Available online at www.sciencedirect.com

ScienceDirect

Discrete Applied Mathematics 156 (2008) 2911–2923

DISCRETE APPLIED MATHEMATICS

www.elsevier.com/locate/dam

Routing to reduce the cost of wavelength conversion☆

Thomas Erlebach[a,*], Stamatis Stefanakos[b]

[a] *Department of Computer Science, University of Leicester, Leicester LE1 7RH, United Kingdom*
[b] *Computer Engineering and Networks Laboratory, ETH Zürich, CH-8092 Zürich, Switzerland*

Received 14 July 2005; received in revised form 27 November 2007; accepted 7 December 2007
Available online 21 February 2008

Abstract

We consider all-optical networks that use wavelength-division multiplexing and employ wavelength conversion at specific nodes in order to maximize their capacity usage. We investigate the effect of allowing reroutings on the number of necessary wavelength converters. We disprove a claim of Wilfong and Winkler [G. Wilfong, P. Winkler, Ring routing and wavelength translation, in: Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '98, 1998, pp. 333–341] according to which reroutings do not have any effect on the number of necessary wavelength converters on bidirected networks. We show that there exist (bidirected) networks on n nodes that require $\Theta(n)$ converters without reroutings, but only $O(1)$ converters if reroutings are allowed. We also address the cases of undirected networks and networks with shortest-path routings. In each case, we resolve the complexity of computing optimal placements of converters.

Keywords: Graph algorithms; Path coloring; Sufficient set; Optical networks; Wavelength assignment

1. Introduction

In optical networks that use wavelength-division multiplexing, connections are established by determining a path from the transmitter to the receiver (the routing phase) and reserving a free wavelength on all the fibers of that path (the wavelength assignment phase). In a network without wavelength converters, each connection must be assigned the same wavelength on all the fibers of its path. Since a single wavelength might not be free on all fibers along a connecting path, this can result in considerable waste of capacity. To improve the capacity usage of the network, wavelength converters – devices that can modify the wavelength of an incoming connection – are placed in certain nodes (the network design phase).

Because of the high costs involved, one usually seeks a small subset of nodes that, when equipped with wavelength converters, will allow the network to run at maximum capacity. Since this problem is dealt with in the network design phase, the placement of wavelength converters has to support any communication pattern that is likely to emerge in the future. The model used until now in the literature requires that any admissible routing with congestion L (i.e., at most L paths share a fiber) can be accommodated with L wavelengths when wavelength converters are present. Under

☆ Research partially supported by the Swiss National Science Foundation (Project AAPCN).
* Corresponding author.
E-mail addresses: t.erlebach@mcs.le.ac.uk (T. Erlebach), stefanak@tik.ee.ethz.ch (S. Stefanakos).

0166-218X/$ - see front matter
doi:10.1016/j.dam.2007.12.001

this condition, a single "bad" routing for a set of requests can necessitate additional wavelength converters for the network. If such routings can be avoided for some sets of requests by using an appropriate routing mechanism, it is likely that the network would need fewer converters to operate at maximum capacity.

In this paper, we address the question whether the number of necessary wavelength converters can be reduced by an appropriate tuning of the routing phase. We therefore require that, for each possible set of connection requests that can be routed with congestion L, there is at least one routing that can be accommodated using at most L wavelengths. In contrast to previous claims in the literature, we exhibit instances where the ratio of the numbers of the required wavelength converters in the two models can be arbitrarily large. We also show that the same holds for the case where the routing is restricted to shortest paths. In each case, we resolve the complexity of computing optimal placements of converters.

1.1. Preliminaries

We model the network by an undirected or (bi)directed graph $G = (V, E)$ (in a bidirected graph, $(u, v) \in E$ implies $(v, u) \in E$). We restrict G to be a *simple* graph (i.e., a graph without parallel edges; pairs of anti-parallel edges (u, v) and (v, u) are allowed in the directed case, however). A *connection request* is a pair of vertices (u, v) and is routed through a (simple) $u - v$ path. The *load* $L(\mathcal{P})$ of a set $\mathcal{P}$ of paths is the maximum number of paths that share a common edge; the load of an edge is the number of paths that use that edge. Regarding wavelengths as colors, a *wavelength assignment* for a set of connections routed through a set of paths $\mathcal{P}$ is an assignment of colors to the paths in $\mathcal{P}$ such that all paths that share an edge get different colors. In the presence of wavelength converters, a coloring is an assignment of a color to every edge of each path. If the converters are placed in the vertices in $S \subseteq V$, we say that a coloring is valid with respect to S if it satisfies the additional constraint that the color assignments to two consecutive edges of a path differ only if their incident vertex is in S.

In order to achieve maximum capacity usage by placing wavelength converters at the nodes $S \subseteq V$, we need that any admissible routing $\mathcal{P}$ can be colored with $L(\mathcal{P})$ colors with respect to S. If we allow arbitrary routings, a set S with this property is said to be *sufficient*. If we restrict to shortest-path routings, a set with this property is said to be *SP-sufficient*. If we are allowed to reroute, we speak about *weakly sufficient* and *weakly SP-sufficient sets*. A set $S \subseteq V$ is weakly sufficient for G if for any set of requests R on G and any routing $\mathcal{P}$ for R there exists a routing $\mathcal{P}'$ for R that can be colored with $L(\mathcal{P})$ colors with respect to S. Weakly SP-sufficient sets are defined similarly but both $\mathcal{P}$ and $\mathcal{P}'$ must contain only shortest paths. We refer to the corresponding minimization problems by appending the word "minimum" before the sought structure; for example "minimum weakly sufficient set."

We will need the following graph-theoretic notation: We denote an induced cycle on k vertices by C_k. We also refer to a C_3 as *triangle*. We call an induced $K_{1,3}$, i.e. the graph on four vertices with one vertex of degree 3 and all others of degree 1, a *claw*. For a graph $G = (V, E)$ and a subset $S \subseteq V$, we denote by $G(S)$ the graph obtained from G by replacing each vertex $s \in S$ by degree-of-s-many copies, each of which is made adjacent to one of the old neighbors of s. The graph $G(S)$ is also called the *exploded graph* of G (with respect to S). A *spider* is a tree with at most one vertex of degree greater than 2. A *chain* is a tree with no vertex of degree greater than 2 (i.e., a path). The maximum degree of an undirected graph or multigraph G is denoted by $\Delta(G)$. The degree of a vertex in a bidirected graph is defined to be the in-degree (or, equivalently, the out-degree) of that vertex.

1.2. Previous work

Wilfong and Winkler [15] show that the only connected bidirected graphs that admit the empty sufficient set are spiders. They also provide an efficient way of determining whether a set S is sufficient: a set S is sufficient for a graph G if and only if every component of $G(S)$ is a spider. Concerning the problem of finding a minimum sufficient set, Wilfong and Winkler show that it is $\mathcal{NP}$-hard even for planar bidirected graphs. They also claim that in bidirected graphs any weakly sufficient set is sufficient; we will disprove this claim in Theorem 4.

Kleinberg and Kumar [11] give a 2-approximation algorithm and a polynomial-time approximation scheme for the minimum sufficient set problem in arbitrary directed graphs and in directed planar graphs, respectively. The main idea behind their 2-approximation algorithm is to transform the instance of the problem to a feedback vertex set problem (for the special case of bidirected graphs they also show a transformation to vertex cover). Kleinberg and Kumar also show that any ρ-approximation algorithm for the minimum sufficient set problem in bidirected graphs implies a ρ-approximation algorithm for the vertex cover problem.