Time for a "Greener" Internet

Mario Baldi
Dept. of Control and Computer Engineering
Politecnico di Torino (Technical University of Turin)
Turin, Italy

Yoram Ofek
Dipartimento di Informatica e Telecomunicazioni
Università di Trento (University of Trento)
Trento, Italy

Abstract—**It is anticipated that the Internet traffic will continue to grow exponentially for the foreseeable future, which will require ever-growing energy (electricity). Since a lot of the Internet traffic growth comes from predictable services (such as video) there is a huge potential for decreasing future Internet energy requirements by synchronizing the operation of routers and scheduling traffic in advance, thus reducing complexity (e.g., header processing, buffer size, switching fabric speedup and memory access bandwidth speedup). Today, scheduling and synchronizing large-scale data transfer operations can be easily achieved by utilizing a choice of tens of global time sources, freely available on earth and in space. In a way, this manuscript shows how to "trade" global time for electricity utilized by the global Internet.**

Keywords—green Internet, low-power network protocols and architectures, low-power switches, pipeline forwarding, optical networking, hardware complexity, speedup.

I. INTRODUCTION

THE Internet has been growing steadily for several years and with it the amount of electricity required to operate the network itself, the servers offering services available through it, and the devices enabling access to those services virtually everywhere and around the clock. Given that the growth trend is not changing, such huge electricity consumption has become significant, especially when current world environmental issues are taken into account.

Actions to curb power consumption can be taken at various levels ranging from the applications and the operating systems (i.e., software) running on servers, in data centers and on end-devices, to the hardware, such as: storage, displays, electronic components and cooling sub-systems deployed in various parts of the system. One important source of power consumption is the network infrastructure itself with routers and switches processing, switching, and buffering an increasing amount of IP packets and transmitting them through links at very high bit rate.

Moreover, during the next ten to fifteen years the Internet is envisioned to undergo several major transitions with respect to technologies, services and size. Bandwidth demands are expected to grow up to 50-100 times [1]. These transitions are led by the evolution of the Internet into a global *quadruple-any* network supporting *any* service for *anyone* at *anytime* and *anywhere*. Specifically, anyone implies that users of all ages and technical skills will be able to access the network; while anytime implies "on-demand" services that will replace scheduled broadcasting services to home users. Thus, most of the Internet growth (possibly more than 90%) will be "driven" and "consumed" by video-based streaming and downloading (see [1]) to home users who are not willing to pay very much. Consequently, in order to enable proper growth of the Internet it is of key importance to maximize both the *energy-scalability index* and *cost-scalability index* of the future Internet. The former is defined as the ratio between the relative network growth and the corresponding relative energy consumption increase; the latter is defined likewise where the divisor is the corresponding relative cost increase. Maximizing the two scalability indexes leads to the following two questions:

1. How to significantly reduce electricity utilization by the network, currently estimated to be about 1.5% of the overall world electricity consumption [12];
2. How to significantly reduce the cost of construction and deployment of network equipment.

This paper outlines a solution to addresses the problem in various complementary ways: reduce per packet processing, limit memory requirements in all network devices, deploy larger switching units (thus reducing the switch reconfiguration frequency), enable full link utilization (thus enabling traffic to be concentrated on fewer optical channels, while the others can be put in a low-power idle mode), enable realization of dynamic all-optical switching. All the above is enabled by utilizing global time in various aspects of the very packet switching process, which implies that network devices are synchronized to the same time-of-day, to realize *pipeline forwarding* of IP packets [1]. Although pipeline forwarding maximizes both the energy-scalability and cost-scalability index, the focus of this work is on the former.

The necessary condition for pipeline forwarding in having the same clock reference in all routers, which can be realized in a global network by using global time or UTC (coordinated universal time) [2], in order to enable an optimal packet forwarding for periodic traffic, such as video. As widely demonstrated in the literature [9][10], pipeline forwarding facilitates the following optimal complexity switch design

[1] Pipeline forwarding is an optimal method that is widely used in manufacturing, e.g., automotive assembly line, and computing, e.g., micro-processor architecture.

[2] It is worthwhile noting that UTC is globally available via tens of sources on earth and in space—e.g., GPS (USA), GLONASS (Russian), Galileo (EU) and Beidou (China). Furthermore, time can be distributed through the network itself (i.e., in-band) using variety of mechanisms and standard protocols—e.g., ITU-T Study Group 15, recommendations for synchronization over packet switched network, IETF: NTPv4 & SNTPv4, IEEE 1588 – PTP (precision time protocol). In addition, many wireless systems distribute time as part of their protocol—e.g., CDMA, WiMAX, Femto-Cell, etc.

features: (*i*) no buffers, (*ii*) no header processing, (*iii*) optimal speedup of 1, (*iv*) switching complexity on the order of *NlogN* (where *N* is the size of the switch) and (*v*) full optical channel utilization.

Consequently, such an optimal switch design has a 10-100 times larger *energy-scalability index* in comparison with current switch designs. Especially, such design can be easily implemented in the all-optical switching domain with links at more than 100 Gb/s, which further reduces the energy requirements per unit of bandwidth. A pipeline forwarding switch design scalable to 10-100 Tb/s in a single chassis was successfully implemented in our laboratory [11].

The paper also proposes a network architecture based on two independent, yet tightly integrated, parallel sub-networks: (1) the current Internet and (2) "super-highways" deploying pipeline forwarding of IP packets. Such a parallel networks architecture enables (*i*) a gradual deployment of the novel pipeline forwarding technology in the current Internet and (*ii*) engineering the bulk of the traffic (mostly video-based), perhaps more than 90%, on the super-highway where power consumption is minimized.

Power consumption has come to the forefront of the scientific discussion as a major issue for the future Internet that is widely recognized to have to be "green" (see for example [2]-[7]). Most of the existing works are focusing on reducing power consumption in end systems, mainly on ways to put sub-system into sleep mode or slowing clocks in some parts of the end system [6]. Other works have focused on reducing energy consumption by data centers and servers [7]. Using photonic technology for energy reduction in the network was proposed in [3].

However, none of the current works considered global time and pipeline forwarding as method for reducing the Internet energy utilization. Consequently, this work should be viewed as complementary to the on-going "Green Internet" research and development.

The paper is organized as follows; Section II discusses some basic enabling technologies, Section III presents a possible extension to current Internet architecture and a plausible gradual deployment strategy, Section IV analyzes the *energy-scalability index*, while Section V presents some conclusions and future directions.

II. ENABLING TECHNOLOGIES

Pipeline Forwarding of IP packets was introduced in [8] and is used in the context of this work to create an architectural solution to reduce power consumption in the Internet, i.e., to realize a "Green Internet". IP packet switches are synchronized and use a basic time period called time frame (TF). A TF can be viewed as a *virtual container* for IP packets. The TF duration can be either derived from an external source — e.g., the UTC second provided by positioning systems such as GPS and Galileo — or distributed through the network. The transmission capacity during each TF can be partially or totally reserved to one or more flows during a resource reservation procedure. Time frames are grouped into time cycle that provides the basis for a periodic repetition of the reservation. This results in a periodic schedule for packets to be switched and forwarded, which is repeated every time cycle. For example, in Figure 1, 100 time frames of duration T_f are grouped into one time cycle.

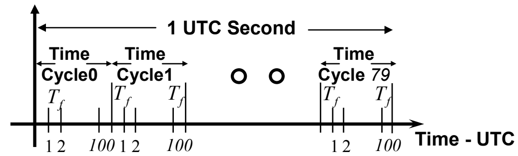

Figure 1: Common time reference structure.

The basic *pipeline forwarding* operation is regulated by two simple rules: (*i*) all packets that must be sent in TF *t* by a node must be in its output ports' buffers at the end of TF *t*-1, and (*ii*) a packet *p* transmitted in TF *t* by a node *n* must be transmitted in TF $t+d_p$ by node *n*+1, where d_p is an integer constant called *forwarding delay*, and TF *t* and TF $t+d_p$ are also referred to as the *forwarding TF* of packet *p* at node *n* and node *n*+1, respectively. The value of the forwarding delay is determined at resource-reservation time, involves solving a scheduling problem, and must be large enough to satisfy (*i*) given the propagation delay between nodes *n* and *n*+1 and the processing and switching time of node *n*+1.

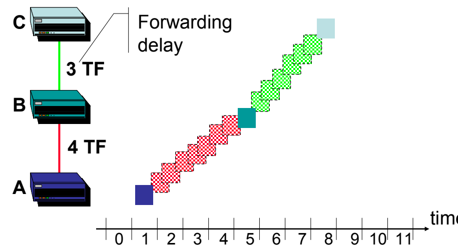

Figure 2: Pipeline forwarding operation.

The periodic scheduling within each node results in a *periodic packet forwarding* across the network so that packets move in the network as in a pipeline (from which the name *pipeline forwarding* or packets), as shown in Figure 2, in an ordered and predictable manner. The quality of the service received by reserved traffic is deterministic since congestion is avoided and the end-to-end delay is consequently known in advance with an uncertainty (i.e., jitter) of less than 1 TF. For example, in the network scenario shown in Figure 2, the delay a packet experiences from node A to node B is 7 TFs.

In *pipeline forwarding*, a *synchronous virtual pipe* (SVP) is a predefined schedule for forwarding a pre-allocated amount of bytes during one or more TFs along a path of subsequent UTC-based switches. The SVP capacity is determined by the total number of bits allocated in every TC for the SVP. For example, for a 10 ms TC, if 20000 bits are allocated during each of 2 TFs, the SVP capacity is 4 Mb/s.

Two implementations of the *pipeline forwarding* have been proposed so far and are described in the following.

A. Time-driven IP switching

Time-driven switching (TDS) was proposed to realize sub-lambda or fractional lambda switching (FλS) in highly scalable dynamic optical networking [9], which requires minimum (possibly optical) buffers. In the context of optical networks, SVPs are called fractional lambda pipes (FλPs).

In TDS/ FλS all packets in the same TF are switched in the same way, i.e., to the same output port. Consequently, header processing is not required, which results in low complexity (hence high scalability) and enables optical implementation. Scheduling through a switching fabric is based on a pre-defined schedule, which enables the implementation of a simple controller. Moreover, low-complexity switching fabric architectures, such as Banyan, can be deployed notwithstanding their blocking features, thus further enhancing scalability. In fact, blocking can be avoided during schedule computation by avoiding conflicting input/output connections during the same TF. Previous results [9][10] show that (especially if multiple wavelength division multiplexing channels are deployed on optical links between fractional λ switches) high link utilization can be achieved with negligible blocking using a Banyan network without speedup.

B. *Time-driven priority*

TDS is suitable for very high speed (optical) backbones, where traffic can be organized in large capacity SVPs handled by high performance switches. Where more flexibility is required, it can be provided by *Time-driven priority* (TDP) [8] that combines *pipeline forwarding* with conventional IP routing. Packets entering a switch from the same input port during the same TF can be sent out from different output ports, according to the rules that drive IP packet routing. Operation in accordance to *pipeline forwarding* principles ensures deterministic quality of service and low complexity packet scheduling. Specifically, packets scheduled for transmission during a TF are given maximum priority; if resources have been properly reserved, all scheduled packets will be at the output port and transmitted before their TF ends.

III. Architecture and Deployment Strategy

Pipeline forwarding seems promising from both a general networking point of view (e.g., deterministic service and scalability), as discussed in the previous section, and an electricity saving point of view, as discussed in Section IV. However, the wide deployment of the current Internet based on asynchronous packet switching makes any changes to its working principles, and consequently to network devices implementing them, impractical. Moreover, the current asynchronous Internet is well suited to traditional data traffic, e.g., e-mail messages, file transfers, web browsing, etc.

The solution is the creation of a *parallel network* (on the same fiber infrastructure with WDM) based on *pipeline forwarding* that coexists with asynchronous IP technology, i.e., currently commercially available IP routers, as shown in Figure 3. Besides carrying typically "best-effort" traffic, such as e-mails, low priority web browsing and file transfers, traditional (asynchronous) IP routers are used to transport the signaling required to setup synchronous virtual pipes in the pipeline forwarding parallel network. The latter carries traffic requiring a deterministic service, such as phone calls, video on demand, videoconferencing, and distributed gaming. As previously mentioned, given the large bandwidth required by most of such video-based services, and especially the ones still to come in the near future, such as 3D video and virtual reality, they are expected to account for more than 90% of future Internet traffic. The current IP infrastructure or its evolution, notwithstanding its limited scalability, is able to support the remaining traffic as it will be only a small fraction (10% of the total) and will not require a service with deterministic quality, i.e., it will basically require the same service (i.e., "best-effort" or differentiated) and capacity currently provided by the Internet. Moreover, the traditional IP network ensures connectivity to the large number of currently existing end-devices.

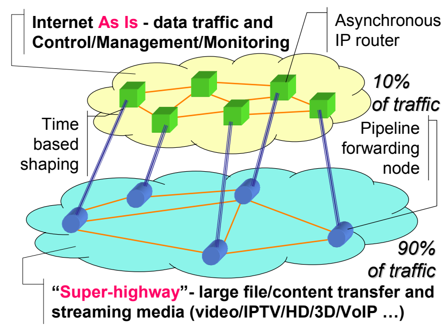

Figure 3: Parallel networks on the same fiber infrastructure with WDM.

The pipeline forwarding parallel network will be a sort of *"super-highway"* as it will carry a large fraction of the traffic with deterministic performance, i.e., packets will flow fast and smoothly through it on an SVP setup for their flow. Moreover, similarly to a system of roads and highways the two parallel networks will have various interconnections through which traffic originated from (directed to) end-devices can enter (exit) the "super-highway". *Time based shaping* is performed on packets entering the pipeline forwarding parallel network in order to schedule their transmission during one of the TFs reserved to the flow they belong to.

As an independent infrastructure, the pipeline forwarding parallel network can be built and deployed incrementally or gradually in islands whenever and wherever it will be most beneficial. For example, deployment of the pipeline forwarding parallel network can start from high density metro areas for then moving to their interconnection and finally to less populated rural areas.

The traditional IP network and the pipeline forwarding parallel network can share the same fiber infrastructure where different DWDM (Dense Wavelength Division Multiplexing) optical channels are deployed for interconnecting traditional (asynchronous) and pipeline forwarding IP routers. As time passes and network equipment is upgraded, pipeline forwarding-unaware routers can be replaced with devices supporting the novel technology; as a result, the two parallel networks are incrementally integrated into a single network infrastructure.

A G-MPLS (Generalized Multi-Protocol Label Switching) control plane can be used for the pipeline forwarding network. Especially in the initial phase of incremental realization of the pipeline forwarding network, the control plane can be run on control nodes within the traditional IP network. This solution allows the deployment of routers that implement only the data plane, i.e., the core pipeline forwarding functionalities, which enables to reduce time to market and cost of pipeline forwarding equipment. A control node, being directly or indirectly interconnected to one or more pipeline forwarding routers, uses a G-MPLS routing protocol, such as OSPF, IS-IS, and BGP, to exchange information with other control nodes through the traditional IP network. RSVP messages among control nodes are deployed to setup and tear down SVPs across the pipeline forwarding network.

IV. Energy-scalability Benefits

There are several ways in which pipeline forwarding provides energy-scalability benefits; the proposed parallel networks architecture enables immediate exploitation of such benefits and brings some additional ones. Each of the following subsections discusses one area in which the energy-scalability index is increased through the deployment of pipeline forwarding.

Before delving into these specifics, it is worth emphasizing once more how each integrated circuit, independently of its function (i.e., whether it is a memory, or a processor, or a switch) is composed of gates, each one draining power and generating heat. Cooling off such heat possibly requires fans and air conditioning that also drain energy[3]. Consequently, any technological and design solution that entails reducing either (*i*) the number of chips or (*ii*) their size in terms of number of gates or (*iii*) the frequency of state changes (i.e., clock speed) contributes to increasing the energy-scalability index (i.e., the amount of data handled with a given amount of hardware) and thereby reducing power consumption and the cost of various networking equipment.

A. Memory requirements

In pipeline forwarding time is deployed to coordinate the forwarding operation in switches and routers so that contention for resources is reduced or avoided, both within routers — e.g., interface processors, switching elements, busses — and between routers — i.e., link transmission capacity. When IP packets contend for a resource, one of them accesses it, while all the others are buffered. Due to the difficulty in controlling contention in asynchronous routers buffers are very large. In general, the higher the amount of traffic routers should handle, i.e., the higher the number of interfaces and their capacity, the larger the amount of buffering required avoiding or reducing packet drop.

[3] It is estimated that in large systems the electricity required for cooling equals the electricity utilized for the actual operation.

As buffers are obviously implemented by using memory chips, pipeline forwarding enables reducing the amount of such integrated circuits on input and/or output interfaces, depending on the router architecture and within the switching fabric.

Moreover, in many traditional router designs such buffers need to operate faster than the interface rate in order to enable multiple functions operating simultaneously (asynchronously) to access packets stored in the same buffer (memory). This feature is called *memory/buffer access speedup*. For example, if the network interface bandwidth is 10 Gb/s and due to the asynchronous nature of the switch/router operation an output buffer should be able to "simultaneously" store packets received from, say, three interfaces, memory will have to operate three times faster so that the three packets can be stored, one after the other, in the same amount time they were received. This memory access speedup of three can be implemented by timing the memory with a three time faster clock or by using a three time wider memory word; in both cases, roughly three times more energy is drained by the memory chip.

The example above shows the huge energy saving benefit pipeline forwarding brings since it does not require memory access speedup above the optimal value of one. Pipeline forwarding requires minimum speedup simply because at any given time there is only one packet contending to access the memory, which is obviously the result of pre-scheduling traffic through the network and having network nodes forwarding packets in a coordinated fashion.

B. Packet processing within TFs

In TDS/FλS IP packets are switched and forwarded inside a TF — that is a virtual container. The content of each TF is switched and forwarded according to its position within the time cycle, and therefore, there is no need for header processing. In TDP packets received during the same TF can be switched to different output interfaces; consequently header processing is required for properly routing packets.

Moreover, pipeline forwarding significantly reduces the hardware complexity, and consequently power requirements, of the schedulers on input/output interfaces and the switching fabric controller, when compared to the ones deployed on asynchronous switches/routers

Since fewer processor cycles are needed to handle each packet, various strategies can be applied to increase the energy-scalability index:

- Processors can be clocked at a lower frequency, thus utilizing less power;
- Processors can be left longer in low power idle/sleep mode.

C. Switching fabric reconfiguration frequency and speed up

Using whole TFs as switching units from input to output enables reducing the reconfiguration frequency of the switching fabric of TDS switches. In fact, the switching fabric is to be reconfigured at the beginning of each TF to setup the connection required for moving the multiple packets received during the TF from scheduled input to scheduled output.

In comparison, an asynchronous router will potentially reconfigure the switching fabric before moving each single packet because packets to be switched from the same input to the same output do not necessarily cluster together before arriving to the switch. Since changing the configuration of switching elements drains more power than leaving them in the same configuration, TDS increases the energy-scalability index of the router.

As discussed in Section IV.A for memory, avoiding contention on output ports, i.e., having more than one packet to be switched to the same output port, eliminates the need for speedup in the switching fabric. Again, lower rate operation implies lower power consumption.

D. Number of switching elements

In order to maximize the switching fabric scalability it is necessary to minimize the switching fabric complexity in terms of number of switching elements, which also minimizes cost and power consumption. The lowest complexity fabric are multistage Banyan interconnection networks, with switching complexity of $O(a \cdot N \cdot \log_a N)$, where N is the number of ports and a is the number of inlets/outlets of each switching element. Banyan is known to suffer from *space blocking*, which means that a connection between an available input and an available output may not be possible because one element on the route through the switch interconnection network is being used for another connection. However, as shown in [9][10], pipeline forwarding of IP packets limits the consequences of space blocking. Intuitively, the multiple TFs in each time cycle provide an additional degree of freedom for scheduling. Namely, if there are K TFs in each time cycle, there are K different possible switching instances, i.e., input/output connection configurations, that can be used to move data through the switch; while a specific input/output connection might be impossible in one instance, it might be in another one. In addition, non-immediate forwarding of TFs can be used [9], which implies that packets in a virtual container may be delayed one or more TFs before being switched and forwarded. In summary, the combination of plurality of TFs in each time cycle together with non-immediate forwarding enable the efficient deployment of Banyan-based switching fabrics with optimal switching complexity.

TABLE I. POWER DISSIPATION OF ELECTRONIC CROSS-POINT SWITCHES

Switching element	Maximum bit rate	Power per input-output	Power per ch. per Gb/s
Vitesse VSC3040 144 ch.	11 Gb/s	0.11 W	0.01 W
Mindspeed M21151 144 ch.	3.2 Gb/s	0.15 W	0.05 W

Figure 4 shows a possible Banyan-topology switching fabric configuration based on a Mindspeed M21151 cross-point switch with an aggregate switching capacity of 10 Tb/s. A Banyan-based switching fabric with aggregate switching capacity of 160 Tb/s (128·128·10 Gb/s) can be built with the architecture shown in Figure 4 using a state-of-the-art cross-point, the VSC3040 by Vitesse with 144 ports at up to 11Gb/s (the M21151 has 3.2 Gb/s ports).

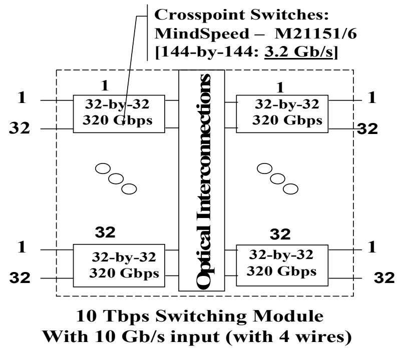

Figure 4: 10 Tb/s Switching Module with M21151

Table 1 shows the power requirement of Mindspeed (M21151) and Vitesse (VSC3040) cross-point switches. The power requirement per Gb/s of switching capacity is orders of magnitude smaller than the one of the top-of-the-line Cisco Systems's router, CRS-1, that requires 10-20W per Gb/s of switching capacity. While of course the latter figure includes the power drained by memory and processors, as discussed in the previous subsection, the buffering and processing requirements of TDS are very small compared to the ones of asynchronous routers, such as CRS-1.

E. Link utilization with efficient provisioning

TDS/FλS enables deterministic service, with negligible packet loss, at high link utilization — more than 90% with negligible blocking. When deploying multiple wavelengths on the same physical link, power savings can be achieved simply by utilizing an optical channel to its full capacity before lighting another one. Moreover, if energy-aware traffic engineering is used, as outlined in Section IV.G, during times of low network utilization (e.g., at night or during weekends) SVPs can be concentrated on limited number of fully loaded channels while turning off the others and possibly entire subnetworks.

In comparison, when DiffServ is deployed to ensure service quality link utilization may be only 20-30%. Also TCP/IP "best-effort" traffic significantly under utilizes the network optical channels due to the congestion avoidance algorithm of the transport protocol. Consequently, aforementioned solutions to reduce power consumption cannot be realized over the current Internet.

F. All-optical switching

Two functions are hard to implement in the all-optical domain: (1) memory/buffer and (2) header processing. However, not requiring such functions, FλS can be realized in the all-optical switching domain. An all-optical FλS switch was realized in our laboratory; although it includes an electronic control system, data are routed without converting them into electrical signals. All-optical switching elements use less power than electric ones. For example, our all-optical FλS implementation requires only 0.0052 W per Gb/s when switching 100 Gb/s optical channels.

G. *Energy-aware Traffic Engineering and Reservation*

As pointed out in [2], coordinated decisions to put nodes and links into a low power sleep mode can result in major power savings. Time has been widely used as a means of distributed coordination. In wireless sensor networks, for example, it is being used to coordinate medium access and data forwarding in a way that nodes can maximize the duration of their idle periods and consequently make it worth to enter a low power mode to minimize energy consumption. In pipeline forwarding time coordinates packet forwarding on a global scale and similarly to wireless sensor networks, it can be leveraged also to reduce power consumption.

For example, when SVPs/FλPs are setup through lightly loaded links, adjacent TFs could be scheduled. As a result, the time cycle will be organized as a sequence of reserved TFs, followed by a sequence of empty TFs. It can be worth powering off links with a large number of subsequent empty TFs until the next reserved TF. Ideally, the network power consumption will be proportional to the traffic actually forwarded.

Moreover, energy-aware traffic engineering might be deployed to route SVPs/FλPs on a limited number of links, which thanks to pipeline forwarding efficiency can be reserved up to 90% of their capacity (see Section IV.E). During low load periods, no traffic travels on parts of the parallel pipeline forwarding network and the corresponding links and nodes can be set into a low power sleep mode. The proposed parallel networks architecture avoids issues stemming from protocols that require periodic packet exchanges, such as routing protocols, thus eliminating the need for sophisticated ad-hoc solutions such as the ones proposed in Section 4.2 of [2]. In fact, while nodes and interfaces in the parallel pipeline forwarding network are possibly set into a low power sleep mode, routing messages and other control data is exchanged among control nodes on the parallel traditional IP network.

V. CONCLUSION

As Internet traffic increases, its power utilization grows. This growth in energy consumption must be considered seriously in any future Internet design, otherwise it may constrain the growth of the Internet itself. By some estimates the current Internet consumes about 1.5% of the world electricity [12]. This energy footprint is expected to grow significantly as Internet traffic is predicted to grow in volume 50 to 100 times in the next decade, as foreseen by some studies, as foreseen by some studies [1]. Thus, the electricity consumption issue brings a new dimension to the Internet design. Various studies have shown that current IP packet switching is not energy efficient. Instead, the time-based IP switching approach realized by pipeline forwarding, and specifically time-driven switching (TDS)/fractional lambda switching (FλS), recently demonstrated by our novel testbed implementation [11] significantly reduces the electricity requirement. The strength of FλS/TDS is its simplicity, while not compromising three most desired properties for the future "green" Internet:

1. Very low energy (electricity) consumption;
2. Sub-lambda switching scalability to orders of 10 and 100 Tb/s in a single chassis (i.e., with limited hardware); and
3. Predictable quality of service (QoS) for streaming media, live (sport and entertainment) events and large file transfers for content distribution and grid computing.

High energy-scalability index implies less hardware per unit of switching capacity, which consequently reduces component and electricity costs. In future studies we will quantify more precisely this electricity reduction in both:

1. Routing and switching, which is manifested by a major reduction in the amount of hardware (a factor of 10-50) due to the elimination of buffers and header processing; and
2. Transmission, which is manifested by the ability to fully utilize each optical channel with sub/fractional lambda switching (FλS), which is not the case with both full-lambda switching (i.e., lambda routing) and traditional (i.e., asynchronous) packet switching.

Furthermore, pipeline forwarding enables to extend the cost/energy efficient time-based IP packet switching all the way to the edges of the network, specifically, where the "last-mile" begins. In fact, several "last-mile" technologies, such as, cable modem and passive optical network (PON), being also time-based, may seamlessly interoperate with the parallel pipeline forwarding network.

REFERENCES

[1] Cisco Systems, Inc., "Approaching the Zettabyte Era," June 16, 2008, available online: http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-81374_ns827_Networking_Solutions_White_Paper.html

[2] M. Gupta and S. Singh, "Greening of the Internet", in Proc. ACM SIGCOMM, Karlsruhe, Germany, August 2003.

[3] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, "Photonic switching and the energy bottleneck," Proceedings of Photonics in Switching, San Francisco, 2007.

[4] K. J. Christensen, C. Gunaratne, B. Nordman, and A.D. George, "The Next Frontier for Communications Networks: Power Management", Elsevier Computer Communications, vol. 27, pp. 1758–1770, June 2004.

[5] E. Pinheiro, R. Bianchini, and C. Dubnicki, "Exploiting Redundancy to Conserve Energy in Storage Systems", In Proc. Sigmetrics, Saint Malo, France, June 2006.

[6] Lawrence Berkeley National Laboratory – US Dept. of Energy, "Energy Efficient Digital Networks," http://efficientnetworks.lbl.gov/

[7] A global consortium dedicated to advancing energy efficiency in data centers, http://www.thegreengrid.org/home

[8] C.-S. Li, Y. Ofek, and M. Yung, "Time-driven priority flow control for real-time heterogeneous internetworking," *IEEE Int. Conf. on Computer Communications (INFOCOM 1996),* San Francisco (USA), Mar. 1996.

[9] M. Baldi and Y. Ofek, "Fractional Lambda Switching - Principles of Operation and Performance Issues," SIMULATION: Transactions of The Society for Modeling and Simulation International, Vol.80, No. 10, 2004

[10] Nguyen V. T., Lo Cigno R. A., Ofek Y., Telek M., "Time Blocking Analysis in Time-driven Switching Networks". In: proc. IEEE INFOCOM 2008, USA, 2008.

[11] IP-FLOW project web site, http://dit.unitn.it/ip-flow/

[12] W. Vereecken et al., "Energy Efficiency in Telecommunication Networks", European Conference on Networks and Optical Communications (NOC), Krems, Austria, 01-03 July 2008, pp. 44-51.