

A STUDY ON THE EFFICACY OF REGULAR VIRTUAL TOPOLOGY DESIGN HEURISTICS FOR OPTICAL PACKET SWITCHING

Olufemi Komolafe¹, David Harle¹, David Cotter²

¹ *Dept. of Electronic & Electrical Engineering, Strathclyde University, Glasgow G1 1XW, U.K.*

² *Corning Research Centre, Adastral Park, Martlesham Heath, Ipswich IP5 3RE, U.K.*

femi@comms.eee.strath.ac.uk

Abstract: This paper studies the design of regular virtual topologies to facilitate optical packet switching in networks with arbitrary physical topologies. The regular virtual topology problem is intractable, hence heuristics must be employed to find (near) optimal solutions expeditiously. Numerous optimisation heuristics are compared in this paper in terms of the quality of the final solution yielded and the efficiency associated with obtaining these solutions. The comparison is conducted for numerous unique randomly generated physical topologies, allowing generic conclusions to be drawn. The heuristics are meritocratically ordered and suggestions made regarding the choice of optimisation heuristic.

Key words: Optical Packet Switching, Virtual Topology Design, Multi-Layer Networks, Combinatorial Optimisation, Artificial Intelligence, Graph Theory.

1. INTRODUCTION

Certain regular multi-processor interconnection architectures are considered attractive for use in future high-speed packet switched optical networks. Such architectures are attractive for numerous reasons including the use of simple and distributed routing schemes, the possibility of the avoidance (or minimisation) of optical buffering, and increased predictability of the network performance. Suitable regular architectures include the Shufflenet [1,2], the de Bruijn Graph [2,3], the Hypercube [2,4], the Manhattan Street Network [5,6] and the Kautz Graph [7].

The multi-processor interconnection architectures may be deployed (or embedded) as regular virtual topologies within arbitrary physical topologies [4,8,9]. Hence, the architectures offer a means to facilitate optical packet switching over arbitrary physical topologies by reducing the packet routing complexity and minimising the optical buffering requirement. Nodes of the interconnection architecture are mapped onto nodes of the physical topology and WDM channels, lightpaths, are established between the appropriate nodes to realise the desired connectivity. It is evident that the length of a particular lightpath is dependent upon which physical topology nodes its end-nodes in the virtual topology are mapped onto. The mapping of a section of an exemplar regular virtual topology onto a physical topology is illustrated in Figure 1. The problem of designing a (regular) virtual topology optimally is known to be *NP*-hard, hence heuristics must be used to find (near) optimal solutions expeditiously [4,9,10,18].




Figure 1. Example of mapping section of regular virtual topology onto physical topology

The design of *irregular* virtual topologies for wavelength routed WDM networks has been studied extensively [18]. In comparison, relatively little work has been done on the design of *regular* virtual topologies to carry packet traffic. Furthermore, arguably most of the work considering regular virtual topology design has been of an *ad hoc* nature; specific optimisation heuristics have been applied when designing regular virtual topologies for specific physical networks [4,8,9,10,11]. Additionally, in many cases, a “black box” view of the virtual topology design process is adopted; attention is paid to the initial and final results and not to the evolution of the design

process. Therefore, it may be said that little is known about the relative performance and progress of different design techniques in a generic sense. Consequently, this paper studies the efficacy of a number of different regular virtual topology design heuristics when applied to numerous randomly generated networks. The objective is to compare the heuristics in terms of the quality of the final solution obtained and the efficiency associated with obtaining this solution. Additionally, the impact that particular graph theoretic properties of the physical topology have on the heuristics' performance is assessed.

The exemplar multi-processor interconnection architecture considered in this paper is the Manhattan Street Network (MSN) featuring a novel routing scheme; *Clockwork Routing* [12,13]. The MSN with the *Clockwork Routing* scheme is well-suited for use in optical packet switched networks for a number of reasons: the routing processing is extremely simple and suitable for optical implementation, no optical domain buffering is required, no resequencing is needed at the destination nodes and throughput is comparable with conventional store-and-forward packet switching [12,13]. The cost considered in this paper is the mean lightpath length. (Throughout this paper, distance is in terms of concatenated hops over a topology, and not in terms of physical length.) The number of physical hops traversed by a lightpath is important for several reasons: it affects the number of wavelengths needed, impacts the dimensions of the wavelength-selective devices, indicates the number of optical cross-connects a packet must traverse between any two adjacent nodes in the MSN and would impact the deployment of optical amplifiers and the consumption of other network resources. Furthermore, lightpath length corresponds to a figure of merit often used in graph theory [14]. The goal, therefore, is to find the mapping of MSN nodes onto physical topology nodes that minimises the mean lightpath length.

This paper compares the efficacy of different heuristics when applied to the problem of deploying the MSN in randomly generated arbitrary networks, with the cost being the mean lightpath length. The remainder of the paper is organised as follows. Section 2 describes the optimisation techniques used and the criteria used to evaluate their efficacy. Section 3 discusses the experimental approach adopted. Section 4 presents and discusses the results obtained. Section 5 concludes the paper.

2. REGULAR VIRTUAL TOPOLOGY DESIGN TECHNIQUES

2.1 Optimisation Heuristics

Given the *NP*-hard nature of the regular virtual topology design problem, heuristics must be employed to find (near) optimal solutions expeditiously. Different optimisation techniques have been applied to regular virtual topology design in isolation; this paper seeks to compare some of these techniques for numerous randomly generated networks, allowing more generic conclusions to be drawn regarding their efficacy. Five different optimisation heuristics, ranging from sophisticated artificial intelligence (AI) based approaches to more intuitive simpler techniques, are considered in this paper and are now discussed.

Simulated Annealing (SA) [15] is a popular and reputable combinatorial optimisation tool which has been successfully applied to the regular virtual topology design problem [4,9]. SA is modelled on the physical cooling of molecules to form crystals. A salient property of SA is that uphill moves are allowed with a regulated and decreasing probability, providing a means of escaping local minima. A previously developed SA implementation, which was found to obtain favourable results [9], was adopted.

Genetic Algorithms (GA) [16] are another AI-based optimisation tool, modelled on evolution in nature. GA have been successfully applied to a plethora of mathematical problems. GA have also been used to tackle the regular virtual topology design problem [10]; the two implementations of GA which were found to be best, *Cycle Crossover* (CX) and *Partially Mapped Crossover* (PMX), are used in this paper.

Hill Climbing (HC) is a relatively simple optimisation technique. An initial solution is (typically randomly) generated. A new solution in the vicinity of the current solution is evaluated. The move is accepted if, and only if, the new solution is better than the current solution. This procedure is repeated a prescribed number of times and the best solution found stored.

Random Search (RS), as its name suggests, is when the search space is explored randomly and the best result found is output. The random search is typically conducted for a predetermined number of points.

2.2 Evaluation of Efficacy of Heuristics

The relative efficacy of simulated annealing, the two implementations of genetic algorithms, hill climbing and random search is to be assessed. Two key metrics used to compare the different optimisation heuristics are:

- The quality of the final solution, i.e. the final mean lightpath length obtained.
- The efficiency associated with obtaining solutions.

The first metric is rather straightforward; simply apply the different optimisation heuristics and note which one obtains the best result. Determining the efficiency associated with finding solutions for each heuristics is less straightforward. The method used in this paper is to consider the number of solutions considered in obtaining the final solution. Evaluating the cost function is relatively time-consuming, hence the number of solutions considered is indicative of the time and computational effort required by the heuristic. Additionally, the number of solutions considered in obtaining a final solution corresponds to the area of the search space explored. Furthermore, monitoring the number of solutions considered means the progress of the optimisation technique can be studied. Lastly, it may be argued that the number of solutions is a more accurate and generic measure of the efficacy than say, the absolute time taken, as it is independent of any idiosyncrasies in the implementation of the heuristic and operating system or central processing unit (CPU) constraints. Hence, the efficacy of the optimisation techniques is studied by observing the mean lightpath length obtained as the algorithm progresses.

3. EXPERIMENTAL APPROACH

In this paper, a 6x6 MSN is deployed in numerous randomly generated 36-node physical topologies. (Methods to resolve a mismatch in the node numbers include the introduction of fictional physical topology nodes [4] or virtual topology node co-location [17].) A fairly large network size of 36 was chosen so that the MSN deployment problems tackled are relatively difficult. Each of the five optimisation techniques was used to seek the MSN node mapping that minimised the mean lightpath length and their progress studied. When generating randomly connected networks, disconnected networks, self-loops and double links were all avoided; in graph theory parlance, only simple connected graphs were considered. Bi-directional links are assumed. The degree of a node is the number of links terminating at the node. The degree of the nodes in the physical topology were uniformly randomly distributed. (The impact of non-uniformly distributed nodal degrees has been previously investigated [17].) Random networks with different mean degrees were studied; mean degrees of 2, 5 and 8 were considered to ascertain the impact, if any, that the mean degree has on the performance of the different optimisation heuristics. In order to make the findings in the paper more generic, over 50 randomly generated unique

networks were used. The networks were tested for uniqueness; two randomly generated networks with the same mean degree are said to be unique if any of the following parameters have different numerical values:

- maximum degree
- degree variance
- mean inter-nodal distance
- maximum inter-nodal distance (network diameter)
- variance of inter-nodal distance
- number of bridges (a bridge is a link whose removal disconnects the network).

The performance of each optimisation technique was evaluated for every unique physical topology randomly generated. The results obtained are presented in the next section.

4. RESULTS

4.1 Presentation

For each unique physical topology, each optimisation heuristic was used to seek the node mapping which minimised the mean lightpath length, with the evolution of the mean lightpath length monitored. Figure 2 shows the example of the simulated annealing results for a sample of 10 arbitrarily chosen different random physical topologies, each with a mean degree of 2. (For clarity, only a fraction of the total number of randomly generated networks with degree of 2 are shown.) Each line in the Figure 2 corresponds to a unique physical topology and shows the progress of the SA algorithm as it seeks to minimise the mean lightpath length. The complete number of physical topologies of degree 2 was generated and the average results for SA calculated. As an example, the standard deviation (STD) and 95% confidence interval (CI) were also computed at selected points in the evolution of the SA algorithm and are presented in Table 1. From Table 1, it may be seen that spread in the mean lightpath length obtained for the different topologies decreases as the SA implementation progresses, hence the confidence interval also decreases. When the random generation of the initial solution in the SA implementation is allied with the fact that unique physical topologies are considered in each case, the difference between the initial mean lightpath lengths will be expected to be relatively large and arbitrary, as found in Table 1. Towards the end of the SA algorithm, when the mean lightpath has been minimised, it is evident that the spread in mean

lightpath length is lower. It has been found that there is a relationship between the optimised mean lightpath length and the mean degree [17]. Hence, it is unsurprising that similar final mean lightpath lengths are obtained for these randomly generated physical topologies, since they all have the same mean degree of 2.




Figure 2. Performance of SA for 10 sample randomly generated networks

Table 1. Statistics for SA for all randomly generated networks with degree of 2

No. of solutions considered (x10,000)	0	40	80	120	160	200	240
Mean	5.810	5.721	5.458	3.906	3.352	3.345	3.343
STD	0.734	0.707	0.548	0.288	0.236	0.238	0.240
95% CI	0.243	0.234	0.182	0.096	0.078	0.079	0.079

Figure 3 shows the progress of the five virtual topology design heuristics as they seek to minimise the mean lightpath length when the MSN is deployed in all the randomly generated networks with a mean degree of 2. Similarly, Figures 4 and 5 show the results for networks with mean degrees of 5 and 8 respectively.

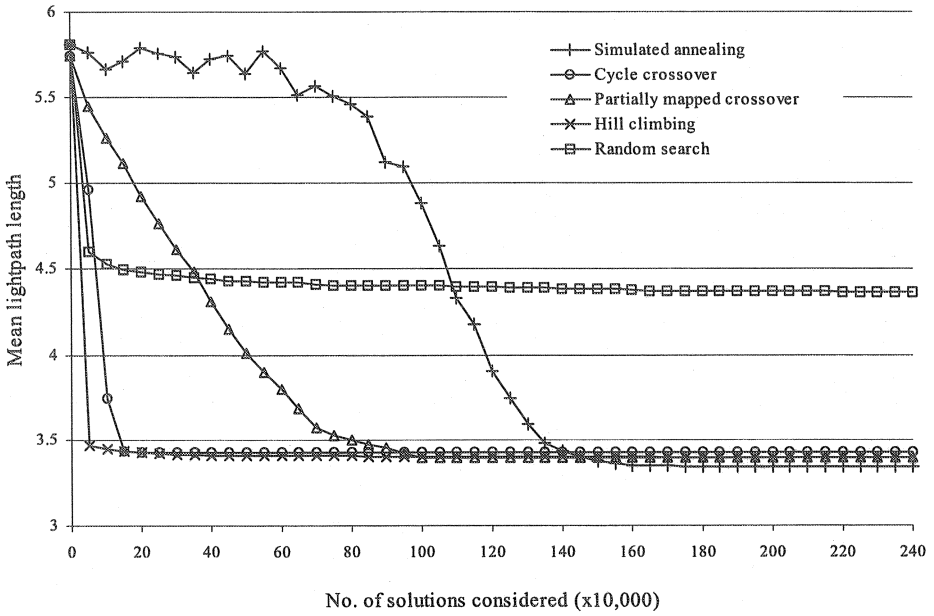


Figure 3. Performance of heuristics for randomly generated networks with mean degree of 2

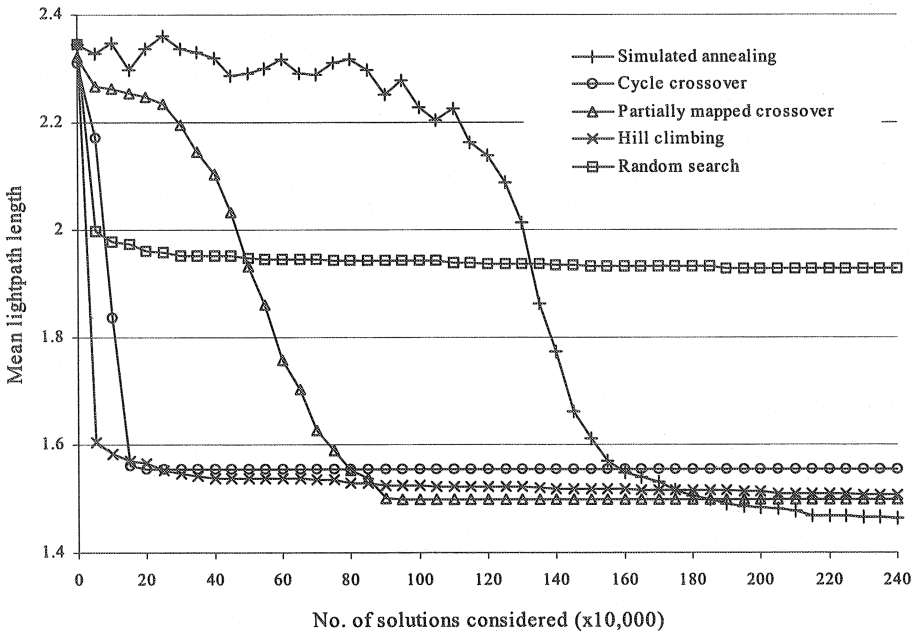


Figure 4. Performance of heuristics for randomly generated networks with mean degree of 5

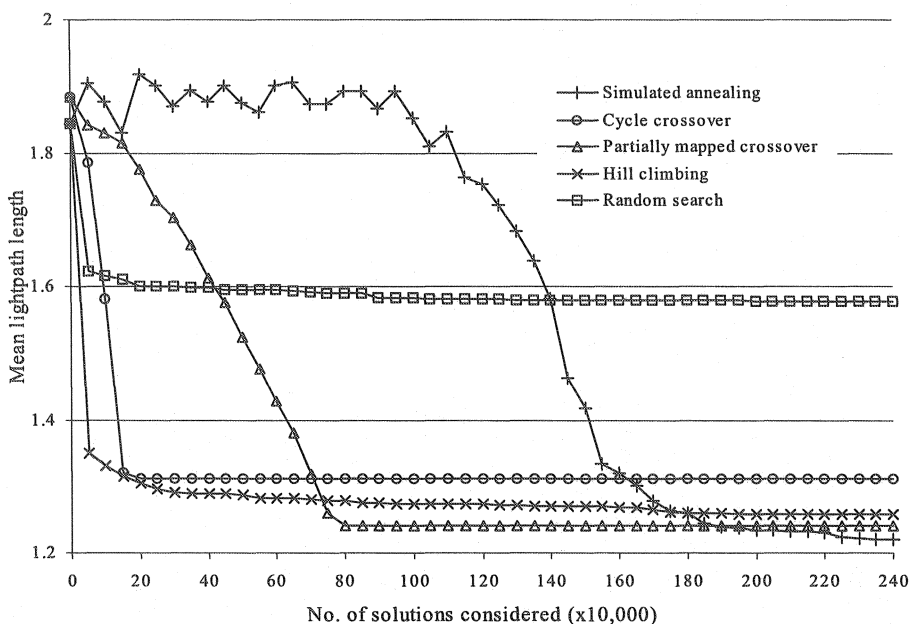


Figure 5. Performance of heuristics for randomly generated networks with mean degree of 8

4.2 Observations

Before commenting on the performance of the individual heuristics, numerous interesting general observations may be made from the results presented in Figures 3 to 5. All the optimisation heuristics reduce the initial mean lightpath length. Furthermore, the initial and final mean lightpath length obtained falls as the mean degree of the networks increases, i.e. the initial and mean lightpath lengths are greater in Figure 3 than in Figure 5. This finding is intuitive as it would be expected that since nodes will be on average closer together in networks with larger mean degrees, the mean lightpath length obtainable will be lower. The amount of improvement obtained on the initial mean lightpath length falls as the mean degree of the physical topology increases. For example, SA manages to reduce the initial mean lightpath length by 42.46% for a mean degree of 2 (Figure 3) whereas the corresponding reduction is 33.81% for a mean degree of 8 (Figure 5). With the exception of random search, the heuristics all seem to converge to a similar value. The fact that these optimisation techniques all work in fundamentally different ways, yet are asymptotic to the same value, suggests that this value may very well be the theoretical minimum mean lightpath length obtainable. Encouragingly, the best final mean lightpath length obtained are of a reasonable magnitude, and suggest that MSN deployment

is a feasible way to facilitate optical packet switching over these arbitrary physical topologies.

4.3 Discussions

It is evident from Figures 3 to 5 that the mean degree of the physical topologies does not appear to have a significant impact on the relative performance of the different optimisation techniques; the same trends are observed regardless of the mean degree.

Firstly, it can be seen that random search (RS) obtains the worst result; there is a slight improvement on the initial mean lightpath length but thereafter, only negligible improvements are obtained. Hill climbing (HC) shows somewhat dramatic improvement on the initial mean lightpath length. However, only marginal improvements are obtained thereafter. It appears that hill climbing finds a local minimum rapidly (hence improves the initial cost quickly), however the algorithm struggles to escape from this local minimum throughout the rest of the implementation. Hill climbing did not obtain the best result found in any case, nevertheless, results which are only slightly poorer than the best results are obtained relatively quickly. A fundamental trade-off between exploration of the search space and exploitation of the information gained so far in the search exists in most optimisation techniques. The relative performance of HC and RS suggest that exploitation is more important than exploration when seeking optimal mappings of the MSN onto the physical topology in terms of the mean lightpath length.

Simulated annealing's performance is extremely interesting. Of all the optimisation heuristics, SA is the only one that explicitly allows uphill moves. Uphill moves are allowed with a probability that decreases as the SA algorithm progresses. Hence, in Figures 2 to 5, it can be seen that SA does indeed allow uphill moves. The frequency with which uphill moves occur diminishes as the algorithm progresses. Such observations are encouraging as they suggest that simulated annealing has been implemented correctly since the experimental results comply with theoretical expectations. Of all the optimisation heuristics, SA takes the longest time to obtain its solution; once again this is in keeping with expectations as a commonly known disadvantage of simulated annealing is the long (and sometimes excessive) time taken to obtain solutions [19]. SA was found to obtain the best results; it obtained the lowest mean lightpath length, regardless of the mean degree of the physical topologies.

The performance of the two genetic algorithm implementations are encouraging. Cycle crossover (CX) obtains good results quickly. This fact has already been observed [10]. The relative performance of CX appears to

decrease as the mean degree of the physical topologies increases, suggesting that CX works better for sparsely connected networks. Partially mapped crossover achieves favourable results. PMX was found to obtain results that were only marginally worse than the SA results, and typically did so considering less than half the number of solutions.

The optimisation heuristics may be ordered meritocratically according to the two efficacy metrics suggested in Section 2. In terms of the final mean lightpath length, the order is SA, PMX, HC, CX and lastly, RS. However, in terms of the number of solutions considered before the final solution is obtained, the ordering is CX, HC, RS, PMX and SA. (Convergence is judged to have occurred at a point after which only negligible improvements on the mean lightpath length are observed.). Hence, the choice of optimisation heuristic depends on the relatively importance associated to the quality of the final solution and the efficiency with which it is obtained. In most practical cases, the quality of the final solution is the more important factor, hence simulated annealing and partially mapped crossover will be the more favoured regular virtual topology design heuristics.

5. CONCLUSION

Different multi-processor interconnection are attractive for optical packet switching as they offer simple and distributed routing schemes and minimise the optical buffering requirement. Such architectures may be deployed as regular virtual topologies in arbitrary physical networks. Using the Manhattan Street Network as an example, this paper has studied the regular virtual topology design problem. The regular virtual topology problem is intractable, hence heuristics must be employed to find (near) optimal solutions in reasonable time. Several optimisation heuristics are compared in terms of the quality of the final solution yielded and the efficiency associated with obtaining this solution. Numerous unique randomly generated physical topologies were used in the studies, allowing generic conclusions to be drawn. The heuristics are meritocratically ordered and recommendations made regarding the choice of heuristic. With the exception of random search, all the heuristics significantly improved on the initial result. However, it was found that either simulated annealing or one of the implementations of genetic algorithms, partially mapped crossover (PMX), was the most efficient virtual topology design heuristic.

REFERENCES

- [1] M. G. Hluchyj and M. J. Karol, "Shufflenet: an application of generalized perfect shuffles to multihop lightwave networks," *IEEE/OSA J. Lightwave Technol.*, vol. 9, pp 1386-1397, Oct. 1991.
- [2] B. Mukherjee, "WDM-based lightwave networks part II: multihop systems," *IEEE Network*, pp 20-31, July 1992.
- [3] K.N. Sivarajan and R. Ramaswami, "Lightwave networks based on de Bruijn graphs," *IEEE/ACM Trans. Networking*, vol. 2, pp.70-79, Feb. 1994.
- [4] B. Mukherjee, S. Ramamurthy, D. Banerjee and A. Mukherjee, "Some principles for designing a wide-area optical network," in *Proc. IEEE Infocom*, pp. 110-119, 1994.
- [5] N. F. Maxemchuk, "The Manhattan street network," in *Proc. IEEE Globecom*, 1985.
- [6] N. F. Maxemchuk, "Routing in the Manhattan street network," *IEEE Trans. Commun.*, vol. com-35, pp. 503-512, May 1987.
- [7] C.P. Ravikumar, T. Rai, and V. Verma, "Kautz graphs as attractive logical topologies in multihop lightwave networks," *Computer Commun.*, vol. 20, pp. 1259-1269, 1997.
- [8] I. Chlamtac, A. Ganz and G. Karmi, "Lightnets: topologies for high-speed optical networks," *IEEE/OSA J. of Lightwave Technol.*, vol. 11, pp. 951-961, May/June 1993.
- [9] O. Komolafe, D. Harle and D. Cotter "Ultrafast optical packet switching over arbitrary physical topologies using the Manhattan street network," in *Proc. IEEE Int. Conf. Commun.*, June 2001.
- [10] O. Komolafe, D. Harle and D. Cotter, "Optical packet switching over arbitrary physical topologies using the Manhattan street network: an evolutionary approach," in *Proc. IFIP Conf. Optical Network Design Modelling*, Feb. 2001.
- [11] O. Komolafe, D. Harle and D. Cotter, "Designing a multi-hop regular virtual topology for ultrafast optical packet switching: node placement optimisation and/or dilation minimisation?," in *Proc. 6th European Conf. on Networks & Optical Commun.*, June 2001.
- [12] F. Chevalier, D. Cotter and D. Harle, "A new packet routing strategy for ultra-fast photonic networks," in *Proc. IEEE Globecom*, 1998.
- [13] F. Chevalier, D. Cotter and D. Harle, "Routing and fairness issues associated with a self routing control strategy in regular mesh photonic networks," in *Proc. Int. Teletraffic Conf.*, pp. 1301-1310, June 1999.
- [14] A. Rosenberg, "Data encodings and their costs," *Acta Informatica*, vol. 9, pp. 273-292, 1978.
- [15] S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671-680, 1983.
- [16] J. H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan Press, 1975.
- [17] O. Komolafe, "High-speed optical packet switching over arbitrary physical topologies using the Manhattan street network," Ph.D Thesis, Dept. of Electronic & Electrical Engineering, University of Strathclyde, Nov. 2001.
- [18] R. Duttra, G. N. Roukas, "A survey of virtual topology design algorithms for wavelength-routed optical networks," *Optical Networks Mag.*, vol. 1, pp. 73-89, Jan/Feb. 2000.
- [19] R. W. Eglese, "Simulated annealing: a tool for operational research," *European J. of Operational Research*, vol. 46, pp. 271-281, 1990.