Available at www.ComputerScienceWeb.com
POWERED BY SCIENCE @ DIRECT®

ELSEVIER

Computer Networks 42 (2003) 251–260

COMPUTER NETWORKS

www.elsevier.com/locate/comnet

Optimal node placement in an optical packet switching Manhattan street network

Olufemi Komolafe [a,*,1], David Harle [b]

[a] Department of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK
[b] Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow G1 1XW, UK

Received 27 March 2002; received in revised form 24 October 2002; accepted 1 February 2003

Responsible Editor: A. Fumigalli

Abstract

The judicious allocation of traffic matrix sources/destinations to nodes of an exemplar optical packet switching architecture, the Manhattan street network, is studied in this paper. The node placement optimisation (NPO) problem is intractable, hence simulated annealing is used to find (near) optimal solutions expeditiously. A number of non-uniform traffic patterns are considered, facilitating the study of the innate parameters of a traffic matrix that affect the optimisation results. A salient finding is that, regardless of the underlying traffic pattern, the efficacy of NPO is correlated to the coefficient of variation of the traffic matrix.

Keywords: Manhattan street network; Optical packet switching; Combinatorial optimisation; Multi-processor interconnection architectures

1. Introduction

1.1. Motivation

The proliferation of the Internet has provided significant impetus to develop improved packet switching architectures. It is widely accepted that optical packet switching architectures will be crucial in future high-speed networks. Two key restrictions associated with optical packet switching are the need to minimise the use of optical domain buffering (due to the non-availability of static random access optical memory and the limitations of travelling wave optical buffers) and to reduce the routing complexity (especially if optical logic devices are to be used to process packet headers). Hence, architectures that offer simple and distributed routing schemes coupled with a contention-free optical domain are attractive for use in optical packet switching networks.

1.2. The Manhattan street network

The Manhattan street network (MSN) [1,2] is the architecture considered in this work. A number

* Corresponding author. Tel.: +44-141-330-6297; fax: +44-141-330-4913.

E-mail addresses: femi@dcs.gla.ac.uk (O. Komolafe), d.harle@eee.strath.ac.uk (D. Harle).

[1] Previously with Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK.

1389-1286/03/$ - see front matter
doi:10.1016/S1389-1286(03)00192-0

of routing schemes have been proposed for the MSN including store and forward routing [2], deflection routing [3] and hot potato routing [4]. However, these different routing schemes are generally ill-suited for high-speed optical packet switching; they typically need routing computations that are too complex for simple optical logic devices, require buffers at intermediate nodes within the MSN or achieve unacceptably low network performance. However, the MSN with *Clockwork Routing* [5,6], a recently proposed packet routing scheme, is an ideal candidate for optical packet switched networks for a number of reasons. The routing processing is extremely simple, 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 [5,6]. *Clockwork Routing* requires that all the nodes in an $n \times n$ MSN be synchronised to a global clock. The timeslots are arranged in a modulo-n sequence of frames. Each node consists of a simple 2×2 crossbar switch. All nodes are in the cross state for the first $n - 1$ timeslots in each timeframe, and in the bar state for the last timeslot in the timeframe. Fig. 1 shows a 4×4 MSN. By appropriately inserting a packet into a timeslot on a particular output link, it is automatically routed to the destination. No additional processing is required at intermediate nodes; these nodes need only determine whether the packet has reached its destination or not. Such a simple "for me or not for me" evaluation may be readily implemented in the optical domain [7]. Optical domain contentions are avoided through the use of peripheral electronic buffers.

1.3. Node placement optimisation

Due to the aforementioned reasons, the MSN with the *Clockwork Routing* is the optical packet switching architecture considered in this paper. That said, different multi-processor interconnection architectures have also been considered for use in optical networks. Suitable architectures include the Shuffle-Exchange (or Shufflenet) [8,9], the de Bruijn Graph [9,10], the Hypercube [9,11, 12] and the Kautz Graph [13]. The ability of such multi-processor interconnection architectures to contend with arbitrary physical topologies and non-uniform traffic are key prerequisites to their use in practical networks. Typically, the architectures may be deployed as regular virtual topologies in arbitrary physical topologies; a relatively well-understood problem [11,12,15–17]. This paper considers the second issue; the problem of carrying non-uniform traffic efficiently over multi-processor interconnection architectures. The objective is to allocate traffic matrix sources/destinations to nodes of the multi-processor interconnection architecture optimally; a problem conventionally referred to as *node placement optimisation* (NPO).

The NPO problem has been studied for a number of architectures, including the Shuffle-

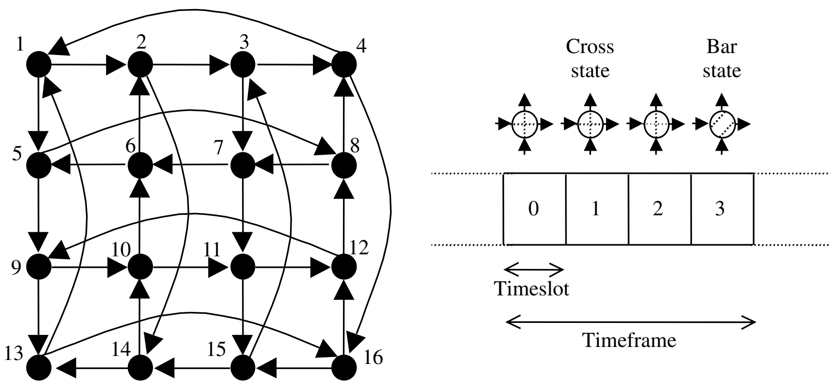

Fig. 1. 4×4 Manhattan street network.

Exchange [18,19,21,23], bus [20], ring [22,23] and MSN [23]. The two costs typically considered are the maximum traffic flow over any link [18,20,22] or the mean traffic-weighted inter-nodal distance [19–21,23]. NPO is known to be NP-hard [20, 21,23], hence heuristics must be used to find (near) optimal solutions expeditiously. The typical approach is to design and apply customised heuristic algorithms to NPO [18–22]. However, it may be argued that this is a somewhat ad hoc approach as the heuristics typically are peculiar to the particular multi-processor interconnection architecture and as such, are not readily applicable to other architectures. Only recently has a more generic optimisation technique, simulated annealing (SA), been applied to NPO [23]. A novel NPO problem encoding is proposed in this paper allowing a different implementation of SA to be applied to NPO.

The conventional approach to NPO is to use instances of traffic matrices derived from selected underlying traffic patterns as optimisation inputs. However, in this paper, a different approach is taken; the impact of the evolution of certain traffic matrix parameters on NPO results is investigated. Such studies are significant as they provide insight into the salient traffic matrix parameter(s) that affects the efficacy of NPO; what statistic(s) of a traffic matrix make it expedient to conduct NPO?

The paper is organised as follows: Section 2 introduces the optimisation technique employed. The optimisation cost and different non-uniform traffic patterns considered are presented in Section 3. Section 4 presents and discusses the results obtained for the selected traffic patterns. Properties of the traffic matrix that affect the efficacy of NPO are studied in Section 5. Section 6 concludes the paper.

2. Simulated annealing

SA [24,25] is a popular and reputable combinatorial optimisation tool that has been successfully applied to the problem of deploying multi-processor interconnection architectures in physical topologies [11,14] and to NPO [23]. SA is an iterative and stochastic search algorithm modelled on the physical cooling of molecules to form crystals, annealing. Initially, the temperature of a solid is increased to a maximum value at which the solid melts. At this high temperature, the molecules can move around freely. The collection of molecules at a given temperature may be simulated using the Metropolis algorithm [26]. The Metropolis algorithm always permits changes that decrease the system's energy but incremental changes are accepted with a probability that diminishes as the temperature falls. The energy of the system corresponds to the cost being optimised. Evidently, therefore, a salient feature of SA is that apparently non-expedient changes are possible, providing means of escaping local minima. The temperature is lowered gradually, with most time being spent at temperatures near the freezing point. In such ground states, provided the temperature has been slowly reduced, molecules are arranged in a highly structured manner; a crystal has been produced. Ideally, the state of minimum energy corresponds to the global optimum.

Each SA implementation is defined by four key design choices [24], upon which success depends:

- *A concise description of a configuration of the system* refers to the encoding of the NPO problem in a form appropriate for SA. The encoding adopted in this work is similar to that used for the Travelling Salesman Problem, a problem tackled successfully using SA [24]. A configuration is a permutation of the number $1, \ldots, N$ and represents the location of traffic matrix sources/destinations in the MSN.
- *A random generation of rearrangements of elements in a configuration* defines the manner in which new solutions are generated or the search space explored. Randomly selected portions of the permutation are swapped to generate new node placements. It was found empirically that swapping $\sqrt{N}/2$ node-pairs each time satisfies the conflicting demands to *explore* as large an area of the search space as possible and to *exploit* any favourable regions discovered.
- *A quantitative objective function*, the cost function, is evaluated to ascertain the quality of solutions. The cost used in this work is discussed in Section 3.1.

- *An annealing schedule of the temperatures and the length of times for which the system is to be evolved* comprises of several vital discretionary parameters. Ideally, the initial temperature should be large enough to ensure that the search space is thoroughly explored, hence a high initial temperature of 1000 was used. Thermal equilibrium is reached once a sufficient number of rearrangements are generated at a given temperature. The traditional view is that equilibrium ought to precede temperature reduction. However, it has been argued that equilibrium need not be reached provided the cooling is sufficiently slow [27]. The pragmatic approach taken is to select a sufficiently large number of rearrangements per temperature step (10 000), such that it is safe to assume equilibrium will be reached at each temperature [28], and additionally, to conduct the cooling very slowly ($T(t+1) = 0.95\,T(t)$, where $T(t)$ is the temperature at step t). Ideally, the SA algorithm ought to be stopped only after the system freezes. Consequently, the SA algorithm was only stopped once the temperature of the system falls below 10^{-4}.

3. Cost and traffic patterns

3.1. Cost

The mean traffic-weighted inter-nodal distance [19–21,23] is the cost used in this paper. The mean traffic-weighted inter-nodal distance is an intuitive cost and, if the links of the multi-processor interconnection architecture are equal, is indicative of the mean packet delay [19,21]. This paper uses SA to find the arrangement of traffic matrix nodes that minimises the mean traffic-weighted inter-nodal distance in the MSN.

3.2. Traffic patterns

Initially, the variation of properties of four different traffic patterns are considered. In each case, the average traffic between N nodes are represented by an $N \times N$ matrix $\Lambda = [\lambda_{ij}]$, where λ_{ij} is the average traffic flow from node i to node j. $\lambda_{ii} = 0$ and the traffic matrix is asymmetrical; i.e. $\lambda_{ij} \neq \lambda_{ji}$.

In order to study the impact of the evolution of the traffic matrix amplitude, λ_{ij} is chosen from a uniform random distribution between 1 and A. A was gradually increased from 10 to 100. Hence, different traffic matrices with the desired amplitudes may be generated and used as inputs to the NPO problem. In the next traffic pattern considered, the mean traffic level is kept constant but the traffic matrix range is varied. Therefore, λ_{ij} is chosen from a uniform random distribution in the range $100 \pm R$, with R being varied from 10 to 90 in steps of 10. The existence of a server node is the third considered traffic pattern. A server node s is (randomly) selected such that the traffic emitted or sunk by s (i.e. λ_{si} and λ_{is} $\forall i, i \neq s$) is S times greater than would otherwise be the case. The variation of S, the server traffic intensity, between 5 and 50 was investigated. Lastly, a traffic hot spot exists between a (randomly chosen) node-pair, a, b. The traffic exchanged by this node-pair (i.e. λ_{ab} and λ_{ba}) is M times greater than they would otherwise exchange. M, the hot spot traffic intensity, was evolved from 5 to 50. Comparing the results from these different traffic patterns will provide insight into the salient properties of a traffic matrix that affect NPO.

4. Results for selected traffic patterns

For each traffic pattern, numerous different traffic matrices were generated for each stage in the evolution of the traffic matrix. Each time, SA was used to find the arrangement of the traffic matrix nodes onto the MSN that minimised the mean traffic-weighted inter-nodal distance. The optimisation results were averaged for all the traffic matrices and these results presented. The results shown are for a 4×4 MSN and a 16×16 traffic matrix. For verification, NPO was performed for larger sized MSNs. Similar trends to those presented here were observed and so these results have been omitted for brevity. Additionally, a genetic algorithm (GA) [29] solution to NPO was implemented and used to verify the SA results. It was found that, in all cases, GA performed comparable to SA and so the GA results have been omitted from this paper for brevity. The fact that

GA and SA seek the optimal solution in radically different ways and yet were found to obtain similar results gives confidence in the veracity of these results.

Previous studies on NPO have typically compared a number of heuristic optimisation results for selected traffic matrices [18–21]. Thus, it has been tacitly assumed that it is always productive to conduct NPO and the properties of a traffic matrix that affect the effectiveness of NPO are somewhat taken for granted. Some preliminary investigations into the properties of multiprocessor interconnection architectures and traffic matrices that impact NPO have been undertaken [23]. It was suggested that the efficacy of NPO is related to the variability and relationship between inter-nodal distances in the architecture and the traffic matrix. (Since a single architecture and routing scheme is being considered in this work, only the variability of the inter-nodal traffic is appropriate.) This paper builds on these ideas by identifying the exact statistical properties of traffic matrices that affect NPO's efficacy and investigating the impact that these have on the optimisation results.

The efficacy of NPO is defined as the improvement obtained by SA on the initial (or random) mean traffic-weighted inter-nodal distance. Numerous traffic matrix parameters were considered to ascertain their relationship to the improvement on the initial costs. Parameters considered included the mean traffic level, the variance, standard deviation and coefficient of variation (CV; ratio of standard deviation to the mean) of the traffic matrix. Additionally, the variance and CV of traffic exchanged by node-pairs, traffic emitted by sources and traffic sunk by destinations were all considered. It was found that a relationship exists between the CV of a traffic matrix and the improvement obtainable on the initial costs. Hence, for each traffic pattern considered, the mean improvement obtained by SA is plotted alongside the mean CV of the traffic matrices.

4.1. Evolution of traffic matrix amplitude

Fig. 2 shows the improvement obtained by SA alongside the mean CV of the traffic matrices, as

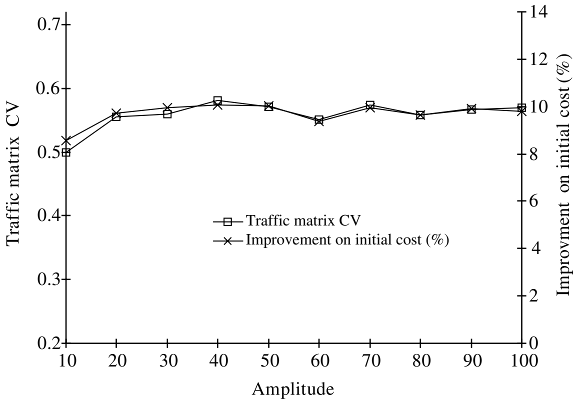

Fig. 2. Traffic matrix CV and efficacy of NPO as function of amplitude.

the traffic matrix amplitude is increased. It is apparent that the improvement on the initial solution and the traffic matrix CV are relatively unchanged with changing traffic matrix amplitude. This is an expected result since, for a uniformly distributed traffic matrix, the amplitude of the traffic matrix ought to have negligible impact of the efficacy of NPO. The results in Fig. 2 are the mean results for 100 different traffic matrices. It is evident that the improvement on the initial solution is proportional to the traffic matrix CV.

4.2. Evolution of traffic matrix range

Fig. 3 shows the impact of varying the traffic matrix range on the optimisation gain and the traffic matrix CV. The mean traffic level was kept constant at 100 units. The results in Fig. 3 are the average of a 100 unique traffic matrices. From Fig. 3 it may be seen that, unsurprisingly, the traffic matrix CV increases linearly with the traffic matrix range. Additionally, the improvement obtained by SA also linearly rises with traffic matrix range. Once again, the traffic matrix CV is proportional to the improvement obtained on the initial random solution for this traffic pattern. Comparing Fig. 2 with Fig. 3, it is evident that, even though the traffic matrix is uniformly randomly distributed in both cases, varying the traffic matrix range has a profoundly different effect to varying the amplitude.

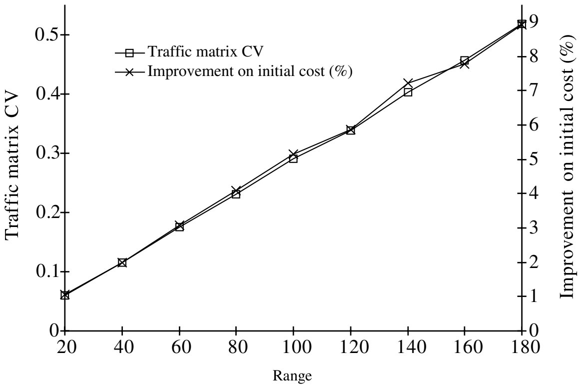

Fig. 3. Traffic matrix CV and efficacy of NPO as function of range.

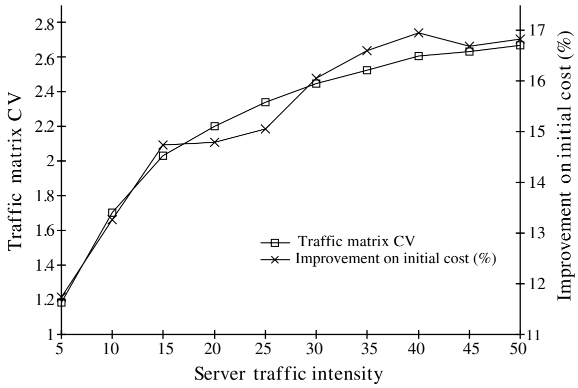

Fig. 4. Traffic matrix CV and efficacy of NPO as function of server traffic intensity.

The trend in Fig. 3 may be readily explained. If the traffic is uniform, then regardless of the arrangement of traffic matrix sources/destinations onto the MSN, the same traffic-weighted internodal distances will always be obtained. Hence, it will be expected that the initial and optimised costs will all have the same numerical value, i.e. no improvement is obtained on the initial solution. Now, when the range of the traffic matrix is small, say 20, relatively speaking, this range is not excessively dissimilar to the uniform traffic case and so the optimisation gain is small. As the range increases, the traffic matrix becomes increasingly dissimilar to the *uniform traffic* case and increasingly similar to the *uniformly randomly distributed traffic* case between 1 and an upper limit. This hypothesis may be readily verified by observing that the values for a range of 180 (Fig. 3) are similar to the values obtained when the amplitude is varied (Fig. 2).

4.3. Evolution of server traffic intensity

Different traffic matrices, with the location of the server node randomly chosen each time, were used to study the effect of the evolution of the server traffic intensity on NPO. The mean results when these 200 traffic matrices are arranged onto an MSN are shown in Fig. 4. Fig. 4 suggests that, once again, the traffic matrix CV is indicative of the improvement obtainable on a typical random initial solution.

Somewhat surprisingly, Fig. 4 shows that, as the server traffic intensity increases, the optimisation gain and the traffic matrix CV have a logarithmic trend. This trend is because for large server traffic intensities, the server's traffic saturates the overall traffic matrix and so negligible improvements are obtained by subsequent increases in the server traffic intensity, i.e. nodes have been positioned optimally relative to the server and scaling the traffic exchanged with the server has no impact on this optimal node arrangement. On the other hand, for lower server traffic intensities, server traffic is less dominant and so improvements are achievable. The results for the server may be summarised as follows: for larger server traffic intensities, the optimal node placement is effectively the optimal placement of nodes relative to the server; an arrangement that does not change with further rises in the server traffic intensity. However, for lower server traffic intensities, the optimal node placement is not necessarily the optimal placement with respect to the server and so improvements may be obtained with small initial increases in the server traffic intensity.

4.4. Evolution of hot spot traffic intensity

The results obtained when the hot spot traffic intensity was varied are presented in Fig. 5. The hot spot locations were randomly chosen in each of the 200 traffic matrices considered.

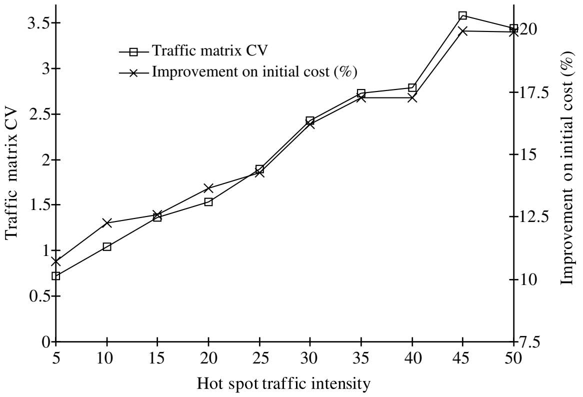

Fig. 5. Traffic matrix CV and efficacy of NPO as function of hot spot traffic intensity.

Fig. 5 suggests that the traffic matrix CV and the improvement obtained on a typical random solution increase with increasing hot spot traffic intensity. In Section 4.3, it was suggested that once the optimal arrangement of nodes relative to the server is realised, subsequent increases in the server traffic intensity will not result in a corresponding increase in the efficacy of NPO, provided the server traffic intensity is large enough to dominate the traffic matrix. Hence, for similar reasons, it would be expected that the hot spot traffic intensity will reach a sufficiently large value and dominate the traffic matrix. After this point, further increases in the hot spot traffic intensity will not significantly improve the solutions. It seems fair to say that the hot spot traffic intensities of 5 to 50 used are not sufficiently large for this phenomenon to be observed. In order to verify this hypothesis, NPO was conducted on a number of different traffic matrices with a hot spot traffic intensity of 500. The mean improvement obtained for hot spot traffic intensities of 5, 50 and 500 are presented in Table 1. From Table 1, it is evident that the improvement

Table 1
Optimisation gain a function of hot spot traffic intensity

Hot spot traffic intensity	Improvement (%)
5	10.698
50	19.879
500	35.937

on the initial cost approximately doubles with each order of magnitude increase in the hot spot traffic factor. Hence, the hypothesis is correct; as the hot spot traffic factor is increased, the improvement on the initial cost obtainable has a logarithmic trend.

An important finding in this paper is that when a non-uniformity (such as a server or hot spot) is present in a traffic matrix, linearly increasing the intensity of the non-uniformity does not lead to a continuous improvement in the optimisation results. This phenomenon occurs because, once sufficiently large, the non-uniformity saturates the traffic matrix and increasing the intensity of the non-uniformity further has negligible impact on the relative NPO result.

5. Efficacy of node placement optimisation

The results presented so far have established that the improvement on the initial cost is proportional to the CV of the traffic matrix, regardless of the underlying traffic pattern; albeit with different constants of proportionality. This is a significant finding as it allows general statements to be made about the efficacy of conducting NPO on a given traffic matrix based on the computation of a relatively straightforward parameter. If the traffic matrix CV is small, little will be gained by conducting NPO. On the other hand, NPO will produce significant yields for a traffic matrix with a more substantial CV. Furthermore, identifying the traffic matrix CV as the salient parameter in determining the efficacy of NPO has significant repercussions for traffic matrix scalability or change. If the traffic changes but there is negligible change to the traffic matrix CV, then it is known that the relative improvement on the initial cost for the new traffic matrix will be similar to that for the old traffic matrix. On the other hand, any change in the traffic matrix that affects the traffic matrix CV is known to change the optimal gain significantly.

In order to investigate the relationship between the traffic matrix CV and efficacy of NPO further, a heterogenous mix of traffic matrices was randomly generated; these included traffic matrices complying to the patterns described in Section 3.2

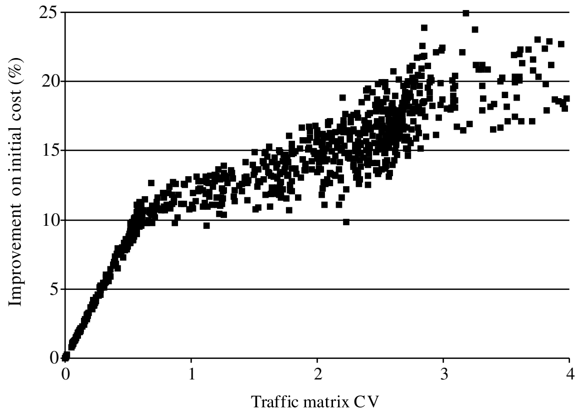

Fig. 6. Traffic matrix CV and efficacy of NPO for heterogenous mix of traffic matrices.

and matrices derived from arbitrary combinations of these patterns. Nevertheless, critically, these matrices were not segregated according to the underlying traffic pattern. Rather, each traffic matrix was characterised by only the CV of inter-nodal traffic. The relationship between the traffic matrix CV and the improvement obtained on the initial solution was studied. Fig. 6 shows the results obtained for a heterogenous mix of over 900 traffic matrices.

Numerous deductions may be drawn from Fig. 6. Firstly, Fig. 6 confirms the previous observations that the traffic matrix CV is broadly indicative of the efficacy of NPO; if the traffic matrix CV is small, little will be gained by conducting NPO and on the other hand, NPO will produce significant yields for a traffic matrix with a more substantial CV.

Somewhat surprising, Fig. 6 suggests that the impact of the traffic matrix CV on the efficacy of NPO may be separated into two distinctive regions. If, for example, best fitting straight lines are drawn, then in the first region, the line will have a steep gradient and the points will have a low variance around this line. In the second region, however, the points are more dispersed around the line and the gradient of the line will be less.

The increased dispersion of points around a best-fitting straight line with increasing traffic matrix CV may be explained by realising that low values in traffic matrix CV means many traffic matrix source–destination pairs will exchange near-identical volumes of traffic, yielding similar traffic matrices, which in turn produce comparable NPO results. However, significantly different traffic matrices may have similar larger values of traffic matrix CV, hence a greater dispersion in the NPO results are obtained. Essentially, there are fewer ways a traffic matrix can have a low CV than have a high CV.

Previous results from Section 4 are recalled to explain the different gradients of best-fitting straight lines. It was noted from Figs. 2–5 that for each traffic pattern, the efficacy of NPO was proportional to the traffic matrix CV. However, critically, the constant of proportionality was observed to be different. Consequently, it will be expected that, if the efficacy of NPO is plotted against the traffic matrix CV for a mix of traffic matrices derived from these underlying patterns, the result will be a superposition of lines, where each line will correspond to an individual traffic pattern. When the traffic matrices in Fig. 6 were segregated according to the underlying traffic patterns, it was found the points could indeed be well represented by a single best-fit line. However, when the traffic matrices are all mixed, it is impossible to distinguish the individual lines (the CV of the traffic matrices derived from the different traffic patterns overlap); with one crucial exception: when the range of the traffic matrix is varied (Fig. 3), the values of the traffic matrix CV (from approximately 0.05–0.55) are unique; no other traffic pattern yields traffic matrices with these values. Hence, the first steep linear region in Fig. 6 corresponds to traffic matrices in which only the range is varied. The less steep and more dispersed region is due to a mix of the other traffic patterns. It may be said, therefore, that it gets harder to predict the optimisation gain as the traffic matrix CV rises.

6. Conclusion

The MSN with the *Clockwork Routing* scheme is an attractive architecture for optical packet switching. The optimal placement of traffic matrix nodes in such an architecture, NPO, has been studied in this paper. Motivated by the intractable nature of the problem, simulated annealing was successfully applied to NPO. A number of traffic patterns were considered and the impact of evolving properties of traffic matrices studied. It was found that a change in the amplitude of a traffic matrix has negligible impact on the relative results of NPO. An increase in the range of a traffic matrix increases the optimisation gain. If a "non-uniformity" such as a server or traffic hot spot exists, it was discovered that an increase in the relative intensity of the non-uniformity leads to an increase in the quality of the optimisation results, up to a certain point. Thereafter, the traffic from the non-uniformity will "saturate" the traffic matrix and only marginal further relative improvements will be obtained.

Conducting NPO for the different traffic patterns provided insight into the innate parameter of traffic matrices that affects the efficacy of NPO. A significant finding of this paper is that the improvement obtainable on a (randomly generated) initial solution is dependent upon the CV of the traffic matrix. In summary, it may be said that for any given traffic matrix, if the traffic matrix CV is low, little is gained by performing NPO but it is possible to estimate the improvement obtainable on an initial random solution beforehand. However, in contrast, if the traffic matrix CV is large, greater gains will be obtained by NPO although it is more difficult to predict these accurately.

References

[1] N.F. Maxemchuk, The Manhattan Street Network, in: Proceedings of IEEE Globecom, 1985.

[2] N.F. Maxemchuk, Routing in the Manhattan street network, IEEE Transactions on Communications 35 (5) (1987) 503–512.

[3] T. Robertazzi, A. Lazar, Deflection strategies for the Manhattan street network, in: Proceedings of IEEE International Conference on Communications (ICC), 1991.

[4] F. Forghieri, A. Bononi, P. Prucnal, Analysis and comparison of hot potato and single buffer deflection routing in very high bit rate optical mesh networks, IEEE Transactions on Communications 43 (1) (1995) 88–98.

[5] F. Chevalier, D. Cotter, D. Harle, A new packet routing strategy for ultra fast photonic networks, in: Proceedings of IEEE Globecom, 1998.

[6] F. Chevalier, D. Cotter, D. Harle, Routing and fairness issues associated with a self-routing control strategy in regular mesh photonic networks, in: Proceedings of International Teletraffic Congress (ITC), 1999.

[7] D. Cotter, J.K. Lucek, High-speed digital optical processing in future networks, Philosophical Transactions of the Royal Society 358 (2000) 2283–2296.

[8] M.G. Hluchyj, M.J. Karol, Shufflenet: an application of generalized perfect shuffles to multihop lightwave networks, IEEE/OSA Journal of Lightwave Technology 9 (10) (1991) 1386–1397.

[9] B. Mukherjee, WDM-based lightwave networks. Part II. Multihop systems, IEEE Network (July 1992) 20–32.

[10] K.N. Sivarajan, R. Ramaswami, Lightwave networks based on de Bruijn graphs, IEEE/ACM Transactions on Networking 2 (1) (1994) 70–79.

[11] B. Mukherjee, S. Ramamurthy, D. Banerjee, A. Mukherjee, Some principles for designing a wide-area WDM optical network, in: Proceedings of IEEE Infocom, 1994.

[12] I. Chlamtac, A. Ganz, G. Karmi, Lightnets: topologies for high-speed optical networks, IEEE/OSA Journal of Lightwave Technology 11 (5/6) (1993) 951–961.

[13] C.P. Ravikumar, T. Rai, V. Verma, Kautz graphs as attractive logical topologies in multihop lightwave networks, Computer Communications 20 (1997) 1259–1270.

[14] O. Komolafe, D. Harle, D. Cotter, Ultrafast optical packet switching over arbitrary physical topologies using the Manhattan street network, in: Proceedings of IEEE International Conference on Communications (ICC), 2001.

[15] O. Komolafe, High-speed optical packet switching over arbitrary physical topologies using the Manhattan street network, Ph.D. Thesis, University of Strathclyde, November 2001.

[16] O. Komolafe, D. Harle, D. Cotter, Impact of graph theoretic network parameters on the design of multi-hop regular virtual topologies for optical packet switching, in: Proceedings of IEEE International Conference on Communications (ICC), 2002.

[17] O. Komolafe, D. Harle, D. Cotter, Impact of non-uniform traffic on the design of multi-hop regular virtual topologies for optical packet switching over arbitrary physical topologies, IEEE/OSA Journal of Lightwave Technology 20 (8) (2002) 1248–1263.

[18] J.P. Labourdette, A.S. Acompora, Logically rearrangeable multihop lightwave networks, IEEE Transactions on Communications 39 (8) (1991) 1223–1230.

[19] S. Banerjee, B. Mukherjee, Algorithms for optimized node arrangements in ShuffleNet based multihop lightwave networks, in: Proceedings of IEEE Infocom, 1993.

[20] S. Banerjee, B. Mukherjee, D. Sarkar, Heuristic algorithms for constructing optimized structures in linear multihop

lightwave networks, IEEE Transactions on Communications 42 (2/3/4) (1994) 1811–1826.
[21] K.L. Yeung, T.P. Yum, Node placement optimization in shufflenets, IEEE/ACM Transactions on Networking 6 (3) (1998) 319–324.
[22] A. Narula-Tam, E. Modiano, Dynamic load balancing in WDM packet networks with and without wavelength constraints, IEEE Journal on Selected Areas in Communications 38 (10) (2000) 1972–1979.
[23] F. Siu, R.K.C. Chang, Effectiveness of optimal node assignments in wavelength division multiplexing networks with fixed regular virtual topologies, Computer Networks 38 (2002) 61–74.
[24] S. Kirkpatrick, C.D. Gerlatt, M.P. Vecchi, Optimization by simulated annealing, Science 220 (4598) (1983) 671–680.
[25] R.W. Eglese, Simulated annealing: a tool for operational research, European Journal of Operational Research 46 (1990) 271–281.
[26] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, E. Teller, Equation of state calculations by fast computing machines, Journal of Chemical Physics 21 (1953) 1087–1092.
[27] B. Hajek, Cooling schedules for optimal annealing, Mathematics of Operations Research 13 (1988) 311–329.
[28] S.P. Brooks, B. Morgan, Optimization using simulated annealing, The Statistician 44 (2) (1995) 241–257.
[29] J.H. Holland, Adaptation in Natural and Artificial Systems, University of Michigan Press, Ann Arbor, MI, 1975.

Olufemi Komolafe was born in Glasgow, Scotland in 1977 and was schooled in Nigeria and Scotland. He obtained a BEng (with First Class Honours) in Electronic and Electrical Engineering (EEE) from the University of Strathclyde, Glasgow in 1998. He received a Ph.D. from the same institution for research in optical network design in 2001. Thereafter, he worked as a research fellow within the EEE Department and then as a researcher at Agilent Laboratories, Scotland. He is currently employed as a research assistant in the Department of Computing Science at the University of Glasgow, Scotland. His current research interests include IP/optical networking, GMPLS, Internet traffic modelling, optical packet switching and network design and management.

David Harle is currently a senior lecturer within the Broadband and Optical Networks Group within the Department of Electronic and Electrical Engineering at the University of Strathclyde, Glasgow. He received his Ph.D. in Integrated Telecommunications from the University of Strathclyde in 1990, having previously been a Research Assistant in the same department. His current research interests within the Broadband and Optical Networks group focus upon performance evaluation, design and management issues associated with current and future broadband and optical communication systems. The author of over 50 research papers and undergraduate texts, He is also a member of the IEE and IEEE.