DYNAMIC ROUTING ALGORITHMS IN TRANSPARENT OPTICAL NETWORKS

An Experimental Study Based on Real Data

Ralf Hülsermann
T-Systems Nova GmbH Technologiezentrum
Ralf.Huelsermann@t-systems.com

Monika Jäger
T-Systems Nova GmbH Technologiezentrum
Monika.Jaeger@t-systems.com

Sven O. Krumke[*,†]
Konrad-Zuse-Zentrum für Informationstechnik Berlin
krumke@zib.de

Diana Poensgen[‡]
Konrad-Zuse-Zentrum für Informationstechnik Berlin
poensgen@zib.de

Jörg Rambau[*]
Konrad-Zuse-Zentrum für Informationstechnik Berlin
rambau@zib.de

Andreas Tuchscherer[*]
Konrad-Zuse-Zentrum für Informationstechnik Berlin
tuchscherer@zib.de

[*] Supported by the DFG research center “Mathematics for key technologies” (FZT86) in Berlin
[†] Supported by the German Science Foundation DFG grant Gr 883/10
[‡] Supported by the Deutsches Forschungsnetz e.V. (DFN)

Abstract Today's telecommunication networks are configured statically. Whenever a connection is established, the customer has permanent access to it. However, it is observed that usually the connection is not used continuously. At this point, dynamic provisioning could increase the utilization of network resources. WDM based Optical Transport Networks (OTNs) will shortly allow for fast dynamic network reconfiguration. This enables optical broadband leased line services on demand. Since service requests competing for network resources may lead to service blocking, it is vital to use appropriate strategies for routing and wavelength assignment in transparent optical networks. We simulate the service blocking probabilities of various dynamic algorithms for this problem using a well-founded traffic model for two realistic networks. One of the algorithms using shortest path routings performs best on all instances. Surprisingly, the tie-breaking rule between equally short paths in different wavelengths decides between success or failure.

Keywords: Dynamic Network Configuration, Routing and Wavelength Allocation, Transparent Optical Networks, Blocking Probability, Simulation

1. Introduction

In recent years, backbone transport network structures and architectures have changed significantly. WDM systems are deployed extensively in today's transport networks. So far, they are only used in static point-to-point connections. However, the WDM technique, applied together with fast reconfigurable Optical Add Drop Multiplexers (OADMs) and Optical Cross Connects (OXCs), allows to establish wavelength-based optical connections very fast. Therefore, the optical layer topology based on optical connections may be reconfigured dynamically. As a result, the virtual optical connection topology provided to the higher layers (e.g., SDH) is not quasi-static anymore. Hence, networking functions like fast provisioning, resilience mechanisms and traffic engineering concepts may be adopted in the optical layer.

Today, broadband leased lines are installed on a long term basis. The observation that customers usually do not use a permanent connection continuously but are "online" only a fraction of the time motivates resource sharing over time. With the fast provisioning capability of optical connections in dynamic Optical Transport Networks (OTNs), broadband leased line services could be offered on demand. We expect an increased number of customers to use the broadband leased line services for much shorter time spans (e.g., hours). Hence, instead of assigning dedicated resources to each leased line customer permanently, the provider may efficiently multiplex many customer services onto a pool of OTN resources. Depending on the amount of resources and the number of connection requests, customers may experience service blocking. Acceptable blocking probabilities, as well as price and profit models, typically depend on the particular service class.

In this paper, we restrict ourselves to the special case of a single service class with uniform price and profit. Our experimental study assumes a transparent optical network without wavelength conversion. We use hypothetical traffic models on benchmark networks based on real-world data. The purpose of our case study is to evaluate how many customers may be served dynamically with a common resource pool at an acceptable blocking probability. The uniform price model already reveals substantial differences in the performance of algorithms. We evaluate variants of several algorithms based on greedy strategies from the literature and compare them to a reference algorithm based on maximum flow computations.

1.1 Related Work

Several static routing schemes have been proposed for defining routes and assigning wavelengths in transparent optical networks at optimizing network performance characteristics like throughput, delay, congestion etc. In the literature this problem is referred to as the Routing and Wavelength Assignment (RWA) problem. Much work has been performed on the optimal design of RWA strategies.

For instance, Yen and Lin investigate a near optimal design of lightpath RWA in purely optical WDM networks by formulating the RWA as a mixed Integer Linear Programming problem (Yen and Lin, 2001). Zang, Jue, and Mukherjee review various routing approaches and wavelength-assignment approaches proposed in the literature and compare them with characteristics of wavelength-converted networks (Zang et al., 2000).

Choi, Golmie, Lapeyrere, Mouveaux, and Su describe different greedy RWA algorithms for static optical networks and compare them in terms of their computational complexity (Choi et al., 2000). Späth and Bodamer evaluate photonic networks under dynamic traffic conditions (Späth and Bodamer, 1998a). They investigate different routing strategies with first-fit wavelength assignment on dynamic Poisson and non-Poisson traffic characteristics in dynamic WDM networks and the influence of limited wavelength conversion on the network performance (Späth and Bodamer, 1998b).

Späth discusses the impact of traffic behavior on the performance of dynamic WDM transport networks (Späth, 2002). Zang, Jue, Sahasrabuddhe, Ramamurthy, and Mukherjee review routing algorithms and investigate their suitability for implementation in systems with distributed control, thereby evaluating the effectiveness of different wavelength assignment strategies and analyzing the influence of convergence times (Zang et al., 2001).

Mokhtar and Azizoglu (Mokhtar and Azizoglu, 1998) propose to classify greedy algorithms by several wavelength ordering schemes. Note that the descriptions in their paper leave some tie-breaking decisions to the implementation. The algorithms in this work are derived from their classification by specification of tie-breaking rules and incorporating current availability of wavelengths in addition to current usage.

1.2 Our Contribution

We provide a comprehensive experimental study of various greedy algorithms for the problem of dynamic routing and wavelength assignment in a transparent optical network. The experiments are based on a well-founded traffic model which enables us to relate the blocking probability to offered traffic. Our results show that greedy algorithms have to be specified *unambiguously*: even changing just a tie-breaking rule may lead to significant changes in performance. Choosing an appropriate greedy algorithm yields blocking probabilities that are on par with a reference algorithm based on much more complicated techniques.

On our benchmark networks we could achieve an offered traffic of 55% (14-nodes network) and 30% (17-nodes network), respectively, at a blocking probability of 0.5%.

1.3 Paper Outline

In Section 1.2, we describe the models that we use for generation of problem instances and input data. In Section 1.3 the algorithms under consideration are explained. Section 1.4 presents the results of the simulation experiments. Section 1.5 is devoted to conclusions. Moreover, we have collected some additional technical information in an appendix.

2. Model

In this section we describe our models for static traffic load, network design and generation of dynamic traffic. We assume the networks to be bidirectional, i.e., every wavelength in a link can be used in either direction. Recall that the service of a connection request requires to establish a fixed lightpath between the corresponding nodes in the network (*circuit switched*). Furthermore, a lightpath uses the same wavelength on every link (*wavelength continuity constraint*). Moreover, on every WDM system on a link, each wavelength can only be used once (*wavelength conflict constraint*).

2.1 Static Traffic Model

Following (Späth, 2002), we assume that the demand for flexible leased line services is approximately proportional to traffic volumes in existing transport networks. Our traffic models are based on US-American and German population data. The population was partitioned into regions which lead to topologies with 17 nodes for Germany and 14 nodes for the US. In order to obtain estimated traffic demands between the regions, we differentiate between three types of traffic (cf. (Dwivedi and Wagner, 2000)): voice (V), transaction data (T, mainly business generated modem and IP traffic), and Internet traffic (I, IP traffic not related to a business environment).

According to (Dwivedi and Wagner, 2000), data traffic between two regions i and j may be estimated by a function depending on the following parameters:

- a constant K_τ depending only on the traffic type $\tau \in \{V, T, I\}$,
- the populations P_i and P_j,
- the numbers of non-production business employees E_i and E_j,
- the numbers of Internet hosts H_i and H_j,
- and the distance D_{ij} between the regions.

Using this notation, the traffics between regions i and j are computed as follows:

$$\begin{aligned} \text{Voice traffic} &= K_V \cdot P_i \cdot P_j / D_{ij} \\ \text{Transaction data traffic} &= K_T \cdot E_i \cdot E_j / \sqrt{D_{ij}} \\ \text{Internet traffic} &= K_I \cdot H_i \cdot H_j \end{aligned}$$

The total traffic between i and j is derived as the sum of these three values.

The values K_τ, $\tau \in \{K, V, I\}$ (for a particular year) are derived from the estimated traffic in a reference year and an estimation of traffic growth (V: 10%, T: 34%, I: 200% per year in our case for 2002). The resulting traffic matrices are displayed in the Appendix 1.5.1.

2.2 Topologies and Dimensioning

We use four different networks based on two topologies. For each of these, the *17-nodes topology* and the *14-nodes topology*, we construct two dimensionings. These will be referred to as the *shortest path dimensioning* and the *low cost dimensioning*, respectively. Both dimensionings are based on the corresponding static traffic matrix.

For the *shortest path dimensioning*, a routing of all static demands is computed using a standard shortest-path algorithm. Then, for each link l of the network, the number $p(l)$ of paths using l is counted, and l is equipped with wavelengths $\lambda_1, \ldots, \lambda_{p(l)}$. Notice that the shortest path dimensioning provides enough capacities for a shortest-path routing of all static demands, if the network is opaque (full wavelength conversion allowed). However, it might not allow for a valid routing of all static demands in the transparent case treated here.

The *low cost dimension* provides enough capacity for all static demands to be routed in the transparent network (not necessarily along shortest paths, though). Such a dimensioning can be computed, e.g., by the software tool described in (Koster et al., 2002).

The resulting four networks are shown in the appendix. Table 1 provides an overview of the total number of wavelength hops for each of them.

	shortest path dim.	low cost dim.
17-nodes topology	166	170
14-nodes topology	828	839

Table 1. Total number of bidirectional wavelength hops for each of the four networks.

2.3 Dynamic Traffic Model

The dynamic arrival of calls is modeled as follows. For each unit of static demand between two nodes u and v of the network, m sources generate connection requests for u and v according to a modified Poisson arrival process (inter arrival times are a constant plus an exponential distribution). The parameter m is called the *multiplex factor*.

Based on an observation of a network provider that a permanent connection is only used $1/12$ of the time we assume the following: every request has a holding time of 1 hour. On average, every source generates one request every 12 hours. This way, a multiplex factor of 1 models a dynamic traffic identical to the actual traffic incurred on permanent connections. Ideally, one could accommodate 12 requests with one hour holding time each within 12 hours. This multiplex factor of 12 corresponds to the traffic that would result from a continuous usage of a permanent line. Since this is the ideal resource utilization of the network we speak of 100% offered load.

Requests of a single source must not overlap: it is unlikely that a single customer requests a connection between two fixed nodes if he already has a valid connection between those nodes. This is modeled by randomly generating the inter arrival times between the requests of a given source according to the exponential distribution with mean 11 hours, then adding the constant holding time of 1 hour.

The multiplex factor serves to manipulate the strain which is put on the network: the higher the multiplex factor, the more requests between two fixed nodes are generated on average within the same time interval. A multiplex factor of 12 corresponds to 100% offered load (as if a permanent connection were used in fact permanently): for each unit of static demand, 12 requests with one hour holding time arrive within 12 hours.

3. Algorithms

In this section we describe the algorithms used in the experimental studies. For the presentation we assume that G is the underlying topology and that $\{\lambda_1, \ldots, \lambda_k\}$ is the collection of all available wavelengths. The collection of algorithms used in the experiments belongs to three classes: two classes of greedy-type algorithms, and one algorithm that is based on maximum flow computations in a capacitated network derived from G.

GENERIC-GREEDY

Input: Two nodes u and v between which a connection should be routed.

1 Let $\lambda_{i_1}, \ldots, \lambda_{i_k}$ be some order on the set of all wavelengths.
Note: The way *how* the order of the wavelengths is chosen leads to different versions of the algorithm, see text.

2 For a wavelength λ, let G_λ be the network restricted to all links where λ is currently still available.

3 Choose the first wavelength λ in the order where there is still a path in G_λ connecting u and v. If no such wavelength exists, reject the request.

4 Compute a shortest u-v-path in G_λ to route the connection (ties broken lexicographically w.r.t. node indices).

Algorithm 1: Generic greedy algorithm for routing connections.

The first six algorithms are based on the greedy-approach of Algorithm 1 and differ in the way how the order of the wavelengths in Step 1 is chosen. Usage and availability are computed w.r.t. the number of wavelength hops.

FIXED1 orders the wavelengths by increasing index.

FIXED2 orders the wavelengths by decreasing index.

SPREAD1 orders the wavelengths by increasing usage .

SPREAD2 orders the wavelengths by decreasing availability.

PACK1 orders the wavelengths by decreasing usage.

PACK2 orders the wavelengths by increasing availability.

Notice that in the empty network, the wavelength with the smallest index corresponds to that one with the highest availability, while the wavelength with the highest index is rarest. This is due to the way the wavelengths were assigned to the links when dimensioning the topology.

GENERIC-EXHAUSTIVE-GREEDY

Input: Two nodes u and v between which a connection should be routed.

1 Let $\lambda_{i_1}, \ldots, \lambda_{i_k}$ be some order on the set of all wavelengths.
Note: The above order is used as a tie-breaking rule which leads to different versions of the algorithm, see text.

2 For a wavelength λ, let G_λ be the network restricted to all links where λ is currently still available.

3 For each wavelength λ compute a shortest path from u to v in G_λ. If no path can be found at all, reject the request.

4 Among all paths choose a globally shortest one to route the connection, breaking ties by choosing the smallest wavelength w.r.t. the order (further ties broken lexicographically w.r.t. node indices).

Algorithm 2: Generic EXHAUSTIVE greedy algorithm for routing connections.

A second set of algorithms is based on the generic algorithm depicted as Algorithm 2. All of these algorithms compute shortest paths for *all* available wavelengths and then choose a globally shortest path. The algorithms differ in the way how Step 4 is implemented, that is, which tie-breaking rule is used in case that more than one globally shortest path is found.

EXHAUSTIVE1 orders the wavelengths by increasing index.

EXHAUSTIVE2 orders the wavelengths by decreasing index.

EXHAUSTIVE3 orders the wavelengths by increasing availability.

Notice that EXHAUSTIVE3 uses the same order on the wavelengths as PACK2.

Finally, as a benchmark, we implemented an algorithm which uses a mathematically more sophisticated cost function to decide which lightpath is the best routing choice, given the current network status. It is called *anticipating disjoint lightpath decrease* (ADLD) and defined as follows. For each request σ_j, let u_j and v_j be the end nodes to be connected and $t_j^{\mathtt{start}}$ and $t_j^{\mathtt{stop}}$ its start and stop time, respectively. Recall that $t_j^{\mathtt{stop}} = t_j^{\mathtt{start}} + 1$ in our setting.

ADLD computes for each available routing choice (P, λ) of request σ_j its cost by the formula

$$c(P, \lambda) := \sum_{s \neq t} d(s,t) \cdot f_j(s,t,\lambda) - \sum_{s \neq t} d(s,t) \cdot f_j^P(s,t,\lambda).$$

Here, $d(s,t)$ is the given static demand between nodes s and t. The values $f_j(s,t,\lambda)$ for each pair (s,t) of nodes are obtained by solving the instance of a *Maximum Flow Problem* (cf. (Ahuja et al., 1993)) defined by the graph G with source s, sink t, and edge capacities

$$\kappa_j(e) := \begin{cases} 1, & \text{if } \lambda \text{ is currently available on } e, \\ \max\{0, 1 - \frac{t^{\texttt{free}}(e,\lambda) - t_j^{\texttt{start}}}{t_j^{\texttt{stop}} - t_j^{\texttt{start}}}\}, & \text{if } \lambda \text{ is currently utilized on } e. \end{cases}$$

Here, $t^{\texttt{free}}(e, \lambda)$ denotes the earliest time at which all connections currently using wavelength λ on edge e will have expired.

Similarly, $f_j^P(s, t, \lambda)$ denotes the value of a maximum (s, t)-flow in G with edge capacities

$$\kappa_j^P(e) := \begin{cases} 0, & \text{if } e \in P, \\ \kappa_j(e), & \text{if } e \notin P. \end{cases}$$

Among all routing choices, a cheapest one with respect to the cost function defined above is selected. If multiple lightpaths incur the same cost, a shortest one is selected.

The idea of ADLD's cost function is to evaluate the decrease in potential future profit caused by the realization of the considered routing choice.

ADLD

Input: The static traffic demand $d(s, t)$ between all pairs of nodes (s, t), and a connection request $\sigma_j = (u_j, v_j, t_j^{\texttt{start}}, t_j^{\texttt{stop}})$, specifying two nodes and a start- and stopping-time.

1 If there is no available lightpath to connect u_j and v_j then reject the request.

2 Compute for each available routing choice (P, λ) its cost
$c(P, \lambda) = \sum_{s \neq t} d(s, t) \cdot f_j(s, t, \lambda) - \sum_{s \neq t} d(s, t) \cdot f_j^P(s, t, \lambda)$.
Here, $f_j(s, t, \lambda)$ and $f_j^P(s, t, \lambda)$ denote the maximum flow in the network G between s and t with edge capacities κ_j and κ_j^P, respectively (see the text for definitions of the capacities).

3 Route the connection on a path with minimum cost, breaking ties by selecting a shortest path with minimum cost.

Algorithm 3: Routing algorithm based on flow computation

4. Experimental Results

For each of the two considered topologies and for each multiplex factor $m = 1, \dots, 12$, we randomly generated a sequence of 20 batches. Each batch contained $10,000$ requests. In addition, one more batch of $10,000$ requests was generated as an onset for each sequence. We assume w.l.o.g. that the simulation starts at time 0 and ends at time T.

Figure 1 displays the resulting *offered traffic values* for both topologies in dependence of the multiplex factor. The shown values are derived as the average of the offered traffic value of all batches, where the traffic value V_o of a batch is computed as the sum of the holding times of all generated requests in the batch, divided by the total duration of the batch. As all requests have the same holding time of 1 hour in our setting, we obtain V_o as

$$V_o := \frac{10,000 \cdot 1h}{T}.$$

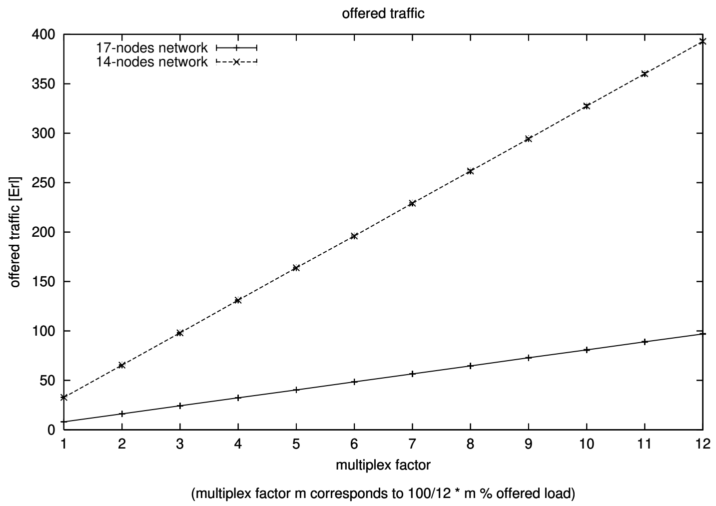

Figure 1. Offered traffic values.

4.1 Blocking Probabilities

For each of the four networks specified in Section 1.2.2, we compare the *blocking probabilities* achieved by the algorithms on the suitable sequence. The blocking probability of an algorithm on a given input sequence is the ratio of rejected and generated requests.

4.1.1 Notes on the diagrams. Since the differences in the performances of some algorithms are neglectable, we only display the results achieved by PACK1, SPREAD1, EXHAUSTIVE1, EXHAUSTIVE3, and ADLD. Among the algorithms whose results are not displayed, SPREAD2 and FIXED1 achieve blocking probabilities very similar to those of PACK1, whereas EXHAUSTIVE2 and EXHAUSTIVE3 perform almost equally. Also FIXED2 and PACK2 achieve nearly identic results, performing slightly better than SPREAD1 but worse than EXHAUSTIVE3.

Figures 2 to 5 display for each of the considered topologies the blocking probabilities achieved by the selected algorithms for different degrees of offered load. Recall that by construction, a multiplex factor of m corresponds to $100/12 \cdot m\%$ offered load.

We have added 95% confidence intervals, which we have computed from the blocking probabilities achieved on each of the 20 batches using a standard method (cf. (Law and Kelton, 2000)). The curve denoted by CONVERT corresponds to an algorithm that may convert the wavelength in each node (opaque routing) and routes a call on a shortest possible path. It serves as a crude estimation of the blocking probability that is unavoidable in the transparent case.

Blocking probabilities are plotted w.r.t. to a logarithmic scale, emphasizing small values. Blocking probabilities above 5% are not acceptable for the customer according to network providers. Blocking probabilities below 0.1% are subject to large relative counting deviations. However, these values are of no major interest because services requiring blocking probabilities smaller than 0.1% will very likely be realized as static connections. Therefore we display only blocking probabilities starting at 0.1%.

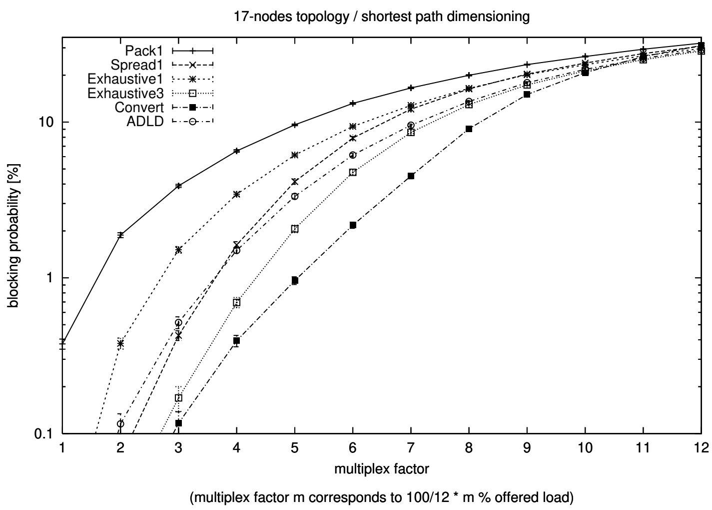

Figure 2. Blocking probability of the algorithms on the 17-nodes network with shortest-path dimensioning

4.1.2 Evaluation. EXHAUSTIVE3 yields the best results in all networks. For instance, at a blocking probability of 0.5%, EXHAUSTIVE3 can handle up to $\sim$ 30% offered load in the topology with 17 nodes and the shortest path dimensioning. Equivalently: it is able to cope almost with multiplex factor 4. In contrast, EXHAUSTIVE1 is only able to deal with multiplex factors up to 2, corresponding to $\sim 18\%$ offered load. In both dimensionings, EXHAUSTIVE3 allows a multiplexing factor of more than 4 at a blocking probability of 1%,

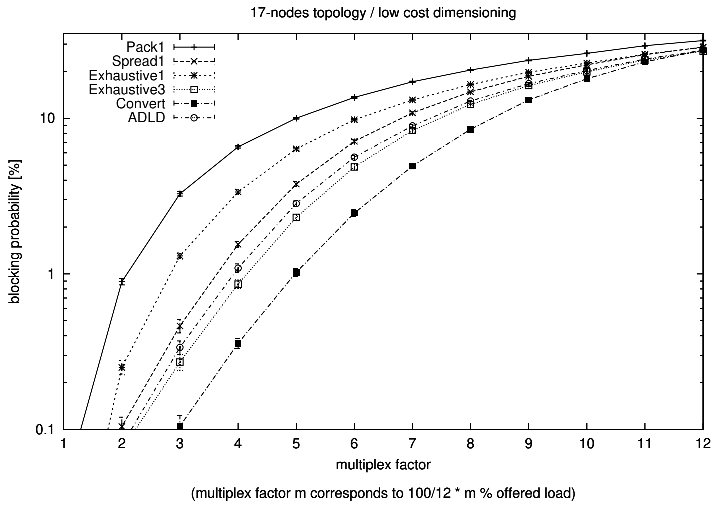

Figure 3. Blocking probability of the algorithms on the 17-nodes network with low cost dimensioning

i.e., more than 4 times the number of customers could be served dynamically compared to a permanent service provisioning with the same network dimensioning.

This remarkable difference between EXHAUSTIVE1 and EXHAUSTIVE3, which differ just in their tie-breaking rule, can also be observed between FIXED1 and FIXED2, SPREAD2 and SPREAD1, as well as between PACK1 and PACK2. The latter effects are plausible because increasing [decreasing] usage is not the same as decreasing [increasing] availability, whenever not all links have the same set of wavelengths.

An algorithm that cannot handle a multiplex factor of 2 at a blocking probability of at most 1% does not lead to any gain in dynamic configuration anymore. This is due to the fact that the profit for dynamic services has to be significantly lower than for permanent services. This applies, e.g., to algorithms PACK1 and—in the case of the 17-node network—EXHAUSTIVE1. Thus, EXHAUSTIVE3 is best whereas EXHAUSTIVE1 is unacceptable: a consequence of changing as little as a tie-breaking rule. We believe these differences to be significant, in particular considering the fact that most standard implementations use EXHAUSTIVE1 and not its counterparts.

Figures 4 and 5 show the simulation results on the 14-nodes network. As before, EXHAUSTIVE3 performs best among the greedy algorithms, while EXHAUSTIVE1 again shows significantly inferior performance. Only for high network load, ADLD achieves slightly smaller blocking probabilities than EXHAUSTIVE3.

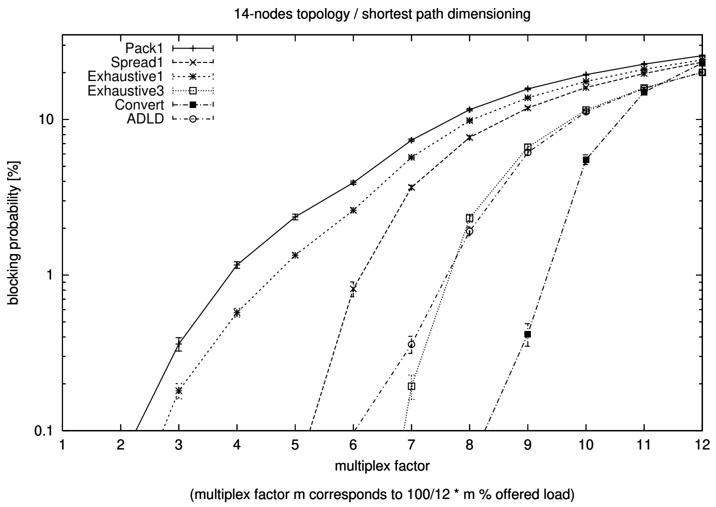

Figure 4. Blocking probability of the algorithms on the 14-nodes network with shortest-path dimensioning

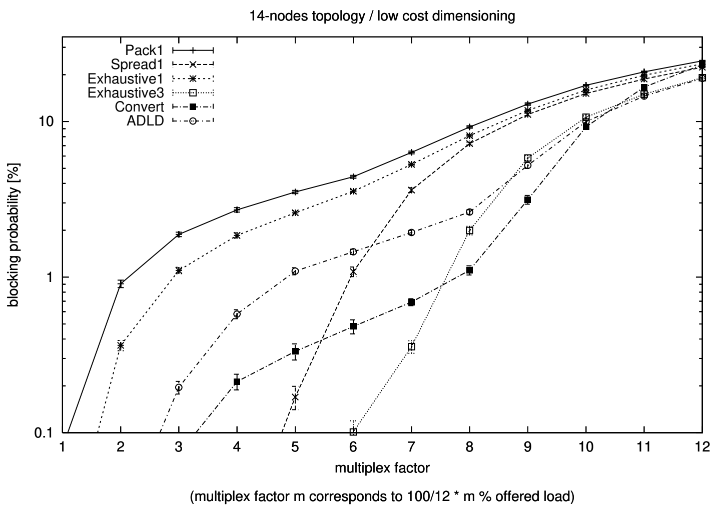

Figure 5. Blocking probability of the algorithms on the 14-nodes network with low cost dimensioning

Here, the routing algorithms have a stronger influence on the blocking probabilities, especially under low load. The range of the multiplex factor given at 1% blocking probability stretches from about 2 for PACK1 to more than 7 for EXHAUSTIVE3 in the low cost dimensioning.

In the 14-nodes network scenario with shortest path dimensioning the blocking probability of CONVERT is again slightly lower than the one of EXHAUSTIVE3. Almost always, full wavelength conversion (CONVERT) yields the smallest blocking probabilities. Rare deviations from this are due to the following: by utilizing conversion, CONVERT can accept some connections that must be rejected by the other algorithms. These, however, require very long paths, blocking the network for future requests.

EXHAUSTIVE3 is relatively close to CONVERT (roughly 20% fewer customers on average). For instance in the 17-nodes network with low cost dimensioning, EXHAUSTIVE3 reaches a multiplex factor of 4 at a blocking probability of 1%, whereas CONVERT permits a multiplex factor of 5. Since converters are expensive, transparent routing with EXHAUSTIVE3 is more cost effective than using full wavelength conversion.

We conclude that EXHAUSTIVE3 is superior to all other algorithms. It even outperforms the more complicated algorithm ADLD in most cases and comes relatively close to full conversion routing. The inferior performance of EXHAUSTIVE1 shows that it is a major issue how to break ties between shortest lightpaths.

4.2 Network Load

In this section, we provide plots that show how much traffic is actually routed in the networks by the algorithms under consideration.

Figures 6 to 9 display the traffic load incurred by the algorithms on each topology for various multiplex factors. The traffic load L of an algorithm measures the capacity utilization of the network. It is defined by

$$L := \frac{V_r \cdot D_{\mathtt{mh}}}{H}.$$

Here, V_r is the realized traffic value of ALG, computed as the sum of the holding times of all accepted connection requests, divided by the duration of the simulation, and $D_{\mathtt{mh}}$ is the mean hop distance of all implemented lightpaths. H is the total number of wavelength hops available in the given topology.

Algorithms with lower blocking probability accept more calls. Therefore, they incur more traffic load. Of course, full conversion allows for a higher degree of network utilization.

Figure 10 shows exemplarily for two algorithms and the 17-nodes network with the shortest path dimensioning the network load which would result if all the accepted requests were routed along a shortest path. The corresponding

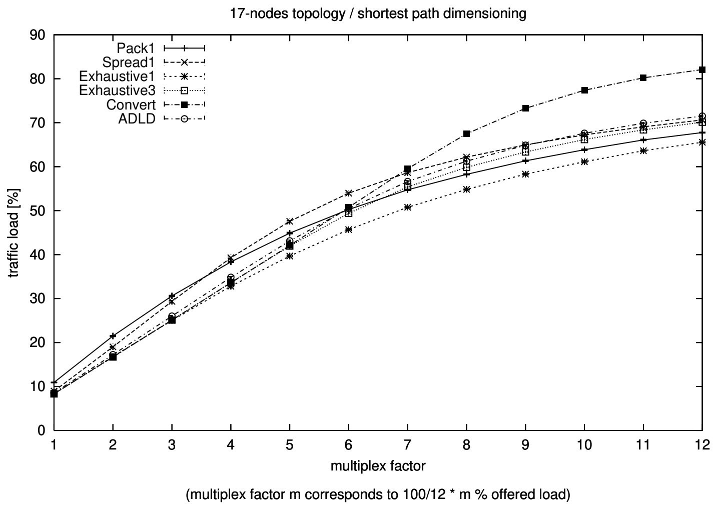

Figure 6. Traffic load of the algorithms on the 17-nodes network with shortest-path dimensioning

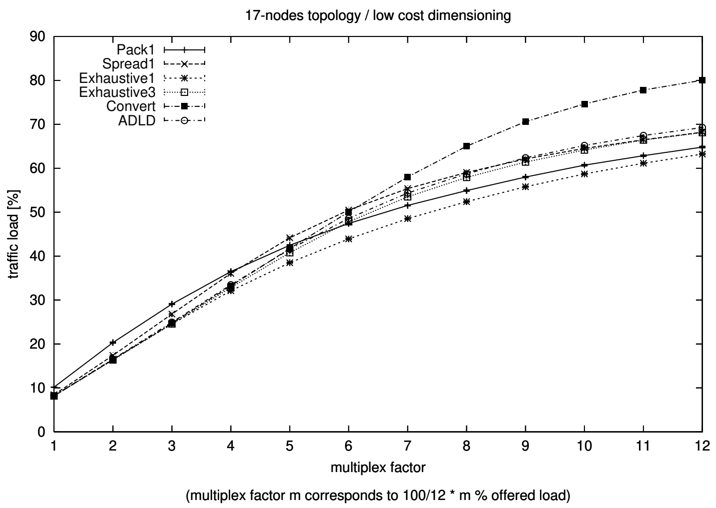

Figure 7. Traffic load of the algorithms on the 17-nodes network with low cost dimensioning

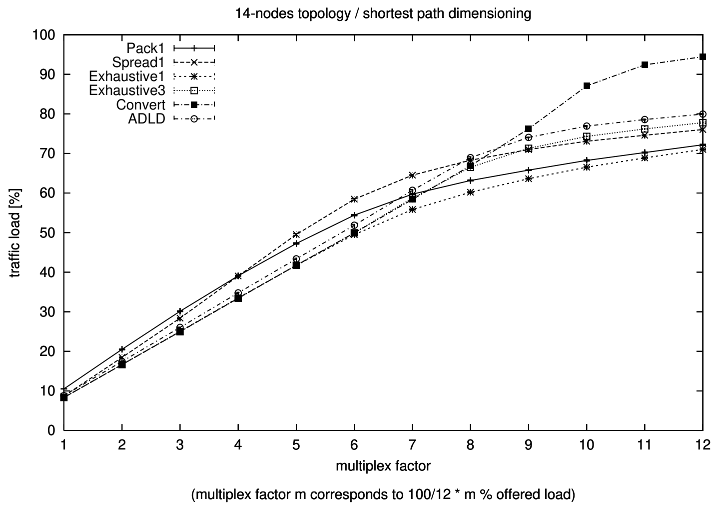

Figure 8. Traffic load of the algorithms on the 14-nodes network with shortest-path dimensioning

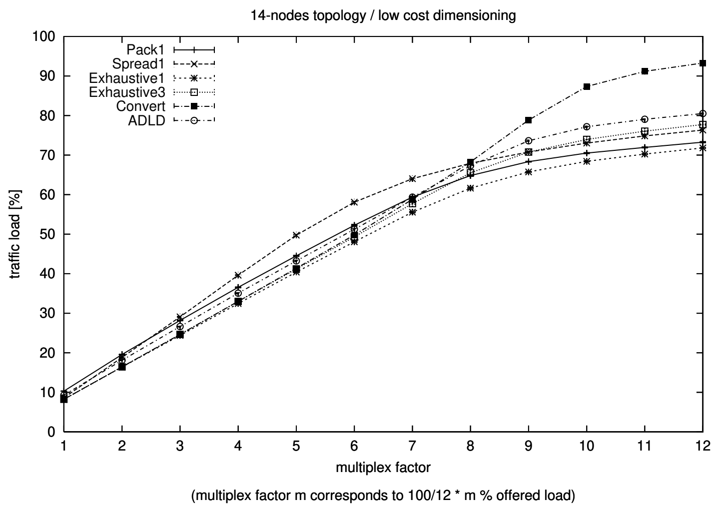

Figure 9. Traffic load of the algorithms on the 14-nodes network with low cost dimensioning

curve which belongs to the set of requests accepted by ALG is denoted by ALG(ideal). The gap between ALG and ALG(ideal) reveals the influence of the path lengths (in hops) on the traffic load.

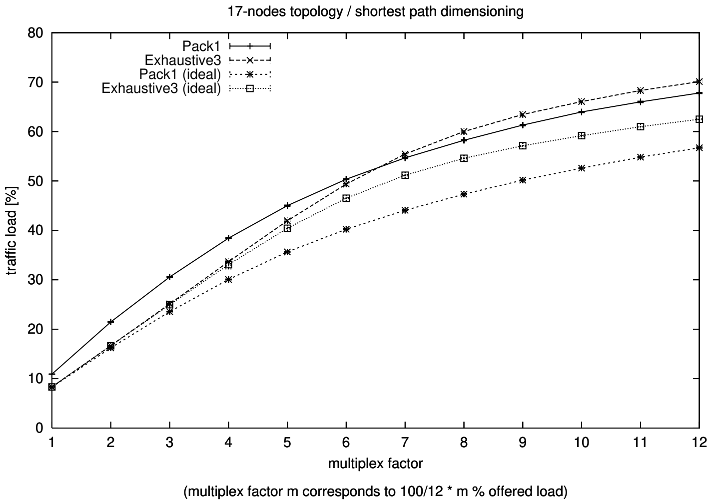

Figure 10. Ideal network loads for two algorithms on the 17-nodes network with shortest path dimensioning

5. Conclusions

We have simulated the behavior of various algorithms for the dynamic routing and wavelength assignment problem on realistic networks under a plausible traffic generation model. We have focused on greedy algorithms that choose a routing and wavelength assignment from a set of routing options by some (predefined or dynamically adapted) order on the wavelengths.

Most surprisingly, the order in which the selection of a wavelength is made substantially influences the performance—*even if only globally shortest paths are considered* (tie-breaking in EXHAUSTIVE): at an identical blocking probability of 0.5%, the best tie-breaking rule may achieve more than twice as much offered traffic as the worst. This results from the fact that there are usually shortest routing options available in many wavelengths, if the optical network is not overloaded. It is a clear warning to the planner not to leave the final decision about wavelength assignment between equally long routing options to chance. Our results suggest that a shortest routing in the currently least available wavelength should be preferred. This trend is apparent in all investigated greedy methods. It results from the fact that only few paths can still be routed in a wavelength with very low availability, and it therefore pays off to use make use of that 'refuse'.

The bright side: the best greedy algorithm compares favorably even with a more sophisticated algorithm (inspired by improved static planning methods). This is a clear indication that a carefully designed greedy algorithm is suitable for the task. Be aware, however, that this conclusion is based on test in a relatively homogeneous environment (unit prices/profits, arrival process with independent inter arrival times). Experiences in other areas of optimization show that the performance of greedy algorithms might severely degrade on very irregular problem instances. Since the development and evaluation of well-funded, more sophisticated algorithms is still in its infancy (no greedy algorithm can take into account varying bandwidth requirements, prices, or profits), further research is needed in this area.

References

Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. (1993). *Networks Flows*. Prentice Hall, Englewood Cliffs, New Jersey.

Choi, J. S., Golmie, N., Lapeyrere, F., Mouveaux, F., and Su, D. (2000). Classification of routing and wavelength assignment schemes in DWDM networks. In *Proceedings of the 7th International Conference on Optical Communications and Networks*, pages 1109–1115.

Dwivedi, A. and Wagner, R. (2000). Traffic model for a USA long-distance optical network. In *Proc. of the Optical Fiber Communication Conference*, volume 1, pages 156–158, Baltimore, USA.

Koster, A. M. C. A., Wessäly, R., and Zymolka, A. (2002). Transparent optical network design with sparse wavelength conversion. This conference.

Law, A. M. and Kelton, W. D. (2000). *Simulation modeling and analysis*. McGraw-Hill, Boston.

Mokhtar, A. and Azizoglu, M. (1998). Adaptive wavelength routing in all-optical networks. *IEEE/ACM Transactions on Networking*, 6(2):197–206.

Späth, J. (2002). Impact of traffic behaviour on the performance of dynamic WDM transport networks. In *Proceedings of the 28th European Conference on Optical Communication*.

Späth, J. and Bodamer, S. (1998a). Performance evaluation of photonic networks under dynamic traffic conditions. In *Proceedings of the 2nd Working-Conference on Optical Network Design and Modelling*.

Späth, J. and Bodamer, S. (1998b). Routing of dynamic Poisson and non-Poisson traffic in WDM networks with limited wavelength conversion. In *Proceedings of the 24th European Conference on Optical Communication*, pages 359–360.

Yen, H.-H. and Lin, F.-S. (2001). Near optimal design of lightpath routing and wavelength assignment in purely optical WDM networks. In *Proceedings of the 5th Working Conference on Optical Network Design and Modelling*.

Zang, H., Jue, J., and Mukherjee, B. (2000). A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. *Optical Networks Magazine*, 1(1):47–60.

Zang, H., Jue, J., Sahasrabuddhe, L., Ramamurthy, R., and Mukherjee, B. (2001). Dynamic lightpath establishment in wavelength-routed WDM networks. *IEEE Communications Magazine*, 39(9):100–108.

Appendix

5.1 Static Traffic Matrices

Node	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
2	1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
3	1	2	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
4	1	1	4	-	-	-	-	-	-	-	-	-	-	-	-	-	-
5	0	0	0	0	-	-	-	-	-	-	-	-	-	-	-	-	-
6	0	0	3	3	2	-	-	-	-	-	-	-	-	-	-	-	-
7	0	1	1	1	1	2	-	-	-	-	-	-	-	-	-	-	-
8	0	0	0	0	0	0	1	-	-	-	-	-	-	-	-	-	-
9	0	0	0	0	0	0	2	0	-	-	-	-	-	-	-	-	-
10	0	1	4	4	2	3	3	1	0	-	-	-	-	-	-	-	-
11	0	0	0	0	0	0	0	0	0	0	-	-	-	-	-	-	-
12	0	1	5	3	1	3	1	0	2	5	1	-	-	-	-	-	-
13	0	1	1	0	0	1	1	0	0	1	1	2	-	-	-	-	-
14	0	0	0	0	0	0	0	0	0	0	1	1	1	-	-	-	-
15	0	0	0	0	0	1	0	0	0	0	0	2	1	0	-	-	-
16	0	0	0	0	1	1	1	0	0	0	0	2	1	1	1	-	-
17	0	0	0	0	1	1	0	0	0	2	0	1	1	0	0	1	-

Table 2. Static demand matrix for the German 17-nodes network

Node	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	-	-	-	-	-	-	-	-	-	-	-	-	-	-
2	13	-	-	-	-	-	-	-	-	-	-	-	-	-
3	2	3	-	-	-	-	-	-	-	-	-	-	-	-
4	4	6	1	-	-	-	-	-	-	-	-	-	-	-
5	6	9	2	4	-	-	-	-	-	-	-	-	-	-
6	3	4	1	2	3	-	-	-	-	-	-	-	-	-
7	5	7	2	3	6	2	-	-	-	-	-	-	-	-
8	1	1	0	1	1	0	1	-	-	-	-	-	-	-
9	4	6	1	3	4	3	4	1	-	-	-	-	-	-
10	9	13	2	6	10	5	8	1	14	-	-	-	-	-
11	6	9	2	4	6	4	5	1	13	16	-	-	-	-
12	11	16	4	7	11	4	8	2	7	15	10	-	-	-
13	1	2	0	1	1	0	1	0	1	2	1	2	-	-
14	3	5	1	2	3	1	2	0	2	4	3	5	1	-

Table 3. Static demand matrix for the US 14-nodes network

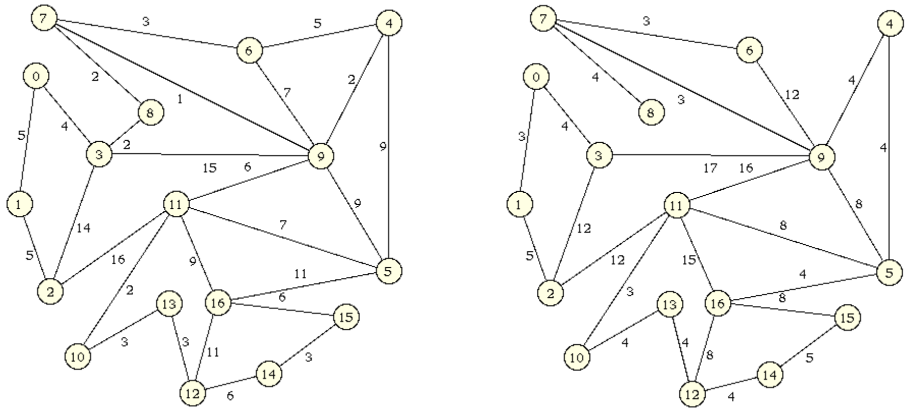

Figure 11. The 17-nodes network with shortest path and with low cost dimensioning

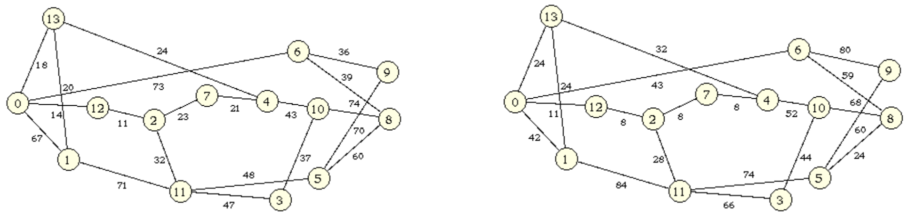

Figure 12. The 14-nodes network with shortest path and with low cost dimensioning