Path allocation for wavelength path sharing

R Friskney

University College London

Abstract: Wavelength path sharing (WPS) was introduced previously as a means of bridging the gap between the bit-rate of wavelengths and the lower bit-rate of anticipated traffic demands while maintaining some of the benefits of wavelength-routed optical networks (WRONs). Work has progressed towards finding more attractive routing algorithms within these networks. In the course of this work, some unexpected observations on the comparison of wavelength routing algorithms have arisen, which are the subject of this paper.

1. Introduction.

This paper starts by briefly describing wavelength path sharing, showing that it provides a method of flexible *optical* bandwidth-sharing of wavelength-routes, rather than using electronic multiplexers or routers. It then discusses how to choose routes for the wavelengths in the system, noting that this is a considerably more sophisticated task than conventional wavelength-routed optical networks - where the needs of only two nodes need be considered. Then, some alternative methods of performing this routing are compared for performance in a sample implementation, using a set of reference networks. Analysis then shows that entirely different answers can be found depending upon the method used for the comparison.

2. Wavelength path sharing (WPS)

Typical wavelengths at the time of writing in modern optical transport systems are 2.5Gbps or – increasingly - 10Gbps. A significant part of the motivation behind this move has been the reduced cost of owning and managing equipment to carry less, but larger channels. However, there are few common end-user applications that require this much capacity – fibre channel at 2Gbps being the fastest single common application – disregarding 10Gb ethernet as it is marketed primarily as an aggregate router-to-router technology. Thus, the oft-discussed vision of agile wavelength-routed optical networks (WRONs) with one application/wavelength seems implausible in the near term – at least one layer of intermediate multiplexing is required for reasonable efficiency. Fixed-time [2] and statistical [3] multiplexing techniques in the electronic domain are well known, but these lose the advantages of a flexible photonic network. Some of these are: reduced optical-electrical-optical conversion costs and greater flexibility with low equipment levels (due to the large number of 'free' dark wavelengths present on a lightly-loaded fibre, all accessible by suitably tunable transmission equipment), bit-rate transparency and format transparency.

Wavelength path sharing was proposed [1] to provide a lower-rate service while still maintaining photonic flexibility. Figure 1 illustrates the sharing of a unidirectional wavelength path; figure 2 illustrates a potential hardware configuration to implement it. Figure 1 shows that a unidirectional wavelength (the straight lines) from A to D that 'calls' at B and C can serve six demands (the arcs). The wavelength is statically photonically routed through Y and Z in an entirely conventional WRON fashion. Figure

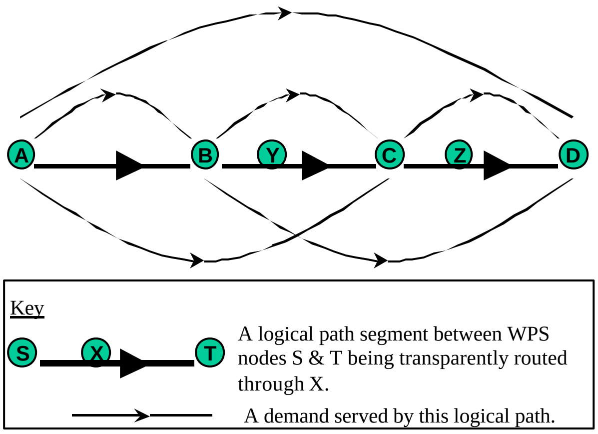

Figure 1 illustrating multiple nodes sharing one wavelength

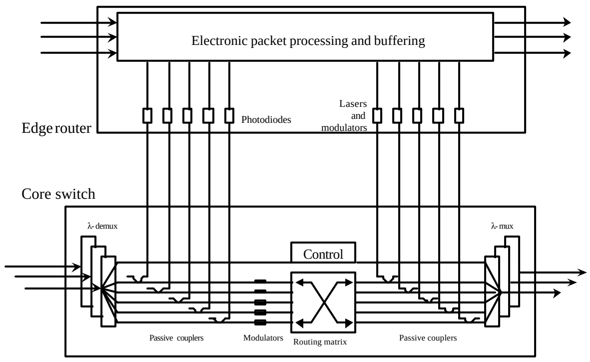

Figure 2 illustrating example (unidirectional for simplicity) hardware for WPS.

2 shows a hardware configuration such as might exist at nodes B or C: a conventional photonic cross-connect/routing matrix doing the slow-moving connectivity to establish the wavelength paths, then lasers, modulators and receivers doing the fast wavelength sharing - allowing this node to drop/inject traffic as it wishes (changing from thru to add/drop potentially as fast as on a packet-by-packet basis), or just allowing a packet/burst to glass straight through to the next participating node without interfering with it – e.g. allowing a 40G burst to pass through optically even if this node only has 2.5G transmission equipment.

3. Wavelength routing for WPS

Routing and wavelength assignment for WRONs using point-to-point demands has been the subject of much work – [4] provides a review. It will be clear that it is a more complex problem to find efficient routes between three or more nodes. However routing for WPS is even more complex than that: It is finding a wavelength that will simultaneously address the needs of (using figure 1 as an example) 6 un-served/under-served demands.

Myers provides an integer linear programming (ILP) formulation in [1] that minimises the maximum congestion on any link. As with most ILP approaches, this scales poorly in execution time with network size, failing to terminate after periods of days on a basic desktop PC for large North American networks. Thus a heuristic approach (not described here) was considered. Further, although levelling out congestion is a known approach to allow the most flexibility in later traffic, when considering capital equipment cost, the number of unique wavelengths – Baroni's N_λ in [4] – will have a more direct effect: dictating required laser tunability ranges, flat optical amplifier bandwidth etc. The heuristic gives fairly close results to the ILP formulation, but clearly inferior to it, and giving minimal efficiency gain over a non-WPS system. However, considering N_λ, using a first-fit wavelength-allocation algorithm the relationship strikingly reverses and the heuristic performs better than the ILP.

$$u_w = \{0,1\} \qquad \forall w \in 1...W$$

$$d_{pw} = \begin{cases} 1 & \text{If path p travels over wavelength } \omega \\ 0 & \text{Otherwise} \end{cases}$$

$$(Eqn.1) \qquad \sum_{p\in P} d_{pw} I(j \in p) \le 1 \qquad \forall j \qquad \forall w \in 1...W$$

$$(Eqn.2) \qquad \sum_{w\in W} d_{pw} \ge 1 \qquad \forall p \in P$$

$$(Eqn.3) \qquad \sum_{p\in P} d_{pw} \le u_w \qquad \forall w$$

Objective:

$$(Eqn.4) \qquad \min\ N_l = \sum_{w\in W} u_w$$

3.1 Wavelength assignment for WPS

As the inversion of performance between considering/not considering wavelength assignment mentioned above was so strong and unexpected, it seemed appropriate to check it. To do this, a new ILP formulation for wavelength allocation was used, derived from Baroni's formulation for static WRONs in [4].

The new formulation is given in equations 1-4, where P is the given set of wavelength routes, and 1…W is a continuous sequence of integer wavelength numbers. A brief explanation: equation 1 says that a wavelength can only be used once on each span; equation 2 says that each path must be assigned one wavelength; equation 3 says that u_w must be set if any path in the network uses wavelength w and equation 4 sets the objective to minimise the number of u_w that are set, i.e. the number of distinct wavelengths used. W must be set to a maximum value of wavelengths just large enough to allow a feasible solution, and thus for the ILP to terminate successfully.

This expands to $|P|.W + W$ variables and $W.E+|P|+|P|.W$ constraints, where E is the total count of edges in the network and $|P|$ is the number of paths being considered. Thus, the choice of W is critical in the scaling. In practical implementation, W was set to the output value of the first fit algorithm – a known feasible solution and thus definitely large enough. It is also a quite practical real-world source because first-fit executes in negligible time compared to the ILP wavelength solution. As will be seen from the results, a more aggressive choice of W would have been both possible and beneficial (in decreasing execution time). This was not investigated further, though.

4. Reference network comparisons

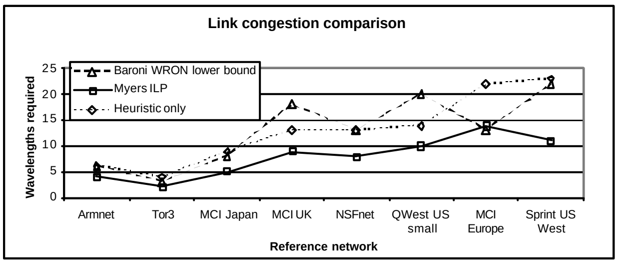

Figure 3 showing worst-case link congestion using both methods – a potential measure of traffic forecast tolerance.

Figure 3 compares the created heuristic to Myers' ILP formulation with regard to the maximum congestion on any link. The Baroni lower bound [4] for WRONs is provided as a reference – as the object of WPS is to achieve increased efficiency over WRONs. Note that this lower bound is of unique wavelengths used within the network (as in figure 4), rather than congestion, and is thus a generous comparison. Reference networks are sorted by number of nodes (increasing from left to right). Figure 3 shows that the Myers ILP does consistently significantly better than the heuristic. The heuristic achieves roughly comparable results to the WRON lower bound. This would seem to indicate that there is little point in using WPS with the heuristic technique, provided a reasonably rapid solution to the much simpler WRON routing problem can be found.

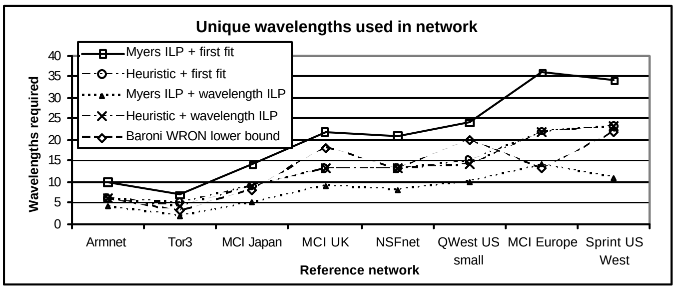

Figure 4 compares the routing techniques in terms of unique wavelength allocation in the network – a perhaps more appropriate metric for minimising capital purchase cost

However, figure 4 provides a rather different perspective – namely of the number of unique wavelengths that are used anywhere in the network. Initial experimentation used the first-fit wavelength allocation algorithm to choose a particular continuous wavelength channel for the length of each wavelength service. The conclusion then inverts: The Myers ILP now consumes considerably more wavelengths than the WRON lower bound, where the heuristic algorithm performs lower or on a par with the WRON lower bound.

Now consider the ILP-based wavelength allocation formulation: The first observation is clear, the λ ILP does drastically better than first-fit for the Myers ILP. However, the λ ILP only barely does better than first-fit for the heuristic. Second observation: The Myers ILP is now performing somewhat better than the heuristic, as would be expected.

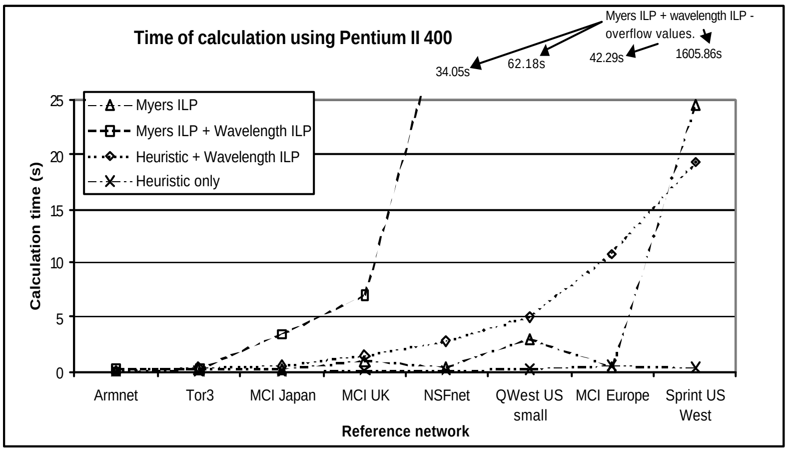

Figure 5 Practical comparison of the computational cost of each method

Previously we have been comparing the benefits of various algorithms in terms of efficiency. Now figure 5 compares the cost – the computation time taken. The heuristic execution time is negligible compared to the Myers ILP implementation, as would be expected. However, note also that the wavelength-allocation ILP running on the output of either approach takes a dramatically greater execution time. The expectation had been that, as it is solving a simpler problem, the wavelength-allocation ILP would add negligible computation time compared to the Myers ILP. In practice, the opposite was observed: as each link in the network was effectively represented *W* times over, the variables/constraints count was much higher. Another interesting phenomenon is that the λ ILP is clearly working much harder to solve the allocation problem for the Myers ILP path allocation versus the heuristic path allocation – despite these being of close size to each other. This begs close future analysis.

5. Conclusions.

This investigation was primarily intended to assess the performance of the proposed heuristic algorithm, but the topic of this paper is the perspective it has given on wavelength routing generally:

1. The first-fit wavelength allocation is potentially deeply misleading as to the potential of routing algorithms – figure 4 shows contradictory conclusions depending upon whether one compares the algorithms using the first-fit algorithm or an ILP approach due to a systematic difference 'error' in result using first-fit. While this is presumably an artefact of the nature of the heuristic algorithm, nonetheless the general point is established that a class of solution may produce a result particularly 'suited' to the operation of the first-fit algorithm. Thus the use of the λ-assignment ILP would seem wise.
2. The λ-assignment ILP has very poor scaling performance, e.g. compared to the Myers ILP. An improved version is highly desirable such that e.g. full-sized US networks can be considered. Note: Given that the wavelength ILP works with fixed paths, it is immaterial whether it is working with WPS networks or WRONs, it would be of general applicability.

Acknowledgments.

Support from Nortel Networks and the EPSRC is gratefully acknowledged. Direction and advice from Prof. Polina Bayvel of UCL is also gratefully acknowledged.

References.

[1] "Architectures for Optical Burst Switching", Alec Myers, M.Res thesis 2001, University College London.
[2] "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria", Telcordia GR-253-CORE. Issue 3, Sept. 2000.
[3] "Internet Protocol", J. Postel, Sept. 1981. Internet Engineering Task Force : RFC0791/STD0005.
[4] "Routing and wavelength allocation in WDM optical networks", Stefano Baroni, University College London, PhD thesis, 1998.