Dynamic Wavelength Routing in WDM Networks via Ant Colony Optimization

Ryan M. Garlick[1] and Richard S. Barr[2]

[1] CSE Dept., Southern Methodist University, Dallas, TX 75275 USA
[2] EMIS Dept., Southern Methodist University, Dallas, TX 75275 USA

Abstract. This study considers the routing and wavelength assignment problem (RWA) in optical wavelength-division multiplexed networks. The focus is dynamic traffic, in which the number of wavelengths per fiber is fixed. We attempt to minimize connection blocking using Ant Colony Optimization (ACO). The algorithm employed quantifies the importance of using length and congestion information in making routing decisions with the goal of reducing total network blocking. The ACO heuristic is shown to be a viable solution to the RWA problem. At high numbers of ants exploring the search space, the ACO algorithm achieves lower blocking rates than an exhaustive search over all available wavelengths for the shortest path.

1 Introduction

As customer demand for telecommunications services continues to grow exponentially, competitive forces have pressured providers to reduce costs and increase operational efficiencies. The current transmission and multiplexing standard for high-capacity U.S. networks is Synchronous Optical Network (SONET), developed by Bellcore in 1985. The primary shortcoming of SONET is the use of a combination of optical and electronic technologies, and the requisite time-consuming conversion from one to another. All-optical networks (AONs) avoid the electro-optical bottleneck by operating exclusively in the optical domain.

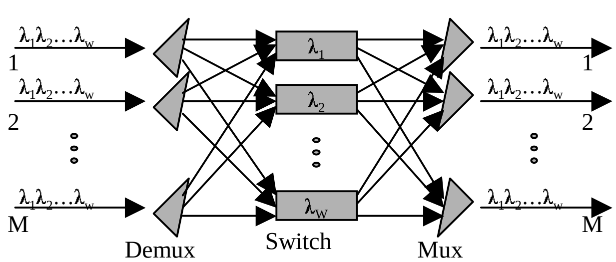

Fig. 1. A Wavelength Crossconnect

A wavelength-routing all-optical network consists of wavelength crossconnect (WXC) nodes interconnected by fiber links. A wavelength crossconnect is shown in Figure 1. This WXC has M incoming fiber links and M outgoing fiber links, with a capacity of W wavelengths (λs) on each link. The higher transmission capacity of all-optical networks is accomplished, in part, by sending multiple signals simultaneously through the same fiber-optic cable. This is achieved through wavelength-division multiplexing (WDM), which transmits multiple data streams simultaneously on different wavelengths. This approach can increase the transmission rate of current fiber spans by a factor of 32, and theoretically by a factor of 10,000, avoiding the expense of deploying new fiber spans [1]. Wavelength-division multiplexing is the key to dramatically increasing the capacity of existing optical networks.

1.1 Motivation

In wavelength-routed WDM networks, lightpaths are established between nodes and can span multiple fiber links in the network. A lightpath is realized by allocating a wavelength on each link comprising a path between two nodes. In the absence of wavelength-conversion equipment, a lightpath must occupy the same wavelength on all links of the path used, a property known as the wavelength continuity constraint. With wavelength converters, a lightpath may change wavelengths, significantly reducing blocking since there need only be an available wavelength on every link in the path. Without converters, *the same* wavelength must be free on every such link.

Once a network is designed, the routing and wavelength-assignment problem (RWA) must be considered. In the static RWA problem, all lightpath requests are known in advance, with the objective of minimizing network resources in order to satisfy all demands.

The focus of this study is the dynamic RWA problem, in which lightpath requests arrive dynamically, and the number of wavelengths is limited. The objective then becomes to minimize connection blocking. The optimal RWA problem is NP-Complete [3], and is thus suited to heuristic methods.

Although the RWA problem has been studied extensively, this paper introduces a new algorithm for evaluating potential routes based on length and congestion information.

1.2 Contribution of the Paper

Previous work in dynamic routing on all-optical WDM networks has progressed from simple shortest path routing to k-shortest path routing, to incorporating network congestion information. Any dynamic routing approach must adhere to the constraints of any other online algorithm, namely the lack of knowledge concerning future connection requests.

Ant-Colony Optimization (ACO) as a method for routing and wavelength assignment on all-optical networks is introduced, and provides valuable insight on the length versus congestion tradeoffs. Ant-Colony Optimization is a technique designed to mimic the ability of a colony of ants to find the shortest path to a food source and route around obstacles. ACO has been applied to the static RWA problem, and is modified to the dynamic RWA problem in this paper. ACO is used to test the hypothesis that occasionally choosing slightly longer paths with less congestion improves blocking performance. ACO provides an effective testing platform for investigating the efficacy of unconstrained dynamic routing, by using ants that prefer paths with lower levels of network traffic.

This study aims to decrease blocked requests through quantifying the importance of using congestion information. When confronted with a shorter path carrying more traffic or a slightly longer path with less congestion, the question of how much additional path length is acceptable to avoid congestion is examined.

1.3 Organization of the Paper

The remainder of the paper is organized as follows: Section 2 examines the existing literature in depth. Existing studies of routing, wavelength selection, and methods of incorporating congestion information are examined. Section 3 introduces the inclusion of path congestion information in fully adaptive routing. The algorithmic implementation of incorporating congestion information into routing and wavelength selection decisions through the Ant-Colony optimization heuristic is introduced. Section 4 presents preliminary results and analysis. Conclusions and future work are discussed in Section 5.

2 Background and Previous Research

It is important to distinguish several terms used in the literature, sometimes in an inconsistent manner. The terms "static" and "dynamic" (or "adaptive") can be applied to the RWA problem and the routing algorithm independently. In this study, we use the following definitions (adapted from [4]), which are common but not universally found in the literature.

The *static RWA problem* involves routing and assigning wavelengths to an entire set of connection requests. All connection requests are known in advance, and wavelengths are assigned to different lightpaths in a manner that attempts to minimize the total number of wavelengths used. The *dynamic RWA problem* is an online problem in which connection requests arrive one at a time, and must be routed as they arrive. The dynamic problem typically has a fixed number of wavelengths per network link, with the goal of routing as many connection requests as possible with the fewest number of blocked connections.

In contrast to referring to the entire problem set, a *static routing algorithm* is one in which the routing procedure does not vary with network state. A connection is blocked if there is no wavelength available on the designated path, even if a different path for the connection exists. Stern and Bala [4] further classify static routing algorithms as fixed if only one path is available on which to route a connection.

Dynamic or adaptive routing algorithms use network state information at the time of connection establishment [8]. Alternate routing refers to routing algorithms in which each connection is assigned a set of admissible paths. The selection of paths is typically made by using network state information, making alternate routing a common complement to dynamic routing strategies.

Unconstrained routing is a subset of dynamic routing in which the set of available choices on which to route the connection request is the set of all available paths from source to destination (barring routes blocked due to traffic.) The reader is referred to [2] for an extensive survey of existing methods for routing and wavelength assignment.

2.1 Static Routing and Wavelength-Assignment Problem

In [5], ACO is applied to the static routing and wavelength assignment problem and provides a background for the ACO approach to the dynamic RWA presented in this paper. In each algorithmic time step, ants move from each source to each destination. Varela introduced backtracking, with each ant keeping a "tabu" list[3] of previously visited nodes. This allows backtracking to avoid dead-ends and cycles – an approach adopted in the ACO algorithms in this paper. When an ant finds itself blocked, it pops its current location from a list of visited nodes and attempts to proceed from the previous location. This ability requires each ant's memory to contain a list of nodes visited in order.

Each ant in [5] maintains its own type of pheromone, and while ants are attracted to their own pheromone, they are repulsed by the pheromone of other ants in order to obtain even loading. The best results are achieved through a global update wherein ants are increasingly repulsed by paths on which more ants have traversed. Maintaining an ant and pheromone type for each connection request is time consuming, however.

Static RWA approaches can provide important insights into the problem and effective routing and wavelength assignment methods. The focus of this paper, however, is the dynamic routing and wavelength assignment problem, as discussed in the next section.

2.2 The Dynamic Routing and Wavelength Assignment Problem

All work in this paper focuses on extensions of work done on the dynamic RWA problem. The previous literature on the topic is extensive, with most studies focusing on a k-shortest-path-based routing scheme [2, 4]. More recent developments have incorporated congestion information into routing decisions [6, 9]. In the discussions in this section, both static and dynamic *routing algorithms* are presented in order to obtain a broad perspective of performance measures.

[3] See Glover and Laguna[7] for an explanation of Tabu Search and related topics

Static Routing The dynamic RWA problem may incorporate static routing. A commonly used static routing strategy is shortest-path routing (SP). The shortest path between each node is statically computed. If this shortest path is blocked, the connection blocks.

Much of the existing literature uses Dijkstra's algorithm [13] to route the demand along the shortest path from source to destination [3]. Static routing is simple and most implementations move the cost of computing routes to an initial set-up step. Due to significant reductions in blocking over static routing, however, most recent research has focused on dynamic routing.

Dynamic Routing Adaptive or dynamic routing algorithms typically provide better performance than static routing [2]. However, they have the added expense of requiring a control node with knowledge of all established connections in order to determine whether specific paths have available capacity. Congestion information must also be maintained to be incorporated into the routing scheme.

Previous research tested the effectiveness of incorporating congestion information by testing a k-shortest-path (ksp) algorithm against a single-shortest-path strategy [8]. This work showed that a *ksp* algorithm will outperform a strategy in which only a single path is available. The single-shortest-path strategy is easy to implement, but leads to higher blocking, particularly under heavy traffic. If the single path is already occupied, the connection will block.

Chan and Yum [9] compare two routing strategies: routing connection requests on the shortest path with available capacity and on the least-loaded route from source to destination. Since the latter paths may be significantly longer, the shortest-routing strategy almost always provided lower blocking. A goal of this paper is to combine these two strategies into a routing algorithm that chooses short paths with low congestion, in an effort to improve performance over each strategy individually.

Typically, a standard shortest path algorithm is used to find routes in the network. In the static routing algorithm, a routing table with one or more paths is computed ahead of time. Dynamic routing algorithms use network state information to determine the feasibility of k shortest paths for routing a connection request.

More recent studies have included congestion information. In its simplest form, a connection request is routed on the path with the least congestion. In [14], the least-loaded routing (LLR) algorithm is proposed and compared to another similar routing scheme using congestion information termed minimum sum routing (MSR). Both algorithms find the k shortest paths and use congestion information in determining the route to select.

Least-loaded routing measures the congestion on each link in each of k shortest paths. The minimum number of wavelengths available on any link on the path is the measure of congestion for the path. The least-loaded path is selected to route the connection request.

Minimum-sum routing employs a metric that reflects the utilization of each link (l) in each path (P). For each of k shortest paths, the metric is calculated as:

$$\sum_{l \in P} \frac{l_c - l_a}{l_c}$$

where l_c is the link capacity and l_a is the number of available wavelengths. The path that minimizes this value is selected as the path to route the connection request.

Test cases of both algorithms produced results supporting the motivations in this paper. With LLR routing, 81.4% of connections are routed over the shortest path, while 91.4% of connections were routed over the shortest path with MSR (at $k = 7$). However, with $k > 1$, and despite choosing slightly longer paths more often, the LLR algorithm significantly outperforms the MSR algorithm. The LLR blocking also improved at a faster rate than MSR as k increased.

The conclusions from these studies is that although selecting the shortest path plays the major role in reducing blocking, incorporating a degree of congestion information will improve blocking percentages even further.

3 Incorporating Congestion Information into Dynamic Routing

While the use of congestion information in online routing has been studied previously, all previous methods route over the least congested path among k shortest paths from source to destination. None of the existing methods quantify the relative weight of congestion information in making a decision to route over a slightly longer path with less congestion.

While scoring paths based on length has been incorporated in almost all previous dynamic routing work, scoring multiple paths based on congestion and length is introduced here. As shown in Section 4, and in [14], the use of congestion information has improved overall blocking percentages across several different routing strategies. It is a goal of this paper to determine the importance of congestion information in relation to the shortest-path model.

3.1 Ant-Colony Optimization

Another approach to the RWA problem incorporates Ant-Colony Optimization, a relatively new category of meta-heuristics. The ACO algorithm developed by Dorigo, et al. [11] mimics the abilities of a colony of ants to find the shortest path to a food source and to re-route that path to circumnavigate obstacles placed in the path. Ants in nature accomplish these relatively complex tasks by depositing small amounts of pheromone wherever they travel. Subsequent ants are attracted to this pheromone and follow the path. If an obstacle is placed in the path and the ants are faced with two alternatives to circumnavigate the obstacle, an equal number of ants will initially follow the longer and shorter paths. The longer path has a greater distance between each ant, and thus gets less pheromone per unit distance. This makes the shorter path more attractive, and as more ants travel the shorter path, it is reinforced with more pheromone.

When an ant completes a route from source to destination, a global update adds pheromone to the entire path, often in inverse proportion to the length of the path, to reinforce shorter routes. ACO is well-suited to path-following, and has been tested on many problems, most extensively the Traveling Salesman Problem

[12].

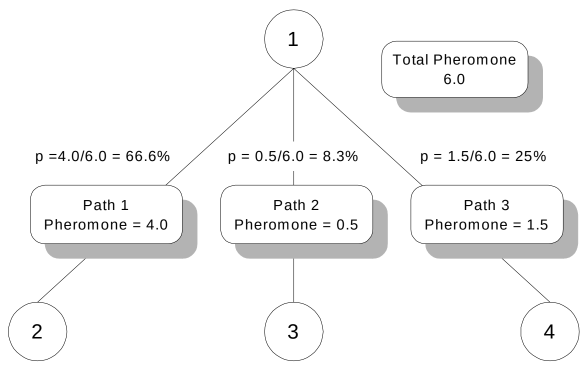

Fig. 2. Example path selection probabilities for an Ant at Node 1

Ant-Colony Optimization is modified here to an unconstrained routing algorithm. Under unconstrained routing, the set of paths comprising potential routes for the connection requests is the set of all paths between a given origin-destination node pair. ACO is an attractive solution for computationally intensive applications because it is inherently parallel, and adding more ant agents (or "ants") generally increases the solution quality. Hence, ACO has the ability to easily improve performance when additional processing units are available. ACO is easily distributed [15], with each processor assigned a different group of ants to execute. Pheromone values on each path and the best results are kept in a scoreboard in one memory space. While this increases communication overhead, the process of extending the algorithm to multiple processors or multiple independent machines, is straightforward and has been successfully applied to the traveling-salesman problem [15].

In the algorithm introduced in this paper, an ant's "life" begins randomly at either the origin or destination node of the demand. It proceeds until it finds the corresponding destination or origin node, using a selected available wavelength. Each ant chooses its wavelength according to parametric rules such as most-used or random selection. At the completion of its search, the ant deposits pheromone along the path. The ACO implementation presented in this paper incorporates the addition of memory in each ant consisting of nodes visited and a perception of surrounding nodes in the network.

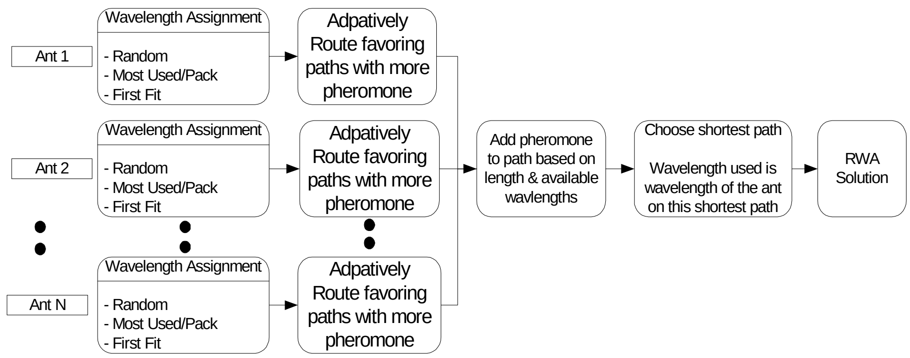

Fig. 3. ACO Flowchart

The next ant then proceeds similarly, choosing each vertex in its search path based probabilistically on the level of pheromone on the link to the next vertex, as shown in Figure 2 and Equation 1. A flowchart for the ACO algorithm applied in this paper is presented in Figure 3.

N	number of ants per connection request
L	set of links available from the current node
ϕ	normalized weight of length vs. weight of # of available wavelengths
l	a link compromising a path P
l_c	total capacity in wavelengths of link l
l_a	available wavelengths on link l
ψ_l	pheromone on link l

Fig. 4. Notation

Throughout our discussions, we use the notations and definitions in Figure 4. The probability γ that an ant will take a path l is the pheromone on that path normalized over the pheromone on all links available from the current node:

$$\gamma_l = \frac{\psi_l}{\sum\limits_{i \in L} \psi_i} \tag{1}$$

Pheromone is deposited on a per-demand basis. The pheromone matrix is reset once the final selection of wavelength and route is made for a connection request. This requires only one type of pheromone and avoids much of the overhead found in the implementation of the static case in [5], which requires running times on the order of hours. Even loading is achieved by having more pheromone deposited on paths that fewer previously routed O-D pairs occupy.

The shortest path found so far is stored as the best path. Subsequent paths with lengths equal to the length of the best path overwrite the best path. As more pheromone is laid down, ants discover paths with equal length but more pheromone. Therefore, the shortest path found with the highest pheromone is selected as the best and final route for the demand. It is important to note that this may not be the shortest available path.

A global pheromone update is performed after each ant completes a route. The shortest path found receives pheromone in inverse proportion to its length. While we want short paths to reduce potential blocking points, we also favor paths that have the fewest conflicts with demands already routed. Therefore, a component of pheromone update includes more pheromone for paths with more available wavelengths. Global update is assigned based on the following equations.
The sum of available lane quantity ratios for a path P is defined as

$$A_P = \sum_{l \in P} \frac{l_a}{l_c} \tag{2}$$

with the mean available lane ratio for a path P of:

$$M_P = \frac{A_P}{|P|} \tag{3}$$

The pheromone value on link l at time-step t, given as ψ_l^t, is updated according to Equation 4 where the scalar parameter ϕ , $0 \leq \phi \leq 1$, controls the emphasis on path length versus available-lane ratio.

$$\psi_l^{t+1} = \psi_l^t + \frac{\phi}{|P|} + M_P(1-\phi), \forall l \in P \tag{4}$$

Although each ant initially chooses a wavelength, the final wavelength selection is not made until all ants have completed a tour from source to destination for this connection request. The best route after N ants is found, and among the contiguous wavelengths available along this path, one is selected based on the most used, first fit, or random wavelength selection policy.

4 Results and Analysis

This study focuses on heuristic methods as opposed to optimization techniques. The complexity of the optimization approach limits its applicability to fairly small networks. For larger networks, particularly when coupled with online routing, computational efficiency is more important, and suggests the use of heuristic methods [4, 8].

For purposes of performance comparison, network blocking is the primary focus. Hereafter, general references to the performance of a particular heuristic refer to the percentage of blocked connection requests calculated during a simulation. All tests were conducted for $5x10^5$ demands at each Erlang.

Throughout this study, certain assumptions are made about reliability and spare capacity on the network. Blocking-probability measurements are based on loading the network to full capacity. The reservation of spare capacity is considered a separate issue.

In all tests, traffic for a source-destination pair arrives according to a Poisson process and the duration of each request is exponentially distributed with mean $1/\lambda$. A single test was the execution of $5x10^5$ connection requests. Load is measured in Erlang for the entire network, as in [2]. If network traffic is modeled at 50 Erlang, and the 51^{st} connection request arrives, a random existing connection is broken and it's resources freed. The ACO algorithm for unconstrained

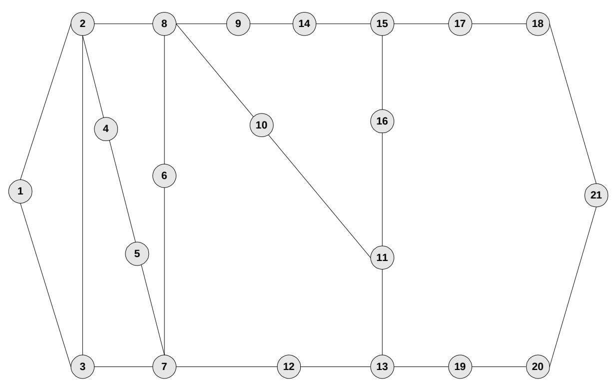

Fig. 5. The ARPA-2 Network

routing was tested using various algorithmic parameters, including number of ants, weight of length and congestion (ϕ), and random, most-used, and first-fit wavelength-selection rules. ACO was then tested against the best-performing algorithm used in [8], in which wavelengths are exhaustively searched for the shortest-available path from source to destination. Testing was conducted on the 21-node, 26-link ARPA-2 network in Fig. 5. Each edge in the network has a capacity of 16 wavelengths for all tests.

4.1 Parametric Tests

The first test conducted concerned the wavelength selection method. Random wavelength selection provided the lowest blocking at 50 ants, but performance peaked at this N for random selection. At 200 ants, most-used was the preferential method, outperforming first-fit and random at all traffic levels. In both tests, differences in blocking were small and details are omitted for brevity. However slight the differences in performance between wavelength selection methods, most-used was the method employed in all subsequent ACO tests.

N (# of ants)

Length Weight (ϕ) = 0.50
l_c = 16 wavelengths
wavelength selection rule = most-used

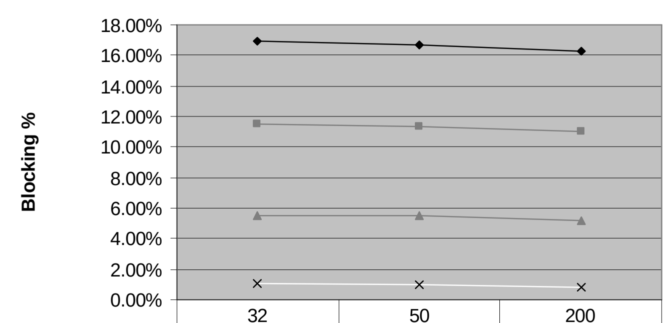

	32	50	200
Erlang 80	16.89%	16.73%	16.27%
Erlang 70	11.52%	11.34%	10.99%
Erlang 60	5.50%	5.48%	5.18%
Erlang 50	1.09%	1.00%	0.82%

# of Ants

Fig. 6. Number of Ants Comparison

Increasing the number of ants per connection request increases the likelihood that more paths are explored and better solutions found. The next set of tests concerned measuring blocking at several levels of N, for a fixed ϕ of 0.50. Performance will only improve up to a certain number of ants, although reductions in blocking percentages were seen at 200 ants. Significant processing is required with $N = 200$, however this increased processing load is easily distributed, a strength of the ACO algorithm. Results of 32, 50, and 200 ants are presented in Fig 6 for 4 traffic levels. 200 Ants provided the best performance in all situations, and is the N value used for comparisons with other algorithms in the next section.

4.2 Computational Comparisons with Published Algorithms

Mokhtar and Azizoglu present several heuristic RWA algorithms in [8]. They achieve the best results through an exhaustive search over all wavelengths for the shortest available path from source to destination. Dijkstra's shortest path algorithm is used [13]. Their exhaustive search provided lower blocking than methods employing static routing, and provides a benchmark for ACO algorithms in this paper.

ϕ and shortest path

N = 200 ants
$l_c = 16$ wavelengths
wavelength selection rule = most-used

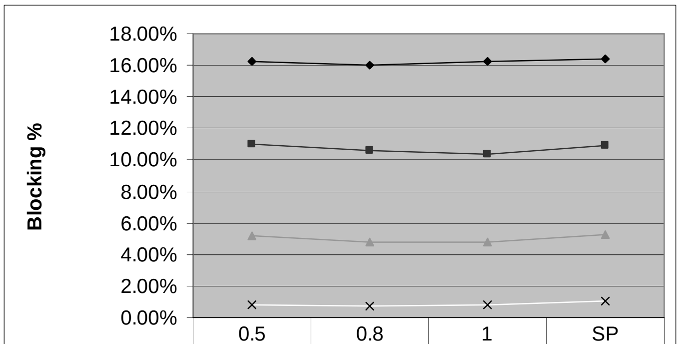

	0.5	0.8	1	SP
Erlang 80	16.26%	16.00%	16.21%	16.37%
Erlang 70	10.99%	10.59%	10.32%	10.88%
Erlang 60	5.18%	4.78%	4.80%	5.26%
Erlang 50	0.82%	0.74%	0.77%	1.04%

Algorithmic parameters

Fig. 7. ϕ and Shortest Path Comparison

The ACO algorithm at various parameters was compared with this algorithm, with the results displayed in Fig. 7. Using 200 ants, ACO provides the best results, outperforming the shortest available path at every traffic load. A ϕ value of 0.80 provided the best results in all but one case. At 70 Erlang, using no congestion information provided the best results by a slight margin. These results seem to indicate that while short paths are important, the best solution incorporates congestion information in selecting the route for each connection request.

4.3 Analysis

As demonstrated in Sec. 4.2, with random wavelength-selection methods, ACO achieves results comparable with the best results for adaptive RWA in [8]. ACO is superior when $N = 200$ and the most-used wavelength-selection rule is used. The value of ϕ in relation to a shortest-path based routing scheme strongly affects both the solution quality and the performance.

With a ϕ value of 1.0, congestion information is ignored, and the algorithm searches exclusively for the shortest path. However, since an initial pheromone value (ψ) of 1.0 is present on all edges in the network, ants may not find the shortest possible path.

5 Conclusions and Future Directions

The methods introduced in this paper demonstrate the validity of a model based on scoring paths for length and congestion. Ant-colony optimization modified for the dynamic RWA problem was shown to outperform the best-performing methods from [8] using large numbers of ants.

Future work includes extending the analysis of congestion quantification in an environment with wavelength grooming. Grooming is the term used to describe the optimization of capacity utilization in transport systems by means of cross-connections or conversions between different transport systems or layers within the same system [16]. In the context of wavelength division multiplexed networks, grooming refers to the addition of wavelength-conversion equipment at the nodes, such that a connection request may avoid the wavelength continuity constraint when passing through a node with a wavelength converter.

References

1. Dutton, H.J., Understanding Optical Communications, Prentice Hall, 1999.
2. Hui, Z., Jue, J., and Mukherjee, B. “A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks,” Optical Networks, January 2000.
3. Zhang, X. and Qiao, C. “Wavelength Assignment for Dynamic Traffic in Multi-fiber WDM Networks,” ICCCN ’98, pp. 479-585, 1998.
4. Stern, T.E. and Bala, K., “Multiwavelength Optical Networks.” Addison-Wesley, 1999.
5. Navarro-Varela, G. and Sinclair, M. “Ant-Colony Optimisation for Virtual-Wavelength-Path Routing and Wavelength Allocation,” Proc. Congress on Evolutionary Computation (CEC’99), Washington DC, USA, July 1999, pp. 1809-1816.
6. Li, L. and Somani, A., “Dynamic Wavelength Routing Using Congestion and Neighborhood Information.” IEEE Trans. Networking, Vol 7, No. 5, Oct. 1999.
7. Glover, F., Laguna M., and Laguna F., ”Tabu Search” Kluwer Academic Publishers, 1997.
8. Mokhtar, A. and Azizoglu, M. “Adaptive Wavelength Routing in All-Optical Networks,” IEEE/ACM Transactions on Networking, Vol. 6, No. 2, April 1998.

9. Chan, K. and Yum, T.P., "Analysis of least congested path routing in WDM lightwave networks." INFOCOM '94. Networking for Global Communications, 13^{th} Proceedings IEEE, 1994. pp. 962-969.
10. Banerjee, D. and Mukherjee, B. "A Practical Approach for Routing and Wavelength Assignment in Large Wavelength-Routed Optical Networks." IEEE Journal on Selec. Areas in Comm., vol 14, No. 5, June 1996.
11. Colorni, A., Dorigo, M. & Maniezzo, V. "Distributed optimization by ant colonies," Proc. First European Conference on Artificial Life, Paris, France, pp. 134-142, 1991.
12. Dorigo, M. and Gambardella, L.M., "Ant-Colony System: A Cooperative Learning Approach to the Travelling Salesman Problem." IEEE Transactions on Evolutionary Computation, pp. 53-66.
13. Dijkstra, E. "A note on two problems in connexion with graphs." Numerische Mathematik, 1959. vol. 1, pp. 269-271.
14. Karasan, E. and Ayanoglu E., "Effects of Wavelength Routing and Selection Algorithms on Wavelength Conversion Gain in WDM Optical Networks." IEEE Trans. Networking, vol 6, pp. 186-196, April 1998.
15. Garlick, R. "A Distributed Ant-Colony Optimization Solution to the TSP," Technical Report, Southern Methodist University.
16. Barr, R. and Patterson, R. "Grooming Telecommunications Networks." Optical Networks, vol. 2, no. 3, May 2001, pp. 20.