

DESIGNING A MULTI-HOP REGULAR VIRTUAL TOPOLOGY FOR ULTRAFAST OPTICAL PACKET SWITCHING: NODE PLACEMENT OPTIMISATION AND/OR DILATION MINIMISATION?

Olufemi Komolafe¹, David Harle¹, David Cotter²

1 – Dept. of Electronic & Electrical Engineering, University of Strathclyde, Glasgow G1 1XW, UK.

2 – Corning Research Centre, Adastral Park, Martlesham Heath, Ipswich, IP5 3RE, UK.

femi@comms.eee.strath.ac.uk

ABSTRACT

This paper studies the design of multi-hop regular virtual topologies to facilitate optical packet switching in networks with arbitrary physical topologies. The inputs to the virtual topology design problem are the physical topology, the traffic matrix and the regular topology. In this paper, this problem is tackled directly and also by decomposition into two sub-problems. The first sub-problem, dilation minimisation, uses only the physical topology and the virtual topology as optimisation inputs. The second sub-problem considers the traffic matrix and virtual topology as optimisation inputs. The solutions of these two sub-problems are compared with each other and against the results obtained when the global problem is optimised (using all three possible input parameters) for a variety of traffic scenarios. This gives insight into the key question of whether the physical topology or the traffic matrix is the more important parameter when designing a regular virtual topology for optical packet switching. Regardless of the approach taken the problem is intractable and hence heuristics must be used to find (near) optimal solutions in reasonable time. Five different optimisation heuristics, using different artificial intelligence techniques, are employed in this paper. The results obtained by the heuristics for the three alternative design approaches are compared under a variety of traffic scenarios. An important conclusion of this paper is that the traffic matrix plays a less significant role than is conventionally assumed, and only a marginal penalty is incurred by disregarding it in several of the traffic cases considered.

1. INTRODUCTION

The Manhattan Street Network [1] is one of several multi-processor interconnection architectures that have been proposed for use in multi-hop optical packet switched networks [2]. Typically, these architectures offer simple and distributed routing schemes, the possibility of ingenious ways to avoid (or minimise) the use of optical buffering, and increased predictability of the network performance. A new routing scheme, *Clockwork Routing*, has been proposed for the Manhattan Street Network (MSN) that is particularly well-suited for optical packet switching [3]. The routing processing is extremely simple and suitable for optical implementation, no optical buffering is required, no resequencing is needed at the destination nodes, and throughput is comparable with conventional store-and-forward packet switching [3].

2. PROBLEM DESCRIPTION

Using the Manhattan Street Network as an example, the paper addresses the deployment of multi-processor interconnection architectures in arbitrary physical topologies. The architectures are deployed (or embedded) as regular virtual topologies. Virtual topology nodes are mapped onto physical topology nodes. WDM channels, lightpaths, are established between the appropriate nodes to realise the desired regular virtual topology connectivity

[4,5,6]. Three inputs to the problem are the regular virtual topology, the physical topology and the traffic matrix. The objective is to find the mapping of regular virtual topology nodes that is optimal with respect to a defined cost. The cost that is considered in the paper is the mean traffic-weighted embedded inter-nodal distance. This cost is important as it embraces two important metrics: the embedded inter-nodal distance, which indicates the number of optical cross-connects packets encounter between source and destination [6], and the traffic-weighted inter-nodal distance, which is indicative of delay [8]. Three alternative design approaches to optimise this cost are studied in the paper and are illustrated in Figure 1. (A 4x4 MSN is also shown in Figure 1.) The first virtual topology design approach is simply to attempt to optimise the traffic-weighted embedded inter-nodal distance directly, using the virtual topology, the physical topology and the traffic matrix as the inputs. The other two approaches are based on decompositions of this problem into different sub-problems.

The first sub-problem, *dilation minimisation* [7], is to try to find the mapping of regular virtual topology nodes onto the physical topology that minimises the mean lightpath length [6]. In this case, the only inputs into the optimisation are the physical topology and the virtual topology (and not the traffic matrix). The second sub-problem, on the other hand, considers only the virtual topology and the traffic matrix (and not the physical topology). This approach is known as *node placement optimisation* [8] and seeks to allocate traffic matrix sources/destinations to virtual topology nodes so that the traffic-weighted inter-nodal distances over the virtual topology are minimised.

In this paper, the virtual topology design problem is decomposed into these two sub-problems, the results of which are then compared against each other and against the global optimisation result. The motivation behind this approach is to ascertain whether and when the physical topology or the traffic matrix is the more important input parameter when designing a multi-hop regular virtual topology to carry packet traffic. Ideally, there are three different input parameters into the optimisation process, but can an (almost) equally good result be obtained by using only two input parameters? Which two? The significance of the results presented in this paper derives from the fact that the traffic matrix typically will be dynamic; determination of the particular conditions when the traffic matrix may be disregarded removes the need for the non-trivial task of trying to find a “representative” traffic matrix. Alternatively, the results presented in this paper may be interpreted as pertaining to the scalability of the solution proposed: if a (near) optimal deployment of a regular virtual topology in a physical topology is obtained without considering the traffic matrix, under what traffic conditions (if any) does this embedding become inferior to an embedding found considering the current traffic matrix? And by how much?

Irrespective of the choice of the design approach adopted, the problem is intractable. Hence, different artificial intelligence based heuristics are used to find (near) optimal solutions in reasonable time. The five optimisation heuristics used in the paper are two implementations of genetic algorithms, simulated annealing, hill-climbing and random search. The two implementations of genetic algorithms have been previously used to deploy the MSN [9] and are known as “Partially Mapped Crossover” (PMX) and “Cycle Crossover” (CX). The use of these diverse heuristics, which explore the search space in such different ways, engenders confidence in the results and trends presented in this paper.

The Manhattan Street Network with *Clockwork Routing* is used as the exemplar multi-hop regular virtual topology, and NSFnet [5] (shown in Figure 1) has been arbitrarily chosen for

the physical topology. For verification, the experiments conducted were also repeated for the deployment of a 6x6 MSN in a 36-node arbitrary physical topology.

3. TRAFFIC SCENARIOS

Dilation minimisation was conducted by using each heuristic to find the mapping of the MSN nodes which minimises the mean lightpath length. Thereafter, the traffic-weighted embedded inter-nodal distances that this mapping produced was computed for different stages in each traffic scenario. This result was compared with the corresponding value when the heuristics were applied to both node placement optimisation and the optimisation of the traffic-weighted embedded inter-nodal distances directly.

In each traffic scenario, nodes send each other a volume of traffic uniformly randomly distributed, before the traffic matrix was modified as required. The four different traffic scenarios studied are:

- Variation of the traffic matrix amplitude.
- Variation of traffic matrix range, with the mean traffic level kept constant.
- Existence of a server within the network, whereby this node sinks/produces X times more traffic than any other node. The impact of X , the server traffic intensity, is investigated.
- Presence of a traffic “hot spot” between a particular node-pair, whereby traffic between a particular pair is Y times greater than they would otherwise send each other. The effect of varying Y , the node-pair traffic intensity, was investigated

4. RESULTS

The performance of the heuristics when designing the virtual topology using each of the three approaches illustrated in Figure 1 are presented in this section. For simplicity it is assumed all fibres have equal physical length, i.e. unity. In the findings presented, the result of optimising the traffic-weighted embedded inter-nodal distance directly is always represented by a continuous line labelled with the optimisation heuristic and $DM+NPO$. Similarly, dilation minimisation results are shown by a dashed line and are labelled with the heuristic and DM . Lastly, node placement optimisation results use a line with longer dashes and are labelled with the corresponding heuristic and NPO . The different heuristics are abbreviated as follows: hill climbing (HC), simulated annealing (SA), first genetic algorithm implementation; partially mapped crossover (PMX), the second genetic algorithm implementation; cycle crossover (CX) and lastly random search (RS). Thus, for example, a series labelled $SA(DM)$ will show the traffic-weighted embedded inter-nodal distance obtained when simulated annealing is used to find the mapping of MSN nodes onto physical topology nodes that minimises the dilation, i.e. the mean lightpath length.

4.1 Variation of Traffic Matrix Amplitude

Figure 2 shows the results obtained as the traffic matrix amplitude was varied. Similar trends could be observed in all five optimisation heuristics used. For clarity, only three are shown in Figure 2. Figure 2 shows, unsurprisingly, that the best results are obtained when the traffic-weighted embedded inter-nodal distances are optimised directly (i.e. the $DM+NPO$ series). However, it can be seen that dilation minimisation (DM) gives significantly better results than carrying out node placement optimisation (NPO), and in fact typically gives results which are comparable with those when the traffic-weighted embedded inter-nodal distance is optimised directly. This suggests that when designing a multi-hop regular virtual topology, if the traffic matrix is uniformly randomly distributed, it makes less of an impact in the design process than the physical topology, regardless of the traffic matrix amplitude.

4.2 Variation of Traffic Matrix Range

The results obtained when the mean traffic nodes transfer is kept constant but the range varied is illustrated in Figure 3. Figure 3 shows that a change in traffic matrix range makes negligible difference to the performance of traffic-weighted embedded inter-nodal distance optimisation ($DM+NPO$) and mean lightpath length minimisation (DM). Furthermore, the results of these two approaches are comparable, regardless of the range of the traffic matrix. Significant fluctuations exist in the results for node placement optimisation. (Similar fluctuations exist in Figure 2 but are more clearly seen in Figure 3 because of the scale of the axis.) The fluctuations are due to the fact that virtual topology designs which are comparably good in terms of traffic-weighted inter-nodal distances may give significantly different result when embedded in the physical topology since, when embedding in the physical topology, links in the Manhattan Street Network are perturbed by different arbitrary amounts. The results for node placement optimisation (NPO) is significantly higher than for the other two approaches. However, the relative performance of node placement optimisation seems to improve with increasing range since, despite the fluctuations, there appears to be a slight downward trend in the results for all the heuristics. This decrease is because as the range of traffic matrix increase, a (disproportionately) large benefit is obtained from placing nodes which send each other large volumes of traffic near each other in the MSN.

In summary, it can be said that the dilation minimisation sub-problem (and not the node placement optimisation sub-problem) dominates traffic-weighted embedded inter-nodal distance optimisation if the traffic matrix range is varied while the mean is kept constant. Furthermore, it has been seen that the dilation minimisation sub-problem is a good approximation for traffic-weighted embedded inter-nodal optimisation.

4.3 Variation of Server Traffic Intensity

A node was randomly selected to act as a server and the impact of the variation of the server traffic factor is shown in Figure 4. Once again, results of only three of the five heuristics have been shown for clarity. Figure 4 indicates that node placement optimisation yields the worst results. Dilation minimisation, on the other hand, gives results which are slightly greater than those for direct traffic-weighted embedded inter-nodal distance optimisation. Consequently, the dilation minimisation sub-problem is more significant than the node placement optimisation sub-problem when a server exists. However, a greater difference exists between the dilation minimisation results and the direct traffic-weighted embedded inter-nodal distance optimisation results than had been observed for variations to the traffic matrix amplitude and range.

4.4 Variation of Node-Pair Traffic Intensity

From Figure 5, it is evident that minimising the traffic-weighted embedded inter-nodal distance directly, i.e. the $DM+NPO$ series, gives the best results for the heuristics, with the sole exception of random search. It can be seen that the lightpath length minimising solutions, i.e. the DM series, tend to give progressively worse results as the node-pair traffic intensity increases. Two important exceptions are the results for random search ($RS(DM)$) and simulated annealing ($SA(DM)$) – the quality of these results does not tend to deteriorate relative to the traffic-weighted embedded inter-nodal distance optimisation results with increasing node-pair traffic intensity. Interestingly, the node placement optimisation embeddings (NPO), initially give the worst results, but the rate of increase appears to be less than the results for dilation minimisation and comparable with traffic-weighted embedded inter-nodal distance optimisation. The trend observed indicates that node placement

optimisation will outperform several of the mean lightpath length minimising results as the node-pair intensity exceeds 50. This hypothesis was found to be true by examining the results when the node-pair intensity was greater than 50. These observations agree with intuition since we would expect that, as the volume of traffic between a particular node pair increases, it becomes increasingly beneficial to position the two nodes next to each other in the MSN rather than to minimise the global mean dilation of the links of the MSN in the physical topology.

An interesting feature of all the results presented in this section is that certain dilation minimising solutions, i.e. $SA(DM)$ and $RS(DM)$, perform significantly better than others and exhibit similar traits to the traffic-weighted embedded inter-nodal distance optimisation. These somewhat anomalous results can be understood by considering what happens within the Manhattan Street Network. Since the node-pair transfer a relatively large volume of traffic, it is evident that to minimise the traffic-weighted embedded inter-nodal distance, the nodes ought to be mapped onto nodes in the MSN whose round-trip inter-nodal distance is minimised and ideally, the links traversed between the two nodes dilated by the minimum amount. The minimum round-trip inter-nodal distance between two nodes in the MSN is four hops. It was found that in a 4x4 MSN a round trip inter-nodal distance of 4 hops exists between 53.33% of all node-pairs. Consequently, there is a greater than 50% chance of the two hot spot nodes being fortuitously placed the optimal distance apart, even if the traffic matrix is not considered in the optimisation. This possibility explains why some of the mean lightpath length minimising embeddings obtain relatively good results. The actual embeddings were examined and it was found that they did indeed position the two nodes within the MSN such as the round-trip inter-nodal distance is four.

Evidently, the dilation minimisation sub-problem dominates traffic-weighted embedded inter-nodal distance optimisation for relatively small node-pair traffic intensities. However, the relative significance of the node placement optimisation sub-problem increases with increasing node-pair traffic intensity.

5. CONCLUSION

The paper addresses the critical issue of designing multi-hop regular virtual topologies for optical packet switching in networks with arbitrary physical topologies. The Manhattan Street Network has been used as an exemplar regular topology. Due to the intractable nature of the virtual topology design problem, five alternative optimisation heuristics are employed. The virtual topology design problem is decomposed into two sub-problems, dilation minimisation and node placement optimisation, which respectively use only the physical topology and virtual topology, or the virtual topology and traffic matrix as optimisation inputs. These sub-problems are compared with each other and the direct optimisation of the global cost under a variety of traffic scenarios. It was found when the amplitude or range of a uniformly randomly distributed traffic matrix was varied, dilation minimisation is the weightier sub-problem as it outperforms node placement optimisation and performs comparable to direct optimisation of the final cost. Similar observations were made when a node is randomly chosen to act as a server. Therefore, in each of these three traffic scenarios, the dilation minimisation sub-problem dominates the virtual topology design problem. In other words, consideration of the physical topology and the virtual topology typically outperforms usage of only the virtual topology and traffic matrix. From the results presented for a traffic hot spot between a randomly chosen node-pair, it can be concluded that the dilation minimisation sub-problem is still dominant for relatively small node-pair traffic

intensities. However, the relative significance of the node placement optimisation sub-problem increases with increasing node-pair traffic intensity.

The obvious answer to the question posed in the paper title is “*and*”. However if a choice is mandated, the answer is then “*dilation minimisation*” for the traffic scenarios considered.

REFERENCES

- [1] “The Manhattan Street Network,” N.F. Maxemchuk, Proc. IEEE Globecom 1985.
- [2] “WDM-Based Lightwave Networks Part II: Multihop Systems,” B. Mukherjee, IEEE Network, July 1992.
- [3] “A New Packet Routing Strategy for Ultra-Fast Photonic Networks,” F. Chevalier, D. Cotter and D. Harle, Proceedings of IEEE Globecom 1998.
- [4] “Lightnets: Topologies for High-Speed Optical Networks,” I. Chlamtac, A. Ganz, and G. Karmi, IEEE Journal of Lightwave Technology, Vol. 11, May/June 1993.
- [5] “Some Principles for Designing a Wide-Area Optical Network,” B. Mukherjee, S. Ramamurthy, D. Banerjee and A. Mukherjee, Proceedings of IEEE Infocom 1994.
- [6] “Ultrafast Optical Packet Switching over Arbitrary Physical Topologies using the Manhattan Street Network,” O. Komolafe, D. Harle and D. Cotter, Proceedings of IEEE ICC 2001.
- [7] “Data Encodings and Their Costs,” A. Rosenberg, Acta Informatica, Volume 9, 1978.
- [8] “Node Placement Optimization in ShuffleNets,” K. Yeung and T. Yum, IEEE Transactions on Networking, Vol. 8, No. 3, June 1998.
- [9] “Optical Packet Switching over Arbitrary Physical Topologies using the Manhattan Street Network: An Evolutionary Approach,” O. Komolafe, D. Harle and D. Cotter, Proceedings of IFIP Working Conference on Optical Network Design and Modelling (ONDM), February 2001.

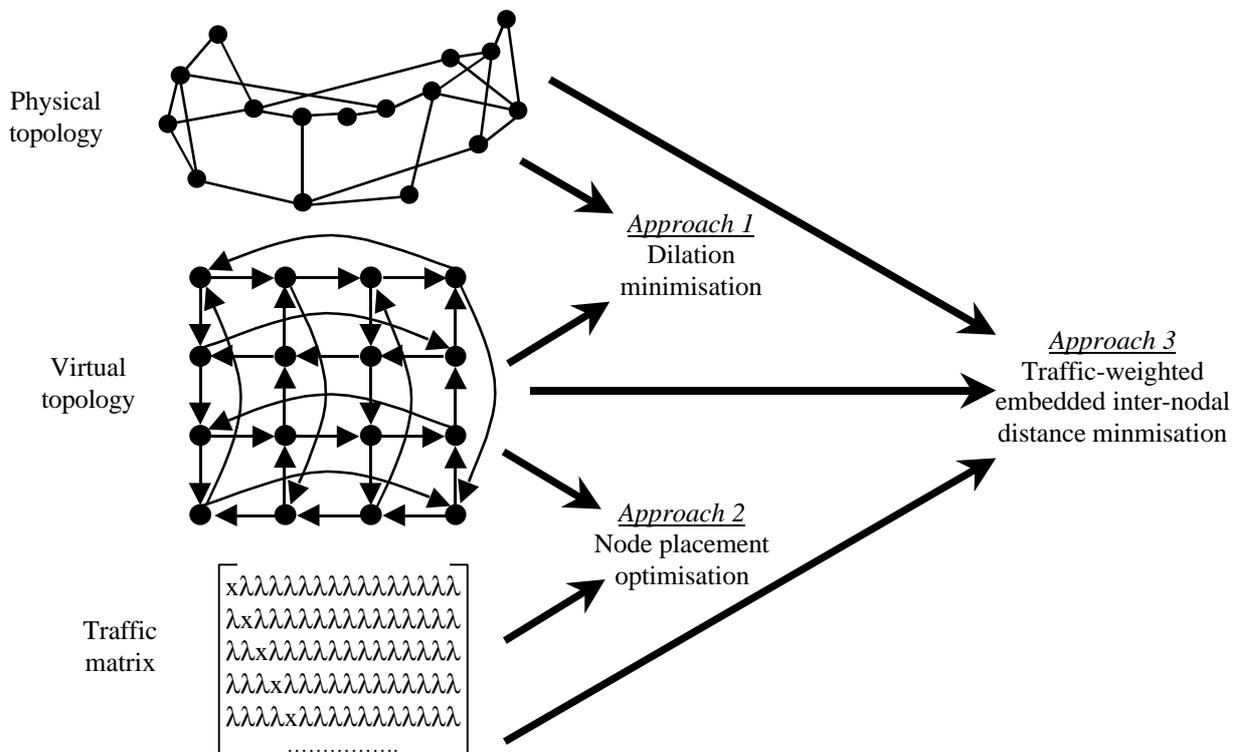


Figure 1 – Three alternative regular virtual topology design approaches

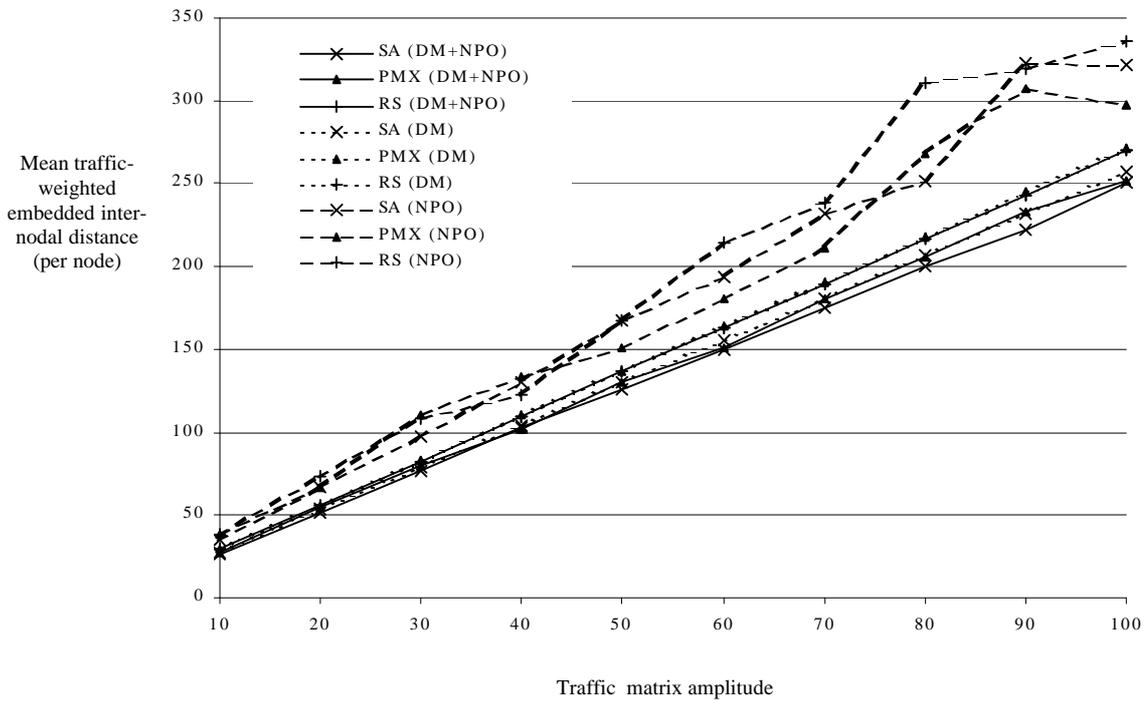


Figure 2 – Variation of traffic matrix amplitude

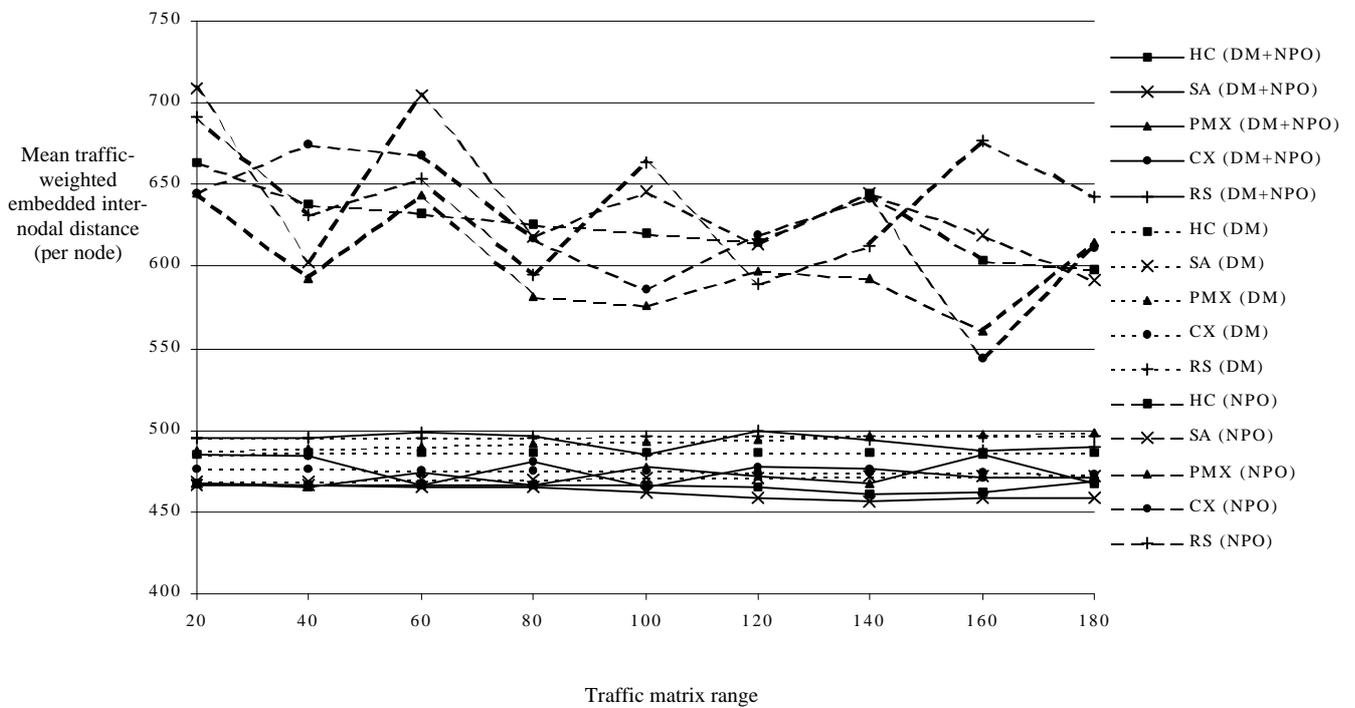


Figure 3 – Variation of traffic matrix range

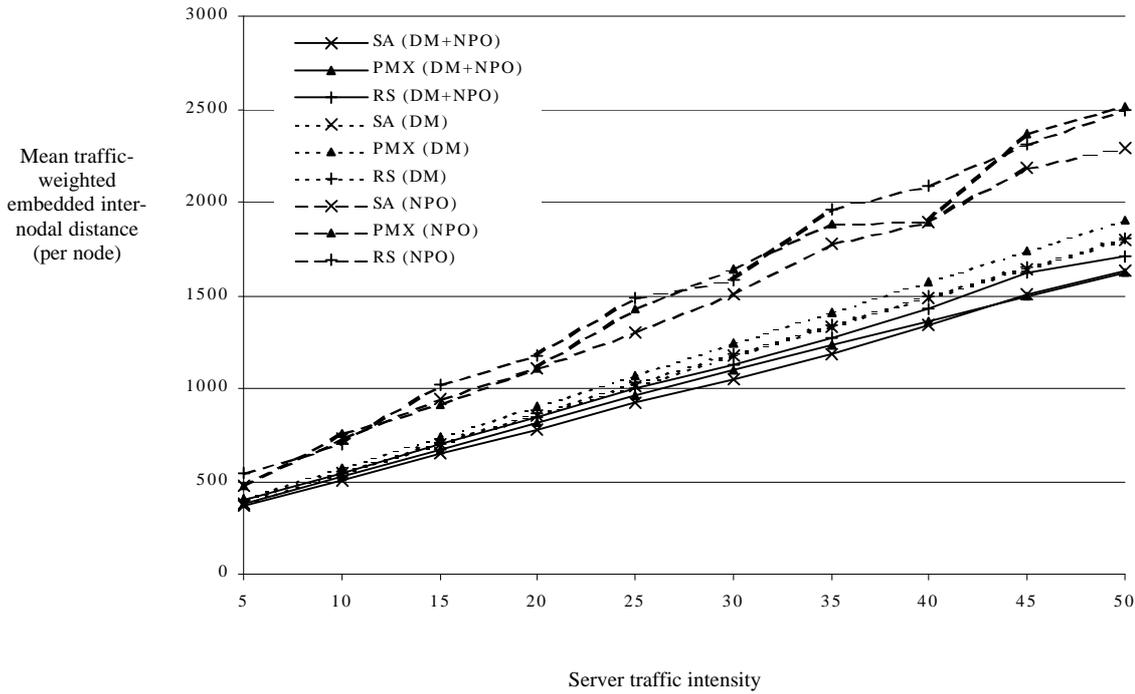


Figure 4 – Variation of server traffic intensity

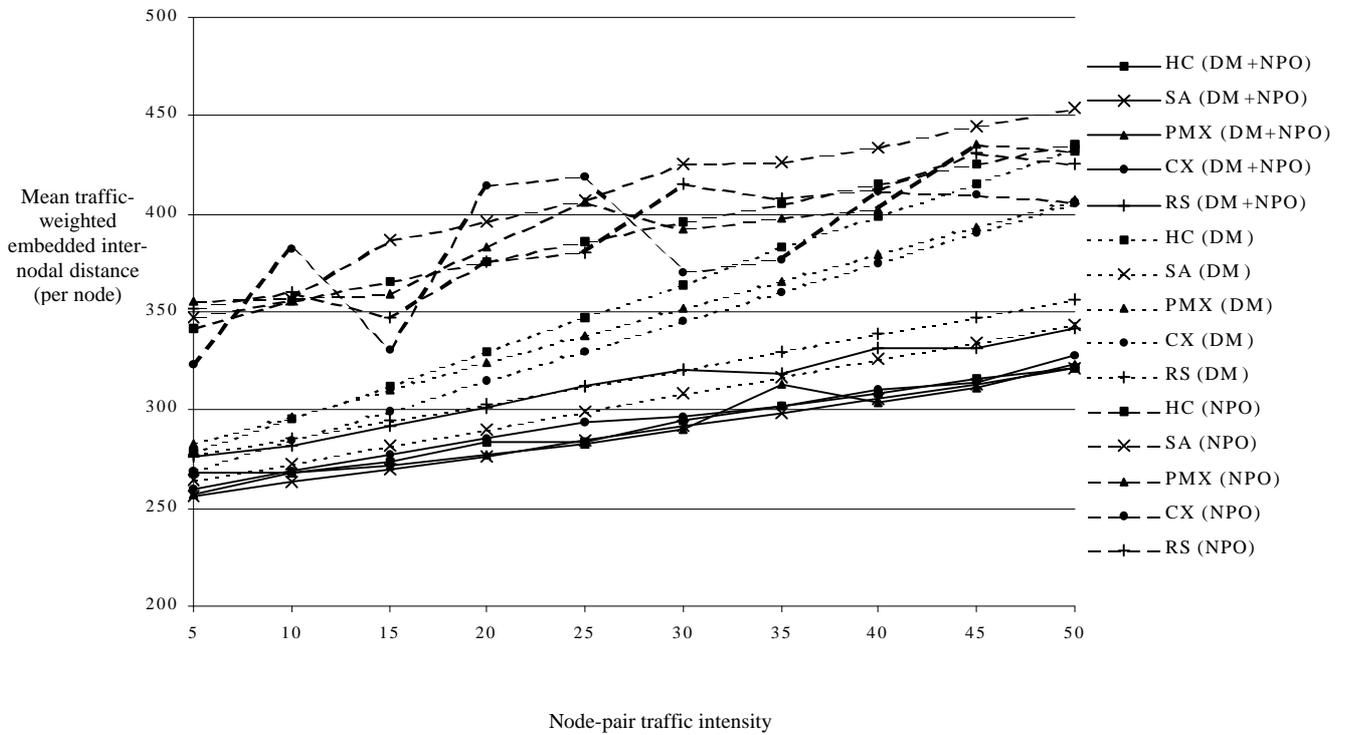


Figure 5 – Variation of node-pair traffic intensity