Assigning routes and wavelengths for collaboration over optical networks

Harzallah, Yosri; Horton, Joseph D.; Spencer, Bruce

Publisher's version / Version de l'éditeur:
https://doi.org/10.1007/s11107-010-0302-y
Photonic Network Communications, pp. 1-12, 2010-11-04

Photon Netw Commun
DOI 10.1007/s11107-010-0302-y

Assigning routes and wavelengths for collaboration over optical networks

Yosri Harzallah · Joseph D. Horton · Bruce Spencer

Received: 28 August 2009 / Accepted: 13 October 2010

Abstract The Routing and Wavelength Assignment problem is investigated in the context of collaboration where tools are shared and used simultaneously over a fiber optic network. Both online and offline versions are discussed, with and without using time as a parameter for scheduling purposes, and including the case where the network is used to carry time-multiplexed traffic. Also, the problem of rescheduling a blocked demand is studied. Several solutions based on Integer Linear Programs (ILP) and heuristics are proposed, implemented and their performance compared. The offline case is solved using two types of ILPs: link and path formulations. ILPs are also proposed for the online problem in addition to heuristic algorithms. While the link formulations give optimal solutions, they take a long time to solve and thus they can only be used for small problems. Path formulations and heuristics scale better but at the expense of optimality. The online approach is recommended when the resort to an offline approach is forbidden by the size of the problem.

Keywords Collaboration · Routing and wavelength assignment · Integer linear programming · Lightpath · Scheduling

Y. Harzallah (✉) · B. Spencer
Institute for Information Technology, National Research Council Canada, Fredericton, New Brunswick, Canada
e-mail: yharzallah@gmail.com

B. Spencer
e-mail: Bruce.Spencer@nrc.gc.ca

Y. Harzallah · J. D. Horton
Faculty of Computer Science, University of New Brunswick, Fredericton, New Brunswick, Canada

Y. Harzallah
École Nationale Supérieure des Mines de Saint Etienne, Saint Etienne, France

1 Introduction

The convergence of software and network technologies has paved the way to virtual organizations and distributed collaboration where scientists and engineers share access to world-class resources without regard to geography. Virtual Organizations (VO) like EarthScope, IPBIR (Integrated Primate Biomaterials and Information Resource), NESS (Network for Earthquake Engineering Simulation) [13] are a few examples that confirm this collaborative trend. The sophistication of the middleware tools, the emergence of the service-oriented architecture paradigm and the huge bandwidth provided by optical networks are all key enablers of a state-of-the-art collaborative cyberinfrastructure. With the wide adoption of fiber optic networking, supported by the Fiber-to-the-Home (FTTH) initiative [4,12], social networks could also share network-intensive applications such as high definition videoconferencing and simultaneous access to large data feeds, such as high definition graphics rendered in real time.

We are using these tools in the smaller scale of a VO called HSVO (Health Services Virtual Organization). A VO can be thought of as a social network where we assume participants work intensely together with some previous knowledge of each other and some common infrastructure. However, allowing people to share resources and work simultaneously together over computer networks raises management difficulties, mainly because of the distance between users and resources, as well as the concurrent use. This outlines the need for a management middleware capable of coordinating the collaboration and ensuring the availability of the resources including the bandwidth. In this paper, we assume participants are interconnected by a tree or mesh network, with lightpath connectivity available via Argia [5].

Published online: 04 November 2010

The National Research Council developed SAVOIR (Service Oriented Architecture for a Virtual Organization's Infrastructure and Resources) for the purpose of managing collaboration. It is a platform that provides a single point of entry to the Web-service enabled shared resources. Even the underlying optical network infrastructure is a resource that can be configured. Optical networks, while they are the answer to bandwidth greedy applications, are an expensive technology. With a solution like Argia, provided by Communications Research Canada and Inocybe Inc, a wavelength-routed network can be virtualized and leased to clients that can remotely configure its switches to create end-to-end connections (lightpaths). The physical topology is then transparent to the user and with remote calls to Web-services one can set up or teardown a private network.

SAVOIR currently does not have a scheduler to reserve resources and bandwidth. In a wavelength-routed optical network, reserving bandwidth means reserving a route and a wavelength. While this problem, commonly known as the Routing and Wavelength Assignment problem, has been the focus of many studies [3,14], less work has been done in the context of collaboration.

In this paper, the problem we address is the feasibility of the network architecture for moderate to large VO. We introduce the collaborative RWA (CRWA) problem, in the case of a non-multicast-enabled network. A specific definition of the traffic demand is introduced, and the availability of the non-network resources is taken into account so that bandwidth is not reserved and wasted between unavailable resources. Both online and offline versions of the problem are considered and compared in this paper. The presented solutions are based on Integer Linear Programming (ILP) formulations from [6] and heuristics.

The paper is organized as follows. In Sect. 2, we define the problem and we compare it to the related previous work in Sect. 3. Section 4 describes the proposed solutions, and Sect. 5 evaluates their performance. Conclusions are drawn in Sect. 6 along with some suggestions of future work.

2 Collaborative RWA problem

Scheduling bandwidth in a wavelength-routed optical network means assigning a route and a wavelength for every lightpath. However, it does not make sense to reserve bandwidth to connect resources that are not available for use, especially in the context of collaboration and virtual organizations where resources are usually under different ownership domains and thus their availability is uncertain. Taking into account, the availability of the non-network resources is then important to avoid wasting bandwidth and eventually blocking other collaborations from being satisfied.

Also, a distributed collaboration implies a hyperconnectivity between the set of users and the set of resources involved. However, when multicast is not available, this means that, at the network level, a collaboration request is actually a request to set up a set of end-to-end simultaneous connections. We define this set as a subsession, that has a start and end time, and define the collaboration session as a set of these subsessions. Indeed, we can imagine a scenario of a collaboration session where resources are brought only when needed throughout the session.

The collaborative RWA problem (CRWA) can then be defined as the problem of assigning a route and a wavelength for every connection of every subsession of session, while taking into account the availability of the resources and users involved in this connection. Solving this problem is crucial to an efficient use of the network.

The problem can have two versions: an offline version (or static) and an online (or dynamic) one. The offline case is when traffic demands are queued and then processed all together while in the online case the problem is solved for one request at a time as soon as it is issued.

Another interesting problem, that naturally follows this, is the problem of finding an alternative start date for a blocked demand. This problem is called the "When" problem.

In this paper, we study the problem with permanent and temporary traffic. Adding the time dimension allows non-overlapping sessions to use the same resources. The Time Division Multiplexing (TDM) capability, where a lightpath uses a certain number of TDM channels on the wavelength, is also studied.

3 CRWA and the classic RWA problem

Several schemes have been proposed to solve the classic RWA problem differing according to the assumptions that can be made on the problem. The RWA problem is known to be hard to solve. The problem is shown in [2] to be NP-hard in the offline case without wavelength converters. Several approaches to solve it are reviewed in [14]. Some of them attempt to tackle the problem as a whole while some break it into two subproblems (a routing problem and a wavelength assignment problem). Integer Linear Programming (ILP), meta-heuristics and heuristic algorithms have been used to solve the problem.

ILP formulations of the RWA problem are, however, NP-complete and thus cannot be used in practical cases with big networks or a big number of traffic demands. The alternatives are, then, heuristics and meta-heuristics such as simulated annealing [10], genetic Algorithms [1] or tabu-search [7]. These algorithms do not guarantee optimal solutions, but their performance can be evaluated with the ILP solution as a bound.

Other studies propose algorithms that handle traffic requests sequentially and assign a path and a wavelength at a time. These algorithms differ according to the scheme used to sort the requests and to define the assignment. Some of these possible schemes are reviewed in [14].

While the RWA problem has been a focus of many researchers, there are few papers studying it in a collaboration context. The few papers that do assume a multicast-enabled network. The communication is then 1–N: one source broadcasts to N destinations using multicast-capable switches that split a signal to N identical signals; or N–N where several sources deliver streams to all the sites. Dynamic multicast is still at its early stage since it is facing several design problems due to fabrication complexity and optical power loss when splitting the signal [15].

Finally, RWA can be seen as a special case of the introduced CRWA where the traffic demand has one subsession with one connection.

4 Proposed solutions

The optical network of interest is a WDM network with no wavelength conversion capabilities: the wavelength continuity constraint applies, a lightpath uses then the same wavelength on every link. The network is represented by an undirected graph where nodes are linked by single bidirectional fibers.

4.1 Notations

$\mathcal{V}$	Set of network nodes: $V_1, V_2, \ldots, V_{\lvert\mathcal{V}\rvert}$;
$\mathcal{E}$	Set of links $E_1, E_2, \ldots, E_{\lvert\mathcal{E}\rvert}$;
C_e	Bandwidth capacity of link E_e measured in number of OC-1 channels;
$G = (\mathcal{V}, \mathcal{E})$	Undirected graph;
$\mathcal{W}$	Set of wavelengths $\lambda_1, \lambda_2, \ldots, \lambda_{\lvert\mathcal{W}\rvert}$;
P_{sd}^{j}	jth element of the ordered set of paths $\mathcal{P}_{sd}$ between V_s and V_d;
L_{sd}^{ep}	Equals 1 if link $E_e \in P_{sd}^{p}$;
$\mathcal{R}$	Set of resources $R_0, R_1, \ldots, R_{\lvert\mathcal{R}\rvert}$;
σ_q	The maximum number of simultaneous connections a resource R_q can be part of;
$\mathcal{U}$	Set of users $U_0, U_1, \ldots, U_{\lvert\mathcal{U}\rvert}$;
$\mathcal{D}$	Set of sessions $D_1, D_2, \ldots, D_{\lvert\mathcal{D}\rvert}$;
w_i	Priority of sessions D_i;
D_{ij}	jth subsession of session D_i ;
$\theta_{ij}^{s}, \theta_{ij}^{e}$	Start and end time of subsession D_{ij};
Θ	Set, of size τ, of relevant time points:
Θ^{r}	Set, of size τ_r, of relevant time points for the online problem;
D_{ijk}	kth source–destination (s–d) connection of subsession D_{ij};
Ψ_{ijk}	⟨node,resource,user⟩ source triple for connection D_{ijk}: $\langle s_{ijk}, q, u\rangle$;
Δ_{ijk}	Destination triple for connection D_{ijk}: $\langle d_{ijk}, p, v\rangle$;
$\mathcal{R}(\Psi_{ijk})$	Resource of Ψ_{ijk};
$\mathcal{U}(\Psi_{ijk})$	User of Ψ_{ijk};
B_{ijk}	Bandwidth requirement for connection D_{ijk} in number of OC-1 channels.

4.2 Problem variables

In order to define the problem we introduce the following variables:

$\delta_i = 1$	if session D_i is accepted, 0 otherwise;
$\delta_{ij} = 1$	if subsession D_{ij} is accepted, 0 otherwise;
$\delta_{ijk} = 1$	if connection D_{ijk} is accepted, 0 otherwise;
$\alpha_{ijk}^{\lambda} = 1$	if wavelength λ is assigned to the s–d connection D_{ijk}, 0 otherwise;
$\gamma_{ijk}^{e\lambda} = 1$	if s–d connection D_{ijk} goes through the link E_e using λ, 0 otherwise;
$\eta_{ijk}^{n} = 1$	if s–d connection D_{ijk} goes through node V_n, 0 otherwise;
$\beta_{ijk}^{p\lambda} = 1$	if s–d connection D_{ijk} uses path $P_{s_{ijk}d_{ijk}}^{p}$ and wavelength λ, 0 otherwise.

4.3 Offline problem

Two ILP formulations are proposed for the offline problem: a link and a path formulation.

4.3.1 Link formulation

In the link formulation, the ILP defines the route by selecting each edge.

The objective is to maximize the number of sessions accepted while keeping the assigned routes short and using a minimum total number of wavelengths.

$$\max\left(10000\sum_{i} w_i\delta_i - 100\sum_{i,j,k,n}\eta_{ijk}^{n} - \sum_{i,j,k,\lambda}\lambda\alpha_{ijk}^{\lambda}\right) \quad (1)$$

In Eq. (1), the coefficients of the three terms are chosen to prioritize the different objectives.

The variables are subject to a set of constraints related to the network traffic and the availability of the resources:

– The wavelength continuity constraint implies that a lightpath enters an intermediate node using wavelength λ and leaves it using the same wavelength λ:

$$\sum_{\lambda \in \mathcal{W}} \sum_{e \in V(n)} \gamma_{ijk}^{e\lambda} = 2\eta_{ijk}^{n} \quad \text{where } n \in \mathcal{V} \setminus \{s_{ijk}, d_{ijk}\} \tag{2}$$

$$\sum_{\lambda \in \mathcal{W}} \sum_{e \in V(n)} \gamma_{ijk}^{e\lambda} = \eta_{ijk}^{n} \quad \text{where } n \in \{s_{ijk}, d_{ijk}\} \tag{3}$$

– An accepted connection is assigned a unique wavelength that is used on every assigned link:

$$\sum_{\lambda \in \mathcal{W}} \alpha_{ijk}^{\lambda} = \delta_{ijk} \tag{4}$$

– On each link a wavelength cannot be shared by two different connections:

$$\sum_{ijk} \gamma_{ijk}^{e\lambda} \leq 1 \tag{5}$$

– $\gamma_{ijk}^{e\lambda}$ has to be equal to zero except for the connections that use λ:

$$\gamma_{ijk}^{e\lambda} \leq \alpha_{ijk}^{\lambda} \tag{6}$$

– A subsession is accepted only if all its connections are accepted and a session is accepted only if all its subsessions are accepted. Later we explain why we do not use a unique variable instead of three.

$$\delta_{ijk} = \delta_{ij} = \delta_{i} \tag{7}$$

– D_{ijk}, if accepted, has to follow a path that goes from s_{ijk} to d_{ijk}:

$$2\delta_{ij} = \eta_{ijk}^{s_{ijk}} + \eta_{ijk}^{d_{ijk}} \tag{8}$$

$$2\gamma_{ijk}^{e\lambda} \leq \eta_{ijk}^{n} + \eta_{ijk}^{m} \text{ where } E_e = (n, m) \tag{9}$$

– A resource has a limited number of simultaneous connections:

$$\sum_{i,j,k} \delta_{ij} \leq \sigma_q \text{ where } \mathcal{R}(\Psi_{ijk}) = q \text{ or } \mathcal{R}(\Delta_{ijk}) = q \tag{10}$$

– Integrity constraints:

$$\delta_i, \delta_{ij}, \delta_{ijk}, \alpha_{ijk}^{\lambda}, \gamma_{ijk}^{e\lambda}, \eta_{ijk}^{n} \in \{0, 1\} \tag{11}$$

The time awareness can be added to this formulation by only changing constraints (5) and (10) as follows:

$$\forall t \in [1, \ldots, \tau - 1] : \sum_{i,j,k | D_{ij} \text{ is active in } [\Theta_t, \ldots, \Theta_{t+1}]} \gamma_{ijk}^{e\lambda} \leq 1 \tag{12}$$

$$\forall t \in [1, \ldots, \tau - 1] : \sum_{i,j,k | D_{ij} \text{ is active in } [\Theta_t, \ldots, \Theta_{t+1}]} \delta_{ij} \leq \sigma_q$$

$$\text{where } \mathcal{R}(\Psi_{ijk}) = q \text{ or } \mathcal{R}(\Delta_{ijk}) = q \tag{13}$$

The TDM support can also be added easily by altering constraint (5) as follows:

$$\sum_{ijk} B_{ijk} \gamma_{ijk}^{e\lambda} \leq C_e \tag{14}$$

4.3.2 Path formulation

In the path formulation, the ILP chooses the path from an ordered set of possible paths. The paths are generated using a k-shortest dissimilar paths algorithm based on the Dijkstra's shortest path algorithm combined with an iterative penalty scheme inspired from [9]. The objective of accepting the maximum number of sessions and the minimum number of wavelengths still stands. In addition, the ILP favors paths with a lower index in the set of paths.

$$\max \left(10000 \sum_{i} w_i \delta_i - 100 \sum_{p,\lambda} \sum_{i,j,k} p \beta_{ijk}^{p\lambda} - \sum_{\lambda} \sum_{i,j,k} \lambda \beta_{ijk}^{p\lambda} \right) \tag{15}$$

– An accepted connection is assigned one wavelength and one route:

$$\sum_{p} \sum_{\lambda} \beta_{ijk}^{p\lambda} = \delta_{ij} \tag{16}$$

– On each link a wavelength cannot be shared by two different connections:

$$\sum_{i,j,k} \sum_{p} L_{s_{ijk} d_{ijk}}^{ep} \beta_{ijk}^{p\lambda} \leq 1 \tag{17}$$

$$\delta_{ijk} = \delta_{ij} = \delta_{i} \tag{18}$$

$$\sum_{i,j,k} \delta_{ij} \leq \sigma_q \text{ where } \mathcal{R}(\Psi_{ijk}) = q \text{ or } \mathcal{R}(\Delta_{ijk}) = q \tag{19}$$

$$\delta_i, \delta_{ij}, \delta_{ijk}, \beta_{ijk}^{p\lambda} \in \{0, 1\} \tag{20}$$

Table 1 Description of ILPs

ILP	Link	Path	Time	TDM
ILP_1	√			
ILP_2		√		
ILP_3	√		√	
ILP_4		√	√	
ILP_5	√			√
ILP_6		√		√
ILP_7	√		√	√
ILP_8		√	√	√

Again the time dimension can be added by modifying Eq. (17) as follows while Eq. (19) is replaced as in the link formulation:

$$\forall t \in [1, \ldots, \tau - 1]: \sum_{i,j,k|D_{ij} \text{ is active in } [\Theta_t,\ldots,\Theta_{t+1}]} \sum_p L^{ep}_{s_{ijk}d_{ijk}} \beta^{p\lambda}_{ijk} \leq 1 \tag{21}$$

The TDM support is defined by changing Eq. (17) with the following constraint:

$$\sum_{ijk} \sum_p B_{ijk} L^{ep}_{s_{ijk}d_{ijk}} \beta^{p\lambda}_{ijk} \leq C_e \tag{22}$$

Time awareness and TDM support can be combined by merging the corresponding constraints.

Table 1 describes the different ILPs that are compared.

4.4 Online problem

The online problem is solved using ILPs and heuristics. In the following D_r represents the new request. The subscript $_r$ is used to refer to this request while the superscript $'$ refers to the previously accepted sessions.

4.4.1 ILP solution

The ILP solutions are based on the offline ILPs (ILP_7 and ILP_8) previously introduced to obtain ILP_9 and ILP_{10}. The link formulation has two similar objective functions. The second favors the use of congested wavelengths when there is a tie on the path length.

In these ILPs, the variables related to the sessions already scheduled are fixed and become input. The variables are then only related to the new requested session. For example, ILP_{10} is formulated as follows:

$$\max \left(10000 \sum_i w_i \delta_i - 100 \sum_{p,\lambda} \sum_{i,j,k} p\beta^{p\lambda}_{ijk} - \sum_\lambda \sum_{i,j,k} \lambda \beta^{p\lambda}_{ijk} \right) \tag{23}$$

subject to:

$$\sum_p \sum_\lambda \beta^{p\lambda}_{rjk} = \delta_{rj} \tag{24}$$

– For every wavelength, the bandwidth used by the new connections has to be less or equal to the available bandwidth on that wavelength:

$$\forall t \in [1, \ldots, \tau_r - 1]: \tag{25}$$

$$\sum_{jk|D_{ij} \text{ is active in } [\Theta^r_t,\ldots,\Theta^r_{t+1}]} \sum_p B_{rjk} L^{ep}_{s_{rjk}d_{rjk}} \beta^{p\lambda}_{rjk}$$

$$\leq C_e - \sum_{ijk|D'_{ij} \text{ is active in } [\Theta^r_t,\ldots,\Theta^r_{t+1}]} B'_{ijk} \gamma'^{e\lambda}_{ijk} \tag{26}$$

$$\delta_{rjk} = \delta_{rj} = 1 \tag{27}$$

$$\forall t \in [1, \ldots, \tau_r - 1]: \sum_{j,k|D_{rj} \text{ is active in } [\Theta^r_t,\ldots,\Theta^r_{t+1}]} \delta_{rj}$$

$$\leq \sigma_q - \sum_{i,j,k|D'_{ij} \text{ is active in } [\Theta^r_t,\ldots,\Theta^r_{t+1}]} \delta'_{ij}$$

$$\text{where } (\mathcal{R}(\Sigma_{ijk}) = q \text{ or } \mathcal{R}(\Delta_{ijk}) = q)$$

$$\text{and } (\mathcal{R}(\Sigma_{rjk}) = q \text{ or } \mathcal{R}(\Delta_{rjk}) = q) \tag{28}$$

$$\delta_{rj}, \delta_{rjk}, \beta^{p\lambda}_{rjk} \in \{0, 1\} \tag{29}$$

4.4.2 Heuristic solutions

The heuristic approach divides the problem into two subproblems: a resource availability problem and a routing and wavelength assignment problem.

The algorithm for the first problem is straightforward. Given the new session request and the scheduled sessions, the algorithm checks if for every resource the availability constraint (28) is satisfied.

Three heuristics are proposed to solve the second problem. The first one is called *Shortest of Longest First* (SLF). It sorts the requests in the decreasing length of the shortest path and then for every connection the shortest lightpath is chosen. If a connection is not the first in the list and is blocked, the algorithm restarts with that request dealt with first. If it is blocked again then the request is rejected. The algorithm is guaranteed to stop either at that step or when all the connections are accepted. The idea behind this algorithm is that it would be harder to satisfy longer connections. The second algorithm, *Most Congested First* (MCF), is based on a different idea. It is that using congested resources first would eventually leave other resources with enough bandwidth to satisfy greedy demands. The version of MCF where the congestion is ignored and the shortest path is chosen is called *Shortest of Shortest First* (SSF) and defines the third heuristic algorithm.

Algorithm 1 Shortest of Longest First (SLF)

Require: Request D_r, $G = (V, E)$, Θ^r
1: **for all** connection D_{rjk} **do**
2: Find the shortest path between s_{rjk} and d_{rjk}
3: **end for**
4: Order the requests using the decreasing length of the shortest path.
5: Create a copy G_λ of G for every wavelength λ.
6: **for all** connection D_{rjk} **do**
7: **for all** λ **do**
8: Find the shortest path in G_λ using a constrained Dijkstra's
9: **end for**
10: **if** no path is found and D_{rjk} is the first in the connections list **then**
11: the request is blocked
12: **return**
13: **else if** no path is found and D_{rjk} is not the first **then**
14: Cancel all the assignments
15: D_{rjk} becomes the head of the connections list
16: Restart at step 6
17: **else if** at least one path is found **then**
18: Select the shortest found path and update the corresponding G_λ
19: **end if**
20: **end for**
21: **for all** λ **do**
22: Commit changes to G_λ
23: **end for**
24: **return** Path and λ assignment for every D_{ijk}

Algorithm 2 Most Congested First Algorithm (MCF)

Require: Request D_r, $G = (V, E)$, Θ^r
1: Create a copy G_λ of G for every wavelength λ.
2: **repeat**
3: **for all** connection D_{rjk} **do**
4: **for all** λ **do**
5: find the shortest path between s_{rjk} and d_{rjk} that has available bandwidth in G_λ.
6: **end for**
7: Select the shortest one of them
8: **end for**
9: **if** this path cannot be found for at least one connection **then**
10: block request.
11: **end if**
12: Route the most congested one, if there is a tie take the shortest one.
13: Update G_λ
14: Remove routed connection from the set of connections
15: **until** All connections are dealt with

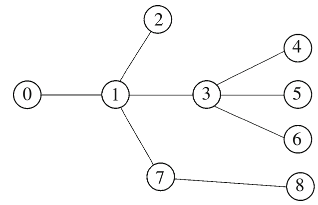

Fig. 1 HSVONET

4.5 "When" problem

The proposed solution for this problem is based on the ILP link formulation of the online problem. A request is rejected because some of its connections use unavailable resources. Relaxing the constraint, $\delta_{rjk} = \delta_{rj} = \delta_r$, binding the connections together, reveals what connections caused the scheduling to fail. Based on this information, one can decide on a candidate time offset by moving the latest start time of the blocked connections to the earliest end time of the already accepted and scheduled connections, which defines the *slideTimes* subroutine in Algorithm 3. Then, the ILP is solved again with these new dates. This is repeated until no connection is blocked.

Algorithm 3 When Algorithm

Require: Blocked Request D_r, Scheduled Sessions D'
1: **repeat**
2: **if** slideTimes(D_r, D') $\neq$ NULL **then**
3: Solve the relaxed ILP9 with D_r and D'
4: **else**
5: **return** Rejected request D_r
6: **end if**
7: **until** No connection is rejected
8: **return** Accepted request D_r

5 Experimental results

In this section, we present experimental results of our proposed solutions. Simulation experiments were conducted with three networks: a tree network (HSVONET) with 9 nodes and two mesh networks as shown in Fig. 1, NSFNET (14 nodes, 19 edges) [8] and EONNET (20 nodes, 38 edges) [11]. Links have 4 and 8 wavelengths. ILPs were solved with the solving framework SCIP and algorithms were implemented in Java. The tests were done on a 3.4 GHz Intel Xeon Dual-Core hyperthreaded quad processor with 32 GB RAM.

The sessions and their requirements were generated randomly. The number of subsessions and the number of connections per subsession are drawn randomly in $[1, \ldots, 3]$. Instances of the Offline problem have a random number of sessions in $[1, \ldots, 20]$ and their solving time limit is set to 4 h.

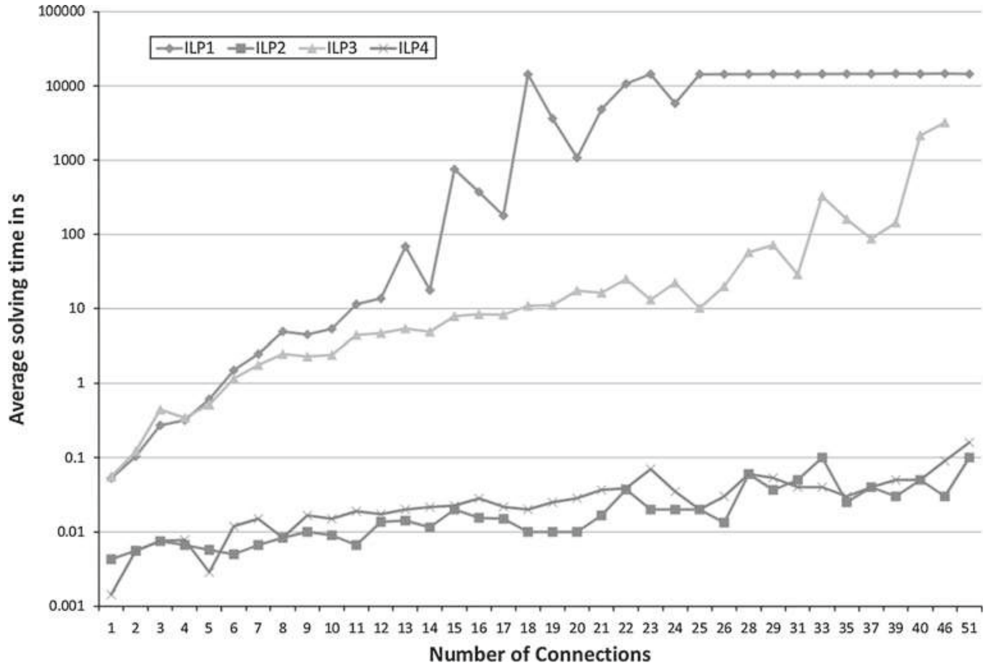

Fig. 2 Comparison of the average solving times of ILP_1, ILP_2, ILP_3 and ILP_4. (NSF, $\lambda = 4$)

5.1 Offline problem

Figure 2 plots, in a logarithmic scale, the average solving times for ILP_1 to ILP_4 according to the total number of connections in the problem. The number of connections represents the size of a problem better than sessions, since sessions can have a random number of connections.

When the size of the problem is small, the time values for the different ILPs are close to each other in the milliseconds range. However, when the problem's size increases, the solving time increases exponentially for the link formulations of the problem, ILP_1 and ILP_3. For the same problem, ILP_3 takes less time than ILP_1 to solve. ILP_1 starts reaching the limit of 4 h for some problems with a number of connections between 18 and 24. For problems with more than 24 connections, ILP_1 always failed, while ILP_3 succeeded in solving every single problem of the set in less than 4 h.

Figure 3 compares the average percentage of blocked connections for ILP_1 and ILP_2, using the NFSNET and 4 wavelengths. When ILP_1 does not fail to solve the problem in less than 4 h, its blocking percentage is never worse than ILP_2. The maximum difference is around 28%.

Figure 4 plots the number of accepted connections by ILP_2 according to the number of accepted connections by ILP_1. Only the solved problems are plotted. A linear regression is also plotted. The trend line shows how far ILP_2 solutions are from the optimal solutions. All the problems under the line have a gap of more than 18.77% with the optimal solution.

The huge difference of scale in the solving times between the link and path ILPs is obvious in Fig. 2. The exponential behavior of the solving time is more evident in the next figure. Figure 5 is a scatter diagram of the solving time according to the number of connections in the case of HSVONET with 4 wavelengths. An exponential trend line is drawn to confirm that. We notice also that the dispersion of the values obtained for the same number of connections increases with that number.

Figure 6 shows the solving time for all the path formulations, ILP_2(♦), ILP_4(■), ILP_6(▲), and ILP_8(×). Adding the TDM support (ILP_8) to the time-aware ILP (ILP_4) does not change the solving time significantly, while adding it to the non-time aware ILP (ILP_2) unexpectedly increases the average solving time.

5.2 Online problem

To compare the solutions to the online problem, we generate sessions randomly and call each of the online implementations (ILPs and algorithms) for processing them sequentially. Table 2 gathers the obtained results.

The solutions based on ILP_9 block fewer requests than ILP_{10} and the heuristic algorithms. The difference is more marked with NSFNET and EONET than with HSVONET. ILP_9 with objective 1 gives better results and blocks fewer sessions than ILP_9 objective 2 in two cases. There is a tie in three cases and in one case the second objective gives better results. Their average solving time is, however, almost

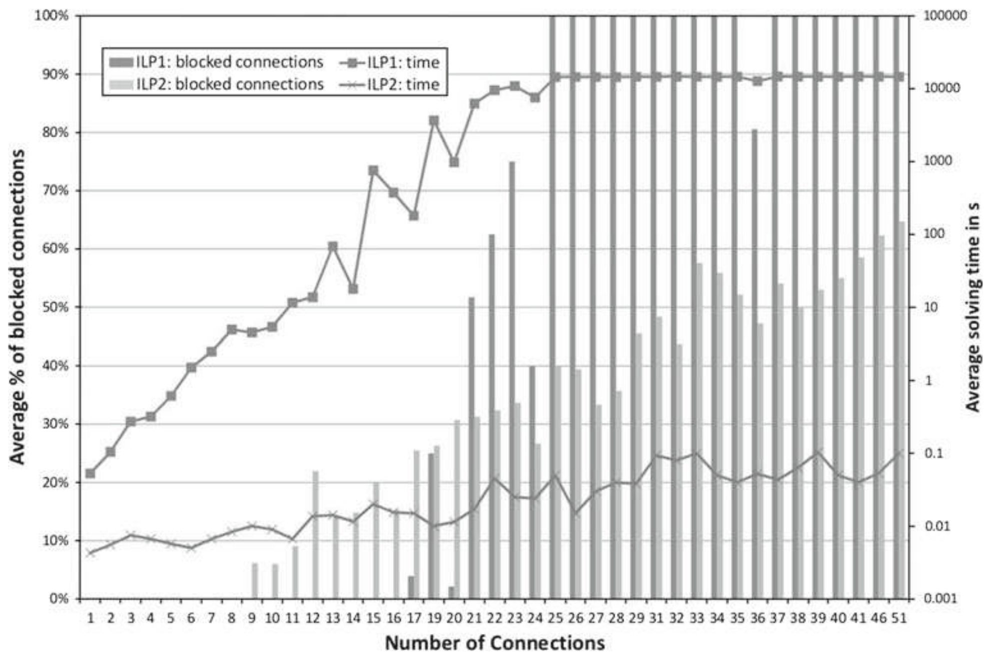

Fig. 3 Comparison of ILP_1 and ILP_2. (NSF, $\lambda = 4$)

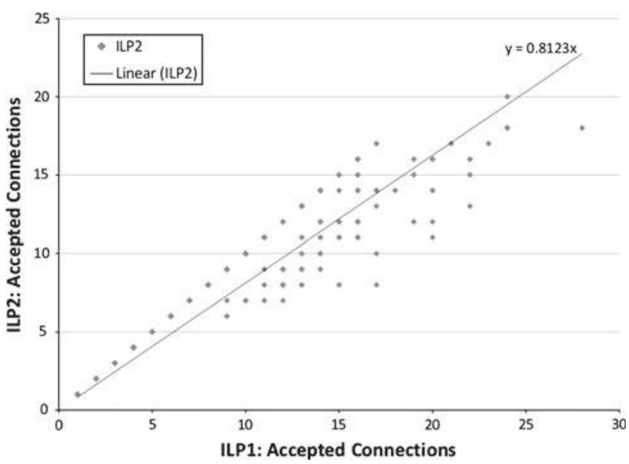

Fig. 4 Comparison of ILP_1 and ILP_2. (NSF, $\lambda = 4$)

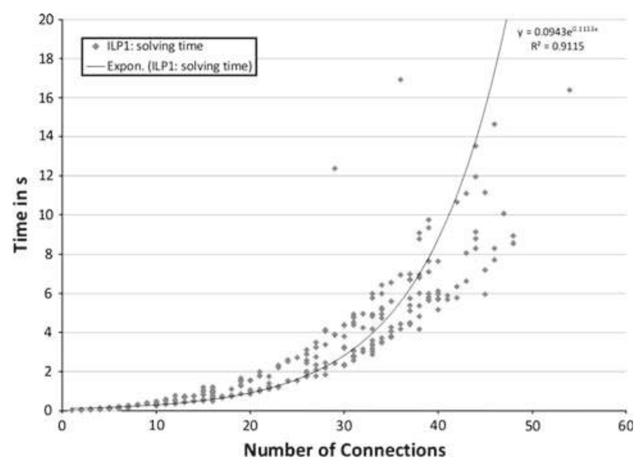

Fig. 5 Exponential solving time. ILP_1. (HSVO, $\lambda = 8$)

the same, since the only difference is the objective function. ILP_{10} has a much lower solving time but has a bad blocking behavior, even worse than the heuristic algorithms for non-tree networks.

By blocking fewer requests, SSF behaves better than SLF and MCF, and MCF better than SLF except for two cases: (NSFNET, 4) where SLF gives the best solution and (EONET, 8) where MCF is best. Among these three algorithms, SSF has the highest average solving time. ILP_{10} has the lowest average solving time among all the solutions as shown in Table 2.

5.3 Online versus offline

In order to compare the online approach to the offline one, we run a modified version of the online ILP_9 (objective 1) with a certain number of sessions and then we run the offline ILP_2 with the same set of sessions and compare the number of blocked sessions in each case. The version of ILP_9 that we use in this experiment does not take into account the time or the bandwidth in order to have more blocked requests with a low number of sessions. Indeed, when time and TDM are considered, the index of the first blocked connection using ILP_9 can be high as reported by Table 2. With a problem of a size higher than the value of that index, the offline

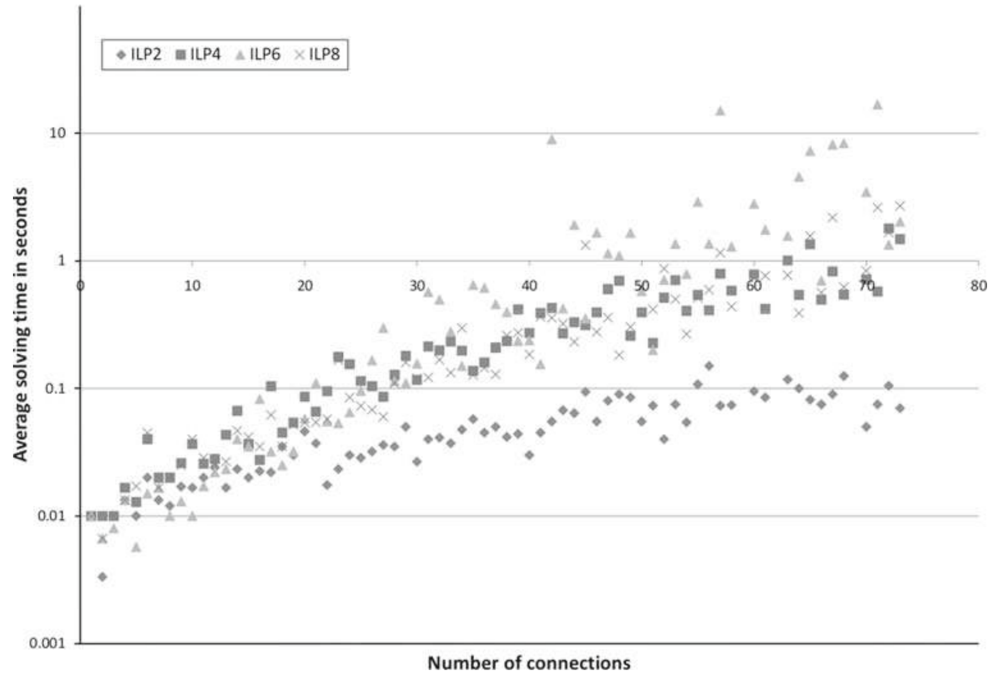

Fig. 6 Comparison of ILP_2, ILP_4, ILP_6, and ILP_8. Average solving time. (EON, $\lambda = 4$)

Table 2 Comparison of the online solutions

	λ	ILP_9 obj1		ILP_9 obj2		ILP_{10}		SLF		MCF		SSF	
		%	Avg. time (ms)	%	Avg. time (ms)	%	Avg. time (ms)	%	Avg. time (ms)	%	Avg. time (ms)	%	avg. time (ms)
NSFNET	4	6.0	281.18	6.0	275.00	19.0	7.92	14.0	17.81	14.5	29.45	14.2	30.44
NSFNET	8	2.0	976.46	3.2	937.56	13.0	11.27	12.7	28.29	11.7	47.86	11.2	49.83
HSVONET	4	58.1	31.92	56.9	32.12	57.1	5.0	67.6	4.09	62.8	5.86	62.1	5.98
HSVONET	8	31.9	118.15	32.2	117.95	32.4	9.45	48.6	10.4	46.4	15.66	42.6	16.15
EONET	4	1.7	1, 377.15	1.2	1, 331.98	14.7	9.22	6.0	37.26	5.5	65.45	5.5	66.94
EONET	8	0.2	4, 685.15	0.2	4, 707.93	14.7	12.67	4.7	54.68	4.5	95.20	5.2	95.25

Table 3 Comparison of the online and the offline (ILP_2) approaches

Network	λ	# sessions	Accepted sessions	
			Online	Offline
HSVONET	4	51	10	16
HSVONET	8	51	12	17
NSFNET	4	51	15	20
NSFNET	8	51	29	34
EONET	4	51	22	30
EONET	8	51	38	44

algorithms do not solve in a reasonable time. Table 3 summarizes the results.

Even though the path formulation is worse than the link one, it still gives better results than the online algorithm by accepting more sessions. However, SCIP had an inconsistent behavior in a few cases with EONET leading to high solving times, almost 21 h with 8 wavelengths.

5.4 The "when" problem

To evaluate the solution to this problem, we generate 200 random sessions sequentially and schedule them with ILP_9 (objective 1). If a session is blocked, we search for the earliest start time. We report on the number of times the Slide Times subroutine is called and the time it takes to find the solution for every blocked session. Figure 7 plots these parameters according to the index of the blocked session.

All the generated sessions are first requested to occur on the same day. The more sessions are accepted, the more computation it takes to find the earliest start time. HSVONET with only 4 wavelengths per link is used on purpose to have

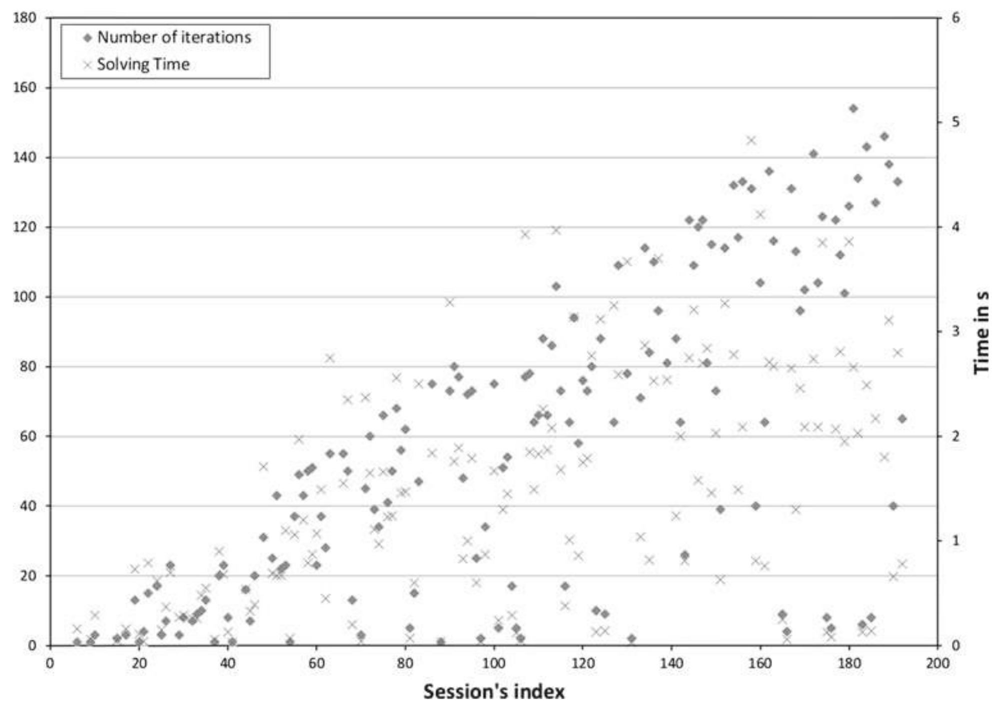

Fig. 7 Number of iterations and time to reschedule blocked sessions. (HSVO, λ = 4)

a worse case where resources are scarce and conflicts are frequent.

6 Conclusion and future work

We studied the routing and wavelength assignment problem in the context of collaboration. A new definition of a traffic demand is introduced, and the availability of the non-network resources used in the collaboration is taken into account to allow efficient advance reservation of all the involved resources. Solving this problem is crucial for an efficient collaborative use of the network and the resources. The solution gives the network enough bandwidth for different kind of social interaction. While this work is currently limited to interactions with participants that have end-to-end fiber connections, it can be relevant in the future to social networks, with the spread of the Fiber-to-the-Home. Also, this study is limited to cases where we have control over dedicated network resources which is enabled by the middleware Argia.

The solution to the offline version was based on Integer Linear Programming. Two types of formulations were developed: a link-based formulation and a path-based formulation. While path formulations lead to small size problems and thus take milliseconds to seconds to solve, they give mediocre blocking percentages comparing to the optimal link formulations. However, the link formulations have an exponential scaling behavior making them inefficient in practice even for tree networks.

It was expected that, for the same number of connections, time-aware link formulations would take more time to solve, but the opposite was observed. For the same number of connections, the solver seemed to take more time to give a solution when more sessions are blocked.

The proposed solutions to the online problem were of two types: ILPs and heuristic algorithms. The ILPs were based on the offline ILPs. A link-formulation ILP with two different objective functions was derived from ILP_7. The first objective enforces the use of shorter paths while the second prefers the use of congested paths. The obtained results do not permit a conclusion regarding the best objective function. However, it is clear that an online path formulation is not recommended for non-tree networks since its blocking probability is even worse than the three proposed heuristics.

Among the heuristics, SSF had unexpectedly the best overall results. It was expected that MCF would be better since preferring congested links would leave more links with more free bandwidth and thus would eventually accept more sessions with high bandwidth demands.

From the comparison between the online and the offline approach, we can assert that an online approach should be used only when an immediate answer to a request is needed or when the size of the problem forbids the resort to an offline approach.

In the case of the online problem, an algorithm is proposed to find an alternative start time for blocked sessions. This algorithm is based on the best online solution which is ILP_9. The main loop of the algorithm can have a high number of iterations especially if the time granularity is small. It is then recommended that a search interval be specified or the heuristic algorithm SSF be used instead of ILP_9.

The use of Integer Linear Programming limits the constraints on the resources to linear ones. While in this work we model the constraint on the number of simultaneous connections a resource can be part of, more and different constraints should be considered. Another limiting factor of ILPs that is relevant to the "When" problem is that, when a problem is infeasible, there is no way to identify precisely the cause. Knowing the cause would certainly reduce the number of iterations in this algorithm.

When developing these solutions, we assumed that we know all about the resources. However, when it comes to a practical use case, one has to consider the important fact that, in a VO, even though the network infrastructure can be controlled in a single domain, the resources connected to it might belong to several domains and thus there would not be a control over its availability. It is important that each resource provides a way to query its availability and usage requirements.

As further research, it would be interesting, in the online case, to study how the problem of rearranging the previously assigned sessions when a new session is blocked, especially if the blocked session has a higher priority than the conflicting ones. In addition, it would be interesting to look for other online heuristics or explore the area of meta-heuristics to improve the blocking probability. The path formulation could probably be improved by looking into other methods to generate the set of possible paths. It would also be interesting to compare suboptimal solutions of the link formulations with the solutions of the path formulations.

References

[1] Banerjee, N., Mehta, V., Pandey, S.: A genetic algorithm approach for solving the routing and wavelength assignment problem in WDM networks. In: 3rd IEE International Conference on Networking, ICN 2004, pp. 70–78. Feb–March (2004)

[2] Chlamtac, I., Ganz, A., Karmi, G.: Lightpath communications: An approach to high bandwidth optical WAN's. IEEE Trans. Commun. **40**(7), 1171–1182 (1992)

[3] Dutta, R., Rouskas, G.N.: A survey of virtual topology design algorithms for wavelength routed optical networks. Opt. Netw. **1**, 73–89 (2000)

[4] Fiber-to-the-Home Council. http://www.ftthcouncil.org

[5] Inocybe Inc. http://www.inocybe.ca

[6] Jaumard, B., Meyer, C., Thiongane, B., Yu, Xiao: ILP formulations and optimal solutions for the RWA problem. Global Telecommunications Conference, 2004. GLOBECOM '04. IEEE 3, 1918–1924 vol.3, Nov.-3 Dec. (2004)

[7] Kuri, J., Puech, N., Gagnaire, M., Dotaro, E., Douville, R.: Routing and wavelength assignment of scheduled lightpath demands. IEEE J. Sel. Areas Commun. **21**(8), 1231–1240 (2003)

[8] Li T., Wang B.: On optimal survivability design under a scheduled traffic model in wavelength-routed optical mesh networks. In: CNSR 06: Proceedings of the 4th Annual Communication Networks and Services Research Conference, pp. 26–33. IEEE Computer Society, Washington, DC, USA (2006)

[9] Lim, Y., Kim, H.: A shortest path algorithm for real road network based on path overlap. J. East. Asia Soc. Trans. Stud. **6**, 1426–1438 (2005)

[10] Mukherjee, B., Banerjee, D., Ramamurthy, S., Mukherjee, A.: Some principles for designing a wide-area WDM optical network. IEEE/ACM Trans. Netw. **4**(5), 684–696 (1996)

[11] O'Mahoney, M.J., Simeonidou, D., Yu, A., Zhou, J.: The design of a European optical network. J. Lightwave Technol. **13**(5), 817–828 (1995)

[12] Reardon Marguerite: Verizon bets big on network infrastructure. October (2008)

[13] The National Science Foundation Cyberinfrastructure Council: Cyberinfrastructure vision for 21st century discovery. March (2007)

[14] Zang, H., Jue, J., Mukherjee, B.: A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. Opt. Netw. Mag. **1**, 47–60 (2000)

[15] Zhou, Y., Poo, G.-S.: Optical multicast over wavelength-routed WDM networks: A survey. Opt. Switch. Netw. **2**(3), 176–197 (2005)

Author Biographies

Yosri Harzallah was born in 1984 and raised in Monastir, Tunisia. He studied for 2 years advanced mathematics and physics at Insitut Préparatoire aux Etudes Scientifiques et Techniques, La Marsa, Tunisia. He received in 2005 a Tunisian government scholarship to pursue a joint master degree in engineering and management at École des Mines de Sainte Etienne in France. He graduated in 2009. The same year he graduated with a Master of Computer Science degree from University of New Brunswick, Canada. Now he is a research associate at the National Research Council—Institute for Information Technology in Fredericton, Canada. His interests are in network scheduling, infrastructure as a service and service oriented architectures.

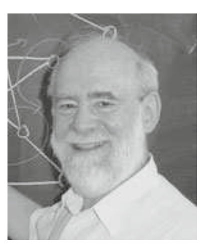

Joseph D. Horton was born and raised in Winnipeg, Canada. He received an honours mathematics degree at the University of Manitoba in 1968, followed by an MA in mathematics at York University the next year, and obtained a doctorate in Combinatorics and Optimization from the University of Waterloo in 1971. His area of research was combinatorial designs, particularly room designs. After 2 years of post doctoral work, he worked for the Regional Municipality of Sudbury, returning to academia in 1980 with Simon Fraser University. He moved to Computer Science at the University of New Brunswick in Fredericton, Canada the next year, where he is now a full professor. He was honored by IBM Canada as a Pioneer of Canadian Computing at the CASCON conference in Toronto in 2005.

Bruce Spencer earned a B.Sc. in Computer Science from Dalhousie University in Nova Scotia, Canada in 1980. He then earned a MMath in 1983 and a Ph.D. in 1991, both from University of Waterloo, Ontario, Canada. He was a Professor of Computer Science at the University of New Brunswick's Faculty of Computer Science from 1990 until 2001 and was a full professor with tenure. He is the founder of the Internet Logic Group at the Institute for Information Technology, National Research Council of Canada. The group's research focus is mainly on Semantic Web techniques, business rules, collaboration systems, procurement systems and reasoning systems for supporting business processes. Bruce is currently a senior research officer at the group and he held the position of group leader several times in his career.