Available online at www.sciencedirect.com

ScienceDirect

Computer Networks 52 (2008) 1105–1122

Computer Networks

www.elsevier.com/locate/comnet

Scheduling bulk file transfers with start and end times

Kannan Rajah, Sanjay Ranka, Ye Xia *

Computer and Information Science and Engineering Department, University of Florida, 301 CSE Building,
P.O. Box 116120, Gainesville, FL 32611-6120, United States

Received 12 May 2007; received in revised form 24 October 2007; accepted 31 December 2007
Available online 8 January 2008

Responsible Editor: A. Kamal

Abstract

The advancement of optical networking technologies has enabled e-science applications that often require transport of large volumes of scientific data. In support of such data-intensive applications, we develop and evaluate control plane algorithms for scheduling bulk file transfers, where each transfer has a start time and an end time. We formulate the scheduling problem as a special type of the multi-commodity flow problem. To cope with the start and end time constraints of the file-transfer jobs, we divide time into uniform time slices. Bandwidth is allocated to each job on every time slice and is allowed to vary from slice to slice. This enables periodical adjustment of the bandwidth assignment to the jobs so as to improve a chosen performance objective: throughput of the concurrent transfers. In this paper, we study the effectiveness of using multiple time slices, the performance criterion being the tradeoff between achievable throughput and the required computation time. Furthermore, we investigate using multiple paths for each file transfer to improve the throughput. We show that using a small number of paths per job is generally sufficient to achieve near optimal throughput with a practical execution time, and this is significantly higher than the throughput of a simple scheme that uses single shortest path for each job.

Keywords: Bulk file transfers; E-science; Optical networks; Scheduling; Multipath; Rerouting; Maximum concurrent flow

1. Introduction

In the past two decades, optical networking technologies have revolutionized communications. The widespread deployment of fiber optic infrastructure has led to low cost, high capacity optical connections. This networking advancement has enabled e-science applications that often require management and transport of large volumes of scientific data [37,9,2]. For instance, the Large Hadron Collider (LHC) facility at CERN [16] is expected to generate petabytes of experimental data every year, for each experiment. In addition to high-energy nuclear physics [9,38,17], a few other e-science applications are radio astronomy [31], geoscience [20], and climate studies [19]. In order to best support the needs of e-science applications, optical research networks are deployed by a consortium

---

[*] Corresponding author. Tel./fax: +1 352 392 2714.
E-mail addresses: krajah@cise.ufl.edu (K. Rajah), ranka@cise.ufl.edu (S. Ranka), yx1@cise.ufl.edu (Y. Xia).

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2007.12.005

of leading research universities, governments and private sector technology companies. Examples include the Internet2-related [30] National Lambda Rail [36] and Abilene [1] networks in the U.S., CA*net4 [15] in Canada, and SURFnet [42] in Netherlands. These networks are small in size (less than 10^3 in the backbone) as compared to the public Internet. This makes it possible to have a centralized network controller for managing the network resources and for providing user service quality guarantee. With the central controller, there is more flexibility in designing sophisticated, efficient algorithms for scheduling user reservation requests, setting up network paths, and allocating bandwidth.

The objective of this paper is to develop and evaluate control plane algorithms for scheduling large file transfers (also known as jobs) over optical research networks. We assume that job requests are made in advance to the central network controller. Each request specifies a start time, an end time and the total file (demand) size. Such a request is satisfied as long as the network begins the transfer after the start time and completes it before the end time. The network controller has the flexibility in deciding the manner in which each file is transferred, i.e., how the bandwidth assignment to each job varies over time on all its allowed paths. The decision process is known as *scheduling*. We call this scheduling problem the *concurrent file transfer problem* (CFTP).

The current research networks generally use routers over optical transmission technologies instead of optical switches. Routers can split or aggregate traffic before putting it into the wavelength channels or the optical Ethernet links. Hence, the problem in this paper is routing and fine-grained bandwidth assignment rather than wavelength assignment, as would be the case in a wavelength-based circuit-switched optical network. It is possible to reserve an end-to-end wavelength path in the current research networks. But, our formulation of the bandwidth assignment problem will be unaffected since we can simply remove the reserved wavelength from the link capacity. We defer the wavelength assignment problem in an all optical network to future research.

We will formulate CFTP as a special type of the multi-commodity flow problem, known as the maximum concurrent flow (MCF) problem [41,26]. While MCF is concerned with allocating bandwidth to persistent concurrent flows, CFTP has to cope with the start and end time constraints of the jobs. For this purpose, our formulations for CFTP involve dividing time into uniform time slices and allocating bandwidth to each job on every time slice. Such a setup allows an easy representation of the start and end time constraints, by setting the allocated bandwidth of a job to zero before the start time and after the end time. More importantly, in between the start and end times, the bandwidth allocated for each job is allowed to vary from time slice to time slice. This enables periodical adjustment of the bandwidth assignment to the jobs so as to improve some performance objective.

Motivated by the MCF problem, the chosen objective is the throughput of the concurrent transfers. For fixed traffic demand, it is well known that such an objective is equivalent to minimizing the worst-case link congestion, a form of network load balancing [41]. A balanced traffic load enables the network to accept more future job requests, and hence, achieve higher long-term resource utilization. We assume that the optical network contains enough IP routers for traffic grooming, which is true for current research networks. Such a network allows fine-grained multiplexing of traffic for better network resource utilization.

In addition to the problem formulation, other contributions of this paper are as follows. First, in scheduling file transfers over multiple time slices, we focus on the tradeoff between achievable throughput and the required computation time. Second, we investigate using multiple paths for each file transfer to improve the throughput. We will show that using a small number of paths per job is generally sufficient to achieve near optimal throughput, and this is shown to be significantly higher than the throughput of a simple scheme that uses single shortest path. In addition, the computation time for the formulation with a small number of paths is considerably shorter than that for the optimal scheme, which utilizes all possible paths for each job.

The rest of the paper is organized as follows. In Section 2, we describe the CFTP and introduce the uniform time slice structure. In Section 3, we formally describe the node-arc and edge-path formulations of CFTP. The latter includes the k-shortest paths and k-shortest disjoint paths variants. In Section 4, we evaluate the algorithm performance for different formulations. In Section 5, we introduce related work. The conclusions are drawn in Section 6.