Future Generation Computer Systems 56 (2016) 284–294

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Scheduling and flexible control of bandwidth and in-transit services for end-to-end application workflows

Mehmet Fatih Aktas *, Georgiana Haldeman, Manish Parashar

Rutgers Discovery Informatics Institute (RDI²), Rutgers University, Piscataway NJ, USA

HIGHLIGHTS

- A flow-based in-transit data staging and processing approach.
- An opportunistic QoS model, based on convex optimization, for resource scheduling.
- SDN controller for enabling in-transit services and end-to-end workflow couplings.
- An evaluation using a Fusion workflow scenario and an emulated virtualized SDN testbed.
- The approach achieves high resource utilization and satisfies workflow QoS requirements.

ARTICLE INFO

Article history:
Received 20 February 2015
Received in revised form
18 August 2015
Accepted 11 September 2015
Available online 26 September 2015

Keywords:
Advance scheduling
Data transport
SDN
In-transit processing
Distributed-area coupled simulation workflows

ABSTRACT

End-to-end scientific application workflows that integrate high-end experiments and instruments with large scale simulations and end-user displays are becoming increasingly important. These workflows require complex couplings and data sharing between distributed components involving large data volumes and present varying hard (in-time data delivery) and soft (in-transit processing) quality of service (QoS) requirements. As a result, supporting efficient data transport is critical for such workflows. In this paper, we leverage software-defined networking (SDN) to address issues of data transport service control and resource provisioning to meet varying QoS requirements from multiple coupled workflows sharing the same service medium. Specifically, we present a flexible control and a disciplined resource scheduling approach for data transport services for science networks. Furthermore, we emulate an SDN testbed on top of the FutureGrid virtualized testbed and use it to evaluate our approach for a realistic scientific workflow. Our results show that SDN-based control and resource scheduling based on simple intuitive models can meet the requirements of the targeted workflows with high resource utilization.

1. Introduction

As scientific discovery is becoming increasingly data driven, scientific applications are moving towards end-to-end workflows that integrate coupled simulations with data sources such as instruments, sensor systems and experiments, and with data analysis and visualization pipelines to facilitate online knowledge extraction. Furthermore, the execution of such workflows often involves geographically distributed resources with runtime interactions, coordination and data exchanges between processes running on these resources [1,2].

Scientific workflows typically involve complex couplings between the workflow components/services, requiring sharing of large data volumes with varying hard (in-time data delivery) and soft (in-transit processing) quality of service (QoS) requirements, and efficient data transport is a key requirement. Specifically, the time between when the data is generated at the producer and when it can be consumed at the consumer can have a significant impact on the execution of the workflow. For example, slower data delivery can throttle the consumer, or faster data delivery may require storing the data at the consumer. Some applications may require that data is delivered at the consumer within a tight time window (for example, data needed to control an experiment), which adds further requirements on the data transport. Additionally, there often exist natural mismatches in the way data is represented at producers and consumers, and data has to be transformed in a timely manner before it can be consumed. As a result, the data transport has to address multiple challenging requirements based on data sizes, data production and consumption rates, strict constraints on data delivery time and data storage, and managing

* Corresponding author.
E-mail addresses: mehmet.aktas@rutgers.edu (M. Fatih Aktas), haldeman@rutgers.edu (G. Haldeman), parashar@rutgers.edu (M. Parashar).

http://dx.doi.org/10.1016/j.future.2015.09.011
0167-739X/

data transformations between producers and consumers. The data transport medium is also typically shared by multiple application workflows with possibly competing application specific coupling requirements. Furthermore, best-effort solutions, which inherently cannot offer service guarantees, will not be able to satisfy these requirements and achieve desired performance over utilized service medium [3].

The goal of this paper is to explore, prototype and evaluate a solution that can address these challenges, and enable the scheduling and control of bandwidth and in-transit services for end-to-end scientific application workflows. Specifically, in this paper, we leverage software-defined networking (SDN) to address issues of data transport service control and resource provisioning to meet varying QoS requirements from multiple coupled workflows sharing the same service medium. The specific contribution of this paper is a disciplined resource scheduling approach for data transport resources that is both application and network aware, and enables flexible control. The proposed control plane achieves flexibility by using SDN abstractions for managing the network, and further extending these abstractions to manage data transport services and schedule data transport resources. We also develop a model for in-transit data staging and data processing using intermediate resources in the data path using the approach outlined in [2,4,5]. Finally, we emulate an SDN testbed on top of the FutureGrid virtualized testbed and use it to evaluate our approach for an end-to-end Fusion workflow. Our results show that SDN-based control and resource scheduling based on simple intuitive models can meet the coupling requirement with high resource utilization.

The rest of the paper is organized as follows. Section 2 defines the problem and provides some background. Section 3 presents our approach for scheduling and control of bandwidth and in-transit services for end-to-end scientific application workflows. Section 4 describes our emulation testbed and presents the experimental evaluation of our approach. Section 5 presents related work. Section 6 concludes the paper.

2. Problem description—enabling end-to-end coupled simulation workflows

Emerging scientific applications integrate simulations with data sources such as experiments and instruments, and analysis and visualization pipelines, into end-to-end workflows. These workflows exhibit different and varying interaction, coordination and data coupling behaviors, such as:

- **Tight coupling**: Coupled processes exchange data very frequently. Therefore, when tight coupling is dominant and processes are coupled over the network, total workflow execution time is typically dominated by the data transfer time. Often involves strict data transport constraints such as delivery time, data integrity and reliability.
- **Loose coupling**: Relatively less frequent, asynchronous and possibly opportunistic data exchanges among the coupled processes. Producer and consumer progress at different rates, and may have different data representations making intermediate data transformation necessary.
- **Dataflow coupling**: Data flows from producers to consumers using publish/subscribe/notify-like semantics, for example, in case of data processing and/or analysis pipelines.

It is essential for the data transport service to be application-aware and self-optimizing to meet requirements of such complex and varying coupling behaviors. Meanwhile, the overheads of service management on the application should be acceptable. In this paper, we focus on control and scheduling of data transfer (across network switches and links) nodes and in-transit service (intermediate staging and processing) hosts to meet varying data transport requirements of differing coupling behaviors, such as those described above. We propose a centralized advance scheduling system consisting of a centralized control layer and scheduler to enable workflow couplings. Enabling a coupling session requires scheduling a network path, and allocating resources on the scheduled path. Data transport resources that we consider are bandwidth and intermediate processing and staging capacity. Scheduler keeps the current state of the service medium, gets the coupling requirements from the applications, and generates scheduling and resource allocation rules upon receiving a scheduling request. These rules are then realized by the controller with a state setup process (Section 3.2).

A scheduling request consists of consumer address, size of data to be streamed and coupling requirements. Differing coupling behaviors translate into different requirements on the data transport service. For example, tight coupling usually has tight delivery time window while loose coupling takes place opportunistically or loose coupling requires more intermediate data manipulation than tight coupling. We formulated resource allocation as a disciplined optimization problem with the objective of maximizing cumulative satisfaction of coupling requirements. Scheduler numerically solves the formulated problem to optimize the resource allocation between coupling sessions (service users) (Section 3.3).

The Software Defined Networking (SDN) initiative aims at making network control simpler and more flexible using well-defined abstractions for forwarding and switch configuration. A key idea behind SDN is replacing the existing distributed control with centralized control. This is achieved by implementing network control programs on top of the network operating system (NOS). NOS provides the control programs with a global view of the network and interfaces for communicating with switches, enabling the control programs to configure the network state.

In this research, we explore how the data transport service control can benefit from SDN, and the capabilities that it provides, such as (1) open and programmable control, (2) faster innovation at the networking layer, (3) easy customization and optimization of network resource scheduling via flexible control. Overall, by leveraging SDN-based networking control, we propose an approach for the control and scheduling of data transport services shared by the end-to-end coupled application workflows.

3. Scheduling and flexible control of bandwidth and in-transit services

In this section, we explore scheduling and flexible control of bandwidth and in-transit services for end-to-end scientific application workflows. We leverage SDN to address issues of data transport service control. We also present a disciplined formulation of resource provisioning to meet varying requirements from multiple coupled applications sharing the same service medium. Specifically, the solution we propose has three key components: (1) A flow-based in-transit service model, (2) Centralized layered architecture for data transport service control, and (3) Application and network aware resource scheduling for workflow couplings. In the following subsections, we will detail these three components.

3.1. In-transit service model

The goal of in-transit service is to offload (1) Opportunistic data transformation, and (2) Data staging for just-in-time data delivery, to available intermediate service hosts. This requires intervening the dataflow between producers and consumers. Packet-based and flow-based approaches are two implementations of intermediate processing/staging.