of bits (telecommunication), while 63% is due to processing of bits (in data centres and telecom equipment) [13]. Therefore addressing the network operating power, both in transmission and signal processing are essential steps in reducing the overall carbon footprint of the Internet.

Efforts to reduce the environmental impact of ICT have mostly concentrated on improving the energy efficiency of individual ICT components or whole ICT systems. In communication networks, equipment level techniques are based on using power saving modes when the equipment is underutilized or improving the energy efficiency of hardware parts and operations. Additionally, network level techniques manipulate the energy profile of components by consolidating traffic, achieving energy savings of the system as a whole.

Substantial advances can be achieved through the innovative use of renewable sources and the development of new architectures, protocols, and algorithms operating on hardware, which will allow significant reductions in energy consumption. BT has announced plans to develop wind farms to generate 25% of its electricity needs by 2016. The BT wind farms will generate a total of 250 MW of electricity (about 2 MW per wind turbine), sufficient power for 122,000 homes or a city the size of Coventry. This will save 500,000 tones of CO_2 each year, the equivalent of a quarter of a million return air trips to New York. To ensure scalability, this problem has to be tackled at a national and European scale and with the activity also examining control plane mechanisms that can provide QoS guarantees in this novel renewable energy ICT scenario, this benefiting from new differentiated QoS and differentiated resilience schemes [14, 15] (where nodes with low renewable power are treated as failing nodes).

In data centers, the primary goal is usually to reduce the ratio of the facility to the IT equipment power usage, called Power Usage Effectiveness (PUE), by increasing the efficiency of the cooling and power distribution systems. Nevertheless, the PUE measures neither the energy efficiency of IT components' operation nor their level of utilization. The former is achieved by technological advancements at the equipment level, while the latter requires the consolidation of resources. Virtualization is a key tool for achieving this goal through the incorporation of multiple operating environments on the same machine and the migration of workloads inside or between data centers [16]. In this way, a significant amount of energy can be saved by putting a fraction of servers to sleep, while increasing the utilization of the active servers, as illustrated in Fig. 1.

Apart from improving the energy efficiency of ICT, its environmental impact can be further reduced by locating communication networks and data centers near clean energy generation sites, so that they can be supplied with both renewable and grid energy (Fig. 1). Supplying remotely generated renewable energy to ICT services rather than to grid workloads in the cities is beneficial for two main reasons. Firstly, renewable sources, such as photovoltaic and wind farms, occupy large areas and are usually built at appropriate remote locations; hence, shipping the generated energy to the cities incurs electricity transmission overheads. On the contrary, information transmission over long distances is highly efficient, allowing the collocation of communication nodes or data centers and renewable energy sites. Locating data centers away from the cities is also beneficial due to the availability of affordable land, low cost electricity and possibly free cooling [17]. Secondly, the intermittent nature of renewable sources combined with the inflexibility of grid workloads and the inability to efficiently store energy, result in substantial energy waste. In contrast to grid workloads, the agility of data potentially allows their processing to be delayed or performed at different geographical locations. For example, the maximization of renewable energy usage in a communication network can be achieved by appropriately routing traffic according to the expected utilization of each node, its energy profile and renewable energy availability. Similarly, data centers utilization can keep up with renewable energy availability by either scheduling non-interactive workloads at later times or migrating appropriate interactive workloads to data centers with abundance of renewable energy. For example, in Fig. 1 the reduction in the solar energy generation at data center A induces workload migration to data center C via a dedicated communication link.

Transporting data instead of energy imposes a number of challenges. Firstly, interactive workload migration has to abide by the Service Level Agreements, which may be violated due to the migration downtime. Therefore, it is essential to characterize workloads and develop migration prediction models and even employ dedicated high-speed low-latency networks to achieve transparent migration. Secondly, the relationship between the energy savings due to workload migration in data centers and the energy consumption of communication networks employed in the migration of resources needs to be established, in order to derive conclusions about the efficiency of the approach and the optimal frequency of re-optimization. Thirdly, the minimization of the nonrenewable energy consumption in accordance with the dynamically varying renewable energy generation and satisfaction of the migration and data center capacity requirements, is a hard optimization problem that needs to be dynamically solved in near real-time. Finally, efficient optimization can only be achieved by building accurate power consumption models of data centers, especially given the large gap between actual and theoretical energy consumption [18].

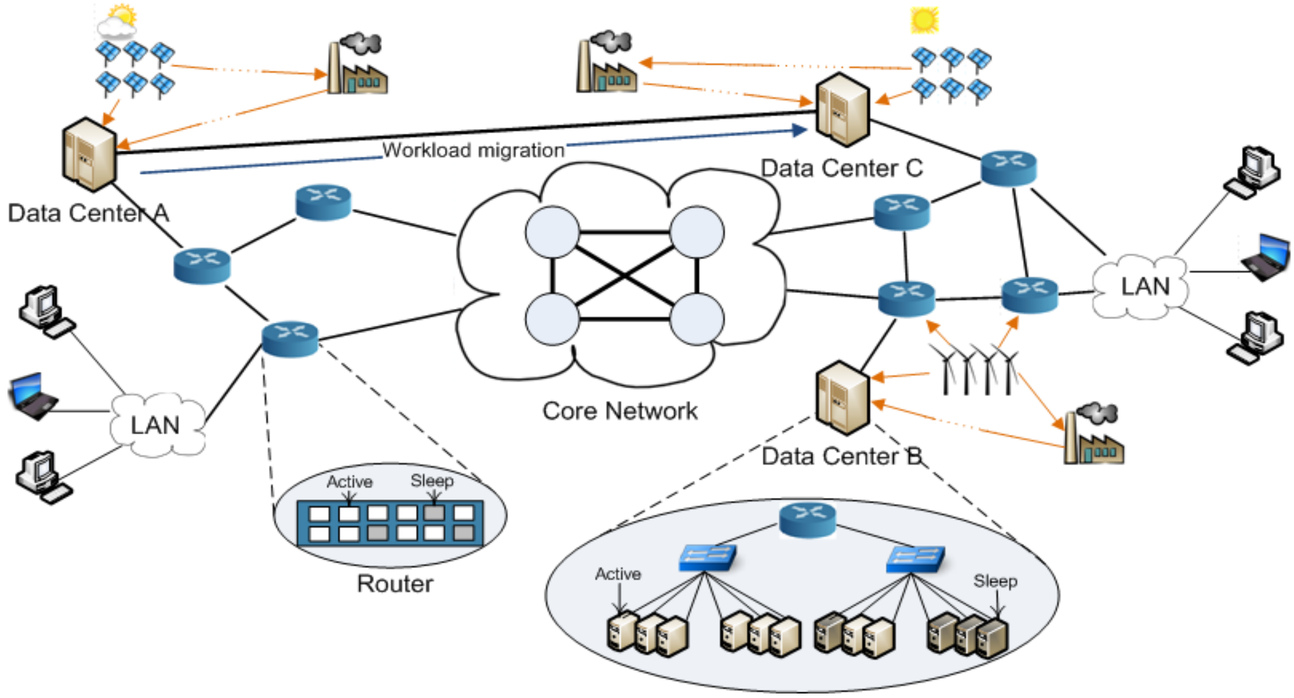

Fig. 1. Example illustrating different opportunities to reduce the energy consumption and environmental impact of ICT

The remainder of the paper is organized as follows: In Section II, we discuss in details the techniques that have been proposed to reduce the energy consumption of communication networks. Section III introduces the hybrid-power IP over WDM network architecture. A Linear Programming (LP) model is developed to minimize the non-renewable energy consumption of the network and a novel heuristic is proposed for improving renewable energy utilization. In order to identify the impact of the number and the location of nodes employing renewable energy on the non-renewable energy consumption of the network, we construct another LP model. We also investigate the additional energy savings that can be gained through Adaptive Link Rate (ALR) techniques where different load dependent energy consumption profiles are considered. In Section IV we expand the discussion of data centers and agile workload migration using virtualization technologies.

In section V we consider some of the photonic systems advances which have the potential to reduce significantly the energy consumption within Ethernet switches and IP routers in the datacenter, showing how integrated photonic switch fabrics are starting to have the performance required for energy efficient high switching applications. Finally, the chapter is concluded in Section VI.

II. Energy efficiency in communication networks

In this section we discuss techniques proposed in the literature for the reduction of energy in communication networks. Specifically, we classify these techniques into two main categories: (a) equipment level, and (b) network level. *Equipment level* techniques aim at reducing the power consumption of individual network devices (e.g. switches, routers) by adjusting the power mode of various components of the devices (e.g. network interfaces, line cards, memories, processors) based on locally collected information. On the contrary, the objective in *network level* techniques is to minimize the power consumption of a network using global state information including the topology, the state of links, traffic demands, QoS requirements and power consumption models of all equipment. This can be achieved either in a centralized manner by collecting all information at a single controller or in a distributed manner through node cooperation.

A. Equipment level

Driven by the idea of reducing power consumption in network components, the research community and industry have made first attempts in the investigation and realization of power-saving technologies. A great deal of work has been done in re-engineering conventional network equipment and network protocols. There are three main methods to reduce the energy consumption of e*quipment level* techniques based on local information. *The adaptive tuning* approach allows the reduction of energy by allocating the optimal amount of equipment resources with respect to the current demand. *Efficient operations* present another alternative to energy saving by re-engineering specific functions of individual components from a performance perspective. Finally, with *proxying techniques* operations of single or multiple network nodes are temporarily performed by a *proxy* while the nodes themselves are sleeping.

1) Adaptive tuning

Adaptive tuning considers a set of approaches that are designed to dynamically adjust characteristics of the network equipment in correspondence to the current service demands. 'Adaptation' is achieved by an extensive management of the onboard system's hardware and brought about by putting various components into a sleep mode or slowing down the hardware operation. These measures result in the reduction of the overall power consumption of the system and are beneficial for energy-conservation inside the network.

In particular, adaptive tuning for switches and routers has been intensively studied within the last decade. These utilize multiple processors, memories and interface cards for data handling and forwarding, which require high-energy expenditure. A number of power consumption solutions based on adaptive tuning of these components are presented in this section.

Sleep/shutdown approach

Recently, a great amount of attention has been devoted to the minimization of the power consumption in network equipment [19] [20]. Sleep/shutdown scheme, being a dominating approach explored in the literature, has been broadly applied in the network interface cards (NIC) of the desktop computers and switching architectures in Ethernet networks. The key principle of sleep/shutdown technique is to power off the system's network cards during the periods of network inactivity. Since the energy demand of an idle interface is almost identical to the one being fully loaded, shutting the device down during the periods of network inactivity can be power-beneficial.

The pioneering study on power minimization in network switches is presented by Gupta et al. in [21]. In particular, the authors did the shutdown of the ports and line cards based on a prediction algorithm, which uses the history of packet inter-arrival times. The maximal power savings of the scheme are achievable only when the packet arrivals to the system are infrequent and transition times between the power states are minimal. However, this is rarely achievable in real networks, where the packet arrivals can be spontaneous and, therefore, hardly predictable for the future. In fact, being in a sleeping mode, the interface cards are not capable of either detecting or receiving arriving traffic, which may introduce additional packet loss and/or delays to the system.

In order to avoid the drawbacks of the previous approach, the schemes with modified NIC's functionality were proposed in [22]-[30]. These schemes introduce additional hardware blocks that are powered on at all times and capable of buffering arriving packets when the system is sleeping. Packet loss minimization is achieved by the combination of different *traffic prediction algorithms*, which trigger a set of timers forcing the system to sleep [23], [24], [29], [30]and [31], *buffer policies* that set up occupancy thresholds to re-initialize the wake up state [23], [24], [26], [28] and [29] and *shadow ports* that cluster multiple ports of the network switches to handle ingress traffic [30]. The coordination of these techniques allows in theory achieving good balance between the performance characteristics of the NIC and its energy savings (40-60%) in underutilized networks (5-10%) [24].

A set of energy-conservation schemes for Passive Optical Networks was introduced in [32], [33] and [34]. Due to the fact that the Optical Line Terminal (OLT) plays the role of traffic distributor and arbitrator in a network, it is possible to switch the modes of Optical Network Units (ONU) from active to sleep, via extensive communication and modifications on the multiple access protocol deployed at upstream direction. To achieve the maximal energy savings under light traffic loads [34], the OLT is capable of storing arriving traffic for inactive ONUs and redistributing the traffic upon the end-system wake up (similar to [22]).

The majority of sleep/shutdown approaches presented at different levels of network equipment can achieve effectiveness in power saving only at low network utilizations. To provide reliable network service, the approaches explored use additional blocks that are powered on continuously to avoid possible losses. With the increase of the traffic load the energy conservation becomes insufficient due to frequent transitions between the power states. Frequent state transitions, in their turn, cause undesirable in-rush currents that can significantly degrade the energy saving of a scheme (due to energy spikes) as well as damage the circuitry [25]. Additionally, the state transition times between the "power on" and "power off" states may be long enough for a packet loss to appear, which will result in retransmissions from the network protocol side. All these factors should be taken into account in practical realization of power-aware network interfaces.

Adaptive link rate

The Adaptive Link Rate (ALR) approach [35] is mostly studied and applied in Ethernet technologies. Contemporary Ethernet provides up to 10Gbps link rate, which consumes a considerable amount of energy irrespective of traffic demand. The Adaptive Link Rate scheme dynamically tunes the link's data rate in correspondence to the varying traffic loads and, as a result, minimizes the energy use.

The Energy Efficient Ethernet (EEE) task force [36] introduced enhancement to the traditional standard by suggesting three main link rates (10Gbps, 1Gbps and 100Mbps) for operation. The rate switching is performed by Rapid PHY Selection (RPS) algorithm that uses a sophisticated frame exchange policy for the negotiation of link speeds. In practice, tuning the transmission rate may require considerable amount of time due to adjustments performed in transmission and synchronization circuits of the Ethernet card. The solution to this issue is a challenge that faces the practical implementation of the ALR approach, since large transition times may introduce performance degradation of the system.

A set of theoretical studies shows the application of the ALR technique in networks with high utilization. The authors in [22], [37] proposed to shape arriving traffic into bursts and, while considering coordination of both ALR and sleep/shutdown approaches, to achieve energy-proportional power consumption [38]. Avoidance of packet loss and packet delay was achieved by over-dimensioning of interface buffers based on the idea that bursts are assembled only at the last-hop of network switches. To address the problem of frequent switching between the transmission rates, the authors of [39]-[43] suggested control of the buffer occupancy of an interface card, while introducing various threshold policies. The thresholds controlling the link utilization were proposed in [42], [43]. The key principle of the approach is to stay at the fixed lower rates (10Mbps or 100Mbps) until the network utilization falls below a certain value, for instance 70%. The implementation of the mentioned threshold policies allows preventing link rate oscillations both for low and high traffic loads as well as for various traffic types [40].

The practical (FPGA-based) implementation of the ALR shown in [44] emphasizes the fact that real physical layer selection times can be almost two orders of magnitude higher than it is proposed for EEE (1ms). The results show that the power consumed by the device is strongly dependent on the transition type and the direction of transition. For instance, the power consumption of an interface, which changes the rate directly from 1Gbps to 10Mbps, is 4 times larger than in a step-by-step transition. In a similar way, the authors show the asymmetric power needs for 1Gbps to 10Mbps (which requires higher energy) and 10Mbps to 1Gbps transitions.

The sleep/shutdown scheme presented in the previous section proved to be energy-efficient only at low network loads. The ALR scheme, in contrast, can be energy beneficial also in networks with high utilization, provided that the rate adaptation time is sufficient to maintain the network services. The latest advances in Ethernet technology show that the standards with higher data rates implement incompatible functionalities to the older designs. Therefore the realization of the ALR approach in practice will imply the complex interaction between the incompatible layers (physical layers) of a device and can be a limiting factor for Quality of Service (QoS) guarantees.

Dynamic Frequency and Voltage Scaling

Dynamic Frequency and Voltage Scaling (DVFS) [45], [46] is a technique used for the power-management of integrated circuits. DVFS exploits the fact that the switching power in a chip decreases in proportion to $CV^2f + P_{stat}$, where V is a core voltage applied, f denotes the clock frequency and P_{stat} is a static leakage power. According to [45], the main contribution to the power consumption in a circuit is attributed to the first term of the formula, so the second term can be usually skipped.

The DVFS scheme proved to be extremely energy-efficient in traditional high-speed CPUs, where the reduction in operating frequency allows a decrease in the supplying voltage level [45], [46]. Modern processors do not show significant power reduction with DVFS due to a set of new advances introduced to the processor's architecture [47]. Since the introduction of DVFS, the transistors' size has been considerably reduced to tens of nanometers. In relation to this, the static leakage power inside the chip has become almost identical to the switching power. The impact of DVFS is further reduced by a small switching power range (transistors being smaller require a lower core voltage to be applied) and more efficient sleep states. The DVFS scheme can be widely applied in the components of network switching and routing architectures, reducing the power consumption of the system during low traffic demands.

2) Efficient Operations

Efficient operations consider a set of approaches that implicitly impact the power consumption of components by improving the performance of network related operations. For instance, line coding can result in energy savings if there is a small transmission overhead and processing requirements. A second example is related to IP lookup in routers. In this case, energy savings are achieved by applying traffic classification techniques that allow the minimization of the computational and storage requirements of IP tables, resulting in the implicit reduction of energy.

Efficient Encoding

The traditional channel coding techniques implemented in current Ethernet standards consider delivery of the prescribed messages through the network with the maximum possible reliability. Unfortunately, the sophisticated implementation of such codes usually requires a considerable amount of energy for the interface operation. The primary investigation on the power consumption of various encoding schemes was presented in [48]. The authors suggest that the encoding circuitry makes the main contribution to the power usage inside the system. While examining the most popular line coding schemes the authors conclude that simplified versions of the codes can be more energy-beneficial but are not always reliable.

Although energy efficient Ethernet encoding schemes are not well explored, this direction has high potential for power conservation at the network equipment level. The Ethernet encoding schemes proposed currently tend to ameliorate the bit-level clock recovery at the receiver side and increase the chances of error detection. On the other hand, the recent implementations of the 10Gbps Ethernet suffer from physical layer compatibility issues with respect to the previous versions of the standard. Therefore, the implementation of advanced encoding schemes should balance the energy consumption of the circuit, the reliability of data transmission and performance. Efficient implementation of the encoding scheme can improve the performance of some of the approaches presented above. For instance, the encoding scheme can complement the Adaptive Link Rate approach reducing the switching time between the transmission rates. In a similar manner the encoding schemes can provide robustness of data transmission, taking advantages of the photonic networks, while dispensing with the latencies of the transport layer protocols (like TCP). Additionally, the redesign and refinement of the physical line codes can promote changes in error detection and correction techniques at the receiver side. While currently the end-hosts utilize the continuous ARQ-based transport protocols to provide correctness in data delivery, Fountain/Raptor codes can provide data recovery without causing network retransmissions. These techniques form a good basis for low latency packet transmissions.

Efficient IP lookup

In high-end routers the speed of IP routing becomes a major performance degradation point that limits the increase in link rates. The traditional IP lookup engines use Ternary Content Addressable Memories (TCAMs) that implement a fully parallel search during a single clock cycle. Unlike Random Access Memories (RAM), TCAMs have more sophisticated onboard functionality featuring comparison circuits for each memory bit. This reduces their scalability in terms of power, storage and density [49],

[50]. A number of studies deal with these shortcomings. Some of them demonstrate that the reduction in the power consumption may be achieved by modification of the TCAM structures. The others utilize alternative low-power memory solutions, like Static Random Access Memories (SRAMs).

A set of modifications to the TCAM's structure was proposed in [51], [52]. TCAM is divided into multiple blocks by distributing equally the number of packet classification rules between each block. Upon a packet arrival, only a set of blocks is searched, while the remaining blocks are disabled.

In contrast, the IPStash approach [53] implements the full associativity of TCAMs with "limited-associativity" of Static Random Access Memories (SRAMs). This provides good performance and lower power consumption. Unfortunately, this approach suffers from extensive power consumption, when the routing table is large.

The implementation of SRAM-based lookup engines in the form of a trie can be used as an alternative [54]. Ties can be further classified into uni-bit tries and multi-bit tries [55],[56]. The packet classification is performed by traversing the trie levels, where nodes represent packet prefixes. Uni-bit tries make correspondence of only one bit to a node and therefore, suffer from frequent memory accesses. On the other hand, multi-bit tries, assigning multiple bits to a single node, may have larger memory occupied due to redundancy.

The large number of prefixes (hundreds of thousands) in a trie may require a large amount of memory to be allocated. In addition, the packet classifications that are performed simultaneously can lead to an extremely large memory access time. The problem is usually resolved by memory pipelining in order to perform multiple lookups in parallel [57]-[59]. In this case, memory arrays are subdivided into sub-arrays, with a possibility to access a specific sub-array directly. The lookup table trie entries can be assigned to the memory sub-arrays in a different way with unbalanced or balanced trie nodes distribution. The unbalanced distribution can degrade the access time, while the balancing approach can influence the system's throughput [58]-[60]. In contrast, the combination of SRAM lookup with external/internal caching can be beneficial in packet classification due to a reduced number of memory accesses for similar traffic arrivals [50], [61].

3) *Proxying*

Currently, desktop computers receive a large amount of network management information, even when they are idle, which prevents them from sleeping. This shortcoming is overcome with the introduction of energy-efficient *proxying* which is capable of managing the arriving network traffic on behalf of sleeping computers and maintaining network presence. Being competent in making traffic classification, the proxy decides the appropriate time for the system to wake up [62]-[64]. Proxying is implemented either locally at the computer (*internal proxying*) or globally within a LAN (*external proxying*).

The onboard proxying adapter adds extra functionality to the NIC and is usually implemented by extra hardware blocks (like microprocessors). The controller implementation presented in [61], [65] shields and processes ARP, ICMP and DHCP requests without any interaction with the rest of the system. On the other hand, due to the limited implementation of the protocol stack TCP segments are processed by the computer and require the system's wake-up. The packet inspection system implemented in [65] performs hardware-based content processing of packets by utilizing partitioned TCAMs and cache memories, which apart from energy-savings improves also the performance characteristics of the system. The application support of the power-proxy was implemented in [66], [67]. The combination of secondary processor and flash memory embedded in the computer's NIC runs the full functionality of the network protocol stack plus a set of applications implemented in the forms of "stubs" (reduced functionality). These allow the support of more levels of functionality and save a considerable amount of energy during the infrequent traffic arrivals.

As the size of a local area network increases, deploying internal proxies at each computer is inefficient due to the complexity and cost of hardware. On the other hand, implementing external proxying in each subnet is capable of managing the arriving traffic, sleeping intervals and transmission rate for a large amount of computers simultaneously [66], [68]-[70]. Two notification routines should be implemented to perform power-management of a network: (a) a sleep management program implemented at the proxy, and (b) a sleep notification program that runs at each client. The sleep management program informs the end-hosts about sleeping possibilities and shows the network presence of sleeping devices for the outer networks. The approaches presented in [68], [69] are capable of responding to some packets on the client's behalf or waking up the sleeping client for certain specified traffic (by means of Wake-on-LAN packets [64]). The possibility of handling user-destined traffic presents a tradeoff between the implementation complexity of the proxy and the amount of energy saved. The SleepServer approach suggests implementing virtual machine instances of the hosts at a proxying machine that provides network presence of devices and efficiently uses the protocol stack deployed at virtual machines [74]. This approach yields significant network power-savings and keeps the connectivity of devices while being in sleeping mode.

4) Problems and Challenges

A great deal of studies has been done to improve the power efficiency of communication networks at various component levels. These studies have contributed to the development of power saving solutions in simulated environments, however their practical application in existing networks is still challenging for two main reasons. Firstly, power mode adaptation is not supported by current hardware, and secondly compatibility has to be maintained with current standards.

Implementing energy efficient and highly reliable encoding schemes is also challenging. Reliability requires the realization of highly complex encoding schemes which consume a large portion of energy for encoding/decoding with respect to the actual transmission energy. Thus, designing highly robust line codes that require low processing are needed. Having robust codes also improves the energy efficiency of reliable transport protocols (e.g. TCP/IP), as the number of required retransmissions is minimized.

According to the discussion in the previous sections, energy-aware equipment may have a large response time and energy consumption during power state transitions. These may result in performance degradation in terms of extensive packet loss and delay. Additionally, it is important not only to minimize the overall transmission energy but also consider the overhead energy due to state transitions. In parallel, due attention is required in building hardware components that efficiently deal with the transition time and energy.

Finally, one of the biggest challenges is to achieve proportionality between energy consumption and utilization of components [38]. Achieving proportionally will not only be beneficial in reducing the energy consumption of network devices, but it will also eliminate the need for sophisticated energy saving approaches.

B. Network level

In the previous section we have discussed techniques that can be used for the energy-efficient operation of various network components such as network interfaces, switches and routers, based on locally collected information and without any node coordination. However, equipment level techniques are not sufficient to guarantee the minimization of the Internet energy consumption. Currently deployed networks are over-provisioned to accommodate more than the maximum expected traffic demand, and over-redundant to deal with link and node failures. As a result, many links are underutilized and all devices are constantly in operation, which provides opportunities for energy reduction. This can be accomplished by disseminating the traffic in a way that minimizes the network's energy consumption by putting specific nodes and links to a power saving mode. Network level optimization requires cooperation between the nodes to collect information about the global network state, which includes the capabilities and energy consumption of all network elements, as well as full information about the traffic demand between nodes and the network topology. These techniques can be employed both at the design stage (*network design*) and during its operation either periodically (*traffic engineering*) or in real-time (*distributed routing*). In these ways, we can achieve energy-efficiency and meet network operation requirements such as user request satisfaction, quality of service (QoS) and reliability. In this section, we discuss research undertaken in these categories in both the electrical (e.g. IP/MPLS routers) layer, as well as multilayer approaches where nodes are comprised of electrical components on top of optical components (e.g. OXCs, ROADMs).

1) Network design

Traditionally, network design has focused on the minimization of the network capital expenditure (CapEx), which accounts for the equipment and installation costs of the network infrastructure [TE4]. However, as Internet traffic and energy costs are exponentially rising the energy consumption is becoming a major issue for network operators for several reasons: (a) the operational expenditure (OpEx) due to the energy cost is significant, (b) ICT is an important contributor to the global energy consumption and it is essential to cut down CO_2 emissions[1], and (c) energy efficient operations alleviate the heat dissipation problem [72]. As a result, energy minimization has become a main goal in designing communication networks.

Before reviewing the research undertaken in this area, we briefly describe the problem examined. The main goal of a network design problem is to find the design parameters that optimize an objective function associated with the minimization of CapEx (or energy consumption), while satisfying the provisioned network traffic demand and other design specifications (e.g. maximum link utilization, end-to-end delay, reliability). Traffic demand information is usually given in the form of a traffic matrix which represents the peak demands between each source/destination node pair [73]. The topology of the network can either be known, or the PoPs (Points of Presence) can be chosen from a set of candidate locations. Candidate equipment information are also available such as different router models that can be used, chassis specifications, number and types of line cards supported by a chassis, as well as capacity, cost and power consumption information of router components. Hence, the goal is to find the type and quantity of equipment to be installed at each location, in order to satisfy the design specifications and minimize the objective function. Mathematical formulation of the network design problem usually results in Mixed Integer Linear Programs (MILP) which fall into the class of multi-

[1] The UK government is already operating a Carbon Reduction Commitment (CRC) mandatory scheme for large energy consumers

commodity flow problems and are NP-hard.

In the electrical layer, the problem of energy-minimized network design has been investigated in [74], [75]. Chabarek et al. [74] assume a known topology and formulate the problem to minimize the total power consumption of the network when the type/number of line cards and chassis can be selected. A general model of power consumption for a router is proposed which considers the power consumption of chassis and line cards in base configuration and the extra power consumption due to the router traffic utilization, while an empirical measurement study is conducted to obtain values for these costs. Results on realistic and random topologies of small size show power reduction of up to 65%. In [75], the relationship between network power minimization and robustness is investigated using an objective function that combines power minimization with mean network delay. Network robustness is assessed with respect to survivability measures, by simulating link failures on topologies derived from the solution of the network design problem for different weighted versions of the two objective metrics. The results indicate that a good trade-off between the two metrics can lead to better network designs.

In IP-over-WDM networks, energy-minimized network design has been introduced in [72]. The authors assume that the topology is known and formulate the problem by encompassing the energy consumed by the node ports, the EDFA amplifiers and the transponders. In this case, the optimal design parameters of the physical layer (e.g. number of fibers and wavelengths on each physical link) should also be found. Additionally, the problem is very challenging as the flow conservation and capacity constraints need to be simultaneously satisfied at the IP and optical layers, resulting in an NP-hard model that has $O(N^4)$ integer variables and $O(N^3)$ constraints, where N is the number of nodes. As a result, the optimal solution can be obtained in reasonable time only for small problems, while for larger problems the authors proposed heuristics techniques. Their results on various real topologies indicated that (a) lightpath bypass is better than non-bypass (25% - 45% less energy consumption), (b) the network consumption is mostly at the IP routers (more than 90%), and (c) CAPEX minimized design is also energy efficient. The latter result was also supported by other independent studies on different networks [76], [77]. Nevertheless, if the energy consumption of the nodes is assumed to be mostly dependent on its load, rather than the number of active elements, then cost efficient design is not also energy efficient [77]. Moreover, the use of diversified-capacity lightpaths and the architecture of the node (e.g. IP-over-WDM or IP-over-OTN-over-WDM) result in energy savings [76], [77] as well. The effect of the node architecture in network design is also considered in [78] for a ring network topology. It is shown that depending on the energy consumption ratio between the electrical and optical layer of one node, different node architectures can be more energy efficient.

2) Traffic Engineering

The objective of traffic engineering (TE) is the control of network traffic flows, in order to optimize resource utilization and network performance under specific QoS requirements [79]. Traditionally, TE has focused on balancing the traffic among links in order to avoid congestion due to traffic bursts. On the contrary, the main goal of power-aware traffic engineering (PATE) is to reduce the energy consumption of the network by consolidating traffic to a few links so as to put idle network components into a power-saving mode. Therefore, performing PATE is an even more challenging problem as energy saving has to be achieved on top of any QoS requirements. As will be discussed in this section, significant energy savings can be achieved due to the over-provisioning and high-redundancy of the network, as well as by periodically performing PATE to exploit periods of low traffic activity.

Modelling the PATE problem usually relies on multi-commodity flow MILP techniques, similar to energy efficient network design. In fact, the two problems have similar mathematical formulations, differing mainly on the decision variables. In the network design problem, the emphasis is placed on selecting the type and amount of equipment to be placed at each PoP; in the PATE problem the decision variables are related to the power mode that has to be adopted for each network component, in order to minimize the total power consumption of the network for the provisioned traffic demand at the considered time period.

Several approaches have been adopted for modelling the power consumption of network components which differ in the type and energy profile of subcomponents considered at each node. In [80], it is assumed that nodes cannot be put to sleep because they are either source or destination nodes; hence, the power consumption metric considered is the number of sleeping links. In [81], the authors proposed to minimise the number of active cables rather than the number of links, by highlighting that each logical link between two nodes is actually a bundle of cables that provide higher availability and linear capacity increase. A common approach adopted in the literature is based on the observation that the power consumption of current internet routers is not significantly affected by the load [74]. As a result, an *on/off model* is adopted that accumulates the power consumption of the active components of each node. Different on/off models have varying degree of detail; in [82] and [83] the consumption of nodes and links is taken into account, in [84] the base consumption of line-cards and their integrated ports is considered, while in [85] the power consumption model includes the contribution from chassis, line cards and ports at different line speeds. An important disadvantage of the on/off model is that it does consider the *energy profile* of a component i.e. the dependency between its power consumption and its utilisation; although this dependency is weak for current communication equipment, future energy-aware components are expected to consume energy proportional to their utilisation [38]. In [86] and [87], a simple energy-profile model is considered where each component has a base consumption (when in idle state) that increases linearly with its load utilisation. Finally, the effect of various energy profiles on the energy consumption of a network is considered in [88].

As already mentioned, PATE targets the consolidation of traffic to a few links and nodes so that many network elements can sleep. Nevertheless, traffic aggregation may create violation of QoS requirements, and for this reason formulations incorporate QoS constraints. To put a threshold on the average link queuing delay, the approach usually taken is to impose an upper bound on the link utilisation which is a fraction of the nominal link capacity (50%-70%) [82]-[85], [87], [88]. Other measures to reduce the overall delay, include the consideration of paths that are smaller than a certain length [84], or the average delay in the network [86]. Load balancing has also been suggested to be performed on top of the PATE solution to provide an even distribution of the traffic among the non-asleep network components [84]. Because PATE routing removes most of the redundant paths, examining the availability of the network is also important. Nevertheless, none of the PATE approaches have incorporated the particular QoS metric into its mathematical formulation; network availability has only been examined in derived simulation results [85], [80].

Because different PATE formulations are NP-hard [82], [80], a number of heuristic approaches have been proposed for obtaining fast and close to optimal solutions. Pruning heuristics have been proposed in [82], [83] and [80]. The idea is that network elements (e.g. nodes and links), that are considered for switching off, are initially ordered according to a desirability criterion such as their power consumption, aggregate traffic or ratio of total to remaining capacity; then, sequentially these elements are orderly switched off, given that the network traffic can be routed with the remaining elements. On the contrary, the heuristics proposed in [81] start from the integer solution of the relaxed MILP problem, and in the process re-compute the desirability criterion in every iteration. These heuristic approaches are of polynomial time complexity that result in fast and relatively accurate solution of the problem. In [84], a non-polynomial heuristic is proposed, which is based on constructing a path-flow formulation of the considered PATE problem for the k-shortest paths for each source destination pair, instead of all the paths. Although the resulting formulation is still MILP, the number of integer variables is considerably smaller so that larger problems can be solved with very high accuracy.

Although an important amount of research has been undertaken in PATE-based routing, little attention has been paid in discussing implementation details using existing protocols and mechanisms. To the best of our knowledge, only Zhang et al. investigated this matter in detail [84]. Specifically, the use of a centralised controller node is proposed that is responsible for the aggregation of information, as well as for the computation and dissemination of the solution, under the assumption that the network runs both OSPF and MPLS protocols. The authors describe how the controller can passively collect all necessary information, and distribute the periodic solution of the problem using TE variations of the

OSPF and MPLS protocols. The implementation of a state transition between on and off and the traffic splitting in flows are also discussed.

A number of interesting results about the energy saving that can be achieved via PATE routing have emerged. However, before starting the discussion, we should emphasize that these results are indicatory, due to the lack of practical implementation in all studies. Additionally, these results are based on assumptions about numerous factors such as the network topology, the traffic demand, the energy model and power consumption values of the network components. This is also indicated by the fact that the energy reduction achieved for different studies ranges between 10% to 70% for realistic network topologies and traffic matrices [87], [85]. The first main result is that the yearly energy reduction that can be achieved for the real topology and traffic data of a specific ISP is around 23%, when an on/off power model is adopted [83]. Another important result is that future energy-aware proportional network devices can result in an order of magnitude energy reduction [88], [87]. Key results can also be inferred about the tradeoff between energy minimisation and QoS satisfaction. Several studies have shown that for low and moderate traffic volumes, energy reduction can be achieved without significantly affecting delay related QoS metrics (e.g. maximum link utilisation, route length increase) [80], [86], [84]. However, the availability of the network, in terms of the average number of disjoint paths between pairs of nodes, drops very quickly to its minimum value, which means that our network is prune to failures [85], [80]. Thus, it is essential to develop network equipment that can switch their power mode very fast. In terms of the solution stability, it has been illustrated in [84] that no abrupt changes occur between successive PATE solution computations.

3) Distributed online energy-aware routing

The PATE approaches discussed in the previous section rely on global information about the system state, including the traffic demand matrix that is not readily available, and require the solution of NP-hard combinatorial optimisation problems; hence, they are centralised and cannot be used for the online dynamic traffic rerouting. In this section we discuss techniques that only rely on information readily available at the servers and can be used for online distributed rerouting of traffic that achieves energy savings.

EATe is the first scalable online technique that has been proposed for energy saving; it is based on adaptive link rate, as well as node and link sleeping [89]. The authors describe techniques that can handle two different power consumption models: (i) when the ratio of idle to maximum component power is low, the authors propose to remove traffic from under-utilised links, in order to switch to a lower link rate, and redistribute it to links that are not underutilised but can support more traffic, and (ii) when the ratio of idle to maximum component power is high, the authors propose to switch off as many links and nodes as possible by redistributing the traffic of under-utilised links. Stability is achieved by explicit feedback from the nodes receiving the redistributed traffic so that only feasible changes are accepted.

Green OSPF protocol relies on aggregating traffic at a neighbourhood of routers in order to put links to sleep [89]. In this protocol there are two sets of routers. “Exporter routers” compute their Shortest Path Tree (SPT) normally, while “importer routers” take as reference a slightly modified version of the SPT of “Exporter router”, rather than computing their own. In this way, a number of routers share the same SPT so that several idle links can sleep. To avoid inconsistencies in routing, a final phase is required in which each router identifies the new topology (thought received LSAs) and re-optimises its routing paths. Nevertheless, this protocol does not address how link overloading can be avoided.

Contrary to green OSPF, GDRP-PS (General Distributed Routing Protocol for Power Saving) considers the possibility of switching of nodes during off-peak hours [90]. In this protocol, node sleeping is coordinated by a central node so that not many nodes sleep concurrently; nonetheless, a node’s decision to ask for sleeping permission is based only on its aggregate link utilisation. Periodically the nodes wake up, reconnect to the network and examine whether they should go back to sleep. Performance evaluation indicates that 20% power saving can be achieved; however, no simulation results or analysis is presented to support whether the network stability is maintained.

4) Problem and challenges

Despite the large number of undertaken network level studies, a number of issues need to be further explored. As already mentioned, network design and traffic engineering problems are NP-hard and researchers often employ heuristic solution techniques. Although a number of different approaches have been proposed in this area, no comparative studies have been performed, while performance evaluation of individual studies is not compared against other similar approaches. For this reason, it is important to build a library of representative test instances for specific problem variations so that benchmarking will be possible. An example of such a library is SNDlib 1.0 for survivable fixed telecommunication in traditional network design [91].

Additionally, the available literature lacks practically implemented approaches; most studies do not even discuss any deployment details about how to collect the require information, when and how often to solve the TE problem, how to disseminate information and how to implement the sleeping of nodes in practice. Furthermore, as TE approaches are performed offline, there is a need for more distributed online routing techniques; so far online approaches have focused on providing intuitive solutions, which provide no optimal results or rigorous guarantees for stability and QoS satisfaction.

Current routing protocols such as OSPF and IS-IS are designed to handle a small number of failures and do not support the simultaneous shutting of a large number of links and nodes, which is the case with energy aware techniques [83]. As a result, modified versions of these protocols need to be designed that will be capable of dynamically adjusting to the constantly changing topology. Moreover, the PATE approaches proposed so far are only suitable for intra-domains, as ISPs are unwilling to share information between them. Consequently, it is important to study inter-domain protocols where energy efficient routing is performed among ISP providers. To this direction, ISPs need to be given incentives to cooperate e.g. by guaranteeing fairly distributed energy consumption and by not requiring exchange of sensitive information.

Finally, it is imperative to investigate the interdependency between equipment level and network-level energy-efficient approaches in order to avoid potential problems. For example, if a PATE technique redirects a large volume of traffic on a link that is running on a lower data rate than usual, this will cause significant delay or packet loss. Additionally, the incorporation of PATE techniques into communication networks will alter the statistical characteristics of traffic, making the efficiency of equipment level techniques doubtful.

III. IP over WDM networks employing renewable energy

An IP over WDM network is composed of two layers, the IP layer and the optical layer. In the IP layer, IP routers aggregate traffic from access networks. IP routers are connected to optical switches in the optical layer. The optical layer provides large capacity and wide bandwidth for the communication between IP routers. Optical switches are connected to optical fiber links where a pair of wavelength multiplexers/ demultiplexers is used to multiplex/demultiplex wavelengths on each fiber. Transponders are used to provide OEO processing for full wavelength conversion at each switching node. For long distance transmission, the EDFAs are used to amplify the optical signal on each fibre. IP over WDM networks are implemented by either lightpath *non-bypass* or *bypass*. Under the lightpath non-bypass, all the lightpaths are terminated, processed and forwarded by IP routers in all intermediate nodes. On the other hand, under the lightpath bypass all lightpaths are directly bypassed through intermediate nodes, i.e. lightpaths are treated as virtual links in the IP layer. Lightpath bypass can significantly save the total number of IP router ports required and as IP routers are the major energy consumption components in an IP over WDM network, minimizing the number of IP router ports can potentially minimize the energy consumption of IP over WDM networks. In [72], the multi-hop bypass heuristic was proposed where the bandwidth utilization is improved by allowing traffic demands between different source-destination pairs to share capacity on common virtual links (lightpaths). Improving the wavelength bandwidth utilization results in fewer virtual links, and therefore, fewer IP router ports and lower energy consumption.

As mentioned in Section I introducing the use of renewable energy in ICT networks will allow significant savings in energy consumption and CO_2 emissions. However, many challenges need to be addressed to develop and deploy renewable energy in ICT networks. In this section, we focus on reducing the CO_2 emissions of backbone IP over WDM networks by introducing renewable energy sources to the network. A hybrid-power IP over WDM network architecture is proposed where the power supply is mixed being composed of non-renewable energy and renewable energy. In this case, the total CO_2 emission of an IP over WDM network will be reduced if a portion of the non-renewable energy consumption is replaced by renewable energy. Therefore, the problem becomes to minimizing the non-renewable energy consumption of the hybrid-power IP over WDM network. A LP optimization model and a new heuristic are set up to minimize the non-renewable energy consumption.

A. A LP model for Hybrid-power IP over WDM Networks

The lightpath bypass was implemented in [72] to reduce energy consumption of the IP over WDM network by reducing the number of required IP router ports. A Mixed Integer Linear Programming (MILP) optimization model to minimize energy consumption was developed. In this section we develop a model focusing on minimizing the non-renewable energy consumption in the hybrid-power IP over WDM network. In this LP model, we assume the network has the topology $G = (N, E)$ with N nodes and E physical links. The nodes that have access to renewable energy can also be powered by non-renewable energy to guarantee QoS when the renewable energy output becomes low. The renewable energy can power the ports, transponders, optical switches, multiplexers and demultiplexers in a node. Fig. 2 shows the hybrid-power IP over WDM network. The total non-renewable energy consumption of the network is composed of:

1) Non-renewable energy consumption of ports without access to renewable energy

$$\sum_{i \in N} PR \cdot \left(Q_i^e + \sum_{p \in P} \delta_{ip} \cdot W_p \right)$$

2) The non-renewable energy consumption of EDFAs

$$\sum_{e \in E} PE \cdot E_e \cdot f_e$$

3) The non-renewable energy consumption of router ports that have access to renewable energy (the non-renewable energy may be used for example to guarantee control at all time)

$$\sum_{i \in N} PRS \cdot \left(Q_i^s + \sum_{p \in P} \delta_{ip} \cdot Ws_p \right)$$

4) The non-renewable energy consumption of transponders that have access to renewable energy (again the non-renewable energy may be used for example to guarantee control at all time) and that of the transponders without access to renewable energy

$$\sum_{e \in E} \left(PT \cdot \omega_e + PTS \cdot \omega s_e \right)$$

5) The non-renewable energy consumption of optical switches that have access to renewable energy (similarly the non- renewable energy may be used for example to guarantee control at all time) and that of the optical switches without access to renewable energy

$$\sum_{i \in N} \left(PO_i \cdot (1 - y_i) + POS_i \cdot y_i \right)$$

6) The non-renewable energy consumption of multiplexers and demultiplexers that have access to renewable energy (here also the non-renewable energy may be used for example to guarantee control at all time) and that of the multiplexers and demultiplexers without access to renewable energy

$$\sum_{i \in N} \left(PMD \cdot DMe_i + PMDS \cdot DMs_i \right)$$

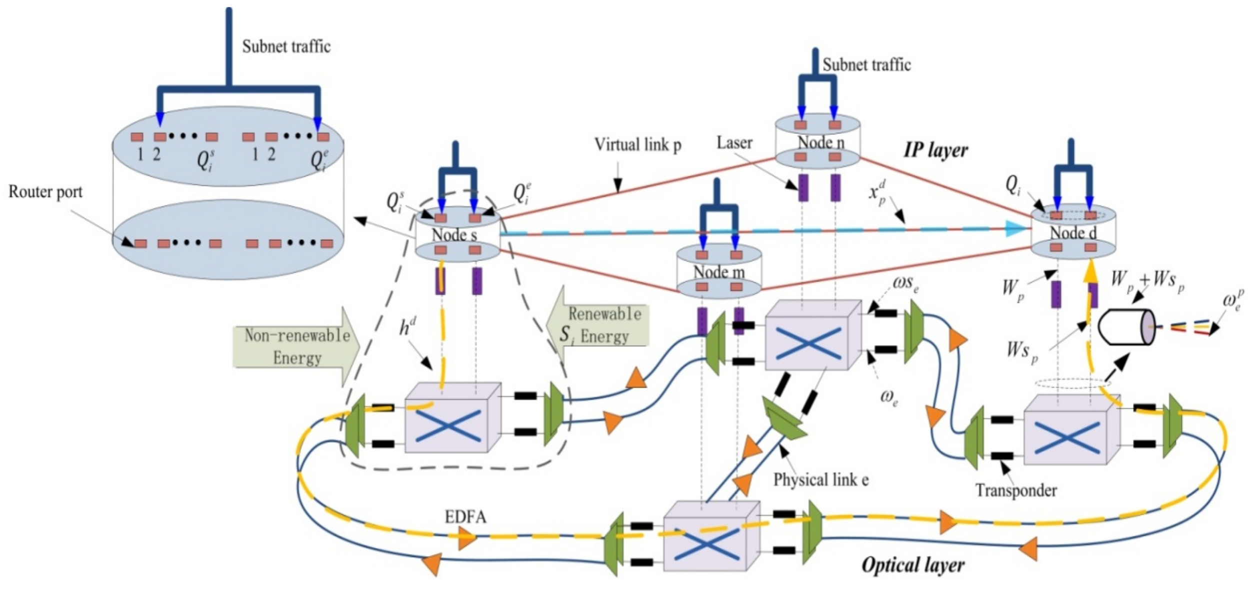

Fig. 2. The hybrid-power IP over WDM network

The LP model that minimizes the non-renewable energy consumed is defined as follows:

Objective: minimize

$$\sum_{i\in N} PR \cdot \left(Q_i^e + \sum_{p\in P} \delta_{ip} \cdot W_p\right) + \sum_{e\in E} PE \cdot E_e \cdot f_e + \sum_{i\in N} PRS \cdot \left(Q_i^s + \sum_{p\in P} \delta_{ip} \cdot Ws_p\right)$$
$$+ \sum_{e\in E} \left(PT \cdot \omega_e + PTS \cdot \omega s_e\right)$$
$$+ \sum_{i\in N} \left(PO_i \cdot (1 - y_i) + POS_i \cdot y_i\right) + \sum_{i\in N} \left(PMD \cdot DMe_i + PMDS \cdot DMs_i\right) \quad (1)$$

Subject to:

$$\sum_{p\in P} x_p^d = h^d \quad \forall\, d \in D, \quad (2)$$

$$\sum_{d\in D} x_p^d \leq \left(W_p + Ws_p\right) \cdot B \quad \forall\, p \in P, \quad (3)$$

$$\sum_{p\in P}\left(\delta_{ip}\cdot W_p+\delta_{ip}\cdot Ws_p\right)+Q_i\le \nabla^i \quad \forall\, i\in N, \tag{4}$$

$$\sum_{e\in E}\delta_{ep}\cdot \omega_e^p = W_p+Ws_p \quad \forall\, p\in P, \tag{5}$$

$$PR^s\cdot\left(Q_i^s+\sum_{p\in P}\delta_{ip}\cdot Ws_p\right)+\sum_{e\in E}PT^s\cdot \omega s_e\cdot\delta_{ie}+PMD^s\cdot DMs_i+PO_i^s\cdot y_i\le S_i$$

$$\forall\, i\in N, \tag{6}$$

$$\sum_{p\in P}\omega_e^p\le W\cdot f_e \quad \forall\, e\in E, \tag{7}$$

$$Q_i^e+Q_i^s=Q_i \quad \forall\, i\in N, \tag{8}$$

$$\sum_{p\in P}\omega_e^p=\omega_e+\omega s_e \quad \forall\, e\in E \tag{9}$$

$$DMe_i+DMs_i=DM_i, \forall\, i\in N \tag{10}$$

The variables and parameters in the equations above are declared as follows:

TABLE I

PARAMETERS (*Pa*) AND VERIABLES (*Ve*) USED FOR LP MODEL

E(Pa)	Physical link set in optical layer,
P(Pa)	Virtual link set in IP layer,
D(Pa)	Traffic demand set between node pairs,
δ_{ip}*(Pa)*	If node *i* belongs to virtual link *p*, δ_{ip} is '1', otherwise it is '0'
δ_{ie}*(Pa)*	If node *i* belongs to physical link *e*, δ_{ie} is '1', otherwise it is '0'
δ_{ep}*(Pa)*	If virtual link *p* starts or ends at physical link *e*, δ_{ep} is '1', otherwise it is '0'

PR(Pa) and PE(Pa)	Non-renewable energy consumption of a router port and an EDFA respectively both use non-renewable energy.
PRS(Pa)	Non-renewable energy consumption of a router port that has access to renewable energy.
PR^s	Renewable energy consumption of a router port that has access to renewable energy.
PO_i*(Pa) and* POS_i*(Pa)*	Non-renewable energy consumption of an optical switch that has access to non-renewable energy only or has access to renewable energy in node *i,* respectively.
PO_i^s	Renewable energy consumption of an optical switch that has access to renewable energy.
PMD*(Pa)and* $PMDS$*(Pa)*	Non-renewable energy consumption of a multi/demultiplexer that has access to non-renewable energy only or has access to renewable energy, respectively.
PMD^s	Renewable energy consumption of a multi/demultiplexer that has access to renewable energy.
DM_i*(Pa)*	The total number of multiplexers and demultiplexers in node *i.*
DMe_i*(Ve)*	Number of multiplexers and demultiplexers in node i which use non-renewable energy.
DMs_i*(Ve)*	Number of multiplexers and demultiplexers in node i which use renewable energy.
y_i*(Ve)*	If the optical switch in node *i* has access to renewable energy $y_i = 1$, otherwise $y_i = 0$.
ω_e *(Ve) and* ωs_e*(Ve)*	Number of wavelength channels on physical link *e* in the optical layer which use non-renewable energy and renewable energy

	respectively.
W_p *(Ve) and* Ws_p*(Ve)*	Number of wavelength channels on virtual link p in the IP layer which use non-renewable energy and renewable energy respectively.
Q_i^e*(Ve) and* Q_i^s*(Ve)*	Number of ports which are powered by non-renewable energy or renewable energy for data aggregation in node *i*.
x_p^d*(Ve)*	Traffic demand *d* between node pairs on virtual link *p*.
ω_e^p*(Ve)*	Number of wavelength channels of virtual link *p* on physical link *e*.
f_e*(Ve)*	Number of fibers on physical link *e*.
E_e*(Pa)*	Number of EDFAs on each fiber on physical link *e*.
PT(Pa) and PTS(Pa)	Non-renewable energy consumption of a transponder that has access to non-renewable energy only or has access to renewable energy respectively.
PT^s	Renewable energy consumption of a transponder that has access to renewable energy.
W(Pa)	Number of wavelengths in a fiber.
Q_i *(Pa)*	Number of ports for assembling data.
∇^i*(Pa)*	Maximum number of ports in node *i*.
S_i*(Pa)*	The maximum output power of the renewable energy source in node *i*.
h^d*(Pa)*	Traffic demand *d* between node pairs.
B(Pa)	Capacity of each wavelength.

The objective function (Equation (1)) aims to minimize the non-renewable energy consumption of the hybrid-power IP over WDM network. Equation (2) and Equation (5) represent the flow conservation constraint in the IP layer and the optical layer. Equation (3) ensures that the traffic flow on each virtual link does not exceed its capacity. The term $(W_p + Ws_p)$ represents the total number of wavelength channels on each virtual link powered by either non-renewable energy or renewable energy. Equation (4) ensures that the limit on the number of router ports in each node is not exceeded. Equation (6) ensures that the renewable energy consumption of router ports and transponders is not larger than the maximum output power of the renewable energy source in each node. Equation (7) and Equation (9) give the limit on the number of wavelength channels in each physical link *e*. Equation (8) ensures that for each node the total number of ports assembling data is equal to the number of router ports using non-renewable energy and the number of ports using renewable energy. Equation (10) gives the limit on the total number of multiplexers and demultiplexers in node *i*.

B. Heuristic Approach

The multi-hop bypass heuristic proposed in [72] improves the bandwidth utilization by allowing traffic demands between different source-destination pairs to share capacity on common virtual links which results in fewer virtual links, and hence fewer IP router ports and lower energy consumption. However, implementing the Multi-hop bypass heuristic which is based on shortest-path routing in the hybrid-power IP over WDM network where renewable energy sources are only available to a limited number of nodes, will only minimize the total energy consumption without considering whether this energy comes from renewable or non-renewable sources. We propose a new heuristic, known as Renewable Energy Optimization hop (REO-hop), that minimize the utilization of non-renewable energy by allowing traffic flows to traverse as many nodes as possible that use renewable energy to ensure that in addition to reducing the total number of IP router ports and transponders, the non-renewable energy consumption is minimized. To maintain QoS, only the two shortest-path routes are considered not to increase the propagation delay. As the traffic pattern is changing and the output power of renewable energy sources varies during different times of the day, the routing paths are dynamic.

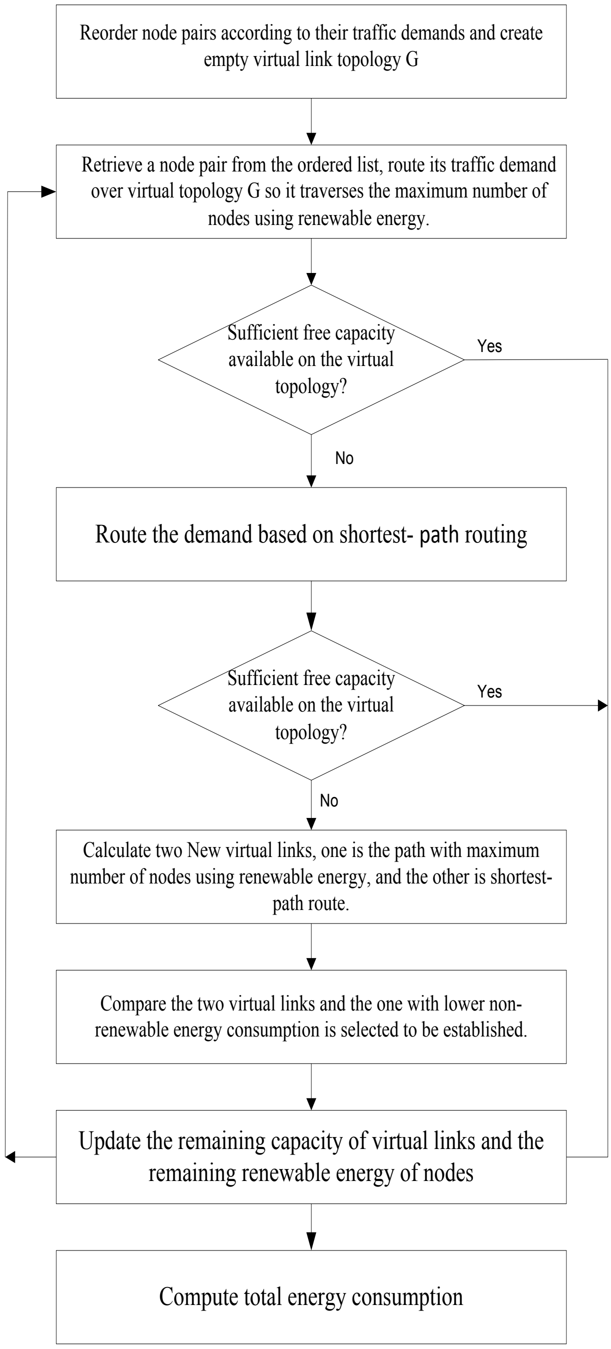

Fig. 3. The REO-hop heuristic Flowchart

The flowchart of the heuristic is given in Fig. 3. The heuristic starts by ordering all the node pairs based on their traffic demands from highest to lowest and creating an empty virtual link topology *G*. A node pair is retrieved from the ordered list, and its traffic demand is routed over virtual topology *G* by comparing the two shortest-path routes and the route with the maximum number of nodes that use renewable energy is selected. If sufficient free capacity is available on virtual topology G, the selected route is accommodated and the remaining capacity on all the virtual links and renewable energy of each node are updated. If the selected route is not available, the other route is selected. If the virtual topology cannot accommodate either routes, a new direct virtual link is established between the source and destination by comparing the non-renewable energy consumption of the route with the maximum number of nodes that use renewable energy, and the shortest-path route. The route with lower non-renewable energy consumption is selected. If the two paths have the same non-renewable energy consumption, the shortest-path route is selected. The new virtual link is added to the virtual topology *G*. and the remaining capacity on all the virtual links and the renewable energy of each node are updated. The above steps are repeated for all the node pairs. After routing all the traffic demands on the virtual topology *G*, the objective function (Equation (1)) is used to calculate the total non-renewable energy consumption of the network.

C. Simulation and Results

To evaluate the non-renewable energy consumption of the architecture and test the REO-hop heuristic, the NSFNET network, depicted in Fig. 4, is considered as an example of a real world network. The NSFNET network consists of 14 nodes and 21 bidirectional links. Solar energy is used as the renewable energy source. Nodes in the NSFNET will experience different levels of solar energy and traffic demands at any given point in time as different parts of the network fall in different time zones. There are four time zones, Eastern Standard Time (EST), Central Standard Time (CST), Mountain Standard Time (MST) and Pacific Standard Time (PST), with an hour time difference between each time zone and the next, we use EST as the reference time.

Real sun rise and sun set data is used to determine the solar energy available at each node during different hours of the day. Fig. 6 shows the solar energy power available to a node [94].

TABLE **II**TABLE II gives the details of the solar energy output power of each node. The solar energy output is non-zero from 6:00 to 22:00 and the maximum output power occurs at 12:00.

The average traffic demand during different hours of the day is shown in Fig. 5 [93]. The average traffic demand between each node pair ranges from 20 Gb/s to 120 Gb/s and the peak occurs at 22:00 in these traffic profiles. We assume that the traffic demand between each node pair in the same time zone is random with a uniform distribution and no lower than 10 Gb/s.

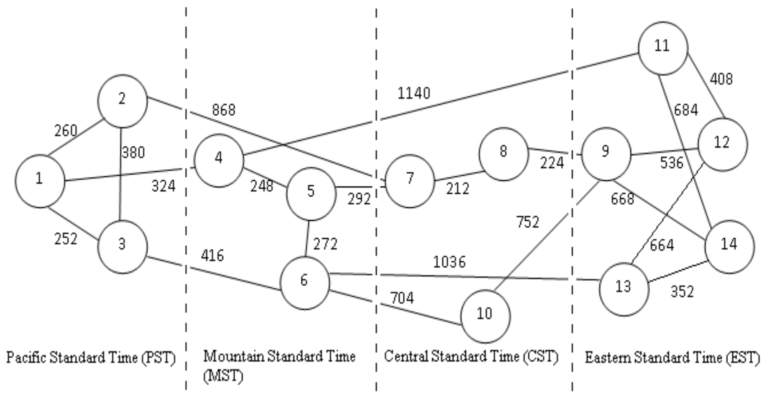

Fig. 4. The NSFNET network with time zones

TABLE II

SOLAR ENERGY OUTPUT POWER AVALIABLE AT EACH NODE (20 kW MAXIMUM OUTPUT POWER)

Node ID / Time (EST)	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	SR: 05:43	SR: 05:14	SR: 05:16	SR: 05:54	SR: 05:27	SR: 05:55	SR: 05:51	SR: 05:39	SR: 05:17	SR: 05:53	SR: 05:18	SR: 05:41	SR: 06:27	SR: 05:30
	SS: 20:04	SS: 21:02	SS: 19:34	SS: 20:40	SS: 20:41	SS: 20:10	SS: 20:53	SS: 20:47	SS: 20:24	SS: 19:56	SS: 20:30	SS: 20:29	SS: 20:30	SS: 20:59
00:00	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW
02:00	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW
04:00	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW
06:00	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	2 kW	0.5 kW	0 kW	0.5 kW
08:00	0 kW	0 kW	0 kW	0 kW	2 kW	0.5 kW	0.5 kW	2.5 kW	6 kW	2 kW	6 kW	6 kW	4.5 kW	6 kW
10:00	1.5 kW	2 kW	2 kW	7 kW	6 kW	6 kW	6 kW	6 kW	13 kW	6 kW	13 kW	13 kW	13 kW	13 kW
12:00	10 kW	10 kW	10 kW	13 kW	13 kW	13 kW	13 kW	18 kW	18 kW	18 kW	20 kW	20 kW	20 kW	20 kW
14:00	18 kW	18 kW	18 kW	18 kW	20 kW	20 kW	20 kW	18 kW	18 kW	18 kW	13kW	13kW	13kW	13kW
16:00	18 kW	18 kW	18 kW	18 kW	18 kW	18 kW	18 kW	13kW	13kW	13kW	7.5 kW	7.5 kW	7.5 kW	7.5 kW
18:00	13kW	13kW	13kW	18 kW	13 kW	13 kW	13 kW	8 kW	8 kW	8 kW	5 kW	5 kW	5 kW	5 kW
20:00	8 kW	8 kW	8 kW	8 kW	8 kW	8 kW	8 kW	3kW	3kW	2 kW	2 kW	2 kW	2 kW	2 kW
22:00	2.5 kW	3 kW	0.5 kW	4 kW	2 kW	0.5 kW	3 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW	0 kW

(SR: Sunrise, SS: Sunset. SR time and SS time are given in the node's local time in June

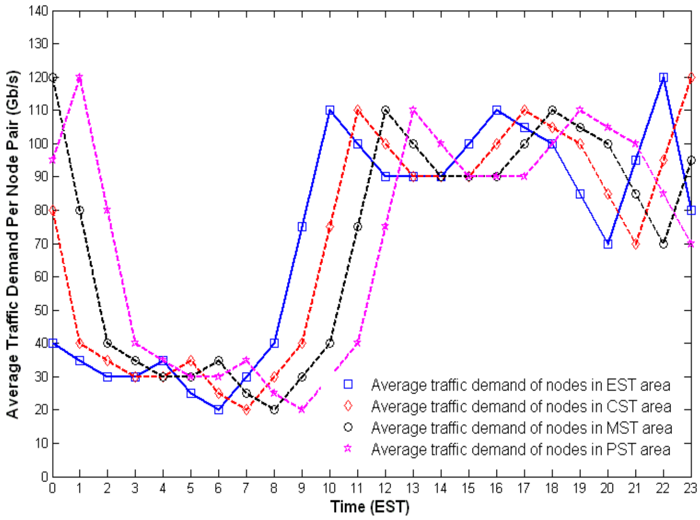

Fig. 5. Average traffic demand in different time zones

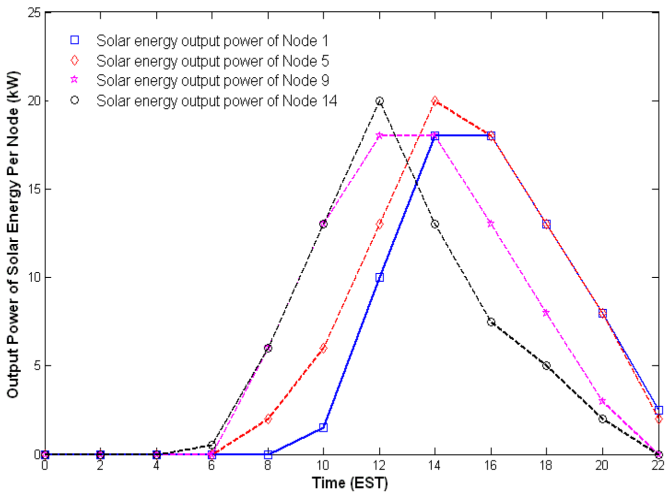

Fig. 6. Output power of solar energy at different nodes in different time zones

1) Non-renewable Energy Consumption of the Network

To maximize the reduction in the non-renewable energy consumption, nodes located at the centre of the network with a high nodal degree are selected to use renewable energy as these are expected to consume more power compared to nodes at the edge of the network as more traffic demands will be passing through them. Nodes 4, 5, 6, 7 and 9 are initially selected to employ solar energy. The impact of the location of nodes using solar energy is studied in Section III.C.2.

We assume that the maximum solar energy available to a node is 20 kW. Typically a one square meter silicon solar cell can produce about 0.28 kW of power [95]. Therefore we require a total solar cell area of about 100 m^2 which can be accommodated in a typical core node. Later we examine the impact of higher solar energy per node.

As mentioned under the Multi-bypass and REO-hop heuristics, some traffic demands are routed in the optical layer by optical switches. We select a suitable size optical switch (the Glimmerglass's 192×192 optical switch) based on the maximum number of wavelengths used in each node. Due to the negligible difference in energy consumption between different optical switch sizes compared to the power consumption of a router port, the same power consumption data for optical switches in the Multi-bypass heuristic and REO-hop heuristic. In Fig. 7, we see that under the multi-hop bypass (the heuristic requiring fewer wavelengths), the maximum number of wavelengths needed is 109.

TABLE III gives the simulation environment parameters. It shows the number of wavelengths, wavelength capacity, distance between two neighbouring EDFAs, and energy consumption of different components in the network. Some of these parameters are similar to those in [72] derived from Cisco's 8-slot CRS-1 data sheets [96 while others are derived from Glimmerglass's 192×192 channels Sytem-600 data sheets [97] and Cisco's ONS 15454 data sheets [98].

The LP problem is solved using the AMPL/LPSOLVE software. We study five different cases: 1) Non-bypass heuristic without renewable energy nodes, 2) Non-bypass heuristic with renewable energy nodes, 3) Multi-hop-bypass heuristic without renewable energy nodes 4) Multi-hop-bypass heuristic with renewable energy nodes, and 5) REO-hop heuristic with renewable energy nodes.

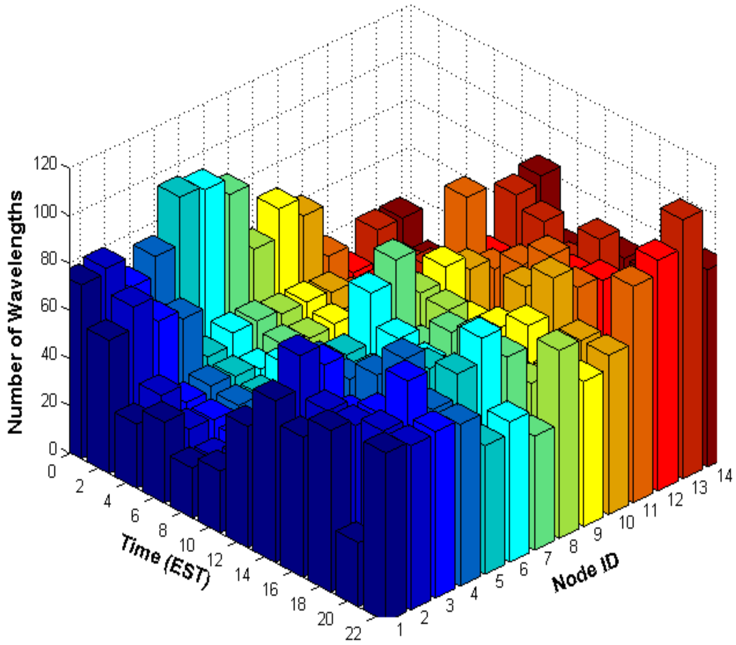

Fig. 7. The number of wavelengths under the Multi-hop-bypass heuristic in throughout the day

Fig. 8 shows the total non-renewable energy consumption of the NSFNET network throughout the day. The curves “Non-bypass without solar energy” and “LP optimal with solar energy” give the upper and lower bounds on the non-renewable energy consumption, respectively. We assume the traffic demands and the output power levels of the solar energy sources given in Fig. 5 and, TABLE II respectively. Compared to the upper bound, both the Multi-hop-bypass and REO-hop heuristics have reduced the non-renewable energy consumption. Reductions introduced by the REO-hop heuristic increase between 6:00 and 22:00 i.e. when the solar energy is significant. Compared with the upper bound, between 12:00 and 18:00, the REO-hop heuristic saves non-renewable energy of about 1000kW. The REO-hop still introduces more savings in the non-renewable energy consumption compared to the Multi-hop-bypass heuristic when there is no solar energy in the network (from 0:00 to 4:00) as REO-hop tries to route demands on virtual links with sufficient capacity rather than using shortest-path routing as with the Multi-hop-bypass heuristic.

The impact of different levels of the maximum solar energy (40 kW, 60 kW and 80 kW) per node is examined in Fig. 9 . A solar cell area of up to 300 m^2 is needed to generate such values. Solar cell cladding with such surface area can be practically built in a typical core routing node location. Increasing the maximum solar energy output linearly reduces the total non-renewable energy consumption under different heuristics.

Fig. 10 shows the reduction in CO_2 emissions under different heuristics. A similar trend to that observed in Fig. 9 is observed. Compared to the non-bypass without renewable energy case, the total CO_2 emissions during a 24 hour period have been reduced by about 47%~52% and 43%~49% under the REO-hop and Multi-hop-bypass heuristics, respectively. In practice, 228 grams of CO_2 approximately are produced by a network component that consumes 1kWh of traditional electricity power [0]; therefore by using the REO-hop heuristic in NSFNET we will be able to save about 1300 tones of CO_2 emissions every year.

Fig. 11 shows the average propagation delay of REO-hop and Multi-hop-bypass heuristics. Compared to the Multi-hop-bypass heuristic based on the shortest path, the REO-hop heuristic has increased the propagation delay. However, the increase in the propagation delay is limited as the REO-hop heuristic routes demands using only the two shortest-path routes (the increase is less than 0.3 ms, i.e. less than 10%) maintaining the QoS. While the propagation delay of the Multi-hop-bypass heuristic based on the shortest path is constant, the propagation delay of the REO-hop heuristic varies throughout the day as the REO-hop heuristic routes demands dynamically based on the output power of solar energy sources in nodes and the available capacity.

TABLE III

SIMULATION ENVIRONMENT PARAMETERS

Distance between two neighboring EDFAs	80 (km)
Number of wavelength in a fiber (*W*)	16
Capacity of each wavelength (*B*)	40 (Gb/s)
Non-renewable energy consumption of a router port (*PR*)	1000 (W)
Renewable energy consumption of a router port (PR^s)	1000 (W)
Non-renewable energy consumption of an optical switch in node *i* (PO_i).	85 (W)
Renewable energy consumption of an optical switch in node *i* (PO_i^s).	85 (W)
Non-renewable energy consumption of an optical switch that has access to renewable energy in node *i* (POS_i)	0 (W)
Non-renewable energy consumption of a router port that has access to renewable energy (*PRS*)	0 (W)
Non-renewable energy consumption of a multiplexer or a demultiplexer (*PMD*).	16 (W)
Renewable energy consumption of a multiplexer or a demultiplexer (PMD^s).	16 (W)
Non-renewable energy consumption of a multiplexer or a demultiplexer that has access to renewable energy (*PMDS*).	0 (W)
Non-renewable energy consumption of a transponder (*PT*)	73 (W)
Renewable energy consumption of a transponder (PT^s)	73 (W)
Non-renewable energy consumption of a transponder that has access to renewable energy (*PTS*)	0 (W)
Non-renewable energy consumption of an EDFA (*PE*)	8 (W)

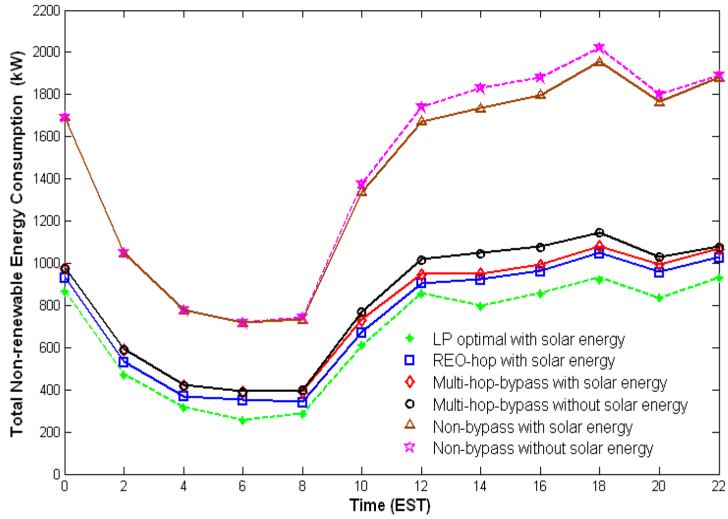

Fig. 8.The total non-renewable energy consumption under different heuristics with and without solar energy

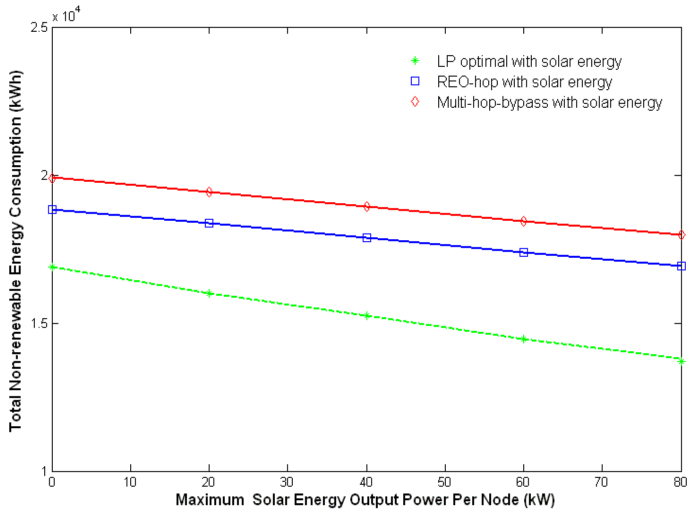

Fig. 9. Total non-renewable energy consumption in 24 hour period for different levels of the maximum solar energy output under different heuristics

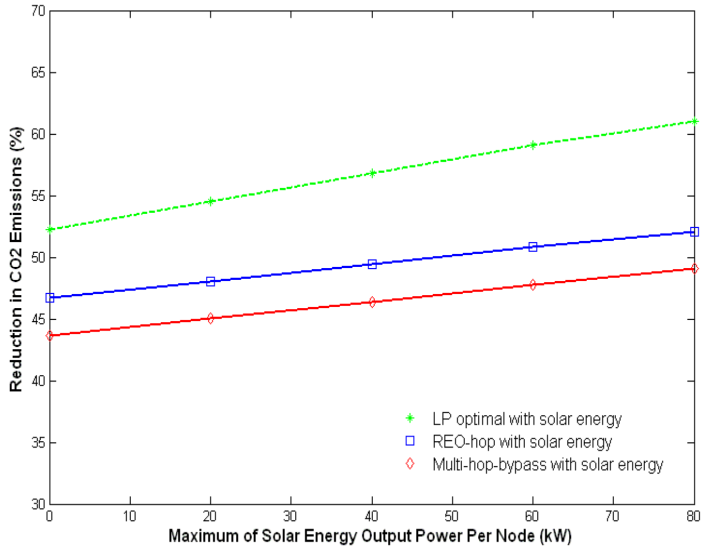

Fig. 10. Reduction in CO_2 emissions in 24 hour period under different heuristics

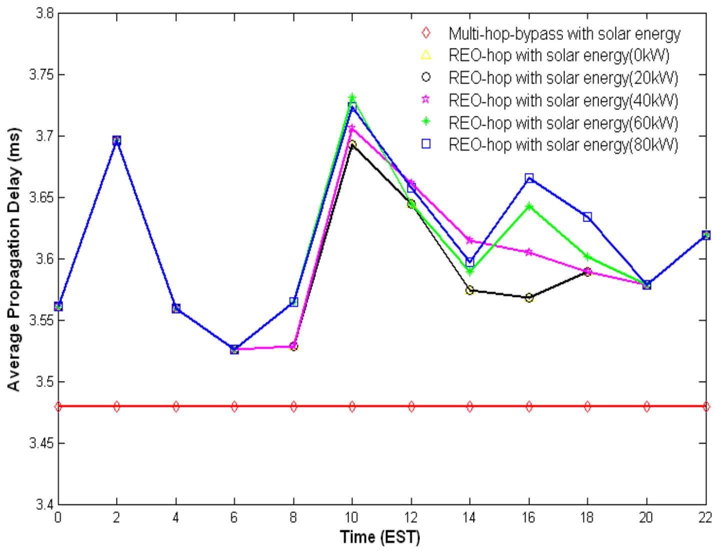

Fig. 11. Average propagation delay under the REO-hop and the Multi-hop-bypass heuristics

2) Number and Location of Nodes that Use Renewable Energy

In Section III.C.1, the five nodes with the highest nodal degree in the NSFNET network core were selected to use renewable energy. In this section we investigate the impact of the number and location of nodes that use renewable energy.

In Fig. 12 we investigate how the non-renewable energy consumption is affected by the maximum output power of renewable energy sources and the number of nodes that use renewable energy. The set of nodes using renewable energy takes the following values: {1}, {1,2}, {1,2,3}, {1,2,3,4}…{1,2,3…14} and the maximum output power of renewable energy sources ranges from 0 kW to 80 kW. It is clear that increasing the number of nodes that use renewable energy and increasing the maximum output power per renewable energy source reduces the total non-renewable energy consumption of network. However, the relation is not linear between the number of nodes and the non-renewable energy consumption which indicates that selecting some nodes to use renewable energy will have more impact on the non-renewable energy consumption than others.

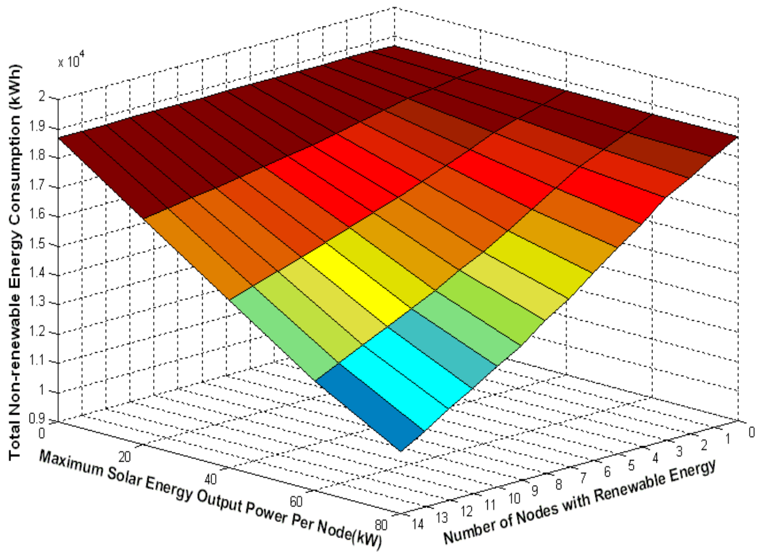

Fig. 12. The total non-renewable energy consumption in 24 hour period under different maximum solar energy power per node and different number of nodes employing solar energy

To investigate the impact of the location of nodes that use renewable energy on the total non-renewable energy consumption, we develop another LP model with the objective of optimizing the selection of nodes that use renewable energy such that the non-renewable energy consumption savings are maximized. The new LP model is subject to the same constraints in section III.A, except that Equation (6) is replaced with Equation (12) and a new constraint is added (Equation (13)).

In this model we consider time as a variable. Therefore, t is added to all the variables in TABLE I where t is the time point of time set T.

The new LP model is defined as follows:

Objective: maximize

$$\sum_{t \in T} \sum_{i \in N} \left(PR^s \cdot \left(Q_{it}^s + \sum_{p \in P} \delta_{ipt} \cdot Ws_{pt} \right) + \sum_{e \in E} PT^s \cdot \omega s_{et} \cdot \delta_{iet} + PMD^s \cdot DMs_{it} + PO_i^s \cdot y_{it} \right) \tag{11}$$

Subject to:

Equations: (2) (3) (4) (5) (7) (8) (9) (10)

(Every variable in the equations above has had the time variable t augmented)

$$PR^s \cdot \left(Q_{it}^s + \sum_{p \in P} \delta_{ipt} \cdot Ws_{pt} \right) + \sum_{e \in E} PT^s \cdot \omega s_{et} \cdot \delta_{iet} + PMD^s \cdot DMs_{it} + PO_i^s \cdot y_{it}$$
$$\leq S_{it} \cdot \varepsilon_i \tag{12}$$
$$\forall\, i \in N, t \in T$$

$$\sum_{i \in N} \varepsilon_i = Ns \tag{13}$$

where ε_i =1 if node i uses renewable energy, otherwise ε_i=0, and Ns is the total number of nodes with access to renewable energy.

Equation (12) ensures that the renewable energy consumption in each node is within the maximum output power of its associated renewable energy source at any time of the day. In practice, the energy produced from solar cells can be stored in batteries; which relaxes the constraints on the availability of solar energy as a function of the day time. The use of energy storage elements is however, not included in the current formulations. Equation (13) implies that the total number of nodes that use renewable energy is limited to Ns which is set in advance.

Fig. 13 shows the optimization results under different levels of the maximum renewable energy output power (20 kW to 80 kW), assuming Ns=5 and the traffic demand shown in Fig. 5. The optimal node selection does not change.

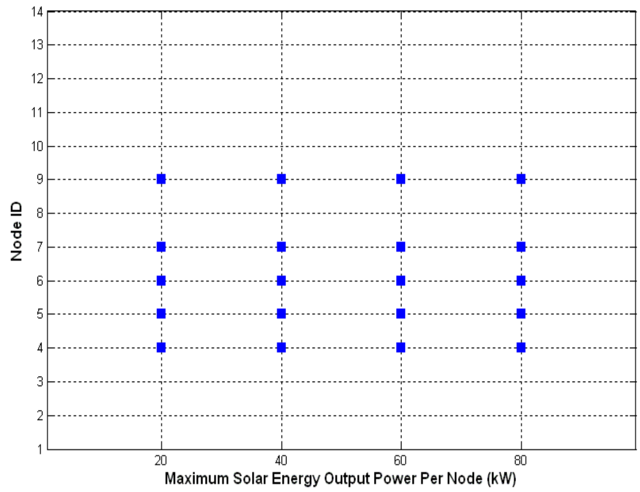

Fig. 13. Optimal location of nodes with access to renewable energy for different values of maximum available solar energy per node (LP-optimal)

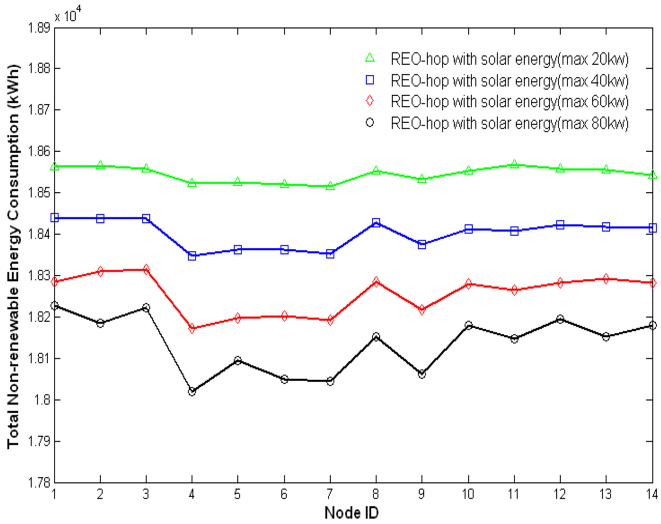

Fig. 14. The total non-renewable energy consumption in 24 hour period with different nodes using renewable energy

In Fig. 14 we verify the optimization results by evaluating the total non-renewable energy consumption under the REO-hop heuristic where we assume that only a single node uses renewable energy. We evaluate the performance under different values of the maximum solar energy per node. Results of the REO-hop heuristic matches the optimization results where the total non-renewable consumption is lower when the nodes in the centre of the network (4, 5, 6, 7, and 9) use renewable energy.

In Fig. 15, the delay and power consumption are evaluated under the REO-hop heuristic with a maximum renewable power of 60 kW per node when nodes at the core {4, 5, 6, 7, 9} or nodes at the edge {1, 2, 10, 11, 14} are selected to use renewable energy. The former node set results in a higher reduction in the non-renewable energy consumption compared to the latter node set. Also selecting the node set {4, 5, 6, 7, 9} results in a lower average propagation delay compared to the node set {1, 2, 10, 11, 14}.

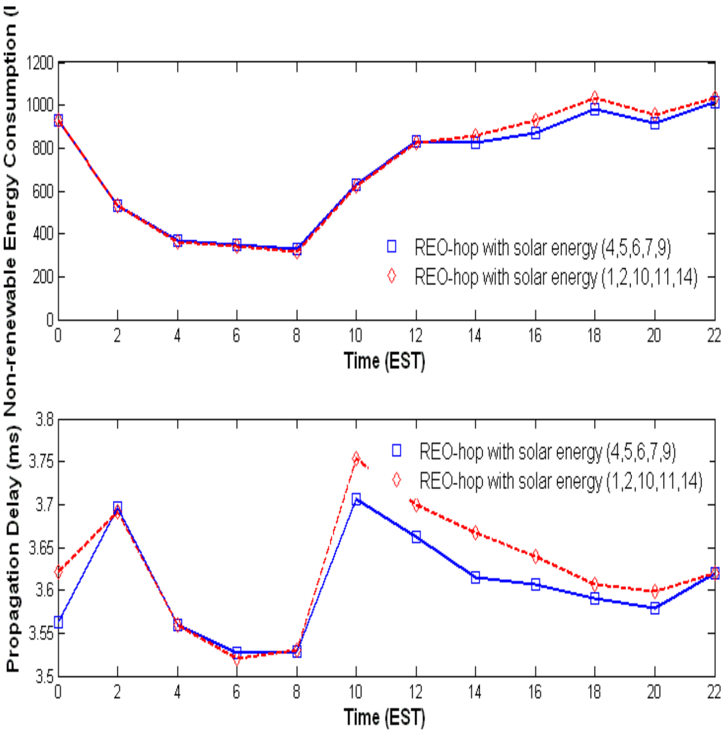

Fig. 15. The non-renewable energy consumption and the average propagation delay under two different node selection scenarios using renewable energy

3) Node Non-renewable Energy Consumption

In this section we show the reduction in the non-renewable energy consumption experienced by each node individually under different heuristics. Assuming the traffic demand in Fig. 5 and that the maximum solar energy output power is 60 kW, we compare the scenarios when no renewable energy sources are used and when some nodes use renewable energy (nodes 4, 5, 6, 7 and 9) in Fig. 16. Under the Non-bypass heuristic, the non-renewable energy consumption of nodes varies significantly between nodes at the center and nodes at the edge of the network as nodes at the centre consume more energy as more traffic flows are routed through them. The Multi-hop-bypass and REO-hop heuristics have significantly reduced the non-renewable energy consumption and its variance. However, nodes at the centre of network still have slightly more non-renewable energy consumption than the nodes at the edge of the network. As expected, the REO-hop heuristic results in further reductions compared to the Multi-hop-bypass heuristic. The non-renewable energy consumption of nodes at the centre of the network, where renewable sources are deployed, has significantly decreased when renewable energy sources are introduced to the network.

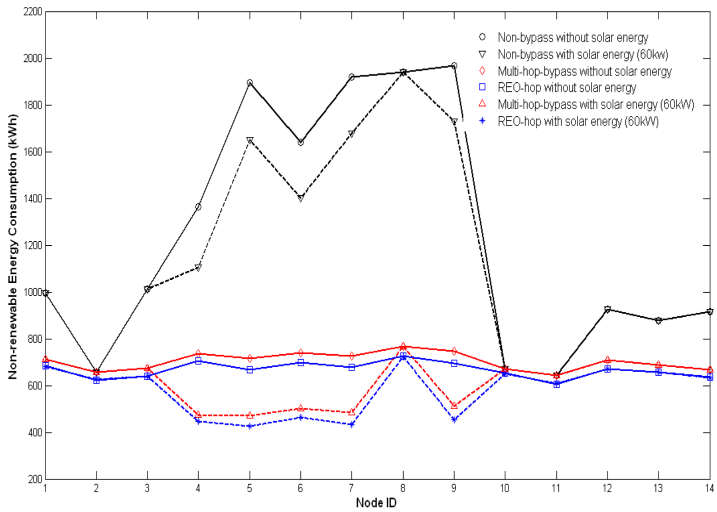

Fig. 16. The non-renewable energy consumption of the nodes in a 24 hour period under different heuristics

4) Non-renewable Energy Consumption under Adaptive Link Rate with the REO-hop heuristic

In this section, we investigate the impact of ALR on the non-renewable energy consumption. Different energy profiles are proposed to define of the dependency between equipment energy consumption and traffic load. Fig. 17 shows different energy profiles for telecommunication equipment [88] where energy consumption is a function of the load on the network component. The latter is expressed as a percentage of the total capacity of the network component. We consider (i) 'On-off' energy profile [88] (ii) 'Linear' energy profile: Here the energy consumption depends linearly on the traffic load, e.g. in switch architectures like Batcher, Crossbar and Fully-Connected [88], [100] (iii) ‘Log10’ energy profile employed in equipment that uses hibernation techniques such as the low-power idle technique for Ethernet [100], [101]. In this approach data is sent as fast as possible to allow the equipment to quickly return to the low-power idle state. (iv) ‘Log100’ energy profile: This profile is considered as a middle function between the ‘On-off’ and the ‘Log10’ profiles. (v) ‘Cubic’ energy profile: Typical in equipment that uses Dynamic voltage Scaling (DVS) and Dynamic Frequency Scaling (DFS) [88].

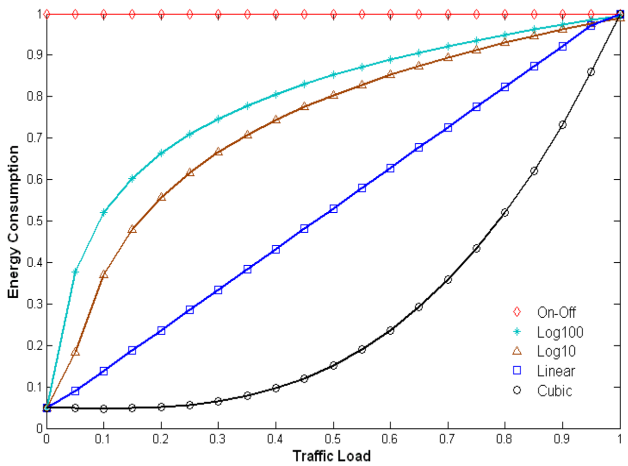

Fig. 17. Different energy profiles

Fig. 18 shows the non-renewable energy consumption of the REO-hop heuristic under different energy profiles for router ports and transponders (the two most energy consuming components in the node) with 20 kW solar power at the five optimal nodes and 80 kW solar power at all nodes. The non-renewable energy consumption curves’ behaviour is subject to the energy profile. The ‘cubic’ profile results in the highest reduction in non-renewable energy consumption. Compared to the ‘on-off’ profile, the ‘cubic’ profile has reduced the non-renewable energy consumption by up to 9% between 12:00 and 20:00. This relatively small reduction in energy consumption associated with ALR is commensurate with the fact that the ALR profiles will only make a difference to partially loaded wavelengths.

Note that when all the nodes have access to 80 kW solar power, the energy consumption continues to decrease beyond 6 am due to the availability of solar power at the nodes. At 22:00 the two sets of curves converge, however the total energy consumption is still lower when all nodes have access to solar power due to solar power availability in more nodes.

The sum of the total power consumed by all nodes, where each node is dimensioned based on the largest number of router ports needed over the 24 hour period is 2010kW which is the peak shown in Fig. 8. A maximum reduction of 85% (Fig. 19) is achieved. The average savings in this case are approximately 65% and vary slightly between the five profiles. Note that these savings are mostly due to architecture design (photonic switching instead of electronic routing) and powering down unused router ports and transponders as the renewable energy is low here and has limited effect.

Fig. 18 also shows the total non renewable energy consumption when all nodes deploy 80 kW renewable energy sources. The maximum reduction in non-renewable energy consumption compared to the peak in Fig. 8 is 97% (Fig. 19). The average saving is about 78% with small variation between different energy profiles.

To appreciate the typical power consumption levels per node, Fig. 20 shows the non-renewable energy consumption of the individual nodes in the network under the REO-hop heuristic and 'cubic' profile when only the 5 optimal nodes have access to 20 kW renewable energy each, and when all the nodes have access to 80 kW renewable energy each.

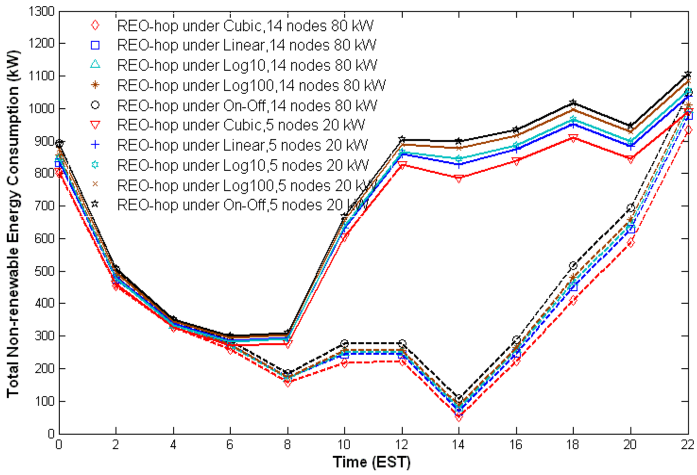

Fig. 18. The total non-renewable energy consumption of the REO-hop heuristic under different energy profiles

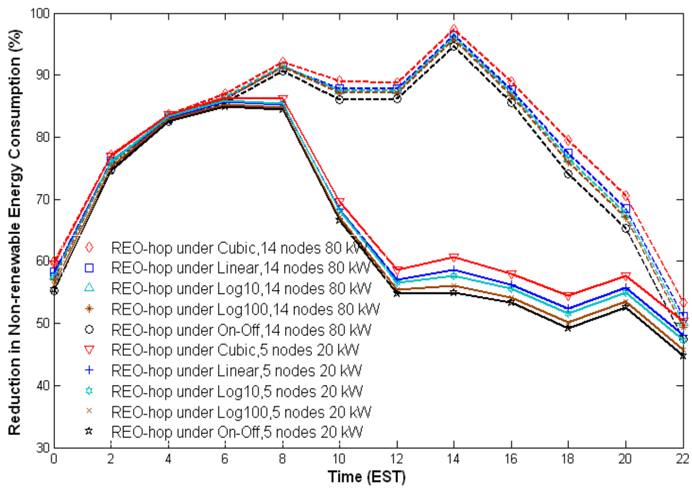

Fig. 19. Reduction in the non-renewable energy consumption of the REO-hop heuristic under different energy profiles

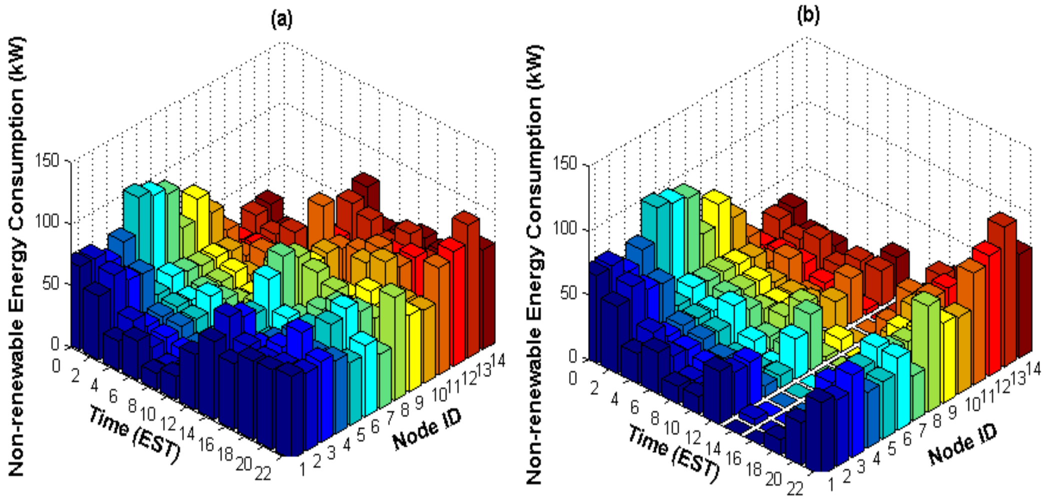

Fig. 20. The non-renewable energy consumption of each node under the REO-hop heuristic and the 'cubic' profile with: (a) 20 kW solar power at nodes (4, 5, 6, 7, 9), (b) 80 kW at all nodes

IV. Data centres and virtualization

Data centres are even more of a contributor to energy bills, and by extension carbon emission, than switching and driving communications lines. As well as housing many power-hungry servers CPUs, data centres have moving parts – in other words disks. Of course, data centres also contain internal interconnects, and these are amenable to many of the power-reduction strategies and techniques as the external network itself. Modern data centres frequently make use of commodity hardware for the servers, since this is low cost and modular. This has some benefit in terms of energy savings since the design of a typical machine for the small office or home installation is often driven by operational as well as capital cost constraints, and so may tend to become more energy-efficient over time.

The granularity of a single commodity server is often mismatched to the distribution of server demands for the many co-located customers in a large data centre. To pool resources more efficiently, while still retaining partitioning or isolation of processing, I/O and storage demands, the typical modern data centre makes heavy use of virtualisation.

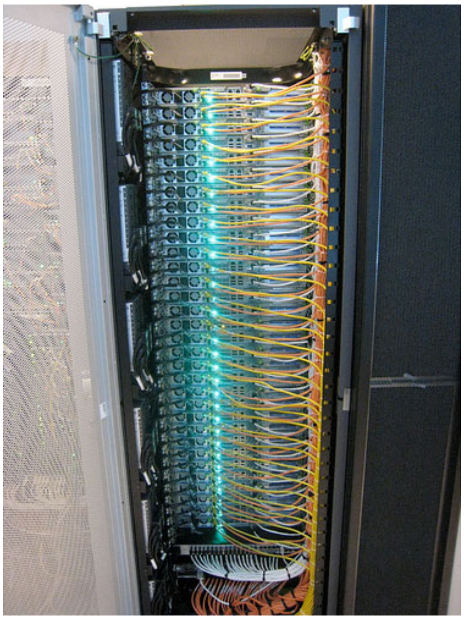

Fig. 21: A Rack with a Top-of-rack switch

A. Virtualisation & Distributed Computing

There are very general lessons to be learned from distributed computing, that can be applied across a range of energy problems, as proposed in Andy Hopper's vision in [103]. More specifically, we have had the capability to determine where to place a job in a distributed system that is planet-wide for many years [104]. Decision procedures for admission and placement of jobs in a system can and do leverage a past work with resource management in networks [105] with real-world workloads [106], 107]. Global workload mobility allows everything from a disk drive [108] to a data centre [109] to power down when electricity demand is high by allowing providers to maintain service levels by moving jobs to alternative locations where demand is low. Intermittent renewable energy can also be absorbed by this mechanism. Computing will be making use of energy that would otherwise be lost through over-generation or use of inefficient generation techniques. Some researchers have considered a dedicated network for connecting a globally distributed set of data centres, so as to provide transparent workload migration [110] although energy conservation during migration will require high-bandwidth, low-latency connections.

B. An Agile, Power-aware API

Provided with measurement and modelling of power utilisation, information provided through a suitable API can permit management and power-waste mitigation. The needs for hardware support with the APIs and the functionality in protocols to exploit such hardware support for energy management, and mitigating the possible negative impact on protocol and service performance has not been examined, and we propose to address these problems. There is also a need to permit agility in the network components; currently there is a highly rigorous determination of the roles for each network subsystem (from host and server through network interface card, switch and router). Work in this area aims to identify and evaluate a flexible interface that permits such extensions as end-point deflection routing (to support migration), network function agility (to accommodate both optimal performance and workload offloading – as determined by power-needs and architectural constraints), and ideas of network and communications subsystems – such as network interface cards and switch subsystems that permit the powering down to sleep of entire systems while providing the lowest level of network response and a fast restart on demand – in the style of Somniloquy [66]. One can envisage such APIs providing the necessary flexibility to allow roles to be moved from host network-stack and application through to being entirely offloaded into core switch components.

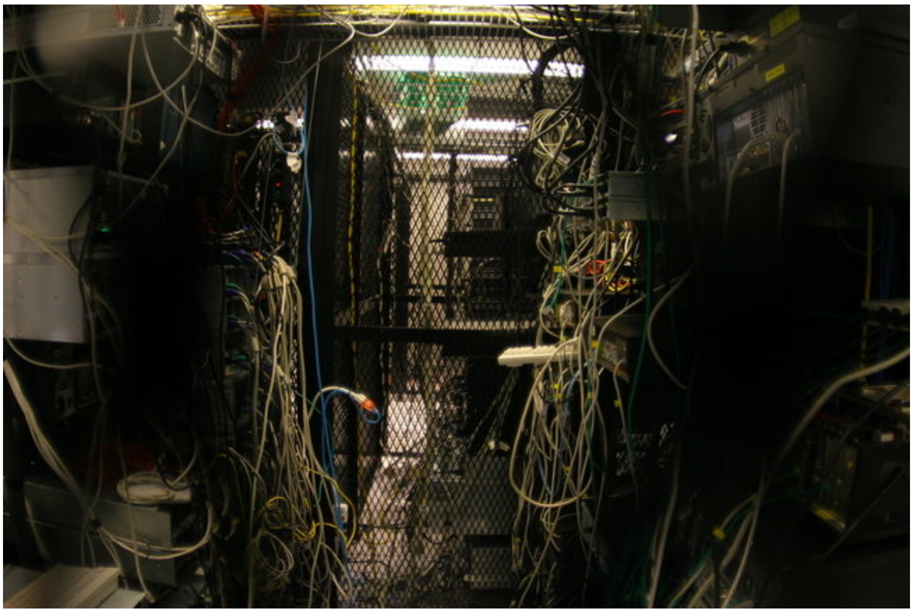

Fig. 22: A Machine Room with typical wiring

C. Agile Virtualization

It is estimated that servers and associated cooling overheads comprise around 1.2% of global energy demand [111]. Service level agreements (SLAs) typically require high levels of availability and so service providers also require highly reliable energy supplies. Some large scale Content Distribution Network (CDN) Providers such as YouTube employ a high-speed, low-latency network to transparently replicate content so that most users perceive low latency, high performance downloads from nearby servers. Others such as Akamai already use energy signals such as price to route requests in their overlays [112]. One can envisage the same networks being used to move workloads around the world [113]. This network would allow the use of computing as a virtual battery ameliorating local changes in power availability by modulating energy consumption in sympathy with local grid demand. This workload mobility will allow continuous service provision, despite variations in power supply, which will allow us to locate computing installations near to remote renewable energy generation sites. This has the potential to improve the viability of remote renewable energy generation sites by allowing one to consume energy on site rather than incurring electricity transmission overheads [103], [114]-[116]. Grid demand varies significantly over short timescales according to many different processes and so the National Grid is continually working to balance generation to demand. This is made possible through reserve generation capacity which is provided either by operating large scale generators at less than full capacity or through the provision of standing reserves which can be activated at short notice. Both of these approaches are less efficient than a large scale plant operating at full capacity. Our vision of global workload mobility will allow data-centres to power down when demand is high by allowing providers to maintain service levels by moving jobs to alternative locations where demand is low. Intermittent renewable energy can also be absorbed by this mechanism. At the lowest level we consider the physical link layer of the network. In particular, one expects future interconnects to be focused on optical networks supporting low latency, low-power switching [117]. However, it takes three years of continuous operation to equal the manufacturing energy embodied in computing hardware [118], so powering down machines to release energy to the general grid is increasing the relative size of this cost. Relative impacts of provisioning fibre optic networks [119] or power transmission systems must also be considered.

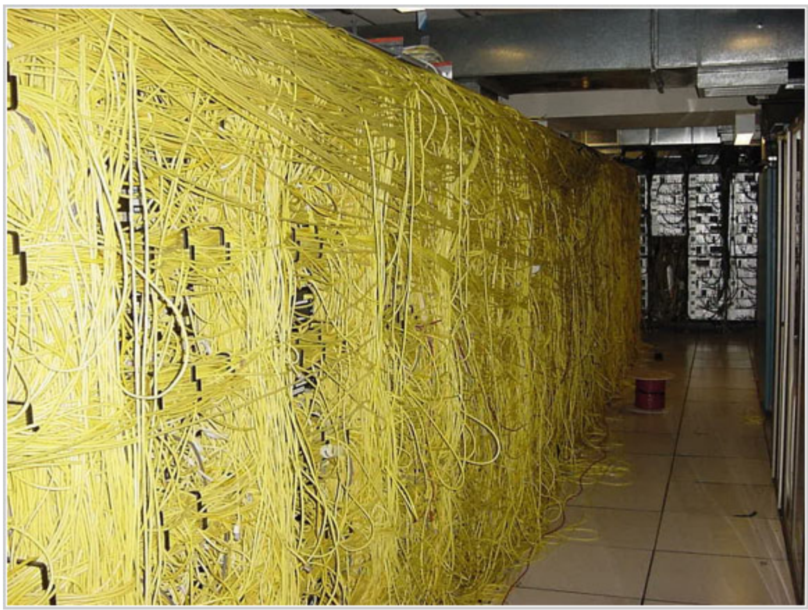

Fig.23: Entropy taking over

D. Workload suitability

Non-interactive or batch-computation jobs are becoming an increasingly significant proportion of computing workload [120]. Examples of this type of workload include indexing web-pages, optical character recognition of books or scientific simulations. These types of workload provide flexibility to our system because they can have weaker constraints on acceptable levels of downtime. For example, one might choose to prioritise the transmission of interactive workloads when shutting down a data centre or even to simply pause non-interactive jobs and wait for power to become available again. Interactive workloads are more challenging because they must continue to be available despite changes in power supply. Very highly responsive workloads might prove unsuitable to this approach because the down time (however minimal) incurred by migration might not be tolerable. However, there is a large class of workloads for which small, occasional periods of downtime will be largely unnoticeable.

Particular examples might include web, email or DNS services, this can benefit from the significant existing Internet-characterisation experience [121]- 123].

Virtualisation and Migration

The recent resurgence in virtualisation provides the ideal software platform for mobile workloads [112]. Virtualisation allows us to run and migrate existing workloads without requiring extensive engineering. There are a number of parameters to virtual machine migration which will affect the performance of our network. The majority of a workload can often be migrated in the background whilst it continues to run on the host machine. This gives two quantities for describing workload movement. Migration time is the total amount of time to fully move the job from one host to another and down-time is the total amount of time for which the job was unavailable. Different workloads (with different patterns of memory use) require different amounts of migration and down time. Another consideration is the location of any storage used by the workload. Conventionally, migration takes place in a local network using a Storage Area Network (SAN). This means that any fixed storage is automatically available to the migrated workload. Suitable guarantees from the network might permit wide-area migration using a SAN. For example, one might imagine a scenario in which the computing workload is moved out of a data centre which then continues to operate but with less power demands providing storage capacity to these workloads. Alternatively, strategies exist for mirroring or live replication of remote storage which can ensure that a workload has an up-to-date image to work from on arrival [124].

E. Dynamic Computation Placement Architecture

At present datacenters are provisioned with worst-case power and cooling capacity in order to enable them to meet user SLAs. However, this provisioning comes at the cost of wasted energy both as stored electrical energy and energy expended in the cooling subsystem. It is our intention to develop an architecture where computation can be moved in accordance with energy policy constraints, e.g. where the cost of energy is cheapest or where there are reserves that would otherwise be wasted, with the goal of controlling the total power consumed by computation. Virtualisation provides the benefit of abstraction – all the components that form the computation (the operating system, binaries, support libraries, etc) are packaged together. This enables us to treat the entire software stack as a single unit for management purposes. Specifically, the migration feature present on most major virtualisation platforms to build a dynamic computation placement architecture. Migration works by physically moving a computation unit from one host to another by moving all state and data information between the hosts. Current migration architectures have been designed for low-latency, low-bandwidth single-hop LAN setups where migrations are carried out occasionally within a small pool of well-known hosts. However, in our architecture all these parameters change. This architecture will require migrations be executed within pools of frequently changing hosts separated by large physical distances connected by high-bandwidth, high-latency networks. Moreover, the underlying network characteristics (both at the routing and physical layers) are likely to deviate substantially from current designs.

This suggests some of the following challenges to meet the goal of a dynamic computation placement architecture work will be required in the following areas:

(i) Characterising computation and workload: As a first step into realising this architecture we must characterise the computation and data footprints of a number of representative workloads for the purposes of determining which computation models are suitable for dynamic computation placement. One could start by leveraging the extensive network monitoring and measurement experience in our team [122], [122121], [125], [126].

(ii) Predicting migration: As a follow-on to the previous step a model will be needed to predict the migration time for any computation in the system thereby creating a theoretical foundation on which we can build the system. Past work on migration-location-identification [104],[127], will assist this phase significantly.

(iii) Network Characterisation: We expect the underlying energy efficient network developed in Section II, will inform us through its models of both network-needs and load and of network-topology, in combination with the opportunities for changes in the underlying physical-network system which we anticipate will contain significantly differing properties compared to current network designs. This will require a complete analysis and characterisation of the network for the purposes of exploiting it for the purpose of system construction, possibly refining recent contributions in techniques for topology modeling and comparison [128].

(iv)Energy Generation Patterns: As we create ever more dynamic patterns of demand, it will be necessary to track the evolution of the generation of power, geographically and temporally. As renewable sources come online, these will be more decentralized and more variable; this might present a challenge for power grid engineers, but for the Internet and the Web, it may represent an excellent opportunity for co-evolution, with small and very large data centers being co-located at all new tidal, wind and solar sites to take up any local over-production as and when needed. This is more a civil engineering challenge than one for computer science and electrical and power engineering.

(v) Recent rapid evolution of smart phone operating systems such as the iPhone and Android Unix based operating system, together with energy aware application design, has led the way. However, we expect to see the same family of techniques applied in OS refinements for Data

Centres and for the Home and Small Office, with an emphasis on enabling efficient and timely migration within the architecture and by changing their design. We anticipate extensions of recent work on disk-energy needs [108], energy-aware whole-data centre storage management, such as in Sierra [129], data-center migration [109], [104] and the vision of a Green ICT [103].

Section V - Advances in Photonic Physical Layer

The network concepts described in I to IV above require an energy efficient physical transport layer. It has been shown [130] that optical point to point transmission data transmission over medium and extended distances has the potential to be extremely energy efficient, with optimal choice of the optical transmission technology, as is currently implemented between ethernet switches and IP routers (for example a 40channel 40Gb/s 1000km transmission system requires 1.1nJ/bit energy consumption). There is, however a significant potential for further energy efficiency improvement by the use of optical technologies within IP routers, as there exists the inherent advantage that optical transmission systems and switches can effectively be bit-rate agnostic, having the potential to switch very high data rate signals at the packet and flow level, rather than at the bit level as in the purely electronically switched IP routers of today. In contrast, energy efficiency presents a real challenge to electronic switching, with current IP routers consuming 20nJ/bit [131], and a 1 Tb/s aggregate IP router consuming 20kW - not only an energy consumption issue, but also a thermal management one.

With optical technologies enabling both energy efficient transport and routing, energy efficient networks can be devised in optimal form both to minimize energy consumption but also reduce related carbon emissions, for example by enabling the ready use of low carbon energy sources.

We therefore concentrate on a discussion of optical switching technologies for high energy efficiency, and describe how optical switch based technologies are starting to achieve the scalability and performance required for implementation within the next generation of switching and routing equipment. In describing such switches, it is important to note developments in transport technologies as any usable switch must be able to operate on the transmission signals. This section therefore first considers the development of energy efficient optical communication links.

From the early 1990's optical transmission systems typically connected telephone exchanges. These telecom systems used temperature controlled semiconductor lasers in single wavelength low data-rate circuits. In addition to the electro-optical conversion energy consumption, the photonic transmitters consumed a greater amount of power in the thermo-electric coolers required to maintain the laser operating temperate than in the light generation and transmission itself. They were also limited to transmission distances of less than 100km before the signal needed to be regenerated electrically, as there were then no effective optical amplifiers in existence.

The invention of the erbium doped fibre amplifier (EDFA) [130,134] enabled efficient optical amplification opening the way to wavelength division multiplexed transmission of many wavelengths over extended distances in an optical fibre, this improving long distance telecommunications. This technology provided a significant boost to IP communications between routers, but was not sufficiently low cost to allow the extension of high data rate communications from the router or switch out to the desktop.

Gigabit Ethernet [135] and 10 Gigabit Ethernet [136] standards drew attention to the growing need for efficient and low cost optical transmission systems over short distances between Ethernet switches and in certain cases from the switch to the user, requiring lasers to operate at high modulation rates at high ambient temperatures, without the need for a thermo-electric cooler. Extensive work on semiconductor laser design enabled operation at 10Gb/s direct modulation at 100ºC at 1.3µm [137-139], providing further reduced energy consumption. As a result, optical links can now be constructed with energy consumptions of 130pJ/bit at rates of 10 Gb/s and operating lengths of up to 10 km at ambient temperatures of 85ºC [132].With the advent of these energy efficient optical links at certain wavelengths it is interesting to develop optical switch technologies which have the potential for energy efficiency.

Broadly optical switch systems can be classified into two types, those whose switching time is longer than a typical packet length, for circuit switching (thermo-optic, and MEMS), and those switches which can switch in several nanoseconds or less, making them suitable for packet switching (lithium niobate and semiconductor).

Micro electro-mechanical switches (MEMS) are capable of being used in large array optical cross-connects with low insertion losses and low power consumption but with relatively slow switching speeds [141,142] (for example a 128x128 way cross connect with 2.6dB average insertion loss has a 57ms switching time [143])

Thermo-optic switches are usually based on Mach-Zehnder modulators with a resistive heating element next to one arm, which is heated to produce an optical phase change and so switch the output of the modulator from one output port to the other.

Examples of this form of switch include silica (e.g. a 2x2 port Mach-Zehnder interferometer based switch, in whch there is an electric heater in each of the arms. In order to switch the output from the interferometer the 45mW heating power is initiated in one arm, causing an optical phase change and output switch in a response time of 3ms [144]) and polymer based switching elements (e.g. a 2x2 port polymer waveguide based Mach-Zehnder modulator, with a heater on one arm. Here to initiate the output switch, the heater is turned on, consuming 4mW, and causing the output to switch from one port to the other in 160µs [145]).

Switches which can operate on the packet timescale are generally of greater interest for applications within IP routers or Ethernet switches. To this end there are several approaches, the use of high speed modulators, silicon photonic and semiconductor optical amplifier based optical switches.

It is possible to create fast optical switches based on lithium niobate Mach-Zehnder modulators – these components offer fast switching speeds, but at the expense of high operating powers and significant losses [146,147]. For example an 8x8 port switch was demonstrated but with 15dB insertion loss [148].

Another potentially attractive technology for nanosecond scale optical switches is that of semiconductor optical amplifiers (SOAs). These have received considerable interest in recent years as they have the potential of combining inherent optical gain, fast switching speed, potential for uncooled operation and the capability of integrating large numbers of components, in order to make functional compact switch sub-systems.

Integrated 2x2 port crosspoint switches have been reported as early as 1990 [149], with designs incorporating low-loss passive waveguides and active SOA switching elements reported in 1993 [150,151]. Fig **24** shows the key parts of such an integrated optical switch, with passive low loss optical waveguides for input and outputs, with turning mirrors and power splitters for compactness and SOA gates, each having sufficient gain to overcome the passive waveguide and splitting losses. In this design two of the SOAs will be active at any time, to allow either straight through or crossed output operation.

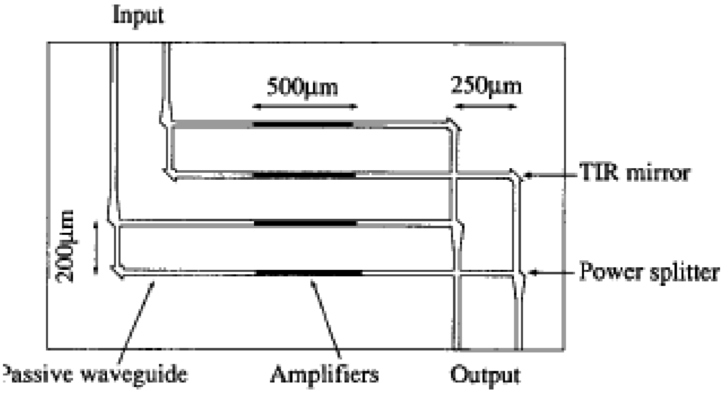

Fig **24**, after [151], showing an early 2x2 port switch, with passive input and output waveguides, beam expander based power splitters, total internal reflection based mirrors and Semiconductor Optical Amplifier switching gates.

Clearly increased port count is of great importance for an optical switch, so considerable research effort has been expended in the design and realization of integrated optical switches. While SOAs have many advantages as optical switching elements, they can introduce excess amplified spontaneous emission noise and also can impart optical patterning distortion to the data signals if they are not designed and operated in the correct way. It is generally desirable to limit the number of cascaded SOAs in any design, to limit the build up of noise and distortion, but this must be balanced against the split-and-combine loss introduced by the number of cascaded splitting stages within any switch design. Bearing this in mind two designs for higher port count switches are presented, one a semi-integrated 8x8 port switch subsystem and the other an integrated 16x16 port switch fabric, together with a route to a switch which can operate uncooled at very high data capacities.

An 8x8 port switch has recently been demonstrateed by Kai et al. of Fujitsu [152]. In this case scheduled multi-wavelength packets (10λ x 10Gb/s) are split in a 1x8 splitter, each output of which is fed to an input of a an integrated array of 8SOAs, which are combined and amplified by a further SOA within the integrated module, shown schematically in fig Fig **25**a. Thus any of the 8 inputs can be directed to any of the 8 outputs. This design uses integration with each of the modules being packaged and controlled, but fibre splitters connect the modules as can be seen in Fig **25**b is contained in a 19” rack mounting enclosure.

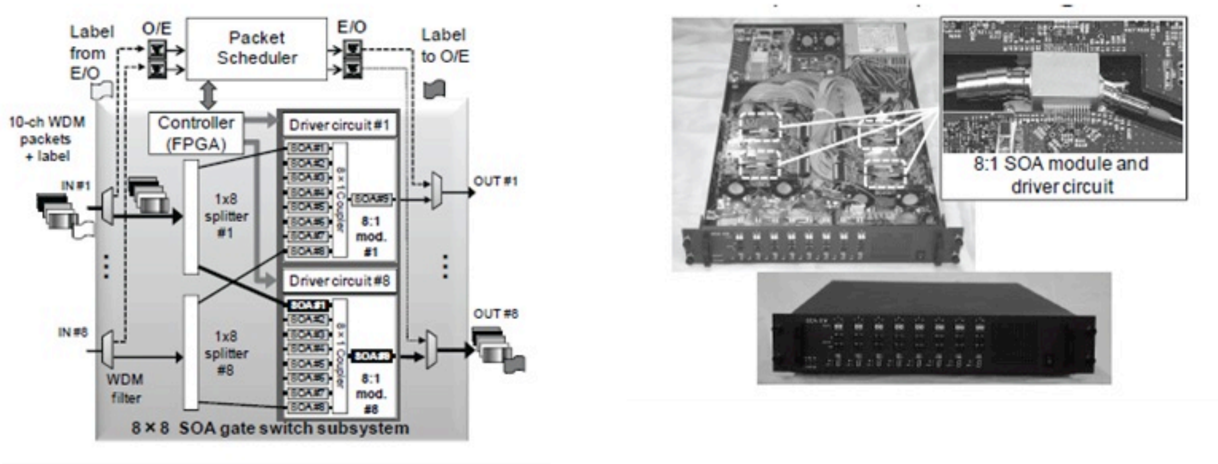

Fig **25**, a. A schematic of the 8x8 switch , &b, the completed switching subsystem from [152]

Switch architecture	**SOA gates**	**cascading stages**	**Max No. of SOA gates in each stage**	**Splitting loss per stage**
Tree	256	1	256	24dB
Benes	224	7	32	6dB
Clos	192	3	64	12dB

TABLE IV SOA optical switch architectures. It can be seen that for the three architectures considered, the Clos is the best compromise between number of required SOA gates, cascaded SOAs and splitting losses.

.

As a route to higher density integration of high speed power efficient optical switches Wang et al. [153] of Cambridge University have produced an integrated 16x16 port optical switch. The architecture for this switch was carefully chosen to optimize the performance of the switch.

TABLE IV shows, for a tree, Benes and Clos architecture, the required number of SOAs, the number of cascaded switching components and the splitting losses between each switching stage.

Thus schematically the switch is made from 12 4x4 port switching elements, each containing an input shuffle network, 16 gating SOAs and an output shuffle network. These 4x4 switching elements are then combined in a Clos architecture, with two additional shuffle networks between them as shown in Fig **26**.

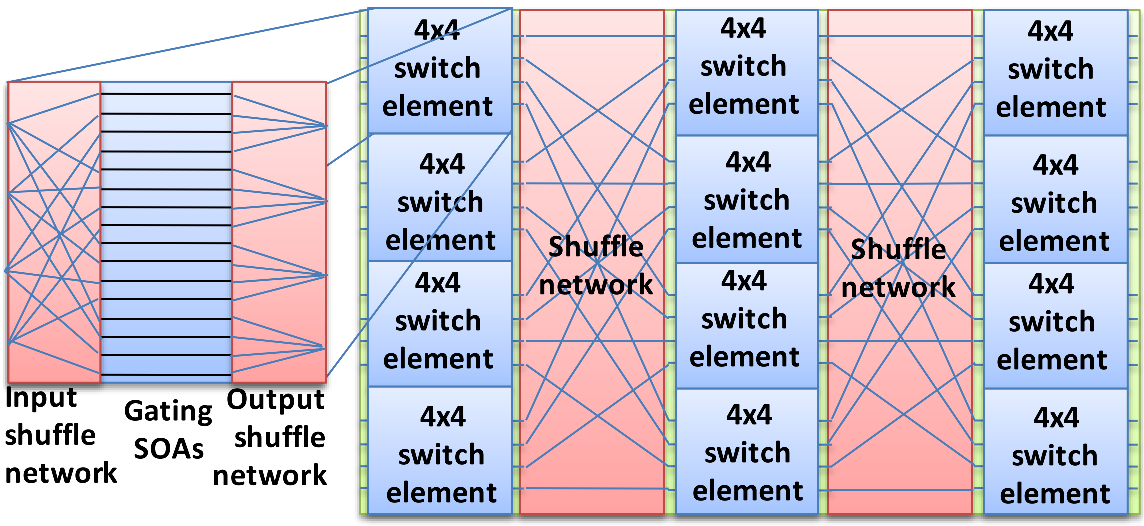

Fig **26** Schematic of the 16x26 port integrated switch in [153]

The switch has been realized on indium phosphide and contains 1114 functional components, including 192 gating SOAs, being the highest component count active photonic integrated circuit yet constructed and is extremely compact, with dimensions of 6.3mm x 6.5mm. The device, shown in Fig **27** operates error free, with a power penalty of between 1.8 and 5 dB. This early demonstration of a 16x16 port integrated switch optical switch is fabricated from all active material, necessitating all of the shuffle networks (most of the switch) to be driven electrically.

It should be emphasized that this first design run of the 16x16 port integrated optical switch fabric is an all active construction, with all of the chip surface being pumped electrically. Improved fabrication with passive waveguides and active SOAs will not only reduce the electrical power consumption by more than half, but will also significantly improve the optical performance of the switch fabric, as the current noise source from the electrically pumped shuffle waveguides will be removed. It is calculated that the active passive version of this switch will require no more than 2pJ/bit of energy in operation in its 16x16 port with 10 wavelengths at 10Gb/s per path operating regime. Although this is by no means all that is required for a router, the power consumption is far less than the 10s of nJ/bit consumed by today's IP routers.

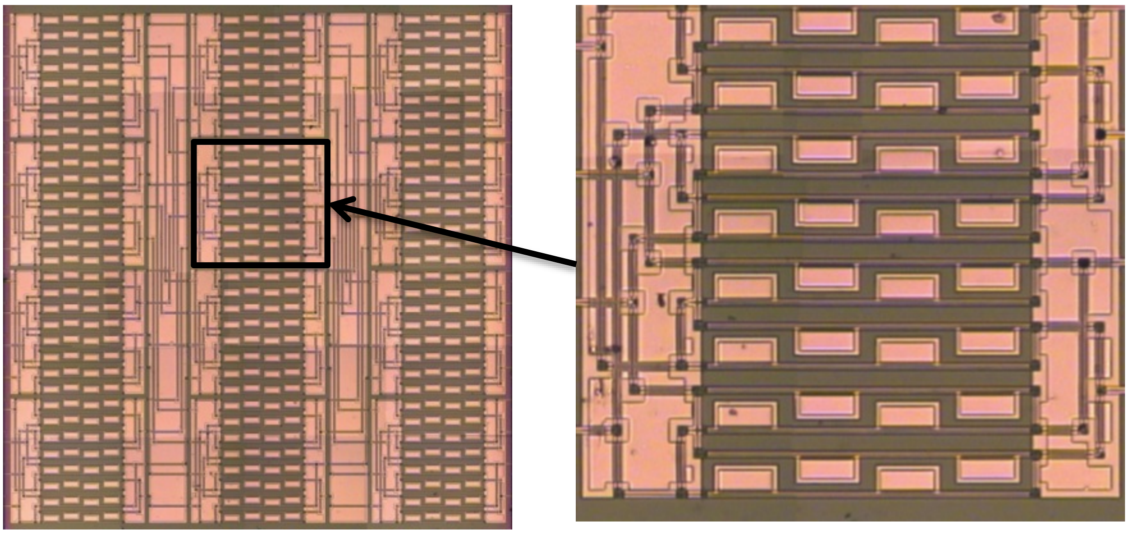

Fig **27**, The fabricated 16x16 port switch, from [153], showing the whole device, and inset, a 4x4 port switching element, with optical splitters, 90º bends, and the 16 horzontal gating SOA switches.

The integrated switch described above is re-arrangeably non-blocking, requiring the switch to be re-configured for all inputs simultaneously. This can easily be achieved in a time-slotted protocol, where all packets are the same length, with centralized scheduling, with buffering being performed at the edges of the network [154]. These issues are inherent to most optical switch designs, as optical buffering is unavailable.

Bergman et al at [155,156] have produced a time-slotted optical switch based on SOAs, which reads WDM Packet headers on the fly and performs deflection routing in a ‘data vortex’. This solution is hard to integrate as it requires both electrical communication from neighbouring nodes and latency between nodes to allow for on the fly processing.

Improvements in energy efficiency and functionality with SOA based optical switch fabrics are anticipated with the next few years. One such key advance is the adoption processing methodologies similar to those which have been learned in CMOS foundries, with the use of standardized ‘building blocks’[157]. This will enable the design of larger switch fabric designs including extensive active/passive material integration, with the SOA gates made from active material with high optical gain at high temperatures and the interconnects and splitters made from low loss passive materials. This will dramatically reduce the energy consumption and cooling requirements of integrated SOA switch fabrics, making cooler-less operation possible, because of the low density of active switching elements on the largely passive switch fabric. In addition, the addition of on-chip power monitors and added functionality like integrated wavelength conversion [158] will greatly reduce the requirement for addition external parts.

The switches described here are constructed from the mature quantum well epitaxial technology, but work has been reported using quantum dots as the active medium for optical amplifiers, with their improved thermal and saturation performance [159]. An integrated 2x2 port switch, operating at 1.3μm has been demonstrated to have stable gain in excess of 10dB from room temperature to over 70ºC [160]. This offers the possibility of large scale integrated optical switches without the need for thermo-electric coolers.

Additionally, recent advances in the field of silicon photonics have enabled optical sources, such as the hybrid silicon laser [161] to be integrated with modulators[162] and detectors to produce high performance low cost optical links for computing applications [163]. Further work based on ring resonators [164] has enabled initial demonstrations [165] of small port count integrated switch fabrics, with the possibility of moderate port count energy efficient switch fabrics in the future.

It is thus envisaged that coolerless optical switch fabrics will scale to larger than 32 or 64 ports, operating at energy efficiencies of 1pJ/bit or better, being good candidates for use within tomorrow's routers and switches.

V. Conclusions

This chapter has presented and discussed different aspects of intelligent energy aware networks. In Section II we discussed the techniques proposed to reduce the energy consumption of communication networks at the equipment and network levels. In Section III we have investigated the use of renewable energy in a hybrid-power IP over WDM network to reduce the non-renewable energy consumption and consequently CO_2 emissions. We have developed an LP optimization model and a novel heuristic, known as REO-hop, to optimize the use of renewable energy in the hybrid-power IP over WDM architecture. The results show that compared to the Multi-hop-bypass, the REO-hop heuristic has reduced the non-renewable energy consumption by

47%~52% while maintaining QoS. Another LP model has been developed to optimize the selection of nodes deploying renewable energy. The results show that compared to nodes at the edge, selecting nodes at the center of the network to deploy renewable energy results in higher reductions in the total non-renewable energy consumption. Also we have investigated ALR where load dependent energy consumption is assumed. The results show that the 'cubic' energy profile gives the highest reduction in the non-renewable energy consumption. Compared to the case where all nodes are statically dimensioned for the maximum traffic in terms of IP ports and optical layer, the REO-hop routing heuristic in the NSFNET network with the optimal node selection, each node has access to 20kW renewable power, and ALR results in a maximum energy saving of 85%, average of 65%. Furthermore, when all the nodes have access to 80kW renewable power each, total energy savings of 97% peak, and 78% average are achieved. In Section IV we have discussed using virtualization technologies for workload migration in data centers to minimize energy consumption. In section V we have considered how photonic technology has the potential to provide energy efficient integrated switch fabrics, reducing the energy consumption of Ethernet switches and IP routers within data centers.

References

1. "Climate Change: Valuing Emissions," Dept. of Environment Food and Rural Affairs (DEFRA), 2008, UK, http://www.defra.gov.uk/ENVIRONMENT/climatechange/research/carboncost.
2. UK Parliament Select Committee on Science and Technology second report, Chapter 3, "The economics of energy efficiency," http://www.publications.parliament.uk/pa/ld200506/ldselect/ldsctech/21/2106.htm.
3. Kenneth J. Christensena, Chamara Gunaratnea, Bruce Nordmanb and Alan D. George, "The next frontier for communications networks: power management," Computer Communications 27 (2004) 1758–1770.
4. C. Bianco, F. Cucchietti and G. Griffa, "Energy consumption trends in the next generation access networks – a teleco perspective," Proc. Of the 29th International Telecommunications Energy Conference, 2007, (INTELEC 2007), Rome, Italy, Sept. 2007, pp. 737-742.

5. Etoh, Minoru; Ohya, Tomoyuki; Nakayama, Yuji, "Energy Consumption Issues on Mobile Network Systems," Energy Consumption Issues on Mobile Network Systems, July 28 2008-Aug. 1 2008 Page(s):365 – 368.
6. Telecom Italia web site, "The Environment," URL: http://www.telecomitalia.it/sostenibilita2006/English/ B05.html.
7. BT Press, "BT announces major wind power plans," Oct 2007, http://www.btplc.com/news/articles/showarticle.cfm?articleid=%7Bdd615e9c-71ad-4daa-951a-55651baae5bb%7D.
8. ITWales, "Green evangelist to call for big changes in computer use to aid environment," ITWales conference, Nov 2007, http://www.itwales.com/997539.htm.
9. "An Inefficient Truth," Global action plan, London, UK, Dec 2007, www.globalactionplan.org.uk/upload /resource/full-report.pdf.
10. K. Kawamoto et al, "Electricity used by office equipment and network equipment in the US, " Energy – The International Journal, vol. 27, No. 3, pp. 255-269, 2002.
11. J. Koomey, "Rebuttal to testimony on 'Koyoto and the Internet: The energy implications of the Digital Economy," Berkeley, CA; Lawerence Berkeley National Laboratory. LBNL-46509, 2000, http://enduse.lbl.gov/projects/ InfoTech.html
12. M. Yamada, T. Yazaki, N. Matsuyama and T. Hayashi, "Power efficient approach and performance control for routers," Workshop on Green Communications at IEEE ICC'09, Dresden, June 2009.
13. SMART 2020 Report, Enabling the low carbon economy in the information age, http://www.theclimategroup.org, 2008.
14. El-Gorashi, T.E.H. and Elmirghani, J.M.H., "Differentiated resilience with dynamic traffic grooming for WDM mesh networks," Proc. 11th International Conference on Transparent Optical Networks (ICTON'09), São Miguel, Portugal, June 28 - July 2, 2009.
15. El-Gorashi, T.E.H. and Elmirghani, J.M.H., "Differentiated resilience for anycast flows in MPLS networks," ICTON 2009, June 28- July 2, Island of São Miguel, Azores, Portugal.
16. S. Nanda and T. Chiueh, "A survey of virtualization technologies," Department of Computer Science, SUNY at Stony Brook, Tech. Rep., 2005.
17. R. Katz, "Tech titans building boom," IEEE Spectrum, vol. 46, no. 2, pp. 40 –54, 2009.
18. X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in Proceedings of the 34th annual international symposium on Computer architecture, ser. ISCA '07. New York, NY, USA: ACM, 2007, pp. 13–23.
19. A. P. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier, "A survey of green networking research," Tech. Rep. arXiv:1010.3880, Oct 2010.
20. R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy- Aware Fixed Network Infrastructures," IEEE Communications Surveys and Tutorials, 2010.
21. M. Gupta and S. Singh, "Greening of the internet," in Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, ser. SIGCOMM '03. New York, NY, USA: ACM, 2003, pp. 19–26.
22. S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation, ser. NSDI'08. Berkeley, CA, USA: USENIX Association, 2008, pp. 323–336.

23. M. Gupta, S. Grover, and S. Singh, “A Feasibility Study for Power Management in LAN Switches,” in Proceedings of the 12th IEEE International Conference on Network Protocols. Washington, DC, USA: IEEE Computer Society, 2004, pp. 361–371.
24. M. Gupta and S. Singh, “Dynamic Ethernet Link Shutdown for Energy Conservation on Ethernet Links,” in IEEE International Conference on Communications, 2007. ICC ’07., 2007, pp. 6156 –6161.
25. H. Tamura, Y. Yahiro, Y. Fukuda, K. Kawahara, and Y. Oie, “Performance Analysis of Energy Saving Scheme with Extra Active Period for LAN Switches,” in IEEE GLOBECOM, 2007, pp. 198–203.
26. M. Rodrıguez-Perez, S. Herreria-Alonso, M. Fernandez-Veiga, and C. Lopez- Garcia, “Improved opportunistic sleeping algorithms for LAN switches,” in Proceedings of the 28th IEEE conference on Global telecommunications, ser. GLOBECOM’09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 35–40.
27. P. Gunaratne and K. Christensen, “Predictive power management method for network devices,” Electronics Letters, vol. 41, no. 13, pp. 775 – 777, 2005.
28. I. Kamitsos, L. Andrew, H. Kim, and M. Chiang, “Optimal sleep patterns for serving delay-tolerant jobs,” in Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, ser. e-Energy ’10. New York, NY, USA: ACM, 2010, pp. 31–40.
29. M. Gupta and S. Singh, “Using Low-Power Modes for Energy Conservation in Ethernet LANs,” in Proceedings of the 26th IEEE International Conference on Computer Communications. IEEE INFOCOM 2007., May 2007, pp. 2451 –2455.
30. G. Ananthanarayanan and R. H. Katz, “Greening the switch,” in Proceedings of the 2008 conference on Power aware computing and systems, ser. Hot- Power’08. Berkeley, CA, USA: USENIX Association, 2008, pp. 7–7.
31. A. Bashar, G. Parr, S. McClean, B. Scotney, M. Subramanian, S. Chaudhari, and T. Gonsalves, “Employing Bayesian Belief Networks for energy efficient Network Management,” in National Conference on Communications (NCC), 2010, pp. 1 –5.
32. T. Smith, R. S. Tucker, K. Hinton, and A. V. Tran, “Implications of Sleep Mode on Activation and Ranging Protocols in PONs,” in 21st Annual Meeting of the IEEE Lasers and Electro-Optics Society (LEOS), 2008, pp. 604–605.
33. R. Kubo, J.-I. Kani, Y. Fujimoto, N. Yoshimoto, and K. Kumozaki, “Sleep and adaptive link rate control for power saving in 10G-EPON systems,” in Proceedings of the 28th IEEE conference on Global telecommunications, ser. GLOBECOM’09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1573–1578.
34. Y. Yan, S.-W. Wong, L. Valcarenghi, S.-H. Yen, D. Campelo, S. Yamashita, L. Kazovsky, and L. Dittmann, “Energy Management Mechanism for Ethernet Passive Optical Networks (EPONs),” in IEEE International Conference on Communications (ICC), May 2010, pp. 1 – 5.
35. B. Nordman and K. Christensen, “Reducing the Energy Consumption of Network Devices,” in Tutorial for the July 2005 IEEE 802 LAN/MAN Standards Committee Plenary Session, ser. IEEE, 2005.
36. “802.3az Energy Efficient Ethernet, meeting materials,” IEEE, Tech. Rep., January 2008. [Online]. Available: http://www.ieee802.org/3/az/index.html
37. P. Reviriego, J. Maestro, J. Herna andndez, and D. Larrabeiti, “Burst Transmission for Energy-Efficient Ethernet,” IEEE Internet Computing, vol. 14, no. 4, pp. 50 –57, 2010.

38. L. A. Barroso and U. Holzle, “The Case for Energy-Proportional Computing,” Computer, vol. 40, pp. 33–37, December 2007.
39. C. Gunaratne, K. Christensen, and B. Nordman, “Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections, and scaling of link speed,” International Journal on Network Management, vol. 15, pp. 297–310, September 2005.
40. C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, “Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR),” IEEE Transactions on Computing, vol. 57, pp. 448–461, April 2008.
41. C. Gunaratne, K. Christensen, and S. Suen, “NGL02-2: Ethernet Adaptive Link Rate (ALR): Analysis of a Buffer Threshold Policy,” in Proceedings of the IEEE Global Telecommunications Conference, 2006. GLOBECOM ’06, 2006, pp. 1 –6.
42. P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, “Energy aware network operations,” in Proceedings of the 28th IEEE international conference on Computer Communications Workshops, ser. INFOCOM’09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 25–30.
43. C. Gunaratne and K. Christensen, “Ethernet Adaptive Link Rate: System Design and Performance Evaluation,” in Proceedings of the 31st IEEE Conference on Local Computer Networks, 2006, pp. 28 –35.
44. B. Zhang, K. Sabhanatarajan, A. Gordon-Ross, and A. George, “Real-time performance analysis of Adaptive Link Rate,” in 33rd IEEE Conference on Local Computer Networks, 2008. LCN 2008., 2008, pp. 282 –288.
45. M. Weiser, B. Welch, A. Demers, and S. Shenker, “Scheduling for reduced CPU energy,” in Proceedings of the 1st USENIX conference on Operating Systems Design and Implementation, ser. OSDI ’94. Berkeley, CA, USA: USENIX Association, 1994.
46. B. Zhai, D. Blaauw, D. Sylvester, and K. Flautner, “Theoretical and practical limits of dynamic voltage scaling,” in Proceedings of the 41st annual Design Automation Conference, ser. DAC ’04. New York, NY, USA: ACM, 2004, pp. 868–873.
47. E. Le Sueur and G. Heiser, “Dynamic Voltage and Frequency Scaling: The Laws of Diminishing Returns,” in Proceedings of the 2010 Workshop on Power Aware Computing and Systems (HotPower’10), 2010.
48. Y. Chen, T. X. Wang, and R. H. Katz. Energy efficient Ethernet encodings. In Proceedings of the 33rd IEEE Conference on Local Computer Networks, Montreal, Quebec, Canada, 14-17 October, pages 122–129. IEEE, Piscataway, NJ, 2008.
49. D. E. Taylor, “Survey and taxonomy of packet classification techniques,” ACM Computing Surveys, vol. 37, pp. 238–275, September 2005.
50. W. Jiang and V. K. Prasanna, “Reducing dynamic power dissipation in pipelined forwarding engines,” in Proceedings of the 2009 IEEE international conference on Computer design, ser. ICCD’09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 144–149.
51. K. Zheng, C. Hu, H. Lu, and B. Liu, “A TCAM-based distributed parallel IP lookup scheme and performance analysis,” IEEE/ACM Transactions on Networking, vol. 14, pp. 863–875, August 2006.
52. F. Zane, G. Narlikar, and A. Basu, “Coolcams: power-efficient TCAMs for forwarding engines,” in Proceedings of the Twenty-Second Annual Joint Conference of the IEEE Computer and Communications, INFOCOM’03, 2003.

53. S. Kaxiras and G. Keramidas, “IPStash: a set-associative memory approach for efficient IP-lookup,” in Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies, INFOCOM’05, vol. 2, 2005, pp. 992 – 1001.
54. W. Eatherton, G. Varghese, and Z. Dittia, “Tree bitmap: hardware/software IP lookups with incremental updates,” SIGCOMM Computer Communication Reviews, vol. 34, pp. 97–122, April 2004.
55. S. Sahni and K. S. Kim, “Efficient construction of multibit tries for IP lookup,” IEEE/ACM Transactions on Networking, vol. 11, no. 4, pp. 650 – 662, 2003.
56. W. Jiang and V. Prasanna, “Architecture-aware data structure optimization for green IP lookup,” in Proceedings of the International Conference on High Performance Switching and Routing (HPSR), 2010, pp. 113 –118.
57. F. Chung, R. Graham, and G. Varghese, “Parallelism versus memory allocation in pipelined router forwarding engines,” in Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures, ser. SPAA ’04. New York, NY, USA: ACM, 2004, pp. 103–111.
58. W. Jiang, Q. Wang, and V. Prasanna, “Beyond TCAMs: An SRAM-Based Parallel Multi-Pipeline Architecture for Terabit IP Lookup,” in Proceedings of the 27th IEEE Conference on Computer Communications, INFOCOM’08, 2008, pp. 1786 –1794.
59. S. Kumar, M. Becchi, P. Crowley, and J. Turner, “CAMP: fast and efficient IP lookup architecture,” in Proceedings of the 2006 ACM/IEEE symposium on Architecture for networking and communications systems, ser. ANCS ’06. New York, NY, USA: ACM, 2006, pp. 51–60.
60. F. Baboescu, D. M. Tullsen, G. Rosu, and S. Singh, “A Tree Based Router Search Engine Architecture with Single Port Memories,” in Proceedings of the 32nd annual international symposium on Computer Architecture, ser. ISCA ’05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 123–133.
61. L. Peng, W. Lu, and L. Duan, “Power Efficient IP Lookup with Supernode Caching,” in Proceedings of the IEEE Global Telecommunications Conference, 2007, 2007, pp. 215 –219.
62. K. J. Christensen, C. Gunaratne, B. Nordman, and A. D. George, “The next frontier for communications networks: power management,” Computer Communications, vol. 27, no. 18, pp. 1758 – 1770, 2004.
63. M. Allman, K. Christensen, B. Nordman, and V. Paxson, “Enabling an Energy Efficient Future Internet Through Selectively Connected End Systems,” in ACM SIGCOMM HotNets. ACM, 2007, pp. 1–7.
64. AMD, “Magic Packet Technology Application in Hardware and Software,” http://support.amd.com/us/Embedded TechDocs/20381.pdf, 2003.
65. K. Sabhanatarajan and A. Gordon-Ross, “A resource efficient content inspection system for next generation Smart NICs,” in Proceedings of the IEEE International Conference on Computer Design, 2008, pp. 156 –163.
66. Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta, “Somniloquy: augmenting network interfaces to reduce PC energy usage,” in Proceedings of the 6th USENIX symposium on Networked systems design and implementation. Berkeley, CA, USA: USENIX Association, 2009, pp. 365– 380.
67. M. Jimeno and K. Christensen, “A Prototype Power Management Proxy for Gnutella Peer-to-Peer File Sharing,” in Proceedings of the 32nd IEEE Conference on Local Computer Networks. Washington, DC, USA: IEEE Computer Society, 2007, pp. 210–212.

68. J. Reich, M. Goraczko, A. Kansal, and J. Padhye, "Sleepless in Seattle no longer," in Proceedings of the 2010 USENIX conference on USENIX annual technical conference, ser. USENIXATC'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 17–17.
69. S. Nedevschi, J. Chandrashekar, J. Liu, B. Nordman, S. Ratnasamy, and N. Taft, "Skilled in the art of being idle: reducing energy waste in networked systems," in Proceedings of the 6th USENIX symposium on Networked systems design and implementation. Berkeley, CA, USA: USENIX Association, 2009, pp. 381–394.
70. Y. Agarwal, S. Savage, and R. Gupta, "SleepServer: a software-only approach for reducing the energy consumption of PCs within enterprise environments," in Proceedings of the 2010 USENIX conference on USENIX annual technical conference, ser. USENIXATC'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 22–22.
71. R. Huelsermann, M. Gunkel, C. Meusburger, and D. A. Schupke, "Cost modeling and evaluation of capital expenditures in optical multilayer networks," Journal of Optical Networking, vol. 7, no. 9, pp. 814–833, 2008.
72. G. Shen and R. S. Tucker, "Energy-Minimized Design for IP Over WDM Networks," Journal of Optical Communications and Networking, vol. 1, no. 1, pp. 176–186, 2009.
73. F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, "Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings," Zuse Institute Berlin, TKN, TU Berlin, Germany, Tech. Rep., 2010.
74. J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power Awareness in Network Design and Routing," in Proceedings of the 27th IEEE Conference on Computer Communications (INFOCOM), Phoenix, AZ, USA, 13-18 April. IEEE, Piscataway, NJ, USA, 2008, pp. 457 – 465.
75. B. Sanso and H. Mellah, "On reliability, performance and internet power consumption," in Proceeding of the 7th International Workshop on Design of Reliable Communication Networks, Washington, D.C., USA, October 25-28,. IEEE, Piscataway, NJ, USA, 2009, pp. 259 –264.
76. E. Palkopoulou, D. Schupke, and T. Bauschert, "Energy efficiency and CAPEX minimization for backbone network planning: Is there a tradeoff?" in Proceedings of the 3rd IEEE International Symposium on Advanced Networks and Telecommunication Systems (ANTS). IEEE, Piscataway, NJ, USA, 2009, pp. 1 –3.
77. W. Shen, Y. Tsukishima, K. Yamada, and M. Jinno, "Power-efficient multilayer networking: design and evaluation," in Proceedings of the 14th Conference on Optical Network Design and Modeling (ONDM). IEEE, Piscataway, NJ, USA, 2010, pp. 1 – 6.
78. I. Cerutti, L. Valcarenghi, and P. Castoldi, "Designing Power-Efficient WDM Ring Networks," in Networks for Grid Applications. Springer Berlin Heidelberg, 2010, vol. 25, pp. 101–108.
79. X. Xiao, A. Hannan, B. Bailey, and L. Ni, "Traffic Engineering with MPLS in the Internet," IEEE Network, vol. 14, no. 2, pp. 28–33, 2002.
80. F. Giroire, D. Mazauric, J. Moulierac, and B. Onfroy, "Minimizing Routing Energy Consumption: from Theoretical to Practical Results," in IEEE/ACM International Conference on Green Computing and Communications (Green- Com), Hangzhou, China, 2010.
81. W. Fisher, M. Suchara, and J. Rexford, "Greening backbone networks: reducing energy consumption by shutting off cables in bundled links," in Proceedings of the first ACM

SIGCOMM workshop on Green networking, New Delhi, India. ACM, New York, NY, 2010, pp. 29–34.

82. L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in Proceedings of the 2009 IEEE international conference on Communications, Dresden, Germany, 14-18 June. IEEE, Piscataway, NJ, USA, 2009, pp. 2298–2303.
83. L. Chiaraviglio, M. Mellia, and F. Neri, "Energy-aware backbone networks: a case study," in First International Workshop on Green Communications, in conjunction with IEEE International Conference on Communications, Dresden, Germany, 14-18 June, 2009, pp. 1 – 5.
84. M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-Aware Traffic Engineering," in Proceedings of the 18th IEEE International Conference on Network Protocols, 2010.
85. P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "Energy Aware Network Operations," in INFOCOM IEEE Conference on Computer Communications Workshops, 2009. IEEE, Piscataway, NJ, 2009, pp. 1 –6.
86. L. Chiaraviglio and I. Matta, "GreenCoop: cooperative green routing with energy-efficient servers," in Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, Passau, Germany, ACM, New York, NY, 2010, pp. 191–194.
87. A. P. Bianzino, C. Chaudet, F. Larroca, D. Rossi, and J.-L. Rougier, "Energy- Aware Routing: a Reality Check," in 3rd International Workshop on Green Communications (GreenComm3), in conjunction with IEEE GLOBECOM 2010, Miami, USA, 2010.
88. J. Restrepo, C. G. Gruber, and C. Machuca, "Energy profile aware routing," in First International Workshop on Green Communications, in conjunction with IEEE International Conference on Communications, Dresden, Germany, 14-18 June, 2009, 2009, pp. 1 –5.
89. N. Vasic and D. Kostic, "Energy-aware traffic engineering," in Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, Passau, Germany, 13 April. ACM, New York, NY, 2010, pp. 169–178.
90. A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An Energy Saving Routing Algorithm for a Green OSPF Protocol," in INFOCOM IEEE Conference on Computer Communications Workshops , 2010. IEEE, Piscataway, NJ, 2010, pp. 1 – 5.
91. K.-H. Ho and C.-C. Cheung, "Green distributed routing protocol for sleep coordination in wired core networks," in Proceedings of the 6th International Conference on Networked Computing (INC), Gyeongju, Korea, 11-13 May. IEEE, Piscataway, NJ, May 2010, pp. 1–6.
92. S. Orlowski, R. Wessaly, M. Pioro, and A. Tomaszewski, "SNDlib 1.0 - Survivable Network Design Library," Networks, vol. 55, pp. 276–286, 2010
93. Y. Chen and C. Chou, "Traffic Modeling of A Sub-network by using ARIMA," Info-tech and Info-net, Proceedings, ICII-2001. vol. 2, pp. 730-735, 2001.
94. H. Wang, H. Yang and H. Wu, "A Fine Model for Evaluating Output Performance of Crystalline Silicon Solar Modules," Photovoltaic Energy Conversion, IEEE 4th World Conference, vol. 2, pp. 2189-2192, 2006.
95. J. Zhao, A. Wang, P. P. Altermatt S. R. Wenham and M. A. Green, "24% Efficient Silicon Solar Cells," Photovoltaic Energy Conversion, Twenty Fourth. IEEE Photovoltaic Specialists Conference, vol. 2, pp. 1477-1480, 1994.
96. Cisco CRS-1 specification data sheet, http://www.cisco.com.
97. Glimmerglass System-600 data sheet, http://www. glimmerglass.com.
98. Cisco ONS 15454 data sheet, http://www.cisco.com.

99. BT Website, “Wind farms to power BT,” 2009, http://www.btplc.com/sharesandperformance/Annualreportandreview/Shareholdermagazine/May2009/windfarms.htm?Terms=380,4305
100. T. T. Ye, “Analysis of power consumption on switch fabrics in network routers,” proceedings of the 39th Design Automation Conference, New Orleans, pp. 524-529, U.S.A, June 10-14, 2002.
101. R. Hays, “Active/Idle Toggling with Low-Power Idle,” presentation at IEEE802.3az Task Force Group Meeting, January 2008.
102. IEEE802.3az, Energy Efficient Ethernet Study Group, http://www.ieee802.org/3/az/index.html.
103. A. Hopper and A. Rice, “Computing for the future of the planet,” Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, vol. 366, no. 1881, pp. 3685–3697, 2008.
104. D. Spence, J. Crowcroft, S. Hand, and T. Harris, “Location based placement of whole distributed systems,” in Proceedings of the 2005 ACM conference on Emerging network experiment and technology, Toulouse, France. ACM, New York, NY, pp. 124–134.
105. A. Moore, “An implementation-based comparison of Measurement-Based Admission Control algorithms,” Journal of High Speed Networks, vol. 13, no. 2, pp. 87–102, 2004.
106. T. Strayer, M. Allman, G. Armitage, S. Bellovin, S. Jin, and A. Moore, “IMRG workshop on application classification and identification report,” ACM SIGCOMM Computer Communication Review, vol. 38, no. 3, pp. 87–90, 2008.
107. H. Jiang, A. W. Moore, Z. Ge, S. Jin, and J. Wang, “Lightweight application classification for network management,” in Proceedings of the 2007 SIGCOMM workshop on Internet network management, Kyoto, Japan. ACM, New York, NY, 2007, pp. 299–304.
108. A. Hylick, R. Sohan, A. Rice, and B. Jones, “An analysis of hard drive energy consumption,” in Proceedings of the IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems. IEEE, Piscataway, NJ, USA, 2008, pp. 103–112.
109. A. Rice, S. Akoush, and A. Hopper, “Failure is an option,” Microsoft Research Technical Report MSR-TR-2008-61, 2008.
110. B. Scheuermann, W. Hu, and J. Crowcroft, “Near-optimal co-ordinated coding in wireless multihop networks,” in Proceedings of the 2007 ACM CoNEXT conference, New York, New York. ACM, New York, NY, 2007, pp. 9:1–9:12.
111. J. G. Koomey, “Estimating total power consumption by servers in the U.S. and the world,” Lawrence Berkeley National Laboratory, Tech. Rep., 2007.
112. A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, “Cutting the electric bill for internet-scale systems,” SIGCOMM Computer Communication Review, vol. 39, pp. 123–134, 2009.
113. B. Cully, G. Lefebvre, D. Meyer, M. Feeley, N. Hutchinson, and A. Warfield, “Remus: high availability via asynchronous virtual machine replication,” in Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation, ser. NSDI’08. Berkeley, CA, USA: USENIX Association, 2008, pp. 161–174.
114. D. J. C. MacKay, Sustainable Energy - Without the Hot Air. UIT, 2008.
115. [Online], “UK Renewable Energy News and Data Centers,” http://www.bbc.co.uk/news/uk-england-kent-11395972, 2010.

116. R. Miller, “Data Centers Heat Offices, Greenhouses, Pools,” http://www.datacenterknowledge.com/archives/2010/02/03/data-centersheat- offices-greenhouses-pools/, 2010.
117. E. Rodriguez-Colina, L. James, R. Penty, I. White, K. Williams, and A. Moore, “TCP Sending Rate Control at Terabits Per Second,” in Proceedings of High-Speed Networking Workshop: The Terabits Challenge, INFOCOM. IEEE Press, Piscataway, NJ, USA, 2006, pp. 1–5.
118. E. Williams, “Energy Intensity of Computer Manufacturing: Hybrid Assessment Combining Process and Economic Input- Output Methods,” Environmental Science & Technology, vol. 38, no. 22, pp. 6166–6174, 2004.
119. E.Wright, A. Azapagic, G. Stevens, W. Mellor, and R. Clift, “Improving recyclability by design: a case study of fibre optic cable,” Resources, Conservation and Recycling, vol. 44, no. 1, pp. 37–50, 2005.
120. D. G. Murray and S. Hand, “Scripting the cloud with skywriting,” in Proceedings of the 2nd USENIX conference on Hot topics in cloud computing, ser. HotCloud’10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–7.
121. A. W. Moore and K. Papagiannaki, “Toward the accurate identification of network applications,” in Passive and Active Network Measurement, ser. Lecture Notes in Computer Science, C. Dovrolis, Ed. Springer Berlin / Heidelberg, 2005, vol. 3431, pp. 41–54.
122. R. Mondragon, A. Moore, J. Pitts, and J. Schormans, “Analysis, simulation and measurement in large-scale packet networks,” IET Communications, Special Issue on Simulation, Analysis and Measurement of Broadband Network Traffic, vol. 3, no. 6, pp. 887–905, 2009.
123. M. Canini, W. Li, M. Zadnik, and A. W. Moore, “Experience with high-speed automated application-identification for network-management,” in Proceedings of the 5th ACM/IEEE Symposium on Architectures for Networking and Communications Systems, Princeton, New Jersey. ACM, New York, NY, 2009, pp. 209–218.
124. R. Bradford, E. Kotsovinos, A. Feldmann, and H. Schioberg, “Live wide-area migration of virtual machines including local persistent state,” in Proceedings of the 3rd international conference on virtual execution environments, San Diego, California, USA. ACM, New York, NY, 2007, pp. 169–179.
125. R. Clegg, M. Withall, A. Moore, I. Phillips, D. Parish, M. Rio, R. Landa, H. Haddadi, K. Kyriakopoulos, J. Aug´e et al., “Challenges in the capture and dissemination of measurements from high-speed networks,” IET Communications, Special Issue on Simulation, Analysis and Measurement of Broadband Network Traffic, vol. 3, no. 6, pp. 957–966, 2009.
126. A. Moore, J. Hall, C. Kreibich, E. Harris, and I. Pratt, “Architecture of a network monitor,” in Proceedings of the Fourth Passive & Active Measurement Workshop 2003, 2003.
127. Sherif Akoush, Ripduman Sohan, Andrew Rice, Andrew W. Moore and Andy Hopper, “Predicting the Performance of Virtual Machine Migration” International Symposium on Modeling, Analysis, and Simulation of Computer Systems, 2010, issue 1526-7539, pp. 37-46, Los Alamitos, CA, USA.
128. D. Fay, H. Haddadi, A. Thomason, A. W. Moore, R. Mortier, A. Jamakovic, S. Uhlig, and M. Rio, “Weighted spectral distribution for internet topology analysis: theory and applications,” IEEE/ACM Transactions on Networking, vol. 18, pp. 164–176, 2010.
129. E. Thereska, A. Donnelly, and D. Narayanan, “Sierra: a power-proportional, distributed storage system,” Microsoft Research Technical Report MSR-TR- 2009-153, November 2009.

130. Tucker, R. S.; , "Green Optical Communications---Part I: Energy Limitations in Transport," Selected Topics in Quantum Electronics, IEEE Journal of , vol.PP, no.99, pp.1-16, 0
131. Tucker, R. S.; , "Green Optical Communications---Part II: Energy Limitations in Networks," Selected Topics in Quantum Electronics, IEEE Journal of , vol.PP, no.99, pp.1-14, 0
132. http://www.avagotech.com/docs/5989-4527EN
133. S.B. Poole, D.N. Payne, M.E. Fermann, "Fabrication of low-loss optical fibres containing rare-earth ions," Electronics Letters , vol.21, no.17, pp.737-738, August 15 1985
134. E. Desurvire, J. Simpson, and P.C. Becker, High-gain erbium-doped traveling-wave fiber amplifier," Optics Letters, vol. 12, No. 11, 1987, pp. 888–890
135. IEEE Std 802.3
136. IEEE Std 802.3ae
137. A.B. Massara, K.A. Williams, I.H. White, R.V. Penty, A. Galbraith, P. Crump, P. Harper, "Ridge waveguide InGaAsP lasers with uncooled l0 Gbit/s operation at 70°C," Electronics Letters , vol.35, no.19, pp.1646-1647, 16 Sep 1999
138. J.K. White, C. Blaauw, P. Firth, P. Aukland, "85°C investigation of uncooled 10-Gb/s directly modulated InGaAsP RWG GC-DFB lasers," Photonics Technology Letters, IEEE , vol.13, no.8, pp.773-775, Aug 2001
139. R. Paoletti, M. Agresti, G. Burns, G. Berry, D. Bertone, P. Charles, P. Crump, A. Davies, R.Y. Fang, R. Ghin, P. Gotta, M. Holm, C. Kompocholis, G. Magnetti, J. Massa, G. Meneghini, G. Rossi, P. Ryder, A. Taylor, P. Valenti, M. Meliga, "100°C, 10 Gb/s directly modulated InGaAsP DFB lasers for uncooled Ethernet applications," ECOC '01., vol.6, no., pp. 84- 85 vol.6, 2001
140. D. Klotzkin, K. Kojima, N. Jordache, N. Chand, P. Kiely, M. Chien, M. Han, E. Michel, S. Ustin, S. Roycroft, D. Melville, R. Kunkel, L. Ketelsen, "High-speed directly modulated Fabry-Perot and distributed-feedback spot-size-converted lasers suitable for passive alignment, unisolated operation, and uncooled environments up to 85°C," Lightwave Technology, Journal of , vol.21, no.1, pp. 69- 78, Jan 2003
141. A.C. O'Donnell, N.J. Parsons, "1*16 lithium niobate optical switch matrix with integral TTL compatible drive electronics," Electronics Letters , vol.27, no.25, pp.2367-2368, 5 Dec. 1991
142. A. Kazama, Y. Itou, M. Horino, K. Fukuda, M. Kanamaru, T. Akashi, T. Ishikawa, T. Harada, R. Okada, "Low-insertion-loss compact 3D-MEMS optical matrix switch with 18 input/output ports," Solid-State Sensors, Actuators and Microsystems, 2005. Digest of Technical Papers. TRANSDUCERS '05., vol.1, no., pp. 972- 975 Vol. 1, 5-9 June 2005
143. M. Mizukami, J. Yamaguchi, N. Nemoto, Y. Kawajiri, H. Hirata, S. Uchiyama, M. Makihara, T. Sakata, N. Shimoyama, H. Ishii, F. Shimokawa, "128×128 3D-MEMS optical switch module with simultaneous optical paths connection for optical cross-connect systems," Photonics in Switching, 2009. PS '09., vol., no., pp.1-2, 15-19 Sept. 2009
144. S. Sohma, T. Goh, H. Okazaki, M. Okuno, A. Sugita, "Low switching power silica-based super high delta thermo-optic switch with heat insulating grooves," Electronics Letters , vol.38, no.3, pp.127-128, 31 Jan 2002
145. N. Xie, T. Hashimoto, K. Utaka, "Ultimate-low-power-consumption, polarization-independent, and high-speed polymer Mach-Zehnder thermo-optic switch," OFC 2009., vol., no., pp.1-3, 22-26 March 2009
146. W. Payne III, "Design of lithium niobate based photonic switching systems," Communications Magazine, IEEE , vol.25, no.5, pp. 37- 41, May 1987

147. P.J. Duthie, M.J. Wale, "16*16 single chip optical switch array in lithium niobate," Electronics Letters , vol.27, no.14, pp.1265-1266, 4 July 1991
148. P. Granestrand, B. Lagerstrom, P. Svensson, H. Olofsson, J.-E. Falk, B. Stoltz, "Pigtailed tree-structured 8×8 LiNbO3 switch matrix with 112 digital optical switches," Photonics Technology Letters, IEEE , vol.6, no.1, pp.71-73, Jan 1994
149. I.H. White, J.J.S. Watts, J.E. Carroll, C.J. Armistead, D.J. Moole, J.A. Champelovier, "InGaAsP 400*200 mu m active crosspoint switch operating at 1.5 mu m using novel reflective Y-coupler components," Electronics Letters , vol.26, no.10, pp.617-618, 1 May 1990
150. G. Sherlock, J.D. Burton, P.J. Fiddyment, P.C. Sully, A.E. Kelly, M.J. Robertson, "Integrated 2*2 optical switch with gain," Electronics Letters , vol.30, no.2, pp. 137- 138, 20 Jan 1994
151. J.D. Burton, P.J. Fiddyment, M.J. Robertson, P. Sully, "Monolithic InGaAsP-InP laser amplifier gate switch matrix," Quantum Electronics, IEEE Journal of , vol.29, no.6, pp.2023-2027, Jun 1993
152. Y. Kai, K. Sone, S. Yoshida, Y. Aoki, G. Nakagawa, S. Kinoshita, "A compact and lossless 8 × 8 SOA gate switch subsystem for WDM optical packet interconnections," ECOC 2008, vol., no., pp.1-2, 21-25 Sept. 2008
153. H. Wang, A. Wonfor, K.A Williams, R.V. Penty, I.H. White, "Demonstration of a lossless monolithic 16×16 QW SOA switch," ECOC '09 Post-deadline, no., pp.1-2, 20-24 Sept. 2009
154. M. Glick, M. Dales, D. McAuley, L. Tao K. Williams, R. Penty, I. White, "SWIFT: a testbed with optically switched data paths for computing applications," Transparent Optical Networks, 2005, Proceedings of 2005 7th International Conference , vol.2, no., pp. 29- 32 Vol. 2, 3-7 July 2005
155. Q. Yang, K. Bergman, G.D. Hughes, F.G. Johnson, "WDM packet routing for high-capacity data networks," Lightwave Technology, Journal of , vol.19, no.10, pp.1420-1426, Oct 2001
156. A. Shacham, B.A. Small, O. Liboiron-Ladouceur, K. Bergman, "A fully implemented 12 × 12 data vortex optical packet switching interconnection network," Lightwave Technology, Journal of , vol.23, no.10, pp. 3066- 3075, Oct. 2005
157. http://europic.jeppix.eu/
158. H. Wang, A. Wonfor, S. Liu, R.V. Penty, I.H. White, "Novel 4×4 port integrated SOA switch and wavelength converter," Photonics in Switching, 2009., pp.1-2, 15-19 Sept. 2009
159. H. Wang, E.T. Aw, M. Xia, M.G. Thompson, R.V. Penty, I.H. White, A.R. Kovsh, "Temperature Independent Optical Amplification in Uncooled Quantum Dot Optical Amplifiers," OFC/NFOEC 2008, vol., no., pp.1-3, 24-28 Feb. 2008
160. E.T. Aw, H. Wang, M.G. Thompson, A. Wonfor, R.V. Penty, I.H. White, A.R. Kovsh, "Uncooled 2x2 quantum dot semiconductor optical amplifier based switch," CLEO/QELS 2008., vol., no., pp.1-2, 4-9 May 2008
161. J.E. Bowers, A.W. Fang, R. Hyundai Park Jones, O. Cohen, M.J. Paniccia, "Hybrid Silicon Evanescent Photonic Integrated Circuit Technology," Lasers and Electro-Optics, 2007. CLEO 2007. Conference on , vol., no., pp.1-2, 6-11 May 2007
162. D. J. Thomson, F. Y. Gardes, G. T. Reed, F. Milesi, and J-M Fedeli, "High speed silicon optical modulator with self aligned fabrication process," Opt. Express 18, 19064-19069 (2010)
163. A. Alduino, L. Liao, R. Jones, M. Morse, B. Kim, W. Lo, J. Basak, B. Koch, H. Liu, H. Rong, M. Sysak, C. Krause, R. Saba, D. Lazar, L. Horwitz, R. Bar, S. Litski, A. Liu, K. Sullivan, O. Dosunmu, N. Na, T.Yin, F. Haubensack, I. Hsieh, J. Heck, R. Beatty, H. Park, J.

Bovington, S. Lee, H. Nguyen, H. Au, K. Nguyen, P. Merani, M. Hakami, M. Paniccia, “Demonstration of a High Speed 4-Channel Integrated Silicon Photonics WDM Link with Hybrid Silicon Lasers,” Photonics in Switching 2010, PDIWI5, 25-28 July 2010

164. M. Lipson, "Silicon photonics: the optical spice rack," Electronics Letters , vol.45, no.12, pp.576-578, June 4 2009
165. B.G. Lee, A. Biberman, J. Chan, K. Bergman, "High-Performance Modulators and Switches for Silicon Photonic Networks-on-Chip," Selected Topics in Quantum Electronics, IEEE Journal of , vol.16, no.1, pp.6-22, Jan.-feb. 2010