Energy-Efficient Wideband Cable Access Networks in Future Smart Cities

Zuqing Zhu and Ping Lu, University of Science and Technology of China
Joel J. P. C. Rodrigues, Instituto de Telecomunicações, University of Beira Interior
Yonggang Wen, Nanyang Technological University

ABSTRACT

The Internet explosion makes energy consumption of the ICT sector a global concern. Nowadays, the power consumption of access networks increases rapidly due to the rollout of broadband services. Cable access networks provide Internet access over the existing cable television systems. Among all wired access networks, they have the second biggest user base, which is still growing fast due to the release of the new industry standard, DOCSIS 3.0. Therefore, we expect that cable access networks will play an important role in the ICT sector of future smart cities. In order to boost up bidirectional throughput, DOCSIS 3.0 introduces the channel bonding technology, which could lead to increased power consumption if not properly managed. In this article, we present a few energy-saving algorithms that could be used by the cable operators to improve the energy efficiency of cable access networks using channel bonding. The discussions cover both the network-wide energy-saving algorithm and the customer-side energy-saving algorithm, and provide investigations on the trade-offs between energy-saving and other network performance metrics, including packet delay and protocol overhead. Our numerical results suggest that effective energy saving can be achieved in wideband cable access networks with the proposed algorithms, and the packet delay and protocol overhead can be reduced if we choose the key parameters properly.

INTRODUCTION

Over the last decade, the Internet has advanced to the era of Web 2.0 with more collaborative and interactive online services. In the near future, the Internet will evolve from numerous isolated websites to an integrated platform to offer various services to end users for communication, news, entertainment, socialization, and so on. Hence, in future smart cities, online information sharing, collaborations, and socialization will become an important part in people's daily life. However, nothing comes as a free lunch. The bandwidth-hungry Web 2.0 applications are pushing the Internet traffic to grow exponentially with an annual rate of more than 30 percent [1]. For instance, high-definition television (HDTV) or video-on-demand (VOD) service usually needs 5–8 Mb/s bandwidth, an interactive virtual conference with telepresence requires more than 24 Mb/s bandwidth and less than 50 ms network latency, and remote surgery needs more than 30 Mb/s bandwidth and less than 1 ms network latency. The consequent bandwidth requirements lead to fast upgrading and expansion of the Internet infrastructure, which make the energy consumption of the information and communication technology (ICT) sector a global concern. According to a recent research report [2], ICT accounts for 2–10 percent of the total carbon emission produced by human activities, and its carbon footprint is already comparable to the aviation industry. The estimation also indicated that 37 percent of ICT's carbon emission is due to the production of the energy for operating the Internet infrastructure.

When looking into the Internet infrastructure, we can see that access networks are among the top carbon emission contributors in it, due to the rollout of broadband services [3]. Moreover, this situation will not change in the short-to-mid-term future, as network operators are investing heavily in deployments and upgrades of broadband equipment to facilitate sufficient capacity for end users. The energy efficiency of wireless access networks can be improved by using smart transmission technologies [4], multipolling mechanisms [5], and so on. According to [2], by 2020, the carbon footprint of wired access network equipment and broadband modems will be 20 times that in 2002. It is known that digital subscriber line (DSL) networks, cable access networks, and passive optical networks (PONs) are the three major types of wired access networks. Green DSL networking was achieved with dynamic spectrum management [6], and energy-efficient network designs have been obtained for PONs by incorporating sleep modes [7]. Cable access networks provide Internet access over the existing cable TV systems. Since these access networks have the second biggest user population among all wired access networks [8], they

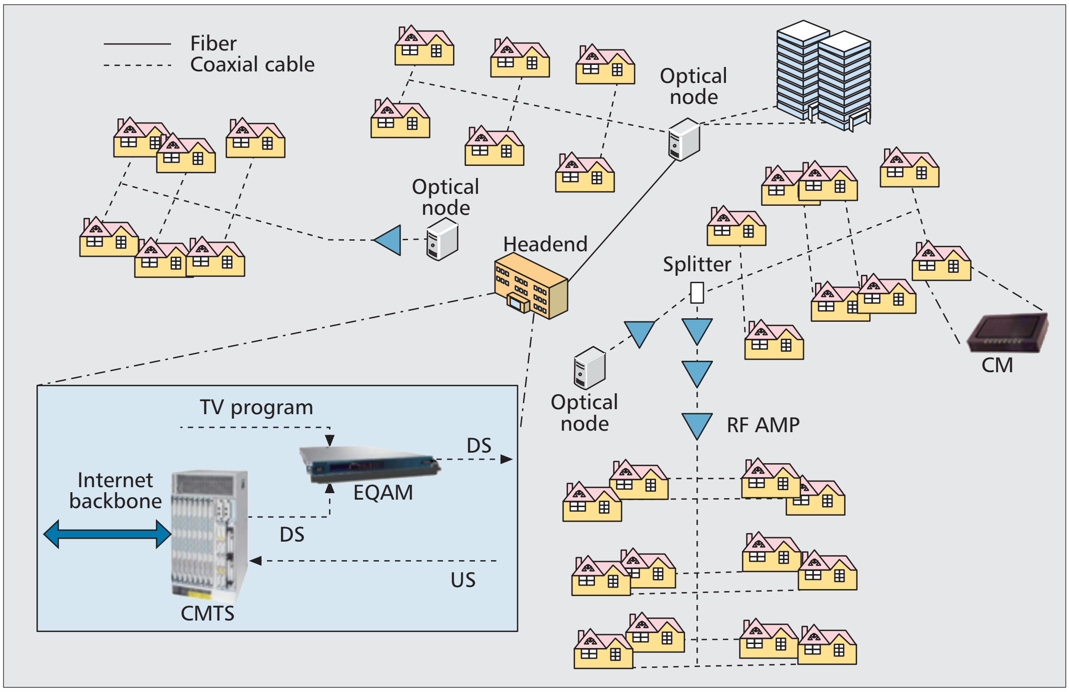

Figure 1. *Typical architecture of a cable access network. CMTS: cable modem termination system, CM: cable modem, DS: downstream, US: upstream, EQAM: edge quadrature amplitude modulation system, RF AMP: radio frequency power amplifier.*

To support bidirectional communication, upstream (US, from customer premises to the operator's headend) transmission is employed and the US and DS transmissions coexist in a frequency-division multiplexing way.

will play an important role in the ICT sector of future smart cities.

In this article, we first review recent technical advances in broadband cable access networks and the energy issues associated with them, and then discuss a few energy-saving algorithms that might be used by the cable operators to improve network energy efficiency. We focus on both the network-wide energy-saving algorithm on a relatively coarse timescale (a few minutes) and the customer-side energy-saving algorithm on a finer timescale (a few milliseconds). We show that the algorithms achieve effective energy saving, and the packet delay and protocol overhead can be reduced if we choose the key parameters properly. Finally, we summarize the article.

Cable Access Networks

Cable networks were originally deployed to deliver TV programs in a downstream-only (DS, from the operator's head-end to customer premises) broadcasting manner. With the rapid development of the Internet, cable operators began to deliver digitally modulated data signals within radio frequency (RF) TV channels over the cable network infrastructure. Figure 1 illustrates the typical architecture of a cable access network. To support bidirectional communication, upstream (US, from customer premises to the operator's head-end) transmission is employed, and the US and DS transmissions coexist in a frequency-division multiplexing way. Based on different standards, US and DS channels are allocated in the frequency ranges of 5–42 MHz (North American and Asian standard) or 5–65 MHz (European standard) and 91–857 MHz (North American and Asian standard) or 112–857 MHz (European standard), respectively [9]. The US and DS frequency ranges are then divided into RF channels with 6 (North American and Asian standards) or 8 MHz bandwidth (European standard) to comply with the existing TV programs.

Located at a cable operator's headend, a cable modem termination system (CMTS) connects the cable access network with the Internet. In the DS direction, the Internet data signal is combined with the TV program signal in the edge quadrature amplitude modulation (EQAM) system and then delivered to cable modems (CMs) in customer premises. The DS bandwidth is shared by subscribers connected to a given cable network segment. With the 256-QAM modulation scheme (the highest modulation level can be used in DS transmission), ~ 40 Mb/s throughput can be achieved in a 6 MHz RF channel for DS transmission. In the US direction, data signals are forwarded from the CMs to the CMTS directly in such a way that the CMs send bandwidth requests for future transmission opportunities, and the CMTS grants the requests based on a scheduling mechanism. The highest modulation level that can be used in US transmission is 64-QAM, and ~ 30 Mb/s throughput can be obtained in a 6 MHz RF channel. The optical nodes in the distribution network bridge physical layer transmissions between the optical and electrical domains using optical-to-electrical (O/E) and electrical-to-optical (E/O) conversions. RF power amplifiers can be employed in the distribution network to maintain the signal level range across the coaxial plant. Since cable access networks use this specific technology solution to distribute RF signals over optical fibers and coaxial cables, they are also referred to as hybrid fiber-coaxial (HFC) networks.

DOCSIS 3.0 redefines the communication between the CMTS and CM by introducing channel bonding. Specifically, channel bonding allows a CM to tune to multiple US and DS channels simultaneously and to groom their capacities for higher data-rate transmissions.

DEVELOPMENT OF DOCSIS STANDARD

Cable operators started to deploy HFC networks globally in the early 1990s, and it has been reported that the worldwide HFC customer number increased faster than that of its major competitor, fiber to the home (FTTH), in the past few years [8]. In order to offer price-competitive broadband services, HFC networks have evolved rapidly since their inception, with a lot of new technologies being developed and standardized in the Data over Cable Service Interface Specifications (DOCSIS) [9]. DOCSIS specifies modulation schemes and the protocol for exchanging information between the CMTS and CM over HFC networks. The European version of DOCSIS is EuroDOCSIS.

Broadband Internet access was first achieved over HFC networks with DOCSIS 1.0, which used MPEG packets with modulation schemes such as 64-QAM or 256-QAM to deliver DS data and specified quadrature phase shift keying (QPSK) or 16-QAM as the modulation schemes for US data transmission. To facilitate quality of service (QoS) provisioning, DOCSIS 1.1 introduced the concept of service flows, and allows network systems to apply different QoS levels to different types of traffic, such as data, voice, and video. Due to the appearance of peer-to-peer (P2P) applications in the Internet, the need for faster US data rates and symmetrical capacities of US and DS led to the development of DOCSIS 2.0. More sophisticated modulation schemes such as 32-QAM and 64-QAM were introduced to US transmission and pushed the US data rate to 30 Mb/s. DOCSIS 3.0 was released to enable cable operators to outpace competitors such as FTTH providers. Table 1 summarizes the development of the DOCSIS standard [9].

CHANNEL BONDING AND RELATED ENERGY ISSUES

DOCSIS 3.0 redefines the communication between the CMTS and CM by introducing channel bonding. Specifically, channel bonding allows a CM to tune to multiple US and DS channels simultaneously, and to groom their capacities for higher-data-rate transmissions. Channel bonding can push the US and DS data rates of a CM to ~120 Mb/s and ~160 Mb/s, respectively. DOCSIS 3.0 specifies the minimum number of transceivers that a CM's hardware must support as four, but does not set a maximum value. Current implementations usually support 4~8 channel bonding. However, when market conditions dictate, cable operators can migrate to larger bonded channel counts (e.g., 12–18). In the channel bonding mode, the management software on a CMTS groups the US/DS ports on its line cards into medium access control (MAC) domains and manages the CMs using the transmitter channel set (TCS) and receiver channel set (RCS). DOCSIS 3.0 defines the TCS and RCS of a CM as the sets of US and DS channels to which it is connected .

Even though channel bonding technology can boost the bidirectional capacity of a CM effectively, it also increases power consumption in HFC networks dramatically. A CM has to equip multiple high-speed RF transceivers to accommodate the bonding requirement, while a CMTS has to be designed using line cards that have an increased number of ports. For instance, a legacy CM usually consumes 6–8 W, while a channel-bonding CM can consume 9–20 W. Also, on the line card of a typical CMTS platform such as Cisco's universal broadband router (e.g., uBR1000), the numbers of US and DS ports are increased from 5 to 20 and from 20 to 60, respectively. It was estimated that the power consumption increase in HFC networks was ~ 66 million W from 2010 to 2011 [10]. Nevertheless, the power consumption can increase even faster with more intensive DOCSIS 3.0 deployments around the world in the near future. The consequent operational expenditure (OPEX) and environmental impacts can become a serious concern.

In most of the current implementations of DOCSIS 3.0, the CMTS determines the sizes of a CM's TCS and RCS (i.e., active US and DS channels) during CM registration and does not change them afterward. By doing so, CM is instructed to support the peak traffic load for customer satisfaction. However, it is well known that network utilization generated by residential services is subject to relatively large hourly fluctuations due to the daily life pattern of human beings [11]. This fact suggests that both the CMs' transceivers and the line card ports on the CMTS can be greatly underutilized in low network utilization hours, such as 4–7 a.m. Hence, if we adjust the operation modes of the CMs and CMTS dynamically according to the traffic fluctuation and deactivate unused channels, energy saving can be achieved.

NETWORK-WIDE ENERGY SAVING WITH TRAFFIC-AWARE DYNAMIC BONDING CHANGES

DOCSIS 3.0 defines a MAC layer operation, dynamic bonding change (DBC), to adjust the TCS and RCS of a CM without introducing significant traffic disruption [9]. With DBC, CMTS can add (DBC-ADD), delete (DBC-DELETE), or replace (DBC-REPLACE) multiple US and DS channels in a CM's TCS and RCS while keeping the CM online. The DBC operation always starts with the CMTS sending a DBC-REQ message to a CM. After receiving the DBC-REQ, the CM performs necessary procedures to prepare for the changes encoded in the message and sends a DBC-RSP back to the CMTS. The CMTS then confirms the change and sends a DBC-ACK message to the CM. Finally, after getting DBC-ACK, the CM commits the changes. DBC was originally designed for the sake of load balancing in DOCSIS 3.0, but a later version of the standard suggests that the CMTS should use DBC to reduce a channel-bonding CM's oper-

DOCSIS version	1.0	1.1	2.0	3.0
Features				
Broadband Internet access	✓	✓	✓	✓
Voice over IP		✓	✓	✓
Videoconferencing		✓	✓	✓
Commercial services			✓	✓
IPv6				✓
DOCSIS upstream				
Channel(s) per CM	1	1	1	≥ 4
Channel capacity (Mb/s)	10	10	30	≥ 120
DOCSIS downstream				
Channel(s) per CM	1	1	1	≥ 4
Channel dapacity (Mb/s)	40	40	40	≥ 160
EuroDOCSIS upstream				
Channel(s) per CM	1	1	1	≥ 4
Channel Capacity (Mb/s)	10	10	30	≥ 120
EuroDOCSIS downstream				
Channel(s) per CM	1	1	1	≥ 4
Channel Capacity (Mb/s)	50	50	50	≥ 200

Table 1. *Development of the DOCSIS standard*.

ation to a single US and DS channel, whenever the CMTS detects the CM operating on backup battery (i.e., loss of A/C power for more than 5 s). This is the first attempt to achieve energy saving with DBC in HFC networks, but it is far from enough. In this section, we show that together with a traffic-aware algorithm, we can implement DBC for more effective energy saving [12].

Due to the asymmetry of US and DS traffic, the US and DS channels need to be handled separately. As the algorithm we discuss here can be applied to an arbitrary transmission direction (i.e., US or DS), we do not have to specify US or DS in the following discussion. The algorithm adjusts the operation modes of the CMs and CMTS in a coordinated way, to reduce the network-wide energy consumption. The energy-saving algorithm can be divided into three phases:

- CM-side channel adjustment
- CM-to-CMTS channel mapping
- CMTS-side channel readjustment [12]

In phase 1, based on how many channels are active in the TCS or RCS of a CM, we define different operation modes for it. For example, for a four-channel bonding CM, we can define the high-power mode with four channels, the moderate-power mode with two channels, and the low-power mode with one channel. The CMTS maintains two traffic thresholds, the high watermark (*HW*) and low watermark (LW). With a fixed time interval (e.g., a few minutes), the CMTS samples the bidirectional traffic loads of a CM. In order to get this done, the CMTS measures the length of the corresponding DS packet queue and calculates the pending US data requests. When the instant traffic sample X_i is obtained, the CMTS gets the average load Y_i by averaging the current and previous traffic samples, $\{X_j, j = i - N - 1, \ldots, i\}$, where N is the number of traffic samples that are used for averaging. The CMTS then determines the CM's operation mode by using $Z_i = \max(X_i, Y_i)$, as the larger one between X_i and Y_i. Specifically, the CMTS compares Z_i with *HW* and *LW*. When Z_i is greater than or equal to *HW*, the CM will operate in the high-power mode; when Z_i is between *LW* and *HW*, the CM will be in the moderate-power mode; otherwise, when Z_i is less than *LW*, the CM will be in the low-power mode. This "slowly-shrinking-and-quickly-restoring" design can achieve minimal impact on the CM's legacy traffic, and undesired long packet delay can be avoided. When the CM's traffic load is decreasing, the CMTS only shrinks its channel set size when the average load Y_i is below *LW* or *HW*, while the CM's channel set size gets restored immediately when the traffic load increases above the thresholds.

Figure 2a illustrates the typical traffic fluctuation (either DS or US) of a CM in a realistic HFC network [11], and Fig. 2b shows simulation results on channel operations obtained with the energy-efficient operation discussed above (*LW* = 25 percent, *HW* = 50 percent, *N* = 5). We assume that there are four channels and plot the status of their transmitters (TXs) under energy-efficient operation. At a time instant, a channel is deactivated if the output of its TX is in white; otherwise, it is active. It can be seen that significant energy can be saved during low traffic hours. For this particular traffic case, the algorithm achieves 35.97 percent energy saving if we assume that the power consumption of a channel is zero when it is deactivated.

In phase 2, when the size of the channel set is obtained, the algorithm determines the actual connection mapping between the CM and the CMTS. A CMTS port (either US or DS) on a line card usually operates on a single frequency channel. Channels from multiple CMs can be tuned to the same CMTS port, and the bandwidth resource of the port is shared by those CMs. The energy-saving algorithm greedily groups CM connections with as small a number of CMTS ports as possible. Specifically, when we have to add a channel to a CM's channel set, we direct the new connection to the busiest CMTS port with a DBC-ADD; and when we remove a channel from a CM's channel set, we take off a connection from the lightest port with a DBC-DELETE. As a result, the CM connections are supported with as few CMTS ports as possible, and the rest of the idle CMTS ports can be shut down for energy saving.

The algorithm adjusts the operation modes of the CMs and CMTS in a coordinated way, to reduce the network-wide energy consumption. The energy-saving algorithm can be divided into three phases: CM-side channel adjustment; CM-to-CMTS channel mapping; and CMTS-side channel readjustment

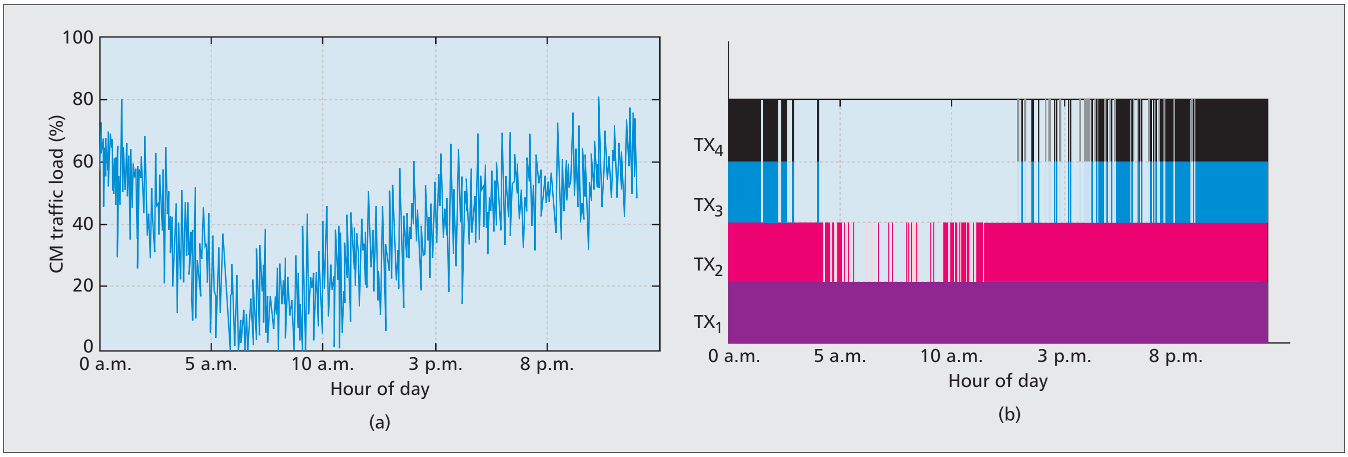

Figure 2. *a) Customer's traffic variation in a day; b) status of the channel-bonding transmitters (TXs) under energy-efficient operation.*

In phase 3, the CMTS readjusts its connections with the CMs to achieve more energy-saving. We define a readjustment threshold (*TH*) in terms of the number of CM connections on a CMTS port. After the CM-to-CMTS connection mapping is done, the CMTS redistributes connections on ports whose number of CM connections is below *TH* to busier ports with DBC-REPLACE, until no change can be made. To evaluate the energy saving on a CMTS achieved by phases 2 and 3 in the algorithm, we simulate a CMTS line card with 16 ports (either DS or US) and 1024 CMs. Each CM has independent traffic input, and the traffic fluctuation follows a similar overall trend as illustrated in Fig. 2a. The simulation results indicate that 31.08–32.61 percent energy saving can be achieved on the CMTS line card over a 24-h period, depending on the choice of *TH*. Similarly, we assume that the power consumption of a CMTS port is zero when it is shut down.

When designing the energy-saving algorithm, we need to pay attention to the energy-delay trade-off as when we shrink the transmission capacity to reduce power, the packet delay can increase due to the mismatch between the required and offered capacities. In [12], we suggested that the energy-delay trade-off can be optimized by choosing the right CM-side parameters for the algorithm, including *LW*, *HW*, and *N*. When the numbers of channels in high-, moderate-, and low-power modes of a CM are determined, packet delay can be reduced by increasing *LW*, *HW*, and *N*. Another trade-off we need to consider in the algorithm design is the number of DBC operations vs. energy saving, as invoking DBC too frequently can result in significant protocol overhead. Moreover, recent research on IEEE 802.3az suggests that switching network devices between active and sleep modes too frequently can lead to high power consumption even at low traffic load due to the transition overhead [13]. The results in [12] suggested that choosing a *TH* at 10 percent of the maximum capacity of a CMTS port can obtain 32.05 percent energy saving on the CMTS, while limiting the average number of DBC operations to 7.4/CM/h.

CM-Side Energy Saving with Energy-Efficient Packet Scheduling

The algorithm discussed above can achieve network-wide energy saving, but in a coarse way, since the operation modes of the CMs and CMTS are only updated every few minutes due to the protocol overhead of DBC. As the CMs typically contribute to more than 50 percent of the per-user power consumption in HFC networks [3], we would like to push the energy saving on the CM side further. In order to make energy-efficient operation more flexible, we need to develop an energy-efficient packet scheduling algorithm that can adaptively assign data traffic to different channels on a channel-bonding CM. The CM's transmitters (TXs) of the channels can work in a manner similar to that of the Eththe ernet interface under IEEE 802.3az energy-efficient Ethernet (EEE) standard [14]. Hence, energy saving can be achieved on a finer timescale on the CMs. Moreover, QoS differentiation can be obtained with packet scheduling on the channel level.

Figure 3a shows the system model of channel-bonding TXs that use energy-efficient scheduling. DOCSIS 3.0 classifies data packets into different service flows and offers a comprehensive QoS control mechanism [9]. By leveraging this mechanism, our energy-efficient scheduling algorithm puts incoming packets into different priority queues according to their QoS requirements and controls the operation modes of TXs on the CM to transmit them [10]. As illustrated in Fig. 3a, we define the numbers of priority queues and TXs as M_Q and M_{CM}, respectively. The energy-efficient scheduling algorithm adjusts the number of active TXs on a CM in a scheduling cycle *T*. Here, *T* can be defined as the MAP Time (e.g. 1–10 ms) to comply with DOCSIS 3.0. As illustrated in Fig. 3b, we define four operation modes for a TX:

- Work, as it is transmitting data
- Report, as it is transmitting necessary US control messages
- Sleep, as it is saving energy with nothing to transmit
- Warm-up, as it is waking up from the sleep mode

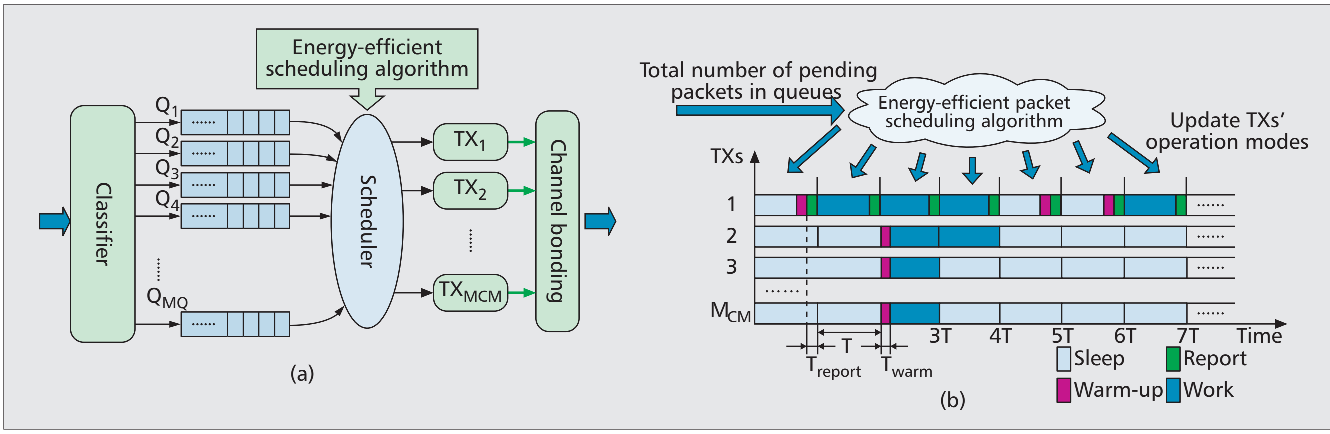

Figure 3. *a) System model; b) time diagram of energy-efficient packet scheduling.*

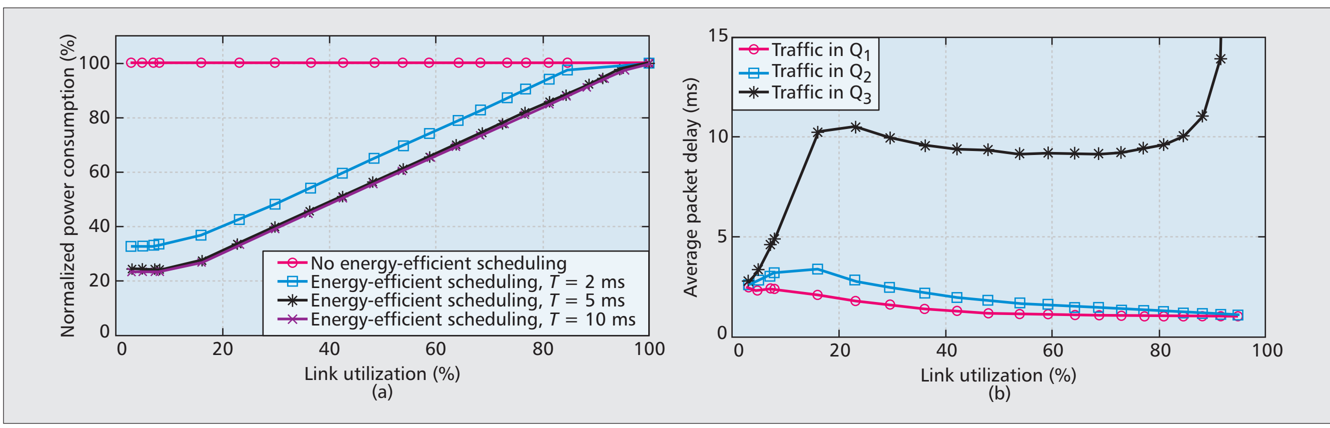

Figure 4. *a) Power consumption vs. link utilization (single packet priority); b) packet delay vs. link utilization (three packet priorities).*

The algorithm determines the operation modes of the TXs on a CM based on the total number of pending packets in all priority queues. Specifically, a series of thresholds are predefined, and the algorithm measures the total number of pending packets in all priority queues, compares the result with the thresholds, and then turns TXs on/off accordingly [10].

Figure 3b shows the detailed operation of the energy-efficient scheduling algorithm. The operation of the TX with the lowest index (e.g., TX1) is special, and it has to transmit necessary US control messages (e.g., bandwidth requests) toward the end of each scheduling cycle T. We assume that TX1 uses a fixed amount of time T_{report} to transmit the control message. We also assume that the warm-up duration of each TX is the same as T_{warm}. When transmitting data packets, the TXs handle those with higher priority earlier according to the non-preemptive rule. To prevent data packets being buffered too long when the traffic load is extremely low, we introduce K_{max} as the maximum number of continuous sleep cycles in which TX1 can stay when there are pending packets in the priority queues.

Figure 4a shows the normalized power consumption as a function of link utilization when running the energy-efficient scheduling algorithm on a four-channel bonding CM with a single packet priority. All packets have a fixed length of 1518 bytes, and the packet arrival follows a Poisson process. We assume that the power consumptions of work and report modes are the same, the power consumption of the sleep mode is 10 percent that of the work mode, and the power consumption of the warm-up mode is 20 percent that of the work mode. The scheduling cycle T is set as 2, 5, and 10 ms, and K_{max} equals 1. Notice that with the energy-efficient scheduling algorithm, the power consumption of the system scales almost linearly with the link utilization, leading to effective energy saving. We also notice that the energy saving starts to show saturation when T reaches 5 ms. This is because when we keep increasing T, TXs can stay in the sleep mode for more time, but at the same time, more packets are buffered in the priority queue, and the TXs also stay in work mode longer after being turned on. Figure 4b plots the average delays of packets when there are three priority queues (Q1, Q2, and Q3) with $T = 4$ ms and $K_{max} = 1$. The packet arrival rates of these three types of packets are the same, and the data rate of a TX is 30 Mb/s. It is interesting to notice that the average delays of higher-priority queues (e.g., Q1 and Q2) can be kept at a small value, while the packets in the lowest-priority queue, Q3, experience a much longer average delay. Hence, when the energy-efficient scheduling is operational, we can assign the data packets from delay-sensitive applications to higher-priority queues and use the lowest-priority queue for best effort traffic.

For the customer-side energy-saving with the channel-level packet scheduling algorithm, we showed that it can achieve successful transmission of delay-sensitive traffic together with energy-saving.

Conclusions

We present a few energy-saving algorithms that can be used by cable operators to improve the energy efficiency of HFC networks using channel bonding. Our discussions cover both the network-wide energy-saving algorithm and the customer-side energy-saving algorithm. We have found that by adjusting the operation modes of CMTS and CMs dynamically, effective energy saving can be obtained with the algorithms. We also verify that additional packet delay and protocol overhead introduced by the energy saving algorithms can be reduced by choosing the key parameters properly. For customer-side energy saving with the channel-level packet scheduling algorithm, we show that it can achieve successful transmission of delay-sensitive traffic together with energy saving.

Acknowledgment

This work was supported by the Fundamental Research Funds for the Central Universities (WK2100060006), by NCET Project (NCET-11-0884), by SRFDP Project (20123402120014), by Instituto de Telecomunicações, Next Generation Networks and Applications Group (NetGNA), Portugal, and by National Funding from the FCT — Fundação para a Ciência e a Tecnologia through the PEst-OE/EEI/LA0008/2011 Project. The authors would like to thank Miss Yabo Yuan from the University of Science and Technology of China for her help on the simulations.

References

[1] *Cisco Visual Networking Index: Forecast and Methodology*, 2009–2014.
[2] http://www.smart2020.org.
[3] J. Baliga *et al.*, "Energy Consumption in Wired and Wireless Access Networks," *IEEE Commun. Mag.*, vol. 49, June 2011, pp. 70–77.
[4] O. Blume, D. Zeller, and U. Barth, "Approaches to Energy Efficient Wireless Access Networks," *Proc. ISCCSP 2010*, Mar. 2010, pp. 1–5.
[5] J. Hsieh, T. Lee, and Y. Kuo, "Energy-Efficient Multi-Polling Scheme for Wireless LANs," *IEEE Trans. Wireless Commun.*, vol. 8, 2009, pp. 1532–41.
[6] P. Tsiaflakis, "Green DSL: energy-efficient DSM," *Proc. ICC 2009*, June 2009, pp. 1–5.
[7] S. Wong *et al.*, "Sleep Mode for Energy Saving PONs: Advantages and Drawbacks," *Proc. GLOBECOM 2009*, Dec. 2009, pp. 1–6.
[8] J. Cioffi, "Lighting Up Copper," *IEEE Commun. Mag.*, vol. 49, May 2011, pp. 30–43.
[9] http://www.cablelabs.com/cablemodem/specifications/index.html
[10] P. Lu *et al.*, "Energy-Efficient Scheduling and Energy-Delay Tradeoff in Green Hybrid Fiber-Coaxial Networks," *Proc. IEEE GLOBECOM 2012*, Dec. 2012, pp. 1–5.
[11] M. Garcia *et al.*, "Analysis and Modeling of Traffic on A Hybrid Fiber-Coax Network," *IEEE JSAC*, vol. 22, Nov. 2004, pp. 1718–30.
[12] Z. Zhu, "Design Energy-Saving Algorithms for Hybrid Fiber Coaxial Networks Based on DOCSIS 3.0 Standard," *J. Opt. Commun. Net.*, vol. 4, June 2012, pp. 449–56.
[13] P. Reviriego *et al.*, "An Initial Evaluation of Energy Efficient Ethernet," *IEEE Commun. Lett.*, vol. 15, May 2011, pp. 578–80.
[14] K. Christensen *et al.*, "IEEE 802.3az: the Road to Energy Efficient Ethernet," *IEEE Commun. Mag.*, vol. 48, Nov. 2010, pp. 50–56.

Biographies

ZUQING ZHU [M'07, SM'12] (zqzhu@ieee.org) received his B.S. degree from the Department of Electronic Engineering and Information Science, University of Science and Technology of China, Hefei, in 2001, and his M.S. and Ph.D. degrees from the Department of Electrical and Computer Engineering, University of California, Davis, in 2003 and 2007, respectively. In July 2007, he joined the Cable Modem Termination System Business Unit (CMTS BU) of Cisco Systems, San Jose, California, as a senior software engineer. At Cisco, he worked on research and development of wideband access routers for HFC cable networks and was involved in the standardization of the DOCSIS 3.0. In January 2011, he joined the University of Science and Technology of China (USTC), where he is now an associate professor. He has published more than 80 papers in peer-reviewed journals and conferences of IEEE, IET, and OSA. He has been in the technical program committees (TPC) of INFOCOM, ICC, GLOBECOM, ICCCN, and others. He is also an Editorial Board member of the *Journal of Optical Switching and Networking* (Elsevier), *Telecommunication Systems Journal* (Springer), *European Transactions on Emerging Telecommunications Technologies* (Wiley), and others. He is a member of OSA.

JOEL J. P. C. RODRIGUES [S'01, M'06 SM'06] (joeljr@ieee.org) is a professor at the University of Beira Interior, Portugal, and researcher at the Instituto de Telecomunicações, Portugal. He is the leader of the NetGNA Research Group (http://netgna.it.ubi.pt), Chair of the IEEE ComSoc Technical Committee on Communications Software, Vice-Chair of the IEEE ComSoc Technical Committee on eHealth, and Member Representative of ComSoc on the IEEE Biometrics Council. He is the Editor-in-Chief of the *International Journal on E-Health and Medical Communications* and *Recent Patents on Telecommunications*, and an editorial board member of several journals. He has been general chair and TPC Chair of many international conferences. He has authored or coauthored over 250 papers in refereed international journals and conferences, a book, and two patents. He was awarded the Outstanding Leadership Award of IEEE GLOBECOM 2010 as CSSMA Symposium Co-Chair and several best papers awards. He is a licensed professional engineer (as a Senior Member), a member of the Internet Society, an IARIA Fellow, and a Senior Member of ACM.

PING LU [S'11] received his B.S. degree from the Department of Optical Information Science and Technology, Xi'an University of Technology, China, in 2011. He is working toward his Ph.D. degree at the School of Information Science and Technology, USTC. His research interest is in energy-efficient networks.

YONGGANG WEN [S'99, M'08] (ygwen@ntu.edu.sg) He received his Ph.D. degree in electrical engineering and computer science (with a minor in western literature) from Massachusetts Institute of Technology (MIT), Cambridge, in 2008; his M.Phil. degree (with honor) in information engineering from the Chinese University of Hong Kong (CUHK), China, in 2001, and his B.En.g degree (with honor) in electronic engineering and information science from USTC in 1999. He is currently an assistant professor with the School of Computer Engineering at Nanyang Technological University, Singapore. His research interests include cloud computing, mobile computing, multimedia networks, cyber security, and green ICT.