Telecommun Syst (2013) 52:1327–1340
DOI 10.1007/s11235-011-9644-8

Lightpath scheduling and routing for green data centres

Jing Wu · James Yiming Zhang · Michel Savoie

Abstract Optical networking technologies enable data centres to be located near sources of green energy (i.e., renewable energy). Since some green energy sources are intermittent and are not always available, we need to dynamically connect distribution networks to the green energy powered data centres. On the other hand, the availability of green energy is reasonably predictable, and we are thus able to schedule connectivity to data centres in advance. We propose a WDM network planning model, which allows lightpaths to slide within their desired timing windows with no penalty on the optimization objective, and to slide beyond their desired timing windows with a deteriorating "green-level". Our simulation results show the tradeoffs between the consumption of brown energy (i.e., energy generated by carbon-intensive means), the capability of providing required connectivity to data centres, network resource utilization, and overall operation objective.

Keywords WDM networks · Lightpath scheduling · Data centre interconnection · Green energy · Network resource utilization

1 Introduction

Optical networks are widely used in connections between data centres, which are facilities that primarily host electronic equipment for data processing, data storage, and communications. In connecting data centres to each other and also to their respective distribution networks, optical networks have attractive advantages, such as high speed, large capacity and energy efficiency [1, 2]. With the advanced optical networking technologies, data centres can be located near sources of green energy (i.e., renewable energy). In this way, instead of the transmission of electricity from sources of green energy to data centres, data can be transmitted to and among data centres over optical networks, which results in less loss in electricity transmission and more efficient use of green energy.

We aim at reducing the overall use of "brown" energy by data centres and on the other hand, maximizing their use of green energy. Green data centres are the data centres that mainly consume the electricity generated by renewable energy, such as wind, solar, tide and hydro energies [3]. Green data centres can avoid the use of brown energy that is generated by carbon-intensive means, such as coal or gas burning power plants. At the same time, green data centres may use brown energy as secondary or backup options. However, using brown energy costs more than using its green counterpart, and more so in social and environmental senses. The cost varies for using different power sources at different times.

Because some sources of green energy are intermittent and are not always available, we need to dynamically change the connectivity of data centres. We use the terminology "connectivity of data centres" to refer to two types of connections: connections between data centres, and connections from data centres to their respective distribution networks for end users [4]. The availability of green energy highly depends on weather and other environmental factors, making it intermittent and fluctuating [5–7]. However, measurement and predication of green energy availability, together with

---

J. Wu (✉) · J.Y. Zhang · M. Savoie
Communications Research Centre (CRC) Canada, 3701 Carling Avenue, Ottawa, Ontario, Canada, K2H 8S2
e-mail: jing.wu@crc.gc.ca

J.Y. Zhang
e-mail: yizhang@site.uottawa.ca

M. Savoie
e-mail: michel.savoie@crc.gc.ca

the use of rechargeable batteries, make green energy reasonably predictable [8–10]. The operation of green data centres must be planned to take into account the predicted green energy supply at different times and sites. Accordingly, we can schedule the change of connectivity of data centres in advance to make better use of greener time slots.

There are daily or weekly patterns of workloads and connectivity requirements of data centres. It has been observed that workloads of transaction-oriented servers vary significantly depending on time of day, day of week, or other external factors [11–14]. Network traffic at the optical layer periodically fluctuates. For core optical networks such as Wavelength Division Multiplexing (WDM) networks, network traffic is observed following daily patterns [15], or weekly patterns [16]. Such traffic patterns can be predicted from historical statistics, which repeat every day (or week) with minor variations in timing and volume. It is expected that a large portion of the traffic that is carried by optical networks are related to data centres. The knowledge about workloads and traffic patterns provides an opportunity to schedule the connectivity of data centres.

Using scheduled lightpaths to connect green data centres has different timing requirements from other lightpath scheduling problems. Existing methods for the lightpath scheduling problems assume that a lightpath should be set up either at a given time, or within a given time window, which makes the lightpath scheduling inflexible for connecting green data centres. For example, in [17–21], network planning was conducted for a set of lightpath requests, each having a pre-specified starting and ending time. In [22–25], static WDM network planning was conducted for fixed holding-time lightpath requests, each one being allowed to slide within its given time window. However, for connecting green data centres, timing flexibility in lightpath scheduling is very important, not only because of the variability of the green power, but also because of the network operations. Network operators are concerned about not only timing violations (related to the "green-level" of data centres), but also the resource utilization, as well as lightpath rejections (related to the connectivity of data centres). Network operators need tools to make wise tradeoffs between these goals, e.g., network operators would rather adjust lightpaths scheduling timing, than reject lightpaths that cannot be accommodated due to their strict timing requirement or impractical timing windows.

To connect green data centres, scheduled lightpaths can be allowed to slightly slide in time, without deteriorating the performance. Sliding-timing scheduling potentially provides better network resource utilization than fixed-timing scheduling. Since lightpaths are scheduled based on the statistical availability of green energy at different data centres and traffic characteristics, minor timing slides should not impact much on the performance of the traffic adaptation, while dramatic timing slides, on the other hand, should be avoided. For example, due to the availability of sunlight at a given solar powered data centre, the network operator needs to provision a lightpath between the data centre to a distribution network during the day time. It usually does not make much difference if the timing slide is far below the variance of green energy availability, e.g., starting the lightpath from 8:00 AM or 10:00 AM. However, setting up the lightpath at 5:00 AM cannot efficiently transfer data from the data centre, due to the lack of sufficient solar power at the data centre.

In the above-mentioned applications, the extent of the timing satisfaction or violation needs to be quantitatively measured. Moreover, a timing window should not be used in a "binary" way, i.e., it can either be satisfied (thus the corresponding lightpath is accepted), or not (thus the corresponding lightpath is rejected). Network operators normally would prefer scheduled lightpaths being centered on their desired timing, with a decreasing tolerance level, which can take into consideration both the "green" level and the timing punctuality, as scheduled lightpaths move away from their desired timing windows. This requires proper modelling of the extent of timing satisfaction or violation, which has not been done by the existing methods for the static lightpath scheduling, and therefore motivates this study.

Our study aims at planning scheduled lightpaths to adapt to relatively stable traffic patterns and predictable availability of green energy. Static scheduled lightpath demands are generated based on forecast traffic patterns and green energy availability at data centres, and are input to our network planning problem. Our problem is different from dynamic lightpath scheduling problems, which generally do not assume any a priori information of how traffic patterns change [26–28]. In our approach, once a lightpath is pre-planned, it becomes available to carry traffic at its scheduled time. In contrast, dynamic lightpath scheduling cannot guarantee the availability of a lightpath. Only when a request arrives, the network operator makes real-time decisions depending on the network resource availability at the moment of the request. Our approach achieves a better coordination of lightpaths than dynamic lightpath scheduling by taking advantage of known traffic patterns, as well as green energy availability for data centres [28].

This paper is organized as follow. In Sect. 2, we outline the energy efficiency problems, challenges, possible solutions and emerging opportunities of data centres. In Sect. 3, we summarize the networking requirements of green data centres, as well as assumptions used in our model. In Sect. 4, we present our model, followed by numeric results in Sect. 5. We conclude this paper in Sect. 6.

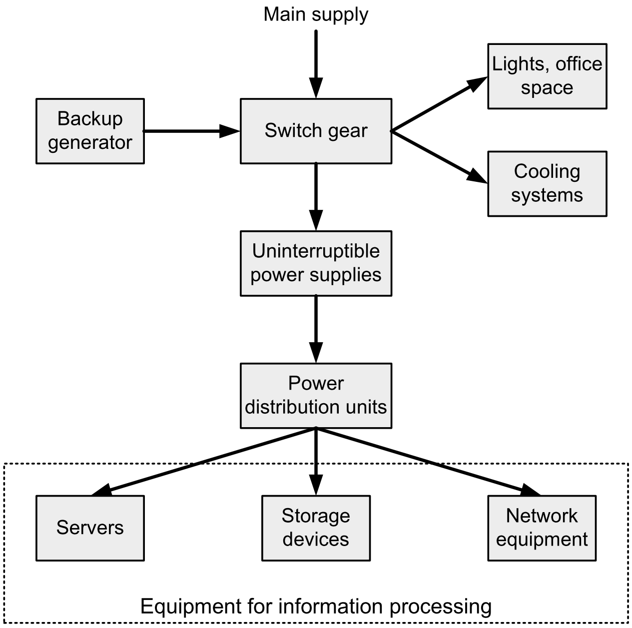

Fig. 1 Typical electrical systems in a data centre [29]

2 Green data centres

The energy efficiency problems, challenges, possible solutions and emerging opportunities of servers and data centres are highlighted in a comprehensive report developed by the United States Environmental Protection Agency (EPA). EPA issued a report to the US congress on energy efficiency of servers and data centres in August 2007 [29]. Some of its key findings are:

- In 2006, servers and data centres in the US consumed 61 billion kWh (kilowatt hours), and their total electricity cost was US $4.5 billion;
- Their energy consumption grew rapidly and the trend continues. From 2000 to 2006, their total electricity consumption doubled;
- On average, their site infrastructure consumed half of the electricity for cooling systems, power delivery, and so on;
- Among the information processing equipment, volume servers consumed 68% of the electricity in 2006; while network equipment steadily consumed approximately 10% of the electricity over 2000–2006.

Typical electrical systems in a data centre include main and backup power supplies, electrical systems for the facility, and power regulatory systems for data processing equipment. Connections of electrical systems in a data centre are illustrated in Fig. 1. Due to the stringent requirements of data processing equipment, power conditioning systems are normally used, such as power distribution units, rechargeable batteries/uninterruptible power supplies.

Data centres suffer from low utilization and low energy efficiency. Normally, the capacity of data centres is designed based on the peak demand. However, the real demand varies over time. Servers operate most of the time at 10%–50% of their maximum utilization levels [12, 13, 30–34]. Unfortunately, when the utilization of servers is low, their power consumption remains high. Currently, servers still consume over 50% of their peak power usage, even when the servers are idle, resulting in low energy efficiency of servers [12, 30, 34, 35]. Since on average data centres consume half of the electricity for cooling systems, power delivery, and so on [29], low energy efficiency of servers results in even lower energy efficiency of data centres. Servers at low utilization still generate significant heat, which requires cooling systems [36]. Even idle servers need power to be delivered, which causes power delivery loss.

The capability of on-demand switching off servers increases the energy efficiency of data centres. In another word, an easy step to increase energy efficiency of data centres is turning off the power of idle servers, and turning them on when necessary [11, 34]. Virtual machine migration allows the relocation of active tasks from one physical host server to another, without any major performance degradation on the active tasks running on servers [37–43]. Active tasks running on low utilized servers may be dynamically consolidated into fewer servers, creating opportunity to scale back the power consumption of idle servers (including powering-off idle servers) [12, 13, 32, 44–49]. When additional service capacity is needed, the servers that are powered off may be powered on and provide service in minutes [50]. Software applications may be intelligently mapped to underlying servers based on the workload, capability and power consumption profiles of these servers [51–58]. As proposed in [59], intelligent distribution of computational workload was explored across geographically distributed data centres to optimize energy consumption and cost. It took advantages of the cost-saving factors, including different and variable electricity prices, peak-demand prices versus off-peak-demand prices due to time zone differences, and green energy versus brown energy.

Making data centres use green energy has different requirements from minimizing the total energy consumption or reducing the peak power consumption of data centres. For example, the total energy consumption may be slightly increased, when we aim at maximizing the usage of green energy. However, our goal is viable due to its long-term social, environmental, and economic benefits. Previous studies addressed the problem of distributing tasks in web servers to minimize their total energy consumption and at the same time to satisfy performance requirements [13, 47, 60–64]. There are also other efforts in reducing peak power requirements at chips, servers, racks and data centres [65].

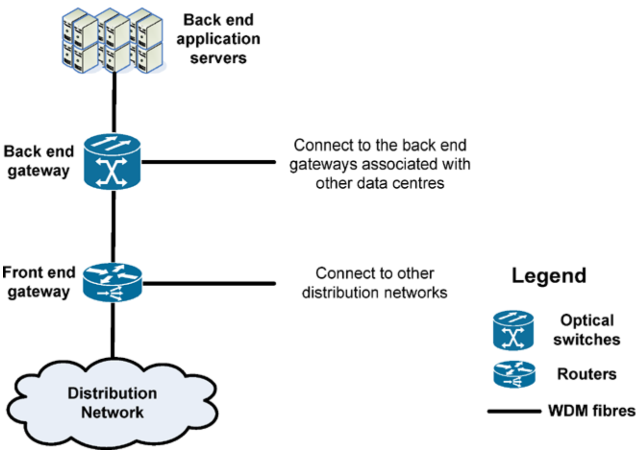

Fig. 2 A data centre interconnected by WDM networks

3 Networking requirements of green data centres

3.1 Networking requirements of green data centres

High-bandwidth networks are used to connect data centres, as well as key systems within a data centre. The data processing equipment within a data centre consists of three key systems (shown in Fig. 2): front and back end gateways, and back end application servers. The front end gateway is connected to distribution networks, e.g., IP service provider networks. The front end gateway supports client-to-server communications and provides the high-speed server-to-server networking with the back end gateway and other front end gateways. The front end gateway is composed of firewalls, and servers for content caching, load balancing, and intrusion detection. The back end gateway is connected to the back end gateways at other data centres, which could be located beyond the same metropolitan domain. In this example, we use WDM networks due to their advantages of protocol transparency, high-bandwidth offered by multiple optical channels, and most importantly their energy efficiency. However, other high-bandwidth networks such as metropolitan area Ethernet networks, and SONET/SDH networks are also viable solutions for data centre interconnection [44].

Optical networks are ideally suited for bridging the gap between the desired locations of data centres from the facility's perspective and from the information processing's perspective. There are benefits to locate data centre facilities in rural and remote areas that provide data security from natural disasters, green energy generation, lower land cost, lower external surrounding temperature, and running water for cooling. However, the information processing should ideally be located close to users to reduce delay and network congestion. Optical networks provide huge bandwidth and point-to-point IP links, while consuming less energy than any other communication networks. By using optical networks, we effectively remove the need of the long-haul transmission of electricity from power generation stations to data centres.

We pre-plan schedules for setting up lightpaths to connect data centres to the requesting distribution networks. We assume that high-volume digital contents are hosted by large data centres. Each piece of content is duplicated to 2 or 3 mirror data centres. Lightpaths originate from the back end gateway of the data centre that store the digital content, and terminate at the front end gateway of the data centre that is located close to the requesting distribution network. Network power cost is considered constant, regardless of carried traffic. We assume that all wavelength channels in all WDM fibres are available for the lightpath scheduling and routing. In reality, optical virtual private networks may be used, where a set of wavelength channels are allocated by the lightpath scheduling and routing algorithm.

Lightpath demands are created based on the predicted green energy availability for data centres and the requests for contents that are hosted by the data centres. If the preferred schedule for a lightpath cannot be satisfied, a lightpath may be shifted in its schedule with a higher cost, due to its increased use of brown energy and the reduced timing satisfactory level. Our problem is to provide the best scheduling and Routing and Wavelength Assignment (RWA) schemes of lightpaths, which are generated off-line by our optimization algorithm, to provide the required connectivity of data centres and at the same time to minimize the use of brown energy.

3.2 WDM network operations, modeling and assumptions

We consider wavelength-routed WDM networks with mesh topologies. We model a general topology WDM mesh network of N nodes interconnected by E links. Each node represents a data centre, including its front and back end gateways. Each link consists of a pair of fibres, each fibre for one direction and having W non-interfering wavelength channels. Two nodes can be connected through a lightpath defined as a concatenated sequence of wavelength channels [66]. We assume no wavelength conversion is used due to its high cost and little benefit for the static WDM planning problem [67]. So a lightpath must use the same wavelength all the way from its source node to destination node.

Lightpaths are scheduled to be set up at the beginning of their starting time slots and be torn down at the end of their finishing time slots. Network-wide synchronous time slots are used for resource allocations and lightpath scheduling. All time slots have the same fixed duration, which should be one hour or larger. We assume the time to set up or tear down a lightpath (i.e., signalling time) is negligible compared to the duration of a time slot. The holding time of a lightpath is fixed and known in advance, measured by the number of time slots. Without losing generality, we number the time slots in our planning time horizon sequentially from 0 to $Z-1$ ($0 \leq t < Z$). The complexity of our scheduling problem increases, as the number of time slots in the planning time horizon increases. A detailed analysis of computation complexity can be found in [68].

We aim at scheduling and allocating network resources to lightpaths, i.e., planning network operations for Scheduled Sliding Lightpath Demands (SSLDs). The network resources primarily include wavelength channels. We provide an accepted SSLD with a schedule (i.e., starting time slot) and an RWA scheme (i.e., a list of allocated wavelength channels). The same RWA scheme is used for the entire holding time of an SSLD, i.e., once an SSLD is accepted, it stays connected from its starting time slot to its finishing time slot. If there is insufficient resource for an SSLD during its holding time, the SSLD is rejected.

4 Problem formulation

4.1 Notations

For the remainder of this paper, the following notations and variables are used:

Network model related parameters:

V	the set of all nodes in the network;
e_{ij}	the fibre between node i ($i \in V$) and node j ($j \in V$);
E	the set of all fibres in the network, i.e., $\{e_{ij}\}$, ($i \in V$, $j \in V$);
W	the number of wavelengths used in the network;
N	the number of nodes in the network.

SSLD related parameters:

l_h	the hth SSLD. If it is accepted, we use the same notation to refer to the lightpath that is provided to the SSLD;
L	the total number of SSLDs, $0 \leq h < L$.

Scheduling related parameters:

Z	the total time slots of our scheduling problem;
t_h	the holding time of SSLD l_h. Based on our assumption, it is the same as the lifespan of lightpath l_h, if SSLD l_h is accepted;
$[b_h, b'_h]$	the desired window of starting time for SSLD l_h, $0 \leq b_h \leq b' < Z$;
y_h	the weight for earliness penalty of lightpath l_h;
r_h	the weight for tardiness penalty of lightpath l_h;
E_h	the overall timing violation penalty of lightpath l_h.

Cost and revenue related parameters:

h_{ijt}	the cost of using a wavelength channel on link e_{ij} ($e_{ij} \in E$) for time slot t;
C_h	the routing cost of lightpath l_h, i.e., the cost of wavelength channels used by lightpath l_h;
D_h	the dual routing cost of lightpath l_h;
P	the penalty for rejecting an SSLD, i.e., the loss of revenue if an SSLD is rejected.

Decision variables:

α_h	the binary integer variable indicating the admission status of SSLD l_h. It is one, if SSLD l_h is accepted. Otherwise, it is zero;
β_h	the starting time slot of lightpath l_h, $0 \leq \beta_h < Z$;
δ^h_{ijct}	the binary integer variable representing the use of the cth wavelength channel on fibre e_{ij} ($e_{ij} \in E$, $0 \leq c < W$) at time slot t by lightpath l_h. It is one, if lightpath l_h uses such wavelength channel at such time slot. Otherwise, it is zero;
A	the admission status of all SSLDs, i.e., (α_h), $0 \leq h < L$;
B	the starting time slots of all lightpaths, i.e., (β_h), $0 \leq h < L$;
Δ_h	the RWA scheme of lightpath l_h, i.e., $(\delta^h_{ijct})_h$;
Δ	the RWA schemes of all lightpaths, i.e., (Δ_h), $0 \leq h < L$.

4.2 Objective function

Our goal is to provide as much required connectivity for data centres as possible for their desired peak operation time slots, so that their usage of green energy is maximized. We model our goal by an optimization objective as minimizing the rejection of requests, the resource usage and the timing violations of lightpaths. We want to accept as many profitable requests as possible, and for the accepted requests, we want to find lightpath schedules that respect their timing preference as much as possible, while at the same time provide them with RWA schemes that use as few resources as possible.

Our objective function is to minimize the function J, i.e.,

$$\min_{A,B,\Delta}\{J\}, \text{ where } J \equiv \sum_{0 \leq h < L} \big[(1-\alpha_h)P + \alpha_h(C_h + E_h)\big].$$

The overall penalty consists of the rejection penalty (i.e., P), the resource usage cost (i.e., C_h) and the timing violation penalty (i.e., E_h). The timing violation penalty could be either an earliness penalty or a tardiness penalty. The shortage of effective service time caused by schedule earliness and tardiness is shown in Fig. 3. When SSLD l_h is scheduled sooner than its desired starting time, the lightpath will be removed sooner than the desired ending time, causing a shortage of the effective service time after the lightpath is removed. This is because we assume the lifespan of lightpath l_h is exactly the same as the holding time of SSLD l_h. On the other hand, when lightpath l_h is scheduled later than its desired starting time of SSLD l_h, there will be a shortage of effective service time before the lightpath starts.

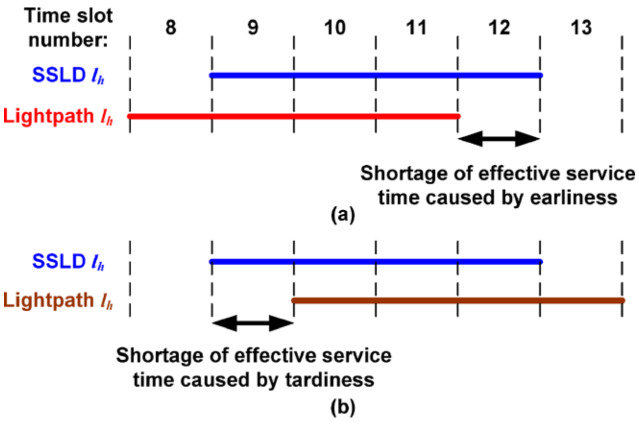

Fig. 3 Shortage of effective service time caused by schedule earliness and tardiness

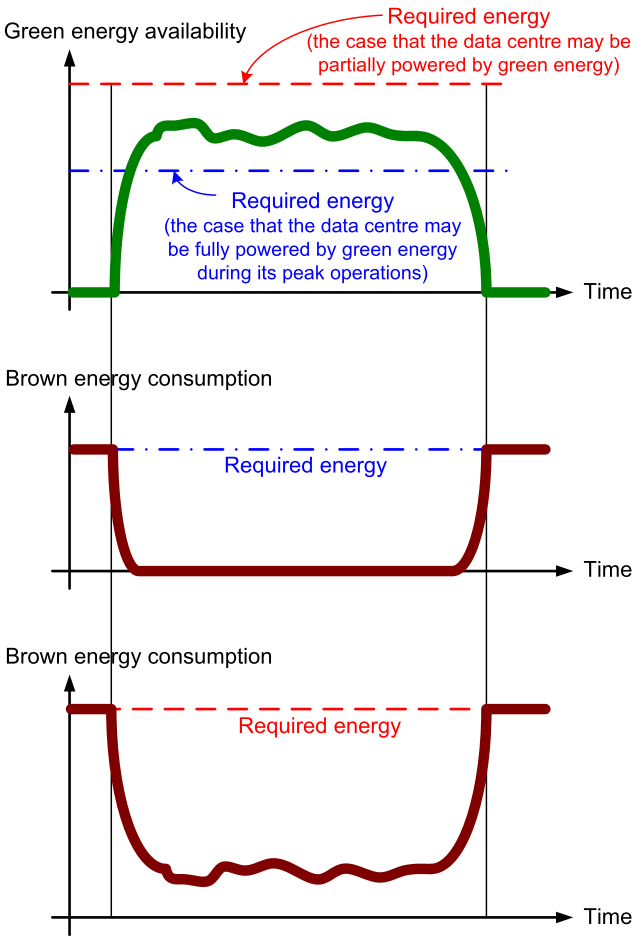

Fig. 4 Green energy availability and brown energy consumption of a data centre for a period of one day

Some sources of green energy are intermittent and are not always available. As an example, we illustrate green energy availability and brown energy consumption of a data centre for a period of one day in Fig. 4. If the data centre operates at its peak capacity, there are two cases of its required energy: either fully or partially powered by green energy. When the data centre operates outside of the green energy availability time window, it has to use brown energy. To maximize its use of green energy, its connectivity needs to be scheduled to match its green energy availability time window as much as possible. We impose higher operational penalty, if its connectivity does not perfectly match its green energy availability.

In this paper, we adopt earliness and tardiness penalties defined in [69]:

$$E_h = \begin{cases} y_h \times (b_h - \beta_h)^2 & \text{if } \beta_h < b_h \\ & \text{(i.e., earliness penalty)} \\ 0 & \text{if } b_h \le \beta_h \le b'_h \\ & \text{(i.e., no penalty, since the} \\ & \text{lightpath starts within} \\ & \text{its desired starting time} \\ & \text{window)} \\ r_h \times (\beta_h - b'_h)^2 & \text{if } \beta_h > b'_h \\ & \text{(i.e., tardiness penalty)} \end{cases}$$

$$0 \le h < L \tag{1}$$

where y_h and r_h are the weights for earliness and tardiness penalties of lightpath l_h. Our earliness and tardiness penalties reflect a high penalty as a data centre is forced to use brown energy, when its connectivity moves away from its desired timing window (shown in Fig. 5). When y_h and r_h are set to infinitely large positive values and the resource costs are set to zero, this formulation then becomes the same as the fixed time-window scheduling problem, which does not allow any timing violations.

The cost of routing lightpath l_h is denoted as C_h and defined as the total cost of using wavelength channels. We use this definition to illustrate that the cost of a routing lightpath is a weighted summation of certain parameters related to links and nodes. Such a definition may be extended to incorporate other parameters without changing the framework proposed in this paper.

$$C_h = \sum_{\beta_h \le t < (\beta_h + t_h)} \left(h_{ijt} \times \sum_{e_{ij} \in E} \sum_{0 \le c < W} \delta^h_{ijct} \right),$$

$$0 \le h < L \tag{2}$$

Our design variables are the admission status of all SSLDs (A), the starting time slots of all lightpaths (B), and the RWA schemes for all lightpaths (Δ). Our design variables are not completely independent. They represent three inter-related sub-problems, i.e., lightpath request admissions, lightpath scheduling, and RWAs.

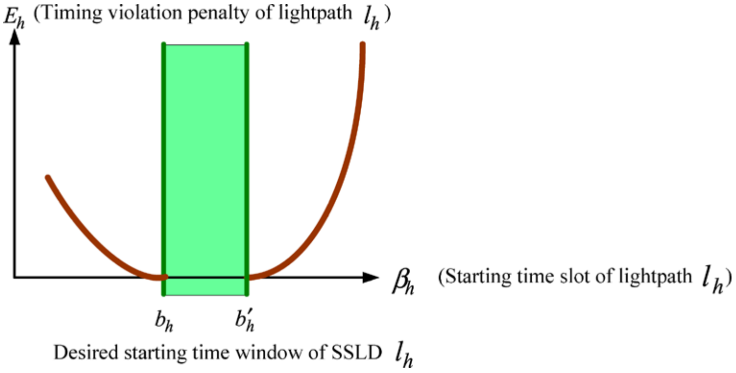

Fig. 5 Example of earliness and tardiness penalties reflecting a decreasing tolerance level as a scheduled lightpath moves away from its desired timing window

4.3 Constraints

4.3.1 Lightpath continuity constraints

If SSLD l_h is admitted, its RWA must be continuous along its path and be terminated at its two end nodes.

$$\sum_{j\in V}\sum_{0\le c<W}\delta^h_{ijct}-\sum_{j\in V}\sum_{0\le c<W}\delta^h_{jict} = \begin{cases}\alpha_h & \text{if } i \text{ is the source node of } s_h\\ -\alpha_h & \text{if } i \text{ is the destination node of } s_h\\ 0 & \text{otherwise,}\end{cases}$$

$$0\le h<L, i\in V, 0\le t<Z \qquad (3)$$

If l_h is accepted (i.e. $\alpha_h = 1$), at its source node, there is one lightpath going out; at its destination node, there is one lightpath coming in; at any intermediate node, this lightpath does not contribute to the number of lightpaths that terminate at this node. For any node that is not related to this lightpath, or when l_h is rejected (i.e., $\alpha_h = 0$), this lightpath does not contribute to the number of lightpaths that terminate at the node. These constraints confine that if and only if $\alpha_h = 1$, during the lifespan of lightpath l_h (i.e., $\beta_h \le t < (\beta_h + t_h)$), there must be a lightpath from its source node to its destination node.

4.3.2 Exclusive wavelength channel usage constraints

$$\sum_{0\le h<L}\delta^h_{ijct}\le 1, \quad e_{ij}\in E, 0\le c<W, 0\le t<Z \qquad (4)$$

Every wavelength channel at any time slot t cannot be used by more than one lightpath.

4.3.3 Lightpath persistency constraints

$$\delta^h_{ijcx}=\delta^h_{ijcy}, \quad 0\le h<L, e_{ij}\in E, 0\le x<Z, 0\le y<Z \qquad (5)$$

During the lifespan of lightpath l_h, its RWA scheme must remain the same for all time slots.

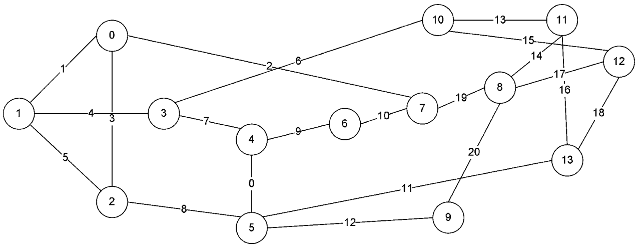

Fig. 6 Example network for performance evaluation (NSFNET)

5 Numeric results

We study the design tradeoffs in an example network operating under randomly generated connectivity requirements of data centres, which represent their best utilization of green energy. Our example network is a mesh topology network (i.e., NSFNET) with 14 nodes and 21 links. Each node represents a data centre, whose internal structure is shown in Fig. 2. The network topology is shown in Fig. 6, which marks the sequence number of nodes and links. In this section, we present results for one particular connectivity requirement of data centres, while the same trends are observed under several other connectivity requirements. The number of SSLDs for all node pairs is listed in Table 1, where the number on the ith row and the jth column represents the total number of SSLDs demands from node i to j. We randomly assign their values between 0 and 3. The total number of SSLDs is 286. Their timing requirements are also randomly generated.

We applied a Lagrangian Relaxation and Subgradient Method (LRSM) to the formulated optimization problem. We aim at obtaining near-optimal solutions to our problem, while providing a tight performance bound that can be used to evaluate the optimality of our solution. Details on LRSM can be found in [70].

In our first example, we study the tradeoffs between multiple design criteria:

- Consumption of brown energy;
- Capability of providing required connectivity to data centres;

Table 1 Number of SSLDs for all node pairs

0	1	3	1	3	1	3	0	2	0	3	2	0	3
0	0	0	2	0	1	0	0	1	0	1	0	0	3
3	0	0	3	0	1	2	3	2	3	1	2	3	0
3	1	0	0	1	1	2	3	2	2	3	2	0	3
1	0	1	2	0	3	3	2	0	3	3	1	1	3
1	2	1	3	2	0	1	3	3	1	2	1	0	2
3	1	2	3	3	3	0	3	3	1	3	2	3	3
0	0	0	0	1	0	3	0	0	1	3	0	2	0
3	0	1	3	3	3	3	0	0	2	3	1	1	2
0	0	0	1	2	0	2	0	1	0	1	0	0	3
1	0	0	2	0	3	1	1	0	3	0	3	0	3
2	3	1	1	3	2	3	2	2	2	3	0	1	3
2	0	0	0	0	1	2	0	3	0	2	0	0	3
1	3	0	2	3	2	3	3	1	2	3	3	3	0

- Network resource utilization; and
- Overall operation objective.

We first set the tardiness penalty to a fixed value and study the impact of the earliness penalty. We vary the weight of the earliness penalty for SSLDs, so that when an SSLD's starting time is earlier than its desired starting time, different earliness penalties are imposed on the operation objective.

The indices of the consumption of brown energy are measured by our timing violation function:

Brown Energy Consumption Earliness Index

$$= \sum_{0 \le h < L} y_h \times (b_h - \beta_h)^2 \tag{6}$$

Brown Energy Consumption Tardiness Index

$$= \sum_{0 \le h < L} r_h \times (\beta_h - b'_h)^2 \tag{7}$$

To better understand the tradeoffs between network parameters, we introduce two measurements of timing violation: the Sum of Earliness Violations (SEV), and the Sum of Tardiness Violations (STV), defined as:

Sum of Earliness Violations (SEV)

$$= \sum_{0 \le h < L} \min\{0, (b_h - \beta_h)\} \tag{8}$$

Sum of Tardiness Violations (STV)

$$= \sum_{0 \le h < L} \min\{0, (\beta_h - b'_h)\} \tag{9}$$

As the weight for earliness penalty y_h increases, fewer earliness violations are observed (shown in Fig. 7). The reason is that with a given total cost budget for each lightpath setting to a fixed value (i.e., $P = 100$), a high weight for earliness penalty quickly uses up the total cost budget.

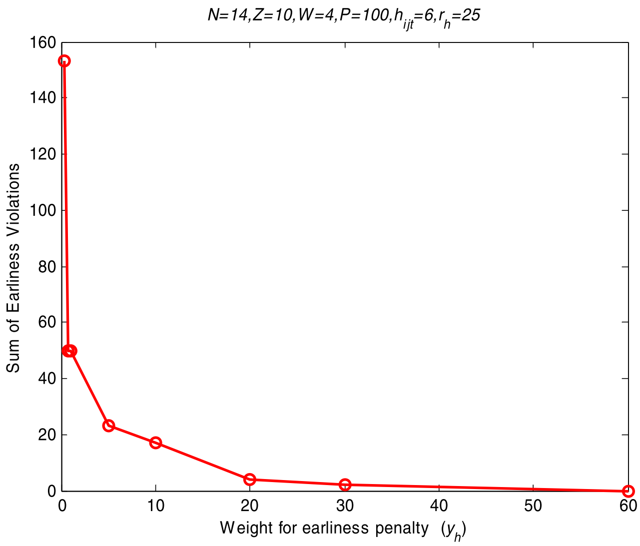

Fig. 7 Sum of Earliness Violations as y_h varies

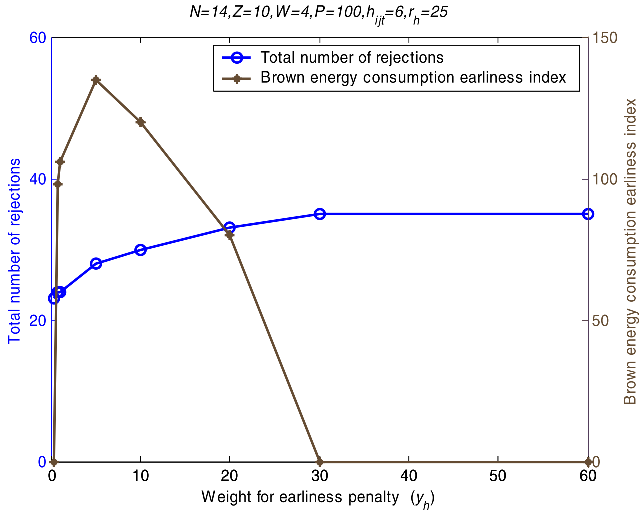

Fig. 8 Total number of rejected SSLDs vs. the Brown Energy Consumption Earliness Index as y_h varies

For the same reason, as the weight for earliness penalty y_h increases, the total number of rejected SSLDs increases, because SSLDs with insufficient cost budget are rejected (shown in Fig. 8). When the weight for earliness penalty y_h reaches a certain value, SSLDs are either scheduled strictly respecting their preferred timing requirements, or rejected due to their insufficient cost budget to cover the high penalty for earliness (shown in Fig. 8).

The Brown Energy Consumption Earliness Index sharply increases, as the weight for earliness penalty y_h increases (shown in Fig. 8). At a certain point, the Brown Energy Consumption Earliness Index gradually decreases (shown in Fig. 8), since the total number of earliness violations decreases at a faster pace (shown in Fig. 7). When all accepted SSLDs are scheduled strictly respecting their timing require-

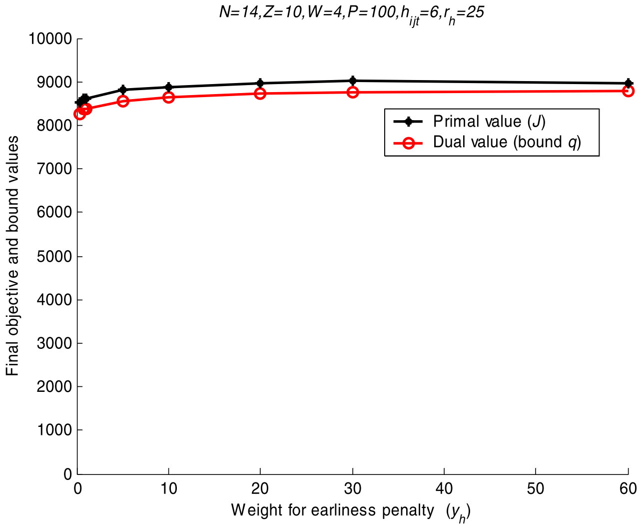

Fig. 9 Achieved optimization objective and its bound as y_h varies

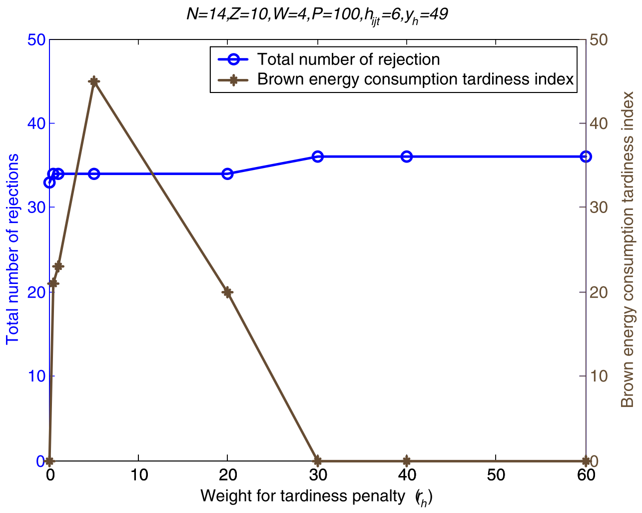

Fig. 10 Total number of rejected SSLDs vs. the Brown Energy Consumption Tardiness Index as r_h varies

ment, the Brown Energy Consumption Earliness Index stays at zero (shown in Fig. 8). No early operation at data centres is scheduled before their green energy approximately reaches peak. However, completely disabling earliness in data centre operation schedule is not optimal for the overall objective, which is evidenced by a higher achieved objective value (shown in Fig. 9). In Fig. 9, we also demonstrate that our results are highly optimal (all within 3% from the bound).

We now set the earliness penalty to a fixed value and study the impact of the tardiness penalty, i.e., we vary the weight of the tardiness penalty for SSLDs. The tradeoff between the consumption of brown energy and the capability of providing required connectivity to data centres is observed in Figs. 10 and 11.

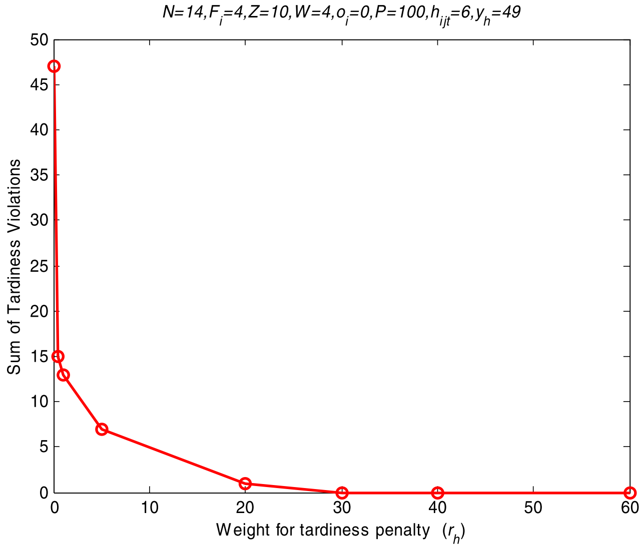

Fig. 11 Sum of tardiness violations as r_h varies

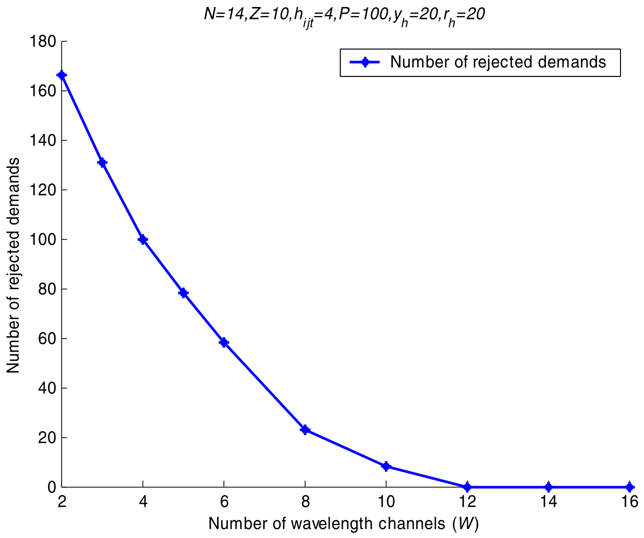

Fig. 12 Number of rejected SSLDs as W varies

In our second example, we study the impact of available network resources on the optimization objective. We vary the number of wavelength channels on each fibre (denoted by W). The impact on the number of rejected SSLDs is shown in Fig. 12. The achieved optimization objective and its bound are shown in Fig. 13. In this example, we set parameters $h_{ijt} = 4$, $P = 100$, $y_h = 20$ and $r_h = 20$. In this way, a lightpath may be scheduled up to two time slots ahead or behind of its desired timing, which causes a penalty of $20 \times 2^2 = 80$ within its 100 total budget. In Fig. 12, we can see that as W reduces, the number of rejected SSLD reduces. We do not observe any obvious trend in changing the timing violations as W changes. In Fig. 13, we can see again that our algorithm consistently produces a near-optimal solution that is very close to the lower bound.

Fig. 13 Achieved optimization objective and its bound as W varies

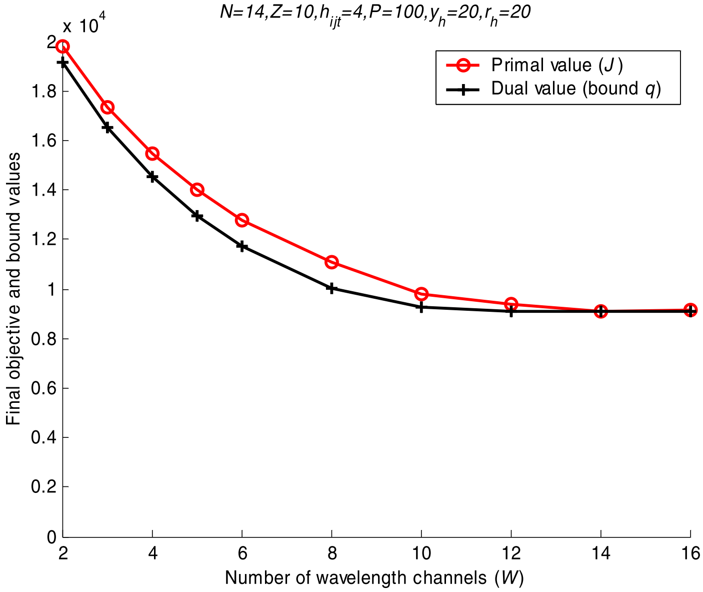

Fig. 14 Impact of h_{ijt} on the average hop count

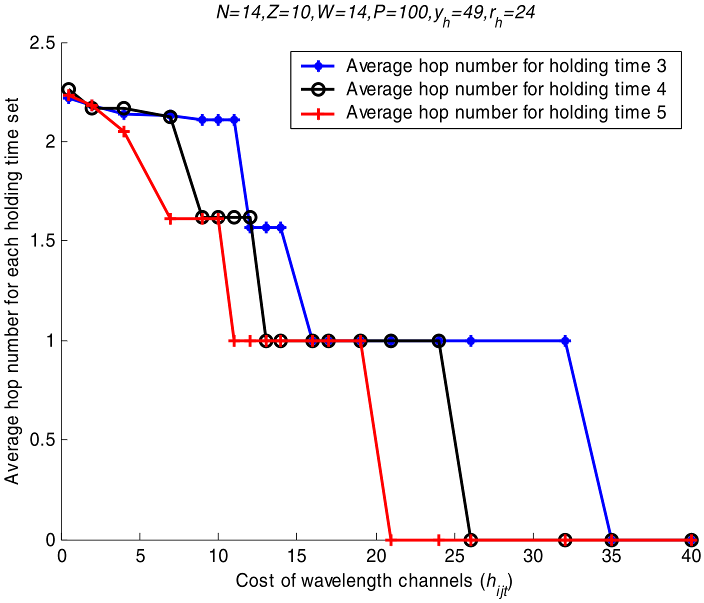

In our third example, we study the impact of the cost of a hop of a lightpath, which includes the cost of one adjacent node and a wavelength channel. We fix the other parameters and vary the cost of wavelength channels (denoted by h_{ijt}). The results are shown in Figs. 14 and 15. We grouped the SSLDs into three groups based on their holding times. We can see in Fig. 14 that as h_{ijt} increases from 0 to 40, the average hop counts of each group drop at different rates. In Fig. 15, we demonstrate that the performance of our scheduling results is mostly optimal for this study. The network operator can thus easily control the hop number of the routings by adjusting the h_{ijt} value. Please note that although for the simplicity of our numerical experiments, we only set the penalty for rejecting an SSLD (i.e., P) to the same value for all SSLDs, our model allows to set different values for individual SSLDs. The penalty for rejecting a given SSLD is an artificial value that operators of data centres are willing to pay for setting up the SSLD.

6 Conclusions

In this paper, we study the benefits and tradeoffs of using scheduled lightpaths to connect data centres for optimizing the use of intermittent renewable energy sources. The prior knowledge of server workload and traffic patterns, as well as

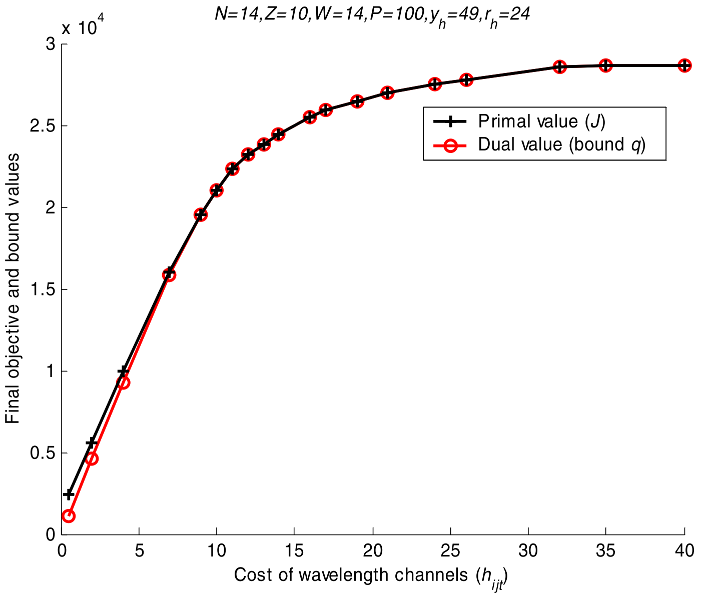

Fig. 15 Achieved optimization objective and its bound as h_{ijt} varies

the green energy sources availability provides an opportunity to schedule lightpaths to serve our goal of maximizing the use of green energy at the network planning stage. We proposed a WDM network planning model, which allows the operator of networks and data centres to set the scheduling and cost related parameters, making tradeoffs between competing operation goals.

We reveal the tradeoffs between the consumption of brown energy and the capability of providing required connectivity to data centres. Our simulation results show how timing flexibility improves network resource utilization and reduces rejections. We also study the impact of available network resources and the cost of network resources on the optimization objective. Our future work includes modeling lightpaths with variable holding times.

References

1. Zhang, Y., Chowdhury, P., Tornatore, M., & Mukherjee, B. (2010). Energy efficiency in telecom optical networks. *IEEE Communications Surveys and Tutorials*, *12*(4), 441–458.
2. Bolla, R., Bruschi, R., Davoli, F., & Cucchietti, F. (2011). Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures. *IEEE Communications Surveys & Tutorials*, *13*(2), 223–244.
3. Rahman, S. (2003). Green power: what is it and where can we find it? *IEEE Power & Energy Magazine*, *1*(1), 30–37.
4. Systems, Cisco (2006). Transport Options for Data Center Optical Interconnection. http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/prod_white_paper0900aecd803e884f.pdf.
5. Wen, J., Zheng, Y., & Donghan, F. (2009). A review on reliability assessment for wind power. *Renewable & Sustainable Energy Reviews*, *13*(9), 2485–2494.
6. Skea, J., Anderson, D., Green, T., Gross, R., Heptonstall, P., & Leach, M. (2008). Intermittent renewable generation and maintaining power system reliability. *IET Generation, Transmission & Distribution*, *2*(1), 82–89.
7. Karki, R., Hu, P., & Billinton, R. (2006). A simplified wind power generation model for reliability evaluation. *IEEE Transactions on Energy Conversion*, *21*(2), 533–540.
8. El-Tamaly, H. H., & Mohammed, A. A. E. (2006). Impact of interconnection photovoltaic/wind system with utility on their reliability using a fuzzy scheme. *Renewable Energy*, *31*(15), 2475–2491.
9. Georgilakis, P. S., & Katsigiannis, Y. A. (2009). Reliability and economic evaluation of small autonomous power systems containing only renewable energy sources. *Renewable Energy*, *34*(1), 65–70.
10. Watson, S. J., & Ter-Gazarian, A. G. (1996). The optimisation of renewable energy sources in an electrical power system by use of simulation and deterministic planning models. *International Transactions in Operational Research*, *3*(3–4), 255–269.
11. Lefurgy, C., Rajamani, K., Rawson, F., Felter, W., Kistler, M., & Keller, T. W. (2003). Energy management for commercial servers. *Computer*, *36*(12), 39–48.
12. Chen, G., He, W., Liu, J., Nath, S., Rigas, L., Xiao, L., & Zhao, F. (2008). Energy-aware server provisioning and load dispatching for connection-intensive internet services. In *5th USENIX symposium on networked systems design and implementation (NSDI 2008)* (pp. 337–350). San Francisco, California, USA, April 16–18, 2008.
13. Chase, J. S., Anderson, D. C., Thakar, P. N., Vahdat, A. M., & Doyle, R. P. (2001). Managing energy and server resources in hosting centers. *Operating Systems Review*, *35*(5), 103–116. Special issue “Proceedings of 18th ACM Symposium on Operating Systems Principles (SOSP 2001), October 21–24, 2001, Banff, Alberta, Canada”.
14. Zhu, Q., Chen, Z., Tan, L., Zhou, Y., Keeton, K., & Wilkes, J. (2005). Hibernator: helping disk arrays sleep through the winter. *Operating Systems Review*, *39*(5), 177–190. Special issue “Proceedings of 20th ACM Symposium on Operating Systems Principles (SOSP 2005), October 23–26, 2005, Brighton, UK”.
15. Gencata, A., & Mukherjee, B. (2003). Virtual-topology adaptation for WDM mesh networks under dynamic traffic. *IEEE/ACM Transactions on Networking*, *11*(2), 236–247.
16. Thompson, K., Miller, G. J., & Wilder, R. (1997). Wide-area Internet traffic patterns characteristics. *IEEE Network*, *11*(6), 10–23.
17. Kuri, J., Puech, N., Gagnaire, M., Dotaro, E., & Douville, R. (2003). Routing and wavelength assignment of scheduled light-

path demands. *IEEE Journal on Selected Areas in Communications*, *21*(8), 1231–1240.

18. Gagnaire, M., Koubaa, M., & Puech, N. (2007). Network dimensioning under scheduled and random lightpath demands in all-optical WDM networks. *IEEE Journal on Selected Areas in Communications*, *25*(9), 58–67.
19. Wang, B., & Li, T. (2010). Survivable scheduled service provisioning in WDM optical networks with iterative routing. *Optical Switching and Networking*, *7*(1), 28–38.
20. Skorin-Kapov, N. (2006). Heuristic algorithms for the routing and wavelength assignment of scheduled lightpath demands in optical networks. *IEEE Journal on Selected Areas in Communications*, *24*(8), 2–15.
21. Lee, S. S. W., Chen, A., & Yuang, M. C. (2010). A Lagrangian relaxation based near-optimal algorithm for advance lightpath reservation in WDM networks. *Photonic Network Communications*, *19*(1), 103–109.
22. Jaekel, A., & Chen, Y. (2009). Resource provisioning for survivable WDM networks under a sliding scheduled traffic model. *Optical Switching and Networking*, *6*(1), 44–54.
23. Wang, B., Li, T., Luo, X., Fan, Y., & Xin, C. (2005). On service provisioning under a scheduled traffic model in reconfigurable WDM optical networks. In *2nd international conference on broadband networks (BROADNETS 2005)* (Vol. 1, pp. 13–22). Boston, Massachusetts, USA, October 3–7, 2005.
24. Su, W., Sasaki, G., Su, C. F., & Balasubramanian, A. (2006). Scheduling of periodic connections with flexibility. *Optical Switching and Networking*, *3*(3–4), 158–172.
25. Nguyen, H. H., Gurusamy, M., & Zhou, L. (2009). Scheduling network and computing resources for sliding demands in optical grids. *IEEE/OSA Journal of Lightwave Technology*, *27*(12), 1827–1836.
26. Shen, L., Yang, X., Todimala, A., & Ramamurthy, B. (2007). A two-phase approach for dynamic lightpath scheduling in WDM optical networks. In *IEEE international conference on communications (ICC 2007)* (pp. 2412–2417). Glasgow, UK, June 24–28, 2007.
27. Tornatore, M., Ou, C., Zhang, J., Pattavina, A., & Mukherjee, B. (2005). PHOTO: an efficient shared-path-protection strategy based on connection-holding-time awareness. *IEEE/OSA Journal of Lightwave Technology*, *23*(10), 3138–3146.
28. Lucerna, D., Tornatore, M., Mukherjee, B., & Pattavina, A. (2009). Dynamic routing of connections with known duration in WDM networks. In *IEEE global telecommunications conference (GLOBECOM 2009)*, Honolulu, Hawaii, USA, November 30–December 4, 2009.
29. Environmental Protection Agency (2007). *EPA report to Congress on server and data center energy efficiency (August)*. http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.
30. Barroso, L. A., & Holzle, U. (2007). The case for energy-proportional computing. *Computer*, *40*(12), 33–37.
31. Tsai, C. H., Shin, K. G., Reumann, J., & Singhal, S. (2007). Online web cluster capacity estimation and its application to energy conservation. *IEEE Transactions on Parallel and Distributed Systems*, *18*(7), 932–945.
32. Pinheiro, E., Bianchini, R., Carrera, E. V., & Heath, T. (2001). Dynamic cluster reconfiguration for power and performance. In *2nd workshop on compilers and operating systems for low power (COLP 2001)* (pp. 75–93). Barcelona, Spain, September 9, 2001.
33. Bohrer, P., Elnozahy, E. N., Keller, T., Kistler, M., Lefurgy, C., McDowell, C., & Rajamony, R. (2002). The case for power management in web servers. In R. Graybill & R. Melhem (Eds.), *Power aware computing* (pp. 261–289). Norwell: Kluwer Academic.
34. Meisner, D., Gold, B. T., & Wenisch, T. F. (2009). PowerNap: eliminating server idle power. *ACM SIGPLAN Notices*, *44*(3), 205–216. Special issue "Proceeding of the 14th international conference on Architectural support for programming languages and operating systems (ASPLOS 2009), March 7–11, 2009, Washington, DC, USA".
35. Fan, X., Weber, W. D., & Barroso, L. A. (2007). Power provisioning for a warehouse-sized computer. *Computer Architecture News*, *35*(2), 13–23.
36. Nathuji, R., Somani, A., Schwan, K., & Joshi, Y. (2008). CoolIT: coordinating facility and it management for efficient datacenters. In *2008 conference on power aware computing and systems (HotPower 2008)* (p. 15). San Diego, California, USA, December 7, 2008.
37. Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., & Warfield, A. (2003). Xen and the art of virtualization. *Operating Systems Review*, *37*(5), 164–177. Special issue "Proceedings of 19th ACM symposium on operating systems principles (SOSP 2003), October 19–22, 2003, Bolton Landing (Lake George), New York, USA".
38. Clark, C., Fraser, K., Hand, S., Hansen, J. G., Jul, E., Limpach, C., Pratt, I., & Warfield, A. (2005). Live migration of virtual machines. In *2nd conference on symposium on networked systems design & implementation (NSDI 2005)* (Vol. 2, pp. 273–286). Boston, Massachusetts, USA, May 2–4, 2005.
39. Sugerman, J., Venkitachalam, G., & Lim, B. H. (2001). Virtualizing I/O devices on VMware workstation's hosted virtual machine monitor. In *USENIX annual technical conference (ATC 2001)* (pp. 1–14). Boston, Massachusetts, USA, June 25–30, 2001.
40. Nelson, M., Lim, B. H., & Hutchins, G. (2005). Fast transparent migration for virtual machines. In *USENIX annual technical conference (ATC 2005)* (p. 25). Anaheim, California, USA, April 10–15, 2005.
41. Figueiredo, R. J., Dinda, P. A., & Fortes, J. A. B. (2003). A case for grid computing on virtual machines. In *23rd IEEE international conference on distributed computing systems (ICDCS 2003)* (pp. 550–559). Providence, Rhode Island, USA, May 19–22, 2003.
42. Milojicic, D. S., Douglis, F., Paindaveine, Y., Wheeler, R., & Zhou, S. (2000). Process migration. *ACM Computing Surveys*, *32*(3), 241–299.
43. Bobroff, N., Kochut, A., & Beaty, K. (2007). Dynamic placement of virtual machines for managing SLA violations. In *10th IFIP/IEEE international symposium on integrated network management (IM 2007)* (pp. 119–128). Munich, Germany, May 21–25, 2007.
44. Nathuji, R., & Schwan, K. (2007). Virtual power coordinated power management in virtualized enterprise systems. *Operating Systems Review*, *41*(6), 265–278. Special issue "Proceedings of 21st ACM Symposium on Operating Systems Principles (SOSP 2007), October 14–17, 2007, Stevenson, Washington, USA".
45. Mahadevan, P., Sharma, P., Banerjee, S., & Ranganathan, P. (2009). Energy aware network operations. In *28th IEEE international conference on computer communications (INFOCOM 2009) workshops* (pp. 25–30). Rio de Janeiro, Brazil, April 19–25, 2009.
46. Bianchini, R., & Rajamony, R. (2004). Power and energy management for server systems. *Computer*, *37*(11), 68–76.
47. Heath, T., Diniz, B., Carrera, E. V., Meira, W. Jr., & Bianchini, R. (2005). Energy conservation in heterogeneous server clusters. In *10th ACM SIGPLAN symposium on principles and practice of parallel programming (PPoPP 2005)* (pp. 186–195). Chicago, Illinois, USA, June 15–17, 2005.
48. Weddle, C., Oldham, M., Qian, J., Wang, A. I. A., Reiher, P., & Kuenning, G. (2007). PARAID: a gear-shifting power-aware RAID. *ACM Transactions on Storage*, *3*(3), 13.
49. Qureshi, A. (2008). Plugging into energy market diversity. In *7th ACM workshop on hot topics in networks (HotNets 2008)*, Calgary, Canada, October 2008.

50. Lu, L., & Varman, P. (2008). Workload decomposition for power efficient storage systems. In *2008 conference on power aware computing and systems (HotPower 2008)* (p. 13). San Diego, California, USA, December 7, 2008.
51. Nathuji, R., Isci, C., Gorbatov, E., & Schwan, K. (2008). Providing platform heterogeneity-awareness for data center power management. *Cluster Computing*, *11*(3), 259–271. Special Issue on Autonomic Computing.
52. Moore, J., Chase, J., Ranganathan, P., & Sharma, R. (2005). Making scheduling "cool": temperature-aware workload placement in data centers. In *USENIX 2005 annual technical conference (ATC 2005)* (pp. 61–75). Anaheim, California, USA, April 10–15, 2005.
53. Verma, A., Ahuja, P., & Neogi, A. (2008). pMapper: power and migration cost aware application placement in virtualized systems. In V. Issarny & R. E. Schantz (Eds.), *Lecture notes in computer science (LNCS) Subseries: programming and software engineering: Vol. 5346. 9th ACM/IFIP/USENIX international conference on middleware (Middleware 2008)*, Leuven, Belgium, December 1–5, 2008 (pp. 243–264). New York: Springer.
54. Bash, C., & Forman, G. (2007). Cool job allocation: measuring the power savings of placing jobs at cooling-efficient locations in the data center. In *2007 USENIX annual technical conference (ATC 2007)* (pp. 363–368). Santa Clara, California, USA, June 17–22, 2007.
55. Burge, J., Ranganathan, P., & Wiener, J. L. (2007). Cost-aware scheduling for heterogeneous enterprise machines (CASH'EM). In *2007 IEEE international conference on cluster computing* (pp. 481–487). Austin, Texas, USA, September 17–20, 2007.
56. Beloglazov, A., & Buyya, R. (2010). Energy efficient resource management in virtualized cloud data centers. In *10th IEEE/ACM international conference on cluster, cloud and grid computing (CCGRID 2010)* (pp. 826–831). Melbourne, Australia, May 17–20, 2010.
57. Berral, J. L., Goiri, I., Nou, R., Julia, F., Guitart, J., Gavalda, R., & Torres, J. (2010). Towards energy-aware scheduling in data centers using machine learning. In *1st international conference on energy-efficient computing and networking (e-Energy 2010)* (pp. 215–224). Passau, Germany, April 13–15, 2010.
58. Kusic, D., Kephart, J. O., Hanson, J. E., Kandasamy, N., & Jiang, G. (2009). Power and performance management of virtualized computing environments via lookahead control. *Cluster Computing*, *12*(1), 1–15. Special Issue on Autonomic Computing.
59. Le, K., Bianchini, R., Martonosi, M., & Nguyen, T. D. (2009). Cost- and energy-aware load distribution across data centers. In *Workshop on power aware computing and systems (HotPower 2009)*. Big Sky, Montana, USA, October 10, 2009.
60. Bertini, L., Leite, J. C. B., & Moss, D. (2010). Power optimization for dynamic configuration in heterogeneous web server clusters. *The Journal of Systems and Software*, *83*(4), 585–598.
61. Rajamani, K., & Lefurgy, C. (2003). On evaluating request-distribution schemes for saving energy in server clusters. In *2003 IEEE international symposium on performance analysis of systems and software (ISPASS 2003)* (pp. 111–122). Austin, Texas, USA, March 6–8, 2003.
62. Rusu, C., Ferreira, A., Scordino, C., & Watson, A. (2006). Energy-efficient real-time heterogeneous server clusters. In *12th IEEE real-time and embedded technology and applications symposium (RTAS 2006)* (pp. 418–428). San Jose, California, USA, April 4–7, 2006.
63. Pakbaznia, E., Ghasemazar, M., & Pedram, M. (2010). Temperature-aware dynamic resource provisioning in a power-optimized datacenter. In *Conference on design, automation and test in Europe (DATE 2010)* (pp. 124–129). Dresden, Germany, March 8–12, 2010.
64. Pakbaznia, E., & Pedram, M. (2009). Minimizing data center cooling and server power costs. In *14th ACM/IEEE international symposium on low power electronics and design (ISLPED 2009)* (pp. 145–150). San Francisco, California, USA, August 19–21, 2009.
65. Nathuji, R., Schwan, K., Somani, A., & Joshi, Y. (2009). VPM tokens: virtual machine-aware power budgeting in datacenters. *Cluster Computing*, *12*(2), 189–203.
66. Chlamtac, I., & Ganz, A. (1993). Lightnets: topologies for high-speed optical networks. *IEEE/OSA Journal of Lightwave Technology*, *11*(5/6), 951–961.
67. Zhang, J. Y., Wu, J., Bochmann, G., & Savoie, M. (2009). A proof of wavelength conversion not improving the Lagrangian bound of the static RWA problem. *IEEE Communications Letters*, *13*(5), 345–347.
68. Zhang, J. Y., Mouftah, H., Wu, J., & Savoie, M. (2010). Lightpath scheduling and routing for traffic adaptation in WDM networks. *IEEE/OSA Journal of Optical Communications and Networking*, *2*(10), 803–819.
69. Jin, B., Luh, P. B., & Thakur, L. S. (1999). An effective optimization-based algorithm for job shop scheduling with fixed-size transfer lots. *Journal of Manufacturing Systems*, *18*(4), 284–300.
70. Zhang, Y., Yang, O., & Liu, H. (2004). A Lagrangian relaxation and subgradient framework for the routing and wavelength assignment problem in WDM networks. *IEEE Journal on Selected Areas in Communications*, *22*(9), 1752–1765.

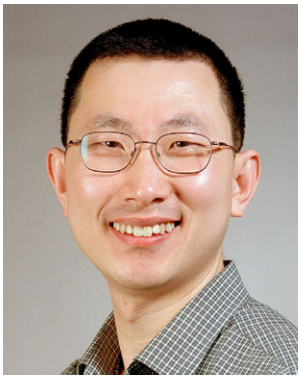

Jing Wu obtained a B.Sc. degree in information science and technology in 1992, and a Ph.D. degree in systems engineering in 1997, both from Xi'an Jiao Tong University, China. He is now a Research Scientist at the Communications Research Centre Canada (Ottawa, Canada), an Agency of Industry Canada. In the past, he worked at Beijing University of Posts and Telecommunications (Beijing, China) as a faculty member, Queen's University (Kingston, Canada) as a postdoctoral fellow, and Nortel Networks Corporate (Ottawa, Canada) as a system design engineer. Currently, he is also appointed as an Adjunct Professor at the University of Ottawa, School of Information Technology and Engineering. He has contributed over 70 conference and journal papers. He holds three patents on Internet congestion control, and two patents on control plane failure recovery. His research interests include control and management of optical networks, protocols and algorithms in networking, optical network performance evaluation and optimization. Dr. Wu is a co-chair of the sub-committee on "network architecture, management and applications" of the Asia Communications and Photonics Exhibit and Conference (ACP), 2009–2011. Dr. Wu is a Senior Member of IEEE.

James Yiming Zhang obtained his Ph.D. degree in electrical engineering from University of Ottawa, Ontario, Canada. He received both his B.S. and M.S. degrees in electronic science from Zhejiang University, China. He has done extensive research work in various fields including optimizations for the communication networks and the manufacturing processes, object detection in images and pattern recognition, machine learning, high-speed digital hardware design and biometrics. He is currently working on natural language processing in Idilia Inc., Montreal, Canada.

Michel Savoie is the research program manager for the Broadband Applications and Optical Networks group of the Broadband Network Technologies Research Branch at the Communications Research Centre Canada (CRC), an Agency of Industry Canada. He maintains expertise in broadband systems and related technologies such as: Infrastructure as a Service (IaaS), Green ICT, Application Oriented Networking (AON), advanced IP and WDM based Optical Networks in order to provide advice on important national initiatives and to demonstrate the application of CRC technologies in a real operational environment. He has managed two projects funded under CANARIE's directed research program on UCLP and was involved with EUCALYPTUS: A Service-oriented Participatory Design Studio project funded under the CANARIE Intelligent Infrastructure Program (CIIP), PHOSPHORUS: A Lambda User Controlled Infrastructure for European Research integrated project funded by the European Commission under the IST 6th Framework and the Health Services Virtual Organization (HSVO) project under CANARIE's Network-Enabled Plaftorms program. He is involved with the High Performance Digital Media Network (HPDMnet) research initiative and the GreenStar Network project under CANARIE's Green IT Pilot program. CRC and its partners are currently engaged in the further development of the IaaS Framework software. Mr. Savoie holds a B.Sc. and M.Sc. in Electrical Engineering from the Univ. of New Brunswick.