Power Consumption and Energy Efficiency in the Internet

Kerry Hinton, Jayant Baliga, Michael Feng, Robert Ayre, and Rodney S. Tucker,
University of Melbourne

Abstract

This article provides an overview of a network-based model of power consumption in Internet infrastructure. This model provides insight into how different parts of the Internet will contribute to network power as Internet access increase over time. The model shows that today the access network dominates the Internet's power consumption and, as access speeds grow, the core network routers will dominate power consumption. The power consumption of data centers and content distribution networks is dominated by the power consumption of data storage for material that is infrequently downloaded and by the transport of the data for material that is frequently downloaded. Based on the model several strategies to improve the energy efficiency of the Internet are presented.

The Internet has become an integral component of the economies of all developed and developing nations. The virtual cycle of improvements in telecommunications supporting economic growth, which, in turn, supports growth in telecommunications infrastructure has served many nations very well. However, this cycle cannot continue without end because all telecommunications networks require resources to function, particularly (electrical) power, to operate. The larger the network becomes (in both capacity and physical size) the more electrical power it consumes. Today the information and telecommunications sector is responsible for approximately 5 percent of the total electrical power consumption in developed national economies [1]. The Internet's infrastructure consumes approximately 1 percent of a developed nation's total electricity consumption in these countries [2–5]. This percentage will grow as higher-speed national broadband access networks are rolled out over the coming years.

The rate of growth of the Internet, in terms of both uptake and capacity increase, means that actually reducing its total power consumption is unlikely to be a realistic goal. The network is growing too fast. A more practical goal is to improve the "energy efficiency" of the Internet. By energy efficiency we mean the amount of data that could be conveyed from end to end per quantum of energy consumed by the network. This measure of energy efficiency is simply the reciprocal of the energy per bit of data transported and/or processed.

Note that although we identify those parts of the Internet that dominate its power consumption (i.e., watts or watts/user), we discuss methods for improving energy efficiency (i.e., reducing Joules per bit). The relationship between these two quantities is power consumption (watts) is equal to energy efficiency (Joules per bit) multiplied by the traffic volume (bits per second). We adopt this approach because the Internet is a complex engineering structure, and any attempt to improve the overall energy efficiency is best focused on those parts that consume the most power. Therefore, a key step in this process is identifying those parts.

In this article, we provide a broad picture of power consumption by Internet infrastructure. We present a "high-level" analysis of the factors that influence the power consumption of the Internet's infrastructure and investigate their relative contributions. We conclude with an overview of some current strategies for improving the energy efficiency of the Internet. We initially focus on the power consumption of the consumer Internet, excluding enterprise networks, data centers, and content distribution. We look at data centers and content distribution later in the article.

Modeling Power Consumption of the Internet

A widely accepted method for modeling the power consumption of Internet infrastructure and related information and communications technology (ICT) infrastructure is based on equipment inventory and/or sales figures [1, 3–6]. Using historical sales data of telecommunications equipment, a broad picture of the quantity of equipment in the network can be estimated. Together with information about the energy consumption of this equipment, this approach can provide a good "order of magnitude" estimate. However, it does not expose the inter-play between demand growth and the consequential power consumption. This is important for estimating how future growth trends may change power consumption patterns as more Internet-based services are taken up.

A complementary approach uses a model based on telecommunications network design principles [2, 7]. In this approach, the Internet is segmented into parts as shown in Fig. 1. For a range of access rates, the energy consumption of each part of the network is calculated using a paper design of the network combined with manufacturers' data on equipment energy consumption for a range of typical types of equipment. This approach enables an overall model of network power consumption to be constructed and provides a platform for predicting the growth in power consumption as the number of users and access rate per user increase.

Figure 1 is a minimalist representation of the network configuration of the Internet. The major components of the net-

0890-8044/11/$25.00 © 2011 IEEE

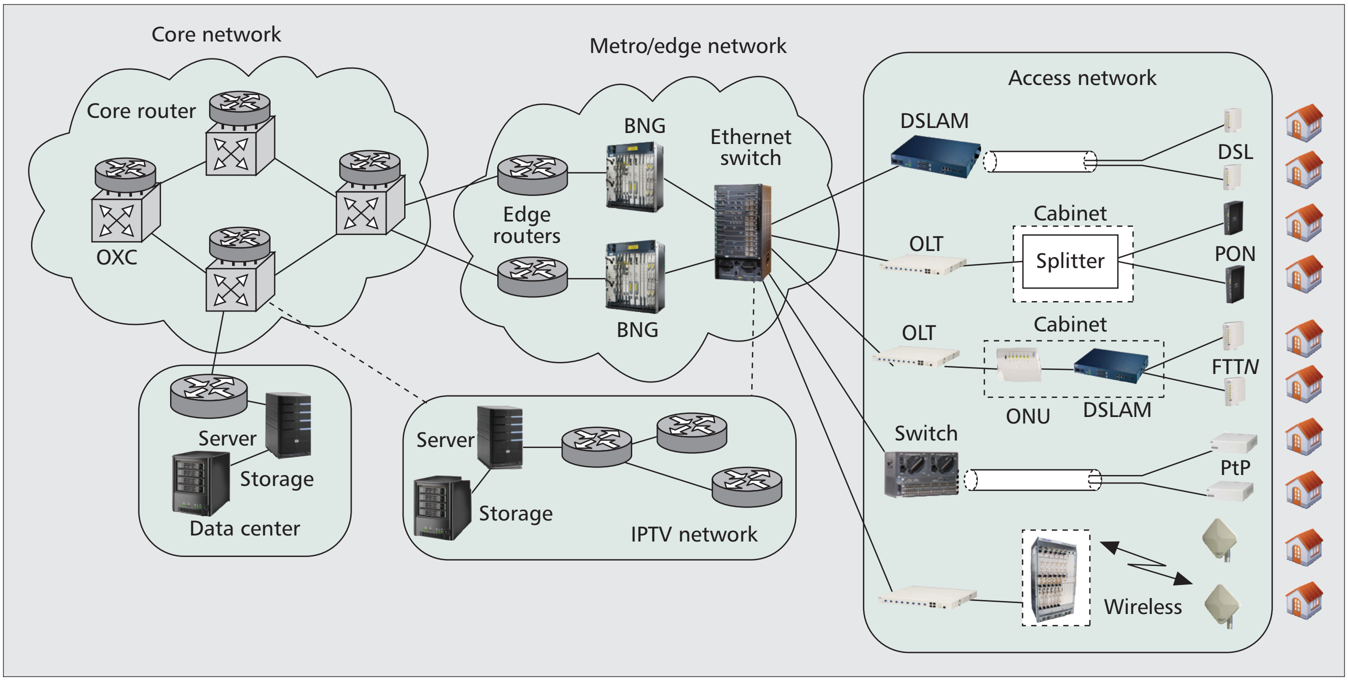

Figure 1. *A high-level network structure with various options for the access network. Also shown are the metro/edge and core parts of the public Internet, and some examples of network data centers and storage networks required to provide web-based services such as IPTV, content distribution, and cloud computing. Power consumption in data centers and content distribution services are not included in the model, and are considered later.*

work are the access, metro, and core networks plus data centers and content distribution networks (e.g., for IPTV). This model is a "first cut" representation of the Internet, and, as such, does not include much of the fine detail of the Internet's true structure and topology. The model does account for the typical hop count for packets that traverse the Internet [8]. The refinement to include a more realistic representation of the Internet's topology is ongoing.

The access network connects individual homes and businesses to their local exchanges. There is a range of technologies in use today and undergoing development. Digital subscriber loop (DSL) uses the copper pairs originally installed to deliver fixed-line telephone service. Fixed-line telephone service, which uses the bandwidth below 3.4 kHz, is left in place, and the higher-frequency bandwidth is used for broadband services. Fibers to the premises (FTTP) installations most commonly use a shared passive optical network (PON) or a point-to-point (PtP) Ethernet connection. In a PON, a single fiber from the network node feeds one or more clusters of customers through a passive splitter. An optical line terminal (OLT) is located at the local exchange, and serves a number of access modems or optical network units (ONUs) located at each customer premises. ONUs communicate with the OLT in a time multiplexed order, with the OLT assigning time slots to each ONU based on its relative demand. In a PtP access network, each ONU is directly connected to the local exchange with a dedicated fiber to the exchange.

In areas where the copper pairs are in good condition, a fiber-to-the-node (FTTN) technology may be used. This technology uses a dedicated fiber from the local exchange to a DSL access multiplexer (DSLAM) located in a street cabinet close to a cluster of customers. A high-speed copper pair technology, such as very-high-speed DSL, is used from the cabinet to the customer premises. In areas where copper and fiber are not available or feasible, wireless can provide Internet access. Technologies for this include WiMAX, High Speed Packet Access (HSPA), and Universal Mobile Telecommunications System (UMTS). For wireless access a wireless modem, located in the customer premises, communicates with a local wireless base station, which, in turn, is connected to the central office.

The local exchanges (or central offices) in a city are connected to each other and to other cities via the metro/edge network. This network also provides connection points for Internet service providers (ISPs). The metro and edge network serves as the interface between the access and core networks. The metro and edge network includes edge Ethernet switches, broadband network gateway (BNG), and provider edge routers. Edge Ethernet switches concentrate traffic from a large number of access nodes uplink to two or more BNG routers. The edge switch connects to two or more BNG routers to provide redundancy. The BNG routers perform access rate control, authentication, and security services, and connect to multiple provider edge routers to increase reliability. The provider edge routers connect to the core of the network.

The core network comprises a small number of large routers in major population centers. These core routers perform all the necessary routing and also serve as the gateway to neighboring core nodes. The core routers of any one network are often highly meshed, but have only a few links to the networks of other providers. High-capacity wavelength-division multiplexed (WDM) fiber links interconnect these routers and connect to networks of other operators.

We initially focus on the power consumption of the consumer Internet, excluding enterprise networks, data centers, and content distribution, returning to these later. We do not address here the energy consumption of equipment within the home network. The home network can take many forms, ranging from a passive cable linking a PC to a modem, through to a multimedia gateway with wired and wireless connections to video, voice, and computing appliances, potentially also including other networking hardware. A study of such networks and their energy consumption is beyond the scope of this article.

Estimating Power Consumption

To estimate the power consumption of the Internet's infrastructure, an access bit rate (in bits per second) is selected. Knowing this access rate and the access technology being used (asynchronous DSL [ADSL], PON, wireless, etc.), and the network design rules, one can calculate the capacity that must be handled by the telecommunications equipment in the access, metro, and core networks. For example, if the access network is a PON (Fig. 1), the design rules may allocate 32 households to be connected to each OLT card port located in the local exchange. Assuming the access rate per user is 10 Mb/s, the total capacity that must be handled by the OLT port is 320 Mb/s. As the access rate increases, so does the capacity that must be handled by the OLT.

The central office (local exchange) houses many OLT cards, which are connected to an Ethernet switch in the metro/edge network, as depicted in Fig. 1. This switch will have a maximum capacity, and the number of OLT cards it can accommodate determines the number of switches required to deal with this capacity. This procedure is repeated to estimate the total traffic within the edge and core networks as well as the amount of equipment required across the whole network. Using representative equipment for each part of the network and employing the manufacturers' specifications for that equipment, one can calculate the numbers of devices required to satisfy the total capacity generated by the chosen access rate. Knowing the power consumption specifications of the equipment provides the information required to calculate the power consumption of the various parts of the network.

Although the results in the plots below are smooth, this is an artifact of the approximations required to make analysis tractable. The model is based on evaluating equipment with adequate capacity to cope with the total demand (plus overprovisioning for redundancy and growth). However, most network equipment has a relatively flat load vs. power profile; thus, deploying the next set of equipment to cope with an increase in traffic causes a step increase in power consumption. The actual power consumption would show a step-wise form close to the smooth lines in these figures. The results presented correspond to averaging out these deployments over the whole network, thereby producing the smoothed traces.

Principal Contributors to Overall Power Consumption of the Internet

Due to space limitations in this article, we now focus on a high-level view of the contributors to the power consumption of the Internet. To do this we identify a number of key contributors.

Network Equipment

The physical network equipment in the network is the major contributor to the power consumption of Internet infrastructure. This includes equipment in the access, metro, and core networks.

- Three technologies dominate the access network. These are fiber (for PON and P2P), copper (for ADSL, VDSL, and hybrid fiber coax [HFC]/cable modem) and wireless. The aggregation of traffic, via statistical multiplexing, from end users is an important function of the access network.
- The metro network includes providing a gateway into the metro and core networks. Local traffic requires routing around city central and suburban areas. The rest is routed into the core network.
- The core network involves core routers and an intercity/international communications system that transport Internet traffic between the core routers.
- Many Internet services provided to access users require exchanging information between the end users and service providers' points of connection to the Internet (often called a point of presence or POP). The transport of this data is "backhaul" and mainly uses wireless or Ethernet transport.
- Network equipment must be powered and cooled. This includes the provision of DC power to the racks that house the equipment and the provision of an uninterruptible power supply (UPS) that ensures continuity of power to the network equipment.

Capacity Planning

Telecommunications network owners need to allow for traffic peaks, future growth, and the protection/restoration of services. This requires some overbuilding of the network which, in turn, increases the power consumption of the network. Sections of the metro/edge and core networks can be 100 percent or more overbuilt depending on the network design policies of the owners.

Services/Cloud Facilities

A significant amount of Internet traffic arises from a wide range of web-based services and resources available to end users via the Internet. Examples of these include cloud services, content delivery, and storage as a service. The data centers that provide these services require significant amounts of equipment and power to function. For example, content services require servers that store the data/content and regulate access to it. Other services require servers hosting, processing, and searching for data. The machines that provide these services are usually networked via a local or wide area network. These networks also consume energy.

Demographics

The Internet is an intra- and international network. The physical distance between population centers has a direct impact on the power consumption of the network. Another important demographic factor is the density of the premises in the population centers. Widely spread premises will require more power to connect to the local exchange.

Service Scenarios

Power consumption is strongly influenced by the type of service being provided. The network will provide the following types of services:

- Shared services: Quasi-real-time or non-real-time shared services such as email, web browsing, and video or audio download, for which short delays are acceptable. These services can be oversubscribed in that many users may share the bandwidth provided without noticing any degradation in speed.
- Dedicated services requiring quality of service (QoS): These include services such as telephony (voice over IP, VoIP), Internet TV, conferencing, and virtual classrooms. These services cannot be oversubscribed. Dedicated capacity for each service must be provided through the access and backhaul networks to the server that is providing the content or service.
- Real-time services delivered to multiple users via multicast. Such services might include broadcast video, near-video-on-demand, and Internet radio. One copy of each requested service is streamed to a switch near the requesting customer(s) and replicated to all requesting customers connected to that switch.

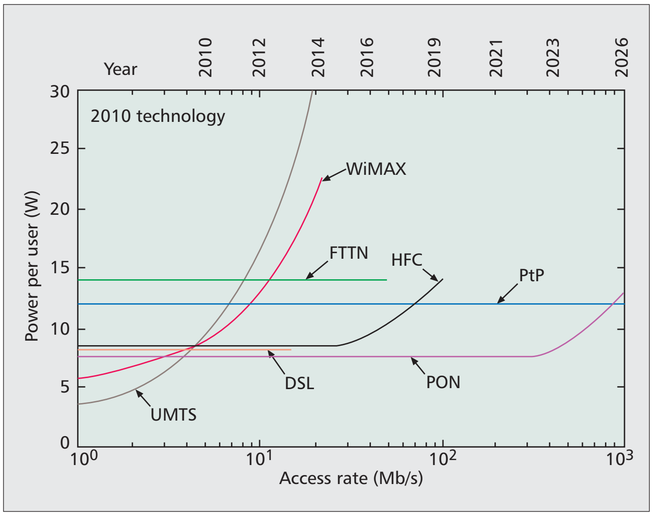

Figure 2. *Power consumption per user for several access network technologies for a range of access rates. Wireless-based (WiMAX, 3G/UMTS) access networks are the most power demanding and fiber based networks the least [9]. Also indicated is the approximate year corresponding to the given access rate assuming 40 percent per annum traffic growth.*

Service Management

All networks and the services they provide must be monitored and managed to ensure they are operating to expectation. These functions add to network power consumption because they require specialized systems and equipment to be inserted into the network.

The focus of this article is on network equipment, not on equipment and networks located within the home. Therefore the analysis does not go beyond the DSL modem/ONU/cable modem/wireless home network gateway or modem.

It is clear that the power consumption of Internet infrastructure is influenced by many factors. These factors often interplay with each other. For example, content distribution centers may communicate with each other and their customers via the public Internet or a private network; depending on demographic factors, this may be over very long distances. These alternatives can have a significant impact on the power consumption of these services.

Where the Power Goes

We have used the analytical methods described above to develop a picture of power consumption in Internet infrastructure and to gauge which parts of the network consume the most power. The details of this work are given in [2, 9–11]. In the following sections we present an overview of some of the key findings of our work. We start with the power consumption of access networks. The details of the analysis are in [9], and some key results are shown in Figs. 2 and 3, which present the power consumed per user as a function of the access rate provided to the user.

We characterize the access rate available to each customer by the access rate advertised and sold to customers by ISPs. However, the metro/edge and core networks are designed by network operators to provide some lower worst-case minimum transmission rate to every customer, taking advantage of the bursty nature of customer Internet traffic. The ratio of the advertised access rate to this minimum per-user rate is referred to as the oversubscription rate. Although the oversubscription rate applied by network providers is typically much higher for wireless access networks than for wired access networks, to facilitate a fair comparison we model the same across all access networks. Note that as the use of the consumer Internet for streaming real-time services increases, high oversubscription ratios will become unsustainable. The results plotted below are based on an oversubscription ratio of 25.

Figure 2 shows that wireless networks such as WiMAX and third-generation (3G)/UMTS consume significantly more power per user than fiber-based access for all but the lowest of access rates. High-speed wireless access is becoming increasingly popular because it provides mobility and ease of access to the Internet. However, unless the energy efficiency of wireless access is improved, its growing popularity may be unsustainable.

Figure 2 shows that fiber-based access networks are the most energy-efficient technologies when high rates are

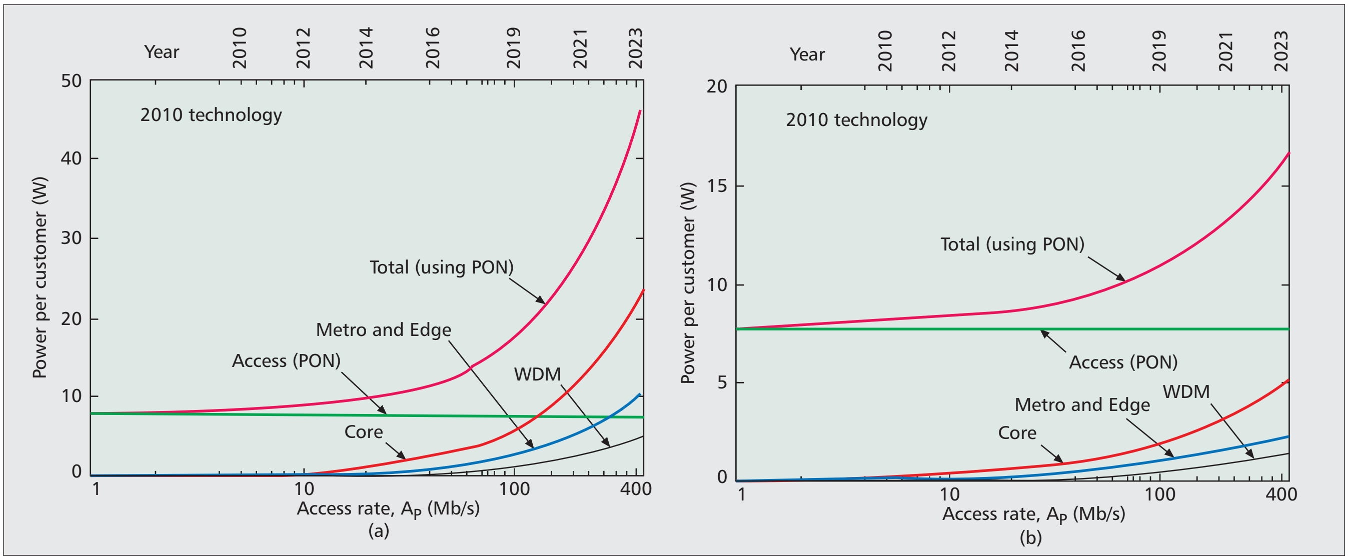

Figure 3. *Power consumption of Internet infrastructure with PON access. The plot includes power consumption of WDM links, core routers, metro and edge network. Plot a) is based on 2010 technology. Plot b) assumes an annual energy efficiency improvement of 10 percent for equipment in the metro/edge and core networks not including data centers and content distribution networks [2]. Both plots include the approximate year corresponding to the given access rate assuming 40 percent per annum traffic growth.*

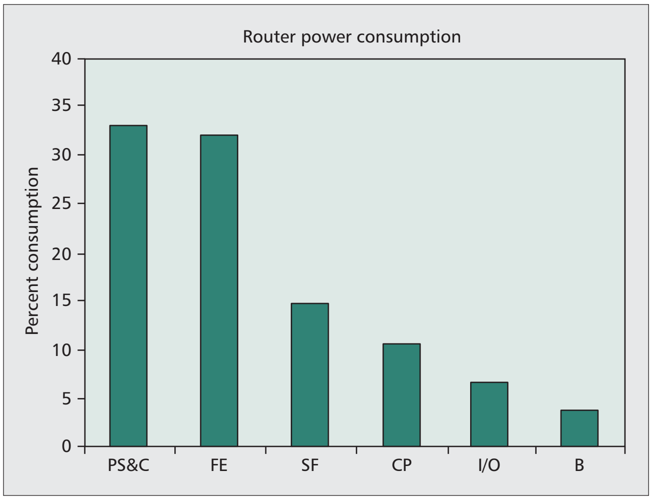

Figure 4. *A Pareto analysis breakdown of power consumed by a core router. The abbreviations are: PS&C: power supply and cooling. FE: forwarding engine. SF: switching fabric. CP: control plane. I/O: input/output cards. B: buffers.*

required. Figure 3 is a plot of the key contributors to power consumption (per customer) of the Internet as a function of the access rate for fiber-based access. The plot is based on 2010 commercially available technology. Over recent years the annual increase in access data rate has been about 40 percent per annum. The calculations used to derive Fig. 3 have included factors to account for power supply, cooling, demographics, service scenarios, capacity planning, and service management contributions that appear in the above list. Cloud and content distribution services are discussed below.

Figure 3 shows that for low access rates, Internet power consumption is dominated by the access equipment (i.e., the equipment used to connect the home to its local exchange); in particular, the ONU located in the home. As access rates increase, the core network power consumption increases and will ultimately surpass access power consumption. Whereas the power consumption of a home ONU or gateway is independent of its access speed, as access rates increase, the volume of traffic in the core must increase. This, in turn, requires a significant increase in the amount of routing equipment and consequential power consumption to such an extent that the core routers dominate power consumption at high access rates. The power consumption growth shown in Fig. 3a assumes 2010 technology in the metro/edge and core networks at all access rates. This plot shows that without ongoing technology improvements, the power consumption of the Internet's infrastructure will grow exponentially toward unsustainable levels due to the demands on the core routers. In reality, there will be improvements in energy efficiency during the time it takes for networks to evolve to higher access speeds. Figure 3b shows the power trends assuming a 10 percent annual improvement in energy efficiency of the metro/edge and core network equipment. This is a realistic improvement rate for networks in which the latest generation of equipment is deployed to accommodate increasing demand [2, 7].

Figure 3 shows that the metro/edge network as well as the optical communications systems that connect between the core equipment do not dominate power consumption. The metro/edge equipment does not have to deal with the volume of traffic that occurs in the core. The WDM optical communications systems that connect the routers are relatively energy efficient in that they can transport substantial capacity at low power.

Because the core routers will dominate power consumption at high data rates, we now turn our attention to these routers. The relative power consumption of subsystems within a core router is shown in Fig. 4 [12, 13]. A fully loaded core router consumes approximately 10 nJ/b when it processes IP packets [2, 13]. The forwarding engine, power supply, and cooling within the router contribute around 65 percent to its total power consumption [12].

Improving Energy Efficiency of the Internet

From Fig. 3 it is clear that the two main areas requiring attention in the context of overall power consumption are the access networks (in particular the home terminal equipment) and the core network routers. The challenge of addressing home terminal equipment has been addressed by the European Union (EU), which has published power consumption guidelines for this equipment. This voluntary code of conduct is designed to improve the energy efficiency of all broadband home equipment sold within the EU [14]. This forms part of the strategies developed in the EU code of conduct.

Three effective strategies to improving equipment energy efficiency are:

- Require equipment to reduce its power consumption when not in use. This low-power state is often referred to as a "sleep" or "idle" state and can be implemented by shutting down those parts of the device that are not needed when the equipment is not communicating. The entire device cannot be turned off because it will lose contact with the Internet. A small amount of power must be used to ensure that the Internet is aware the device is available and is able to awaken the device when required. Because modern electronics can operate at very high speeds, even very short (much less than 1 s) sleep states can be very effective in reducing power consumption [15].
- Reduce the processing rate of a device when its work load is low. Many devices can operate over a range of bit rates. Electronic circuits consume less power when operating at a lower speed. Thus, when the traffic load on a device is low, power consumption can be reduced by lowering the speed at which the device operates. This is often referred to as rate adaptation [15].

  These approaches are being applied to make the Ethernet protocol more energy efficient and can be applied across many parts of the Internet. [16]
- Improving the energy efficiency of core routers. This will require either improving the signal processing technology within the router or changing the function of the router. Also, strategies for dimensioning the core network to improve energy efficiency will become increasingly important in the future as the core network starts to dominate power consumption. These strategies have been the subject of significant research over recent years.
- Deploy the most energy-efficient access network technology available. The dominance of access network equipment in today's network is a clear focus for improving the energy efficiency of the Internet.

One proposal has been to replace the electronic circuitry within a router with photonic circuits. This approach is motivated by the expectation that photonic switching technologies can operate at far higher speeds than electronics. Current trends indicate that the maximum processing speed attainable by electronics in the next few years will be about 100 Gb/s, while photonics holds the promise of attaining speeds over 10 Tb/s. It has been proposed by some researchers that many electronic routers could be replaced with far fewer photonic machines, thereby reducing overall power consumption. Unfortunately, power consumption trends to date for the key photonic signal processing technologies do not support this

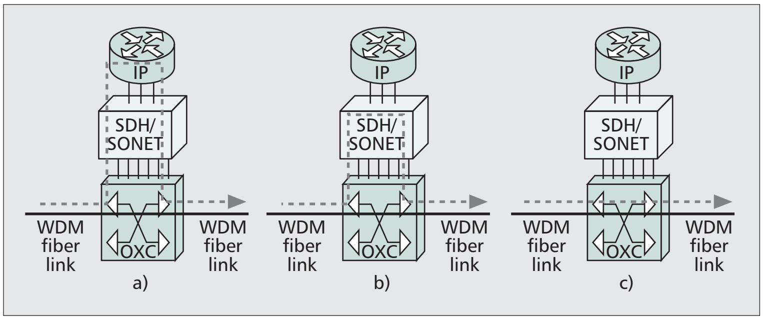

Figure 5. *Traffic flow (grey dashed line) in a node with SDH/SONET and WDM layers: a) all traffic is passed up through the lower layers and processed by the IP router; b) traffic is processed by the SDH/SONET switch, bypassing the IP router; c) Traffic is switched by the optical cross-connect, bypassing both the SDH/SONET and IP layers. Lower layers are progressively more energy efficient.*

scenario. Today complementary metal oxide semiconductor (CMOS) is about five orders of magnitude less energy consuming than the photonic technologies [17]. Furthermore, while CMOS has shown a trend of continually decreasing power consumption, the photonic technologies are showing very little improvement [18]. The net effect of this is that, whenever intensive signal processing or computation is required, electronics is the most energy-efficient technology available.

Another option is to re-architect networks to reduce the traffic processed in the IP routers. In this approach the network would be redesigned so that a large proportion of Internet traffic bypasses routers in the core network [12, 19, 20]. As Internet traffic travels between its source and destination, on average it is processed by about 14 routers. These routers are not directly connected. Rather, they communicate via optical communication systems that use the synchronous digital hierarchy/synchronous optical network (SDH/SONET) protocol. The physical connections are based on WDM in which many independent optical channels propagate through fibers deployed around the globe. IP routers can be considered as sitting at the top level of a multilayer stack of equipment, as depicted in Fig. 5. Processing traffic at the IP level (Fig. 5a) typically requires about 10 nJ/b. Processing at the SDH/SONET layer (Fig. 5b) requires around 1–3 nJ/b and in the WDM layer (Fig. 5c) less than 1 nJ/b [2,19]. Therefore using SDH/SONET and WDM to bypass the routers reduces the size and power consumption of the routers in the core of the network because much of the traffic is processed at the more energy-efficient SDH/SONET and/or WDM layers [19].

Content Distribution and Data Centers

We now consider content distribution networks with a focus on video distribution. The provision of video and TV content via the Internet (e.g., IPTV) is a key driver of Internet growth. As these services grow the power consumption of the equipment required also increases. Adopting the same network-model-based approach, the energy required to download content from a data center can be calculated. A content distribution network can be connected via the Internet or a private network. (Both options are shown in Fig. 1.) The difference is that content distribution via the Internet results in the content traveling through several routers, each adding to the power consumption of the distribution. Using a private network avoids the routers, but requires a privately owned (and paid for) network to connect content directly to the local distribution point.

IPTV, like most content, is typically stored on hard disks in a data center which has its own internal network. When the content is requested it must be retrieved from the appropriate disk within the data center, and transmitted from the data center to the metro Ethernet edge switch and down to the TV set-top box (STB) in the home. To calculate the power consumption of this process, a file size is chosen (e.g., a 2 hour standard definition movie is about 1.8 Gbytes) and the number of request per hour for that movie. Given this and using typical equipment power consumption figures, the network equipment required to provide the movie to the STB is determined, and the power consumption for the download is calculated. For a given file size, the energy per download depends on the number of downloads per hour. The results for a 1.8 Gbyte file (2 hour standard definition movie) are shown in Fig. 6. In this example the movie is replicated in 20 data centers spread across a service area.

Looking at Fig. 6, we see that the power consumed by storage of the content (on hard disks) is constant and independent of the number of downloads per hour; it is set by the power consumption of spinning the disks. The power consumed by the servers that respond to customer requests for content and extract the content from the disks, and the power required to transport the content from the data center to the customer depend on the number of requests for the content. This means the energy per download being dominated by the storage (disk) power when there are only a few downloads per hour. If the content is popular, resulting in many downloads per hour, the energy per download is dominated by servers and the transport network because the storage power is shared amongst many customer requests.

Therefore, popular content should be stored closer to the user, meaning there will be multiple copies of popular content geographically spread across the network. This also reduces the number of routers the content must pass through to reach the user. For less popular content, fewer copies should be stored in centralized sites. The most energy-efficient solution is a compromise between the number/location of the stored copies and transport to the user's home. The precise details of the traces in Fig. 6 and the optimal deployment of the video copies depend on the popularity of the content.

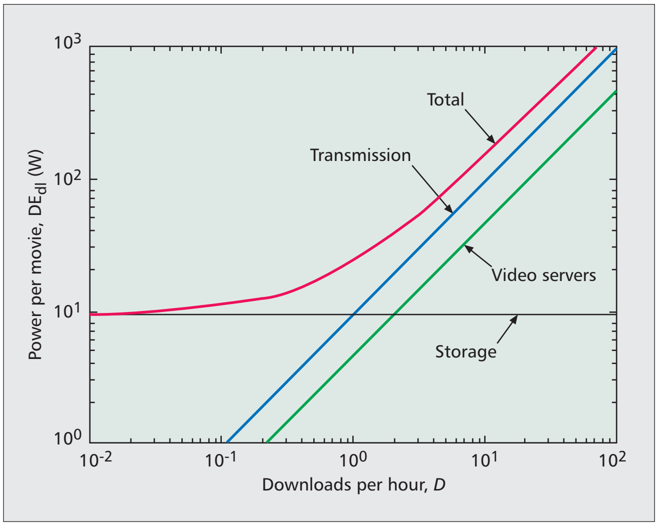

Figure 6. *Power consumption per download for a standard definition 2 hour video that has 20 copies replicated in data centers [10].*

Conclusions

The importance of the Internet and ICT is continually increasing both in terms of economic growth and as a source of greenhouse gas production. Therefore, to manage the power consumption of the Internet, it is important to understand where the energy is consumed in the Internet's infrastructure. This article has described one approach to attaining this understanding and used it to identify those parts of the Internet that dominate its power consumption. This information can then be used to make the Internet more energy efficient.

As the world becomes more energy constrained, humankind will need to develop and refine strategies for improving the energy efficiency of the Internet. A "common rule of thumb" is the "80/20 rule," which states that 20 percent of causes produce 80 percent of effects. In this article the strategies to address this 20 percent have been identified.

References

[1] GeSI, The Climate Group, "Smart 2020: Enabling the Low Carbon Economy in the Information Age," 2008.

[2] J. Baliga *et al.*, "Energy Consumption in Optical IP Networks," *J. Lightwave Tech.*, vol. 27, no. 13, 2009, pp. 2391–403.

[3] K. Kawatomo *et al.*, "Electricity Used by Office Equipment and Network Equipment in the US," *ACEEE Summer Study Conf. Energy Efficiency in Buildings (NLBL-45917)*, 2000.

[4] C. Barthel, S. Lechtenbohmer, and S. Thomas, "GHG Emission Trends of the Internet in Germany," *Int'l. Climate Policy and IT Sector Joint Wksp. with IGES and the Wuppertal Inst.*, 2001.

[5] V. Türk, *Assessing the Resource Intensity of the Internet Infrastructure: Data Analysis for a Material-Flow Oriented Approach and First Results on Electricity Consumption*, IIIEE, M.Sc. thesis, Lund Univ., Sweden. 2001.

[6] S. Lanzisera, B. Nordman, and R. Brown, "Data Network Equipment Energy Use and Savings Potential in Buildings," *ACEEE Summer Study on Energy Efficiency in Buildings*, 2010.

[7] D. Kilper *et al.*, "Power Trends in Communications Networks," *IEEE J. Sel. Topics in Quantum Elect.*, to appear, 2011.

[8] P. Mieghem, *Performance Analysis of Communications Networks and Systems*, Cambridge Univ. Press, 2006.

[9] J. Baliga *et al.*, "Power Consumption in Access Networks," *OFC 2008*, paper OThT6, 2008.

[10] J. Baliga *et al.*, "Architectures for Energy-Efficient IPTV Networks," *OFC/NFOEC 2009*, paper OThQ5, 2009.

[11] J. Baliga *et al.*, "Carbon Footprint of the Internet," *Telecommun. J. Australia*, vol. 59, no. 1, 2009, pp. 05.1–05.14.

[12] R. Tucker *et al.*, "Evolution of WDM Optical IP Networks: A Cost and Energy Perspective," *J. Lightwave Tech.*, vol. 27, no. 3, 2009, pp. 243–52.

[13] O. Tamm, C. Hermsmeyer, and A. Rush, "Eco-Sustainability System and Network Architectures for Future Transport Networks," *Bell Labs. Tech. J.*, vol. 14, no. 4, 2010, pp. 311–28.

[14] European Commission, "Code of Conduct on Power consumption of Broadband Equipment," v. 3, Nov. 2008.

[15] S. Nedevschi *et al.*, "Reducing Network Energy Consumption via Sleep and Rate-Adaptation," *Proc. 5th USENIX Symp. Network System Design and Implementation*, 2008, pp. 232–26.

[16] IEEE 802.3az-2010, "Energy Efficient Ethernet"; http://www.ieee802.org/3/az/index.html.

[17] K. Hinton *et al.*, "Switching Energy and Device Size Limits on Digital Photonic Signal Processing Technologies," *J. Sel. Topics in Quantum Elect.*, vol. 14, no. 3, 2008, pp. 938–45.

[18] K. Hinton *et al.*, "The Future Internet — an Energy Consumption Perspective" invited paper FT1, OECC 2009.

[19] G. Eilenberger *et al.*, "Energy-Efficient Transport for the Future Internet," *Bell Labs Tech J.*, vol. 15, no. 2, 2010, pp.147–68.

[20] M. Feng *et al.*, "Reducing NGN Power Consumption with IP/SDG/WDM," *ACM E-energy Conf.*, 2010.

Biographies

KERRY HINTON (k.hinton@ee.unimelb.edu.au) received an Honors B.E. in 1978, an Honors B.Sc. in 1980, and an M.Sc. in mathematical sciences in 1982, all from the University of Adelaide. He was awarded a Ph.D. in theoretical physics from the University of Newcastle Upon Tyne, United Kingdom, and a DipIR from Newcastle Upon Tyne Polytechnic in 1984. In 1984 he joined Telstra Research Laboratories, Victoria, Australia, and worked on analytical and numerical modeling of optical systems and components. His work has focused optical communications devices, architectures, monitoring, and physical layer issues for intelligent all-optical networks. He was also a laser safety expert within Telstra. In 2006 he joined the ARC Special Centre for Ultra-Broadband Information Networks at the University of Melbourne, Australia, where he is undertaking research into the energy efficiency of the Internet and optical communications technologies.

JAYANT BALIGA (jbaliga@ee.unimelb.edu.au) received a B.Sc. degree in computer science and a B.E. degree in electrical and electronic engineering (with first class honors) in 2007 from the University of Melbourne, Australia. He is currently working toward a Ph.D. degree in electrical engineering at the same university. His research interests include optical network architectures and wireless communications.

ROBERT W. A. AYRE (r.arye@ee.unimelb.edu.au) received his B.Sc. degree in electronic engineering from George Washington University, Washington, DC, in 1967, and B.E. and M.Eng.Sc. degrees from Monash University, Melbourne, Australia, in 1970 and 1972, respectively. In 1972 he joined the Research Laboratories of Telstra Corporation, working in a number of roles primarily in the areas of optical transmission for core and access networks, and in broadband networking. In 2007 he joined the ARC Special Centre for Ultra- Broadband Networks (CUBIN) at the University of Melbourne, continuing work on networking and high-speed optical technologies.

RODNEY S. TUCKER [M'76, SM'81, F'89] (r.tucker@ee.unimelb.edu.au) is a Laureate professor at the University of Melbourne and research director of the ARC Special Research Centre for Ultra-Broadband Information Networks. He has held positions at the University of Queensland, the University of California, Berkeley, Cornell University, Plessey Research, AT&T Bell Laboratories, Hewlett Packard Laboratories, and Agilent Technologies. He is a Fellow of the Australian Academy of Science and the Australian Academy of Technological Sciences and Engineering. He received his B.E. and Ph.D. degrees from the University of Melbourne in 1969 and 1975, respectively. He was awarded the Institution of Engineers Australia Sargent Medal in 1995 for contributions to electrical engineering, and was named IEEE Lasers and Electro-Optics Society Distinguished Lecturer for 1995–1996. In 1997 he was awarded the Australia Prize for his contributions to telecommunications, and in 2007 he was awarded the IEEE Lasers and Electro-Optics Society Aron Kressel Award.

MICHAEL Z. FENG (mzfeng@ee.unimelb.edu.au) received his Bachelor of Engineering degree (with first class honors) in electrical and electronic engineering in 2008 from the University of Auckland, New Zealand. He is currently a Ph.D. student at the Centre for Ultra-Broadband Information Networks (CUBIN), University of Melbourne. His research interests include optical communications technologies and energy efficiency of optical networks.