

Next Generation Content Delivery Infrastructures: Emerging Paradigms and Technologies

Giancarlo Fortino
University of Calabria, Italy

Carlos E. Palau
Universitat Politècnica de Valencia, Spain

Managing Director: Lindsay Johnston
Senior Editorial Director: Heather A. Probst
Book Production Manager: Sean Woznicki
Development Manager: Joel Gamon
Development Editor: Myla Merkel
Acquisitions Editor: Erika Gallagher
Cover Design: Nick Newcomer

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: <http://www.igi-global.com>

Copyright © 2012 by IGI Global. All rights reserved. No part of this publication may be reproduced, stored or distributed in any form or by any means, electronic or mechanical, including photocopying, without written permission from the publisher. Product or company names used in this set are for identification purposes only. Inclusion of the names of the products or companies does not indicate a claim of ownership by IGI Global of the trademark or registered trademark.

Library of Congress Cataloging-in-Publication Data

Next generation content delivery infrastructures : emerging paradigms and technologies / Giancarlo Fortino and Carlos E. Palau, editors.

p. cm.

Includes bibliographical references and index.

Summary: "This book delivers state-of-the-art research on current and future Internet-based content delivery networking topics, bringing to the forefront novel problems that demand investigation"-- Provided by publisher.

ISBN 978-1-4666-1794-0 (hardcover) -- ISBN 978-1-4666-1795-7 (ebook) -- ISBN 978-1-4666-1796-4 (print & perpetual access) 1. Web site development. 2. Information storage and retrieval systems. I. Fortino, Giancarlo, 1971- II. Palau, Carlos E., 1970-

TK5105.888.N4825 2012

004.67'8--dc23

2012002872

British Cataloguing in Publication Data

A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

Chapter 6

Quality Guaranteed Media Delivery over Advanced Network

Zhiming Zhao

University of Amsterdam, The Netherlands

Paola Grosso

University of Amsterdam, The Netherlands

Jeroen van der Ham

University of Amsterdam, The Netherlands

Cees de Laat

University of Amsterdam, The Netherlands

ABSTRACT

Moving large quantities of data between distributed parties is a frequently invoked process in data intensive applications, such as collaborative digital media development. These transfers often have high quality requirements on the network services, especially when they involve user interactions or require real time processing on large volumes of data. The best effort services provided by IP-routed networks give limited guarantee on the delivery performance. Advanced networks such as hybrid networks make it feasible for high level applications, such as workflows, to request network paths and service provisioning. However, the quality of network services has so far rarely been considered in composing and executing workflow processes; applications tune the execution quality selecting only optimal software services and computing resources, and neglecting the network components. In this chapter, the authors provide an overview on this research domain, and introduce a system called NEtWork QoS Planner (NEWQoSPlanner) to provide support for including network services in high level workflow applications.

DOI: 10.4018/978-1-4666-1794-0.ch006

1. INTRODUCTION

The development of a large multi-media application involves raw material acquired from different sources and is often a collaborative effort among several parties. Moving large quantities of semi-finished material between distributed locations is a frequently invoked process during the development phase. IP-routed paths in the Internet are not the most suitable way to transfer high-quality digital media. Streaming content in 4K-format (4096 pixels of horizontal resolutions) or higher formats has two basic requirements: sufficient network capacity and quality of experience for the end user. Uncompressed 4K content requires a network bandwidth of more than 7Gbps. This can be accomplished by ensuring that the whole end-to-end path, from source to destination, is provisioned over a 10Gbps channel. Nowadays this is technically feasible, but it cannot be a priori guaranteed in the Internet where there is very limited control over the segments the data will be routed through. Furthermore, packet loss, reordering and varying jitter cannot be avoided in a best-effort environment as the Internet. These performance hiccups cause severe degradation of the viewing performance.

To investigate solutions to these problems, several research initiatives have started. Notably, a group of researchers and industrial partners started in 2006 the CineGrid collaboration (<http://www.cinegrid.org>). CineGrid is a non-profit organization whose members form an interdisciplinary community focused on the research, development, and demonstration of networked collaborative tools to enable the production, use and exchange of very-high-quality digital media over photonic networks. The basic idea of CineGrid is that network circuits implemented over photonic networks provide the proper guarantees of bandwidth and quality of service for media delivery applications. A challenge is to integrate the network in the overall delivery framework, where also computing nodes and many types of software components

play an important role. The techniques used for digital content delivery are easily ported to support content-delivery networks (CDNs), such as (Fortino, Russo, Mastroianni, Palau, & Esteve, 2007). It is our opinion that these types of networks could fully utilize the advanced network services exposed by the network providers and that they could integrate the workflow planning techniques in their optimizations.

We shall highlight our research focuses: modeling the meta information of network resources and media material, managing operation sequences of data access and movement sequences, and using advanced network infrastructure to provide quality guaranteed connections for moving large quantity data. Workflows are the natural way to address this resource selection and composition problem and they are playing an important role in the daily operations and use of grid and cloud infrastructures. Particularly, in the scientific community workflow systems have gained popularity among researchers to support complex experiments (Zhao, Belloum, & Bubak, 2009). Still the application of workflows to media delivery scenarios is fairly new. The inclusion of network information in the workflow planning phase and the use of a continuous feedback regarding the current status of the network resources during workflow execution are the two main novel aspects of our work.

This chapter is about using these technologies in digital content delivery. First, we will introduce the background for our work, and review the state of the art. We then introduce a system called NEWQoSPlanner, and we discuss how it can be used to enhance the resource description, discovery, network path selection and provisioning for content delivery.

2. BACKGROUND

Advanced network architectures can provide quality guaranteed services for data intensive applications, such as content delivery, which have

a high requirement on the data delivery and on the data operations. We will focus in this section on the technologies involved in developing such kind of new (network) infrastructures.

2.1 Data Intensive Applications over Grid

Grids, and nowadays clouds, provide a suitable environment for the execution of data intensive applications. Data can be processed in parallel at different locations and later on transported back to a single place where the final computation is performed. Data Grids provide also support for distributed storage, allowing users to leverage the infrastructures present in multiple data centers (Maassen, Verstoep, Bal, Grosso, & de Laat, 2009; Venugopal, Buyya, & Ramamohanarao, 2006). The combination of computing Grids and advanced network services, as the ones offered by many Research and Education Networks, has enabled the support of applications from various scientific fields, e.g. high-energy physics, geosciences, bioinformatics, ecology, astronomy. In particular high-definition video and digital-cinema streaming, high performance computing, visualization and virtual reality applications fully exploit this type of infrastructures.

A very illustrative example is the plethora of applications that already in 2005 were showcased during the iGrid2005 conference (Smarr, Brown, de Fanti, & de Laat, 2006). Participants could witness a full range of working demonstrations in the area of visualization and video streaming. Cosmic ray data collected in Tibet was sent to Beijing and later on to all projects partners that needed to process it, using *lightpaths services* (Nan, Ma, Zhang, & Chen, 2006); a high-quality collaborative environment that used HD video provided an enhanced video conferencing system to participants around the globe (Holub et al., 2006); interactive 3D video streams were transported over 10Gbit/s intercontinental dedicated connections (Jo et al., 2006); video transcoding, from high-

resolution broadcast video into MPEG_2 format, in order to reduce file size and resolution, made use of remote computers connected by on-demand paths (Grasa et al., 2006). iGrid2005 also showed for the first time a real-time, international transmission of 4K digital cinema and 4K Super High Definition digital video between Japan and the USA west coast (Shimizu et al., 2006).

Several data intensive applications were also presented, in particular in the area of computational astrophysics and astronomy. For example; 797GB astronomical data from the Sloan Digital Sky Survey was sent to computing centers across the world (Grossman et al., 2006); or, the electronic long base interferometry (e-VLBI) application (Sobieski et al., 2006), which can provide ultra-high resolution images of faint and distant objects in the universe, required the creation of a single computational environment allowing data coming on the network links to many radio telescopes around the globe.

One very interesting use case that well highlighted the fruitful symbiosis of grids and advanced networks was the seamless migration of virtual machines over the wide area network (Travostino et al., 2006). A live virtual machine could be moved from Amsterdam (NL) to San Diego (USA) with just 1-2s of application downtime. Balancing the work load across data centers and disaster recovery are among the most important motivations for this advanced use of grids and networks, and in general to support data intensive applications over lambda grids.

2.2 Advanced Network Services for Data Intensive Applications

The Internet, based on the TCP/IP architecture, has been designed as a best-efforts service, where the intelligence is in the end nodes at the edges. The functionalities related to problems occurred during communication on the networks, such as for example the reordering of packets due to the erroneous order of packets arrival or the detection

of duplicates, are all function of the Transport Layer, *i.e.* the Transmission Control Protocol (TCP). In the core of the Internet the Internet Protocol (IP) '*limits*' itself to perform fast and efficient delivery.

Quality of Service can be defined as the statistical performance guarantees that a network system can make in terms of throughput, delay, jitter and loss. The Internet treats everybody, *i.e.* all packets, as equal. But applications are not equal as they differ in the sensitivity to delays, and the mission critical value of the data being transported.

There are four principles that govern the implementation of Quality of Service (QoS) in the Internet. All of these principles have led to the development of protocols and techniques to mitigate and solve the original design flaw of TCP/IP:

1. Marking (packet classification) allows routers to distinguish between different classes of packets; and router policies to treat packets accordingly;
2. Isolation (scheduling and policing) provides protection for one class from other classes; it ensures sources adhere to bandwidth requirements; scheduling and policing are QoS functions performed at the edges of the network.
3. High resource utilization. While providing isolation, it is desirable to use resources as efficiently as possible.
4. Call admission (signaling) prevents that traffic is ingested in the network beyond link capacity. Application flows need to declare their needs, and the network may block calls if it cannot satisfy the requests.

Packet marking has relied on the use of the Type of Service (ToS) field in the IP packets. The Integrated Services (Intserv) architecture provides QoS guarantees in IP networks for individual application sessions. It was defined in 1994 in RFC 1633 (Braden, Clark, & Shenker, 1994). Its main

characteristics are 1) the support for resource reservation, as routers maintain state info of allocated resources and QoS requests and 2) the possibility to admit/deny new call setup requests. IntServ does not scale well, because maintaining per-flow router state with large number of flows is difficult.

In 1998 RFC 2474 (Nichols, Blake, Baker, & Black, 1998) and RFC 2475 (Blake et al., 1998) defined the Differentiated Services (DiffServ) protocol. DiffServ provides simple functions in the network core, and relatively complex functions at edge routers (or hosts). It does not define service classes, but it provides the functional components to build such service classes. DiffServ supersedes the ToS field in IPv4 to make per-hop behavior (PHB) decisions about packet classification and traffic conditioning functions. The PHB results in a different observable (measurable) forwarding performance behavior, but it does not specify what mechanisms to use to ensure required PHB performance behavior. DiffServ marks packets using the IPv4 ToS and the Traffic Class field in IPv6. Six bits are used for Differentiated Service Code Point (DSCP) and to determine PHB that the packet will receive, while 2 bits are unused. DiffServ has replaced IntServ as the protocol of choice to provide different level of services in the Internet. Still DiffServ fails to guarantee a specified service level, as there is no guarantee that packets marked will receive the expected service. This limitation has led to look for different ways to provide applications with the proper guarantees in terms of throughput, delay, jitter and loss.

The idea of providing data intensive applications with deterministic point-to-point connections was fostered by a community of research networks, later organized in the Global Lambda Integrated Facility (GLIF). This community provides a global network to support data-intensive scientific research, and also supports middleware development for optical networking. The ideas in this community led to the concept of *hybrid networking*, the offering of packet switched (IP)

Figure 1. Visualization techniques shown during the iGrid2005 workshop: high-definition passive stereo displays (bottom right), auto-stereoscopic displays (top right), ultra-high definition tiled projection and LCD displays (left top and bottom)



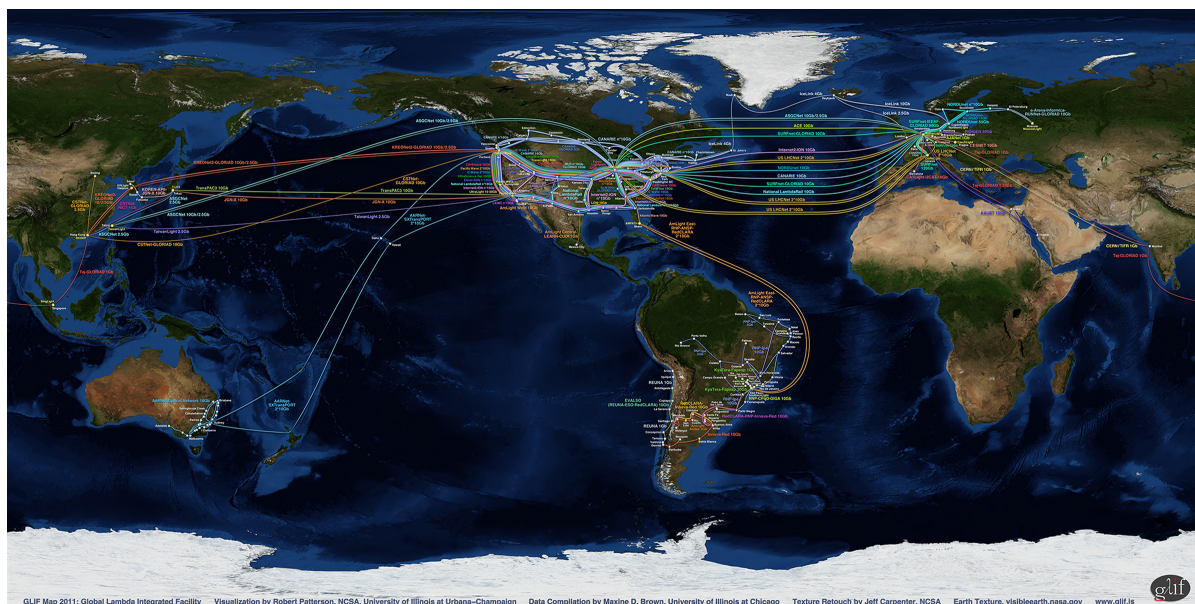
services and circuit switched connections over the same physical network infrastructure (de Laat, Radius, & Wallace, 2003).

Since most data intensive applications operate in a large-scale environment, with collaborators at different locations, the networks required for these applications are nearly always multi-domain networks. De Laat estimated in 2000 that a typical network connection for a physics experiment crosses seven domains (de Laat & Blom, 2000). To achieve inter-domain operations, the different networks have to collaborate. For dedicated network connections, this collaboration is done in the GLIF community. In few years time a number of international network connections have been established to provide the inter-domain

connectivity. Figure 2 shows a collection of the interconnections provided by partners in the GLIF community as of May 2011.

The GLIF community is working hard at improving the lightpath provisioning process by exchanging experiences, documenting processes and developing middleware. In the meantime, the available speeds of lightpaths keep growing. While 10Gbit/sec links were introduced only a few years ago, 40Gbit/sec links are now becoming available to application developers and 100Gbit/sec hardware is just becoming available (Dumitru, Koning, & de Laat, 2010). These kinds of links provide unique opportunities for the transport of high-quality media and the construction of CDN architectures.

Figure 2. GLIF world map of May 2011, with all network connections offered by its participants ¹



2.3 Information Model of Advanced Network Infrastructures

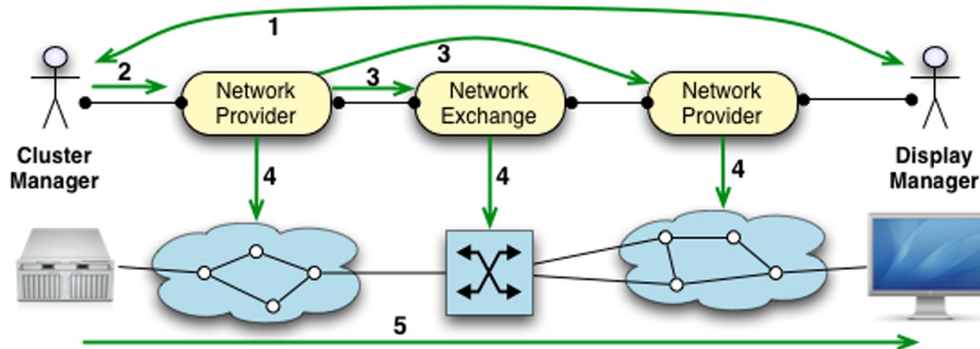
Automatic path finding and provisioning of inter-domain lightpaths is important to facilitate the usage of the advanced network infrastructures mentioned in 2.2. Figure 3 shows the steps that currently need to be taken to establish a network connection for any high level application, in this example between a cluster and a display. If we examine this procedure in more detail, we see that it is broken up in the following underlying steps:

1. The user formulates the requirements, including the end points and the network characteristics like bandwidth, latency, jitter, minimum packet size (if applicable), reliability, etc.
2. These requirements must be communicated to their upstream network provider. The network provider must gather information about available resources, including the resources in other networks, as the two end-points are typically in different networks.
3. The network provider must, in collaboration with the other network providers, determine a valid path that uses available resources, and is within the specs of the user. The resources needed for the path must be reserved in all networks involved.
4. Once the reservations are all confirmed, the reserved resources must be configured in the networks. The end-to-end path must be tested, and in case of faults the faults must be examined and resolved. The network provider informs the user, and the user must configure the end nodes (e.g. configure the IP addresses and set the routing table).
5. The user runs the applications.

Currently, this whole process of acquiring a (working) lightpath across multiple domains can take several weeks, a lot of emails and phone calls and extensive testing. It is clear that the whole process needs to be improved and automated in order to scale.

The example described in Figure 3 shows that the intermediate steps by the network operators

Figure 3. Steps to set up a network connection between a cluster and a display



involve a lot of communication. In order to determine the path, they have to exchange topology and capability information. Once a path has been determined an operator must communicate the specifics to the other operators involved. There are several information models available to describe network topologies. However, these information models are either aimed at a knowledgeable single network operator, since they are mostly suitable for describing monitoring, diagnostics and configuration information, such as SNMP (Case, Mundy, Partain, & Stewart, 2002), NetConf (IETF, 2011) or CIM (DMTF, 2011).

Other network information models are aimed more at describing topologies such as the Network Measurements-WG (Nsi-Wg, 2011) model or G.805 (International Telecommunications Union, March 2000), however these models are not intended to publish topology information to other domains. Another very complete network information model is GMPLS (Farrel & Bryskin, 2006); however, that model is squarely aimed at networking devices, and is not suitable for publishing outside the domain, or for extension to other applications.

Based on the existing work, a Network Description Language (NDL) (Ham, Dijkstra, Travostino, Andree, & de Laat, 2005) has been developed to model the network information, which uses the Resource Description Framework (W3C, 2010) as

its basis. NDL provides generic globally unique identifiers, so that different domains can publish and share network descriptions. In the next section, we will have more discussion on NDL.

2.4 QoS and Workflow Systems in Data Intensive Applications

Delivering large quantity data over advanced networks involves several steps in setting up networks and needs invocation of different services to perform the data movement and processing. Scientific workflow systems are suitable in this context to hide low-level integration details and to automate the management of delivering and sharing digital contents in the applications. But it also requires workflow systems to meet the performance requirements on the data operations.

The development of scientific workflow can be roughly divided into four phases (Deelman, Gannon, Shields, & Taylor, 2009): composition, enactment, execution, and post analysis. Service oriented architecture plays a key role in decomposing workflow processes and in integrating them. The Quality of Services needs to be included in each phase of the workflow lifecycle to optimize the global performance of the application to meet the user's requirements. In the following sections we briefly review the existing work from four

aspects: workflow composition, service selection, execution control and provenance.

A workflow composition process that is QoS-aware must: 1) compose a service of the highest quality and 2) determine the quality of the composition process itself. The first goal is achieved by computing the global quality starting from the QoS attributes of constituting services (Lecue & Mehandjiev, 2009). Graph reduction is a widely used approach (Cardoso, Miller, Sheth, & Arnold, 2002); a pre-defined set of logic patterns defines certain reduction rules which can be used to simplify the logical dependencies among constituting services. From the reduction rules, the quality parameters are computed; for instance the computing time of two sequentially connected services is computed as the sum of the quality of each of them, the computing time of two parallel services is computed as the maximal one. The second goal requires modeling the quality attributes of the semantic links between services, the composition quality of the workflow can then be evaluated by the semantic fit and the reliability of the selected service in the workflow.

Searching for suitable services from available resources is a basic procedure in composing a workflow. QoS aware service selection implies two steps: properly formulating the requirements and selecting resources that meet these requirements. Rosenberg proposed a QoS enabled description language, the Vienna composition language (VCL) (Rosenberg, Leitner, Michlmayr, Celikovic, & Dustdar, 2009), to specify an abstract flow for workflow composition. VCL defines an abstract workflow as four parts: feature definition, feature constraints, global constraints and the business protocol (the desired workflow language). The feature constraints and global constraints include both functional constraints and QoS attributes. The problem of resource selection has been formulated differently. A commonly used formulation is *shortest path finding in a weighted graph*, in which the available services are represented as a directed graph according to the service types, and

the graph nodes are labeled by the quality attributes of the service (Li, Chen, Wen, & Sun, 2008). Well known shortest path finding algorithms include Bellman-Ford and Dijkstra's. These algorithms exhibit optimal performance because of their greedy search strategy and avoid backtracking operations during the search; however, the minimal cost path found by the algorithms is often not the most optimal solution if there are multiple constraints on the quality attributes. Therefore, the problem has also been formulated as a multi constraint optimal path problem (Yu, Kirley, & Buyya, 2007), or multi objective optimization problem. Ant colony optimization (ACO) is a meta heuristic search approach proposed in (Li & Xu, 2003; Alaya, Solnon, & Ghdira, 2007) for discovering minimum cost path in a graph, and for solving NP-hard combinatorial optimization problems. Fang, Peng, Liu and Hu (2009) applied ACO in service selection and proposed a multi objective ACO approach that can simultaneously optimize several objectives. Genetic algorithms in searching optimal paths, and constraint programming or Integer programming methods are also widely used for the multi objective optimization problem.

Workflow execution is the mapping of workflow processes onto underlying computing resources and the scheduling of the execution sequence. Task based scheduling is a straightforward approach, in which the workflow tasks are submitted to the local manager of the computing infrastructure. Several researchers have instead proposed a workflow level scheduling that takes into account future task performance (Harada, Ushio, & Nakamoto, 2007); this approach will achieve higher performance and better resource utilization than only using local resource managers. Multi objective optimizations are widely used to formulate the problem of QoS aware scheduling. Avanes and Freytag (2008) proposed a constraint programming based approach to search for the best match between workflow requirements and the available computing resources. The basic idea is

to describe the quality requirements and resource dependencies as constraints by partitioning the workflow into different parts based on the patterns and QoS requirements. One of the contributions from Avanes work is that the network dynamics has been also included in the procedure of constraint resolving. Resource provision plays an important role to improve the fault tolerance and the performance of the workflow (Juve & Deelman, 2008). Basically, provisioning can be either static or dynamic. Advance reservation is a typical static provisioning mechanism, and several batch based schedulers support it. Based on the quality requirements, the workflow engine reserves computing resources and time slots from the Grid resource manager. One of the disadvantages of static provisioning is its overhead on the total cost for computing the workflow. To improve this, Raicu, Zhao, Dumitrescu, Foster and Wilde (2007) proposed multi level scheduling strategies, in which the application level scheduler is able to interact with the low level resource manager to tune the requirements at runtime. This approach introduces a dynamic component in the provisioning process.

The provenance service tracks the events occurred in the workflow execution, and allows scientists to trace the evolution of data computed in the workflow and to obtain insights in the experiment processes. Moreover, provenance data can also be used to debug errors of the workflow execution and optimize the workflow design. The Open Provenance Model (OPM) (Moreau et al., 2008) emerges as a standard model to represent workflow provenance information. Including QoS information of the workflow processes and the execution in the provenance model allows scientists to analyze the quality of the services and the workflow scheduling. Michlmayr, Rosenberg, Leitner and Dustdar (2009) provide the provenance service using a QoS aware middleware, which records the changes of the service quality as events. Evaluating trust and reliability of the

provenance data itself has also been discussed in the literature (Rajbhandari, Contes, Rana, Deora, & Wootten, 2006). However, research on the provenance model which includes the QoS information of the workflow processes is still in its very early stage.

The above technologies contribute necessary building blocks to enable the delivery of large quantity digital content over distributed environments. However, putting them all together and providing quality guarantee for overall applications in terms of high quality of both media content and the delivery is not trivial; not only the network QoS is not directly included the scheduling loop of high level workflow systems, but also optimizing the usage of network services require real time monitoring of network infrastructures is not an easy task. In the next section, we will formulate the problem and propose an agent based solution to enable quality guarantee for workflows that handle content delivery over advanced network.

3. NETWORK QoS AWARE WORKFLOW PLANNING

In the previous section, we reviewed different technologies involved in delivering content over network, and we argued that including network QoS in high level applications is essential to enable global quality guarantee on applications. In this section, we will discuss an agent based solution for this problem. Our focus is on improving existing workflow systems by adding an extra planner.

We had two alternatives when we looked at the inclusions of QoS aware functionalities in scientific workflow systems: 1) re-engineer the functional components of existing systems to include the QoS support, or 2) consider existing systems as legacy systems, and provide QoS support as pluggable components to the systems. Each alternative has advantages and disadvantages, we chose ultimately the second approach.

3.1 Design Requirements

Network QoS support can be applied to: QoS aware resource selection, resource provisioning and quality assured workflow execution. The designed system thus needs to meet the following functional requirements:

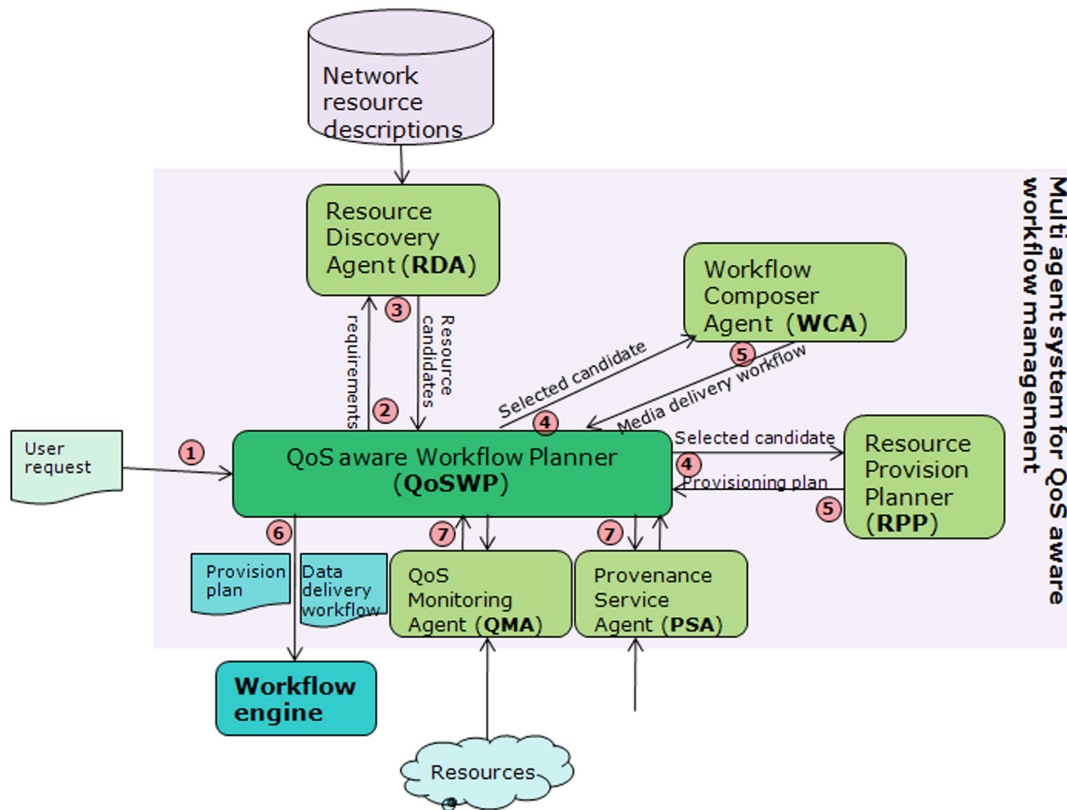
1. The system must include QoS aware resource discovery and selection of network resources. To support this we must have descriptions of the network resources and their quality attributes, we must provide a search tool that checks the suitable resources based on the input requirements.
2. The system should be able to generate a resource provisioning plan for the selected resources based on the input requirements. The plan is made based on the provisioning services that the available network infrastructure provides.
3. The system should be able to generate workflows that handle large data movement between network resources with guaranteed data transfer quality, and wrap the generated workflow as a service, which can be executed standalone or included in a third party workflow.
4. At runtime, the system should provide monitoring services to track the actual state of the network resources. It should also provide interfaces for third party workflows to invoke during their provenance procedure to record the runtime information.

The Agent Oriented (AO) methodology complements the object and component oriented methods with knowledge related notions to manage system complexity (Massonet, Deville, & Neve, 2002), and emerges as an important modeling and engineering approach for constructing complex systems, such as workflow management systems. The concept of *agents* originated in the mid-1950s as a *'soft robot' living and doing its*

business within the computer's world (Kay, 1984). Wooldridge distinguished three types of agent architectures: deliberative, reactive and hybrid (Wooldridge & Jennings, 1995). The difference between the deliberative and reactive architectures is that the former incorporates a detailed and accurate symbolic description of the external world and uses sophisticated logic to reason about the activities, while the latter one only implements a stimulus-reaction scheme. Reactive architectures are easier to implement but lack a subtle reasoning capability. Hybrids of the two schemes are commonly used. During the past two decades, agent based models, in particular reactive models, have been applied as an advanced technology in modeling and constructing complex systems. Agent frameworks, such as FIPA (Bellifemine, Poggi, & Rimassa, 2001), abstract the structure of basic agents and define standardized communication languages to represent interactions between agents, which facilitate the implementation of agent-based applications.

JADE (Java Agent DEvelopment Framework) is a free software and distributed by Telecom Italy (Case et al., 2002). Fully implemented in Java, JADE realizes a FIPA compliant multi agent middleware. In our project, we choose JADE as the implementation framework. Firstly, the JADE platform can be distributed across machines and the configuration can be controlled via a remote GUI; the Java language makes the development portable; the JADE framework allows agents move from one machine to another at runtime. Moreover, being compliant to the FIPA protocol, JADE provides a standard architecture for scheduling agent activities, which makes the inclusion of high level functionality easy, e.g., adding a Prolog module for activity reasoning. Finally, the ontology enabled agent communication between agents promotes seamless integration among the semantic network description, QoS aware searching modules, underlying models of workflow descriptions, and other necessary functional components of our system.

Figure 4. The basic architecture of the NETwork QoS planner



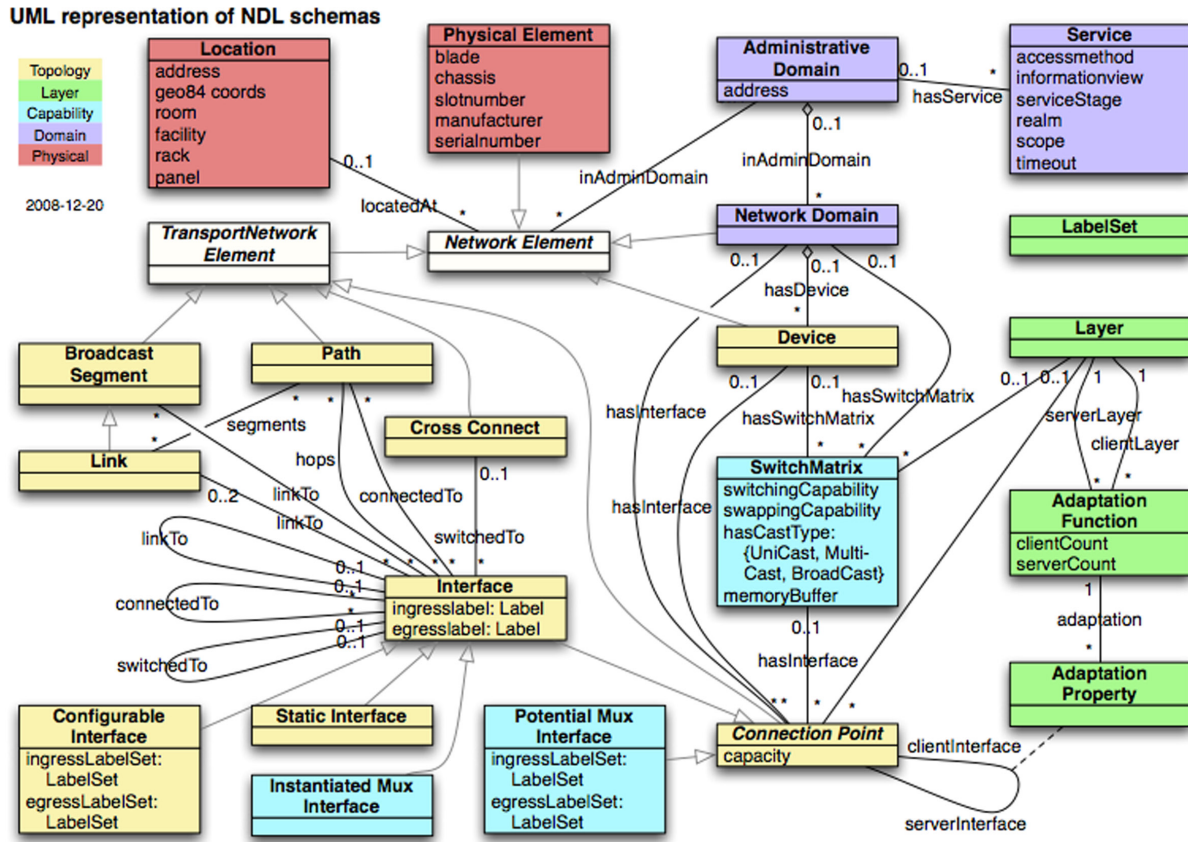
3.2 An Agent Based QoS Workflow Planner

We propose an agent based architecture, composed of a *QoS aware workflow planner (QoSWP)* and five more agents: a *Resource Discovery Agent (RDA)*, a *Workflow Composition Agent (WCA)*, a *Resource Provisioning Planner (RPP)*, a *QoS Monitor Agent (QMA)* and a *Provenance Service Agent (PSA)*. Figure 3 provides a conceptual schema of our agent system.

The QoSWP coordinates the other agents to select suitable services, to propose optimal network connections between the services, and to create the necessary scripts for the workflow engine to invoke the requested services. A typical use case scenario will illustrate the role of each component (see Figure1). The QoSWP receives the request for data process services and the service requirements

from the user (step1). After that, the RDA reads the description of the resources and the network topologies from the registry, and searches suitable data sources and destinations, and network paths between them (step2). The RDA returns a list of qualified candidates, and sorts them based on the quality metrics of each candidate (step3). From the candidates, the QoSWP selects the best one, and requests WCA and RPP to generate a resource provisioning plan and a data transfer workflow (step4 and step5), both of which will be executed by the workflow engine (step6). At run time, the QMA monitors the actual state of the resources and checks whether the global quality required by the workflow is satisfied (step7). Based on the states updated by the QMA, the QoSWP decides if the resources of the workflow should be adapted. The provenance service records events in the resources provisioning, allocation, and combine

Figure 5. The concept schema of network description language



the actual state of the quality attributes with the log data (step7).

In the rest of the section, we will discuss the detailed design issues.

3.3 Semantic Network Description

Semantic web technologies (Berners-Lee, Hender, & Lassila, 2001) provide suitable solutions to describe network topologies, devices, and the QoS requirements for data *and network resource*. We have developed two ontologies for describing CineGrid services and network topologies respectively. The CineGrid Description Language (CDL) describes the services and resources available on top of the network infrastructure. The Network Description Language (NDL) models the different levels of a network infrastructure:

physical, domain, capability, layer and topology (Ham et al., 2008).

NDL (see: <http://www.science.uva.nl/research/sne/ndl>) comprises of a series of RDF schemas that categorize information for network topologies, network technology layers, network device configurations, capabilities, and network topology aggregations. The main use cases so far have been generation of network maps, lightweight offline path finding and more recently multi-layer path finding, and network topology information exchange. NDL has been used primarily in the research community in the Netherlands: UvA, SARA and SURFnet (SURFnet, 2002). It also has been applied to the GLIF Optical Lightpath Exchanges (see: <http://www.glif.is>).

NDL chooses RDF because 1) RDF allows easier exchange of information between independent

domains and 2) it is easily extendible and it allows integration of independent data models developed in other fields, by other researchers. Several tools that consume RDF data are publicly available and make the use of this syntax straightforward. NDL is a modular set of schemata, defining an ontology to describe computer networks. Figure 5 shows the UML diagram of the NDL schemas.

1. The topology schema describes devices, interfaces and connections between them on a single layer. The classes and properties in the topology schema describe the topology of a hybrid network, without detailed information on the technical aspects of the connections and their operating layer. Through this lightweight schema NDL provides an easy toolset for basic information exchange and path finding.
2. The layer schema describes generic properties of network technologies, and the relation between network layers. The topology schema defines network topologies on a single layer. The NDL layer schema allows applications to describe multi-layer networks, like hybrid networks. The NDL layer schema is based on a formal model, which uses ITU-T G.805 functional elements (see: <http://www.itu.int/rec/T-REC-G.805/en>) and the concept of labels as described in GMPLS (see: <http://www.ietf.org/rfc/rfc3945.txt>).
3. The capability schema describes device capabilities.
4. The domain schema describes administrative domains, services within a domain, and how to give an aggregated view of the network in a domain. It allows network operators to provide an aggregated view of their domain to neighboring domains, rather than the full topology. An important concept in the domain schema is that of Service Descriptions.
5. Service descriptions allow domains to point applications to the (web) services they of-

fer. The idea is that domains publish static information in NDL, and provide a web service for dynamic information or more confidential data, like reservation requests. Furthermore, different domains will have different opinions on what is “static” and “non-sensitive”.

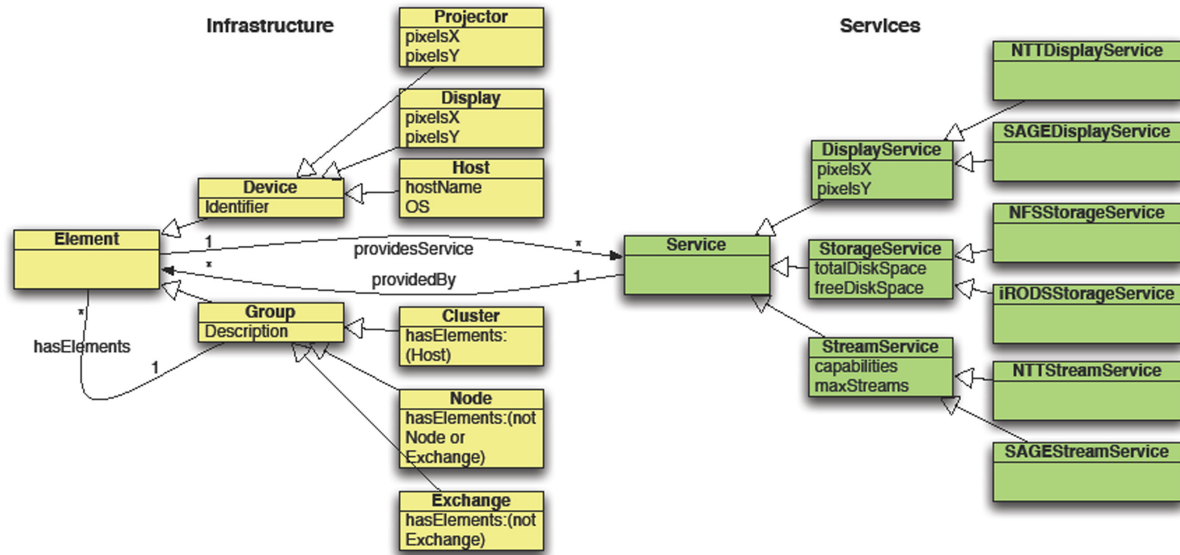
6. The physical schema describes the physical aspects of network elements.

The CineGrid Description Language (CDL) defines an ontology for describing CineGrid resources. CDL consists of two parts, an infrastructure ontology and a service ontology. The service ontology describes the tasks a device can perform for the users of the CineGrid Exchange. Devices in the Exchange nodes perform multiple types of tasks, possibly at the same time. We map these tasks into services; and the user of the ontology deals directly with services.

In order to do resource planning the service ontology had to be mapped onto network infrastructure descriptions. The infrastructure part of the CDL provides a loosely mapped paradigm between the CDL and the underlying network description schemas. In this way, other network description schemas, for example the Network Markup Language currently under development in the OGF (OGF, 2011), can also be integrated with the CDL.

Figure 6 shows the classes used in CDL and their relations to each other. The classes related to the infrastructure are on the left side and the classes representing the different services are on the right side. Let’s begin with the service ontology in the CDL. We have a generic Service class, which contains three sub classes: DisplayService, StorageService and StreamService. The DisplayService defines all the common properties needed for a display service; StorageService defines all the common properties needed by a storage service; and, StreamService defines all the common properties needed by a streaming service. We identify seven specific implementations of these

Figure 6. The concept schema of CineGrid description language (CDL)



three main services. SAGEDisplayService is a service to display video on tiled panels running SAGE (Sage, 2010); this specific service inherits all the properties of DisplayService, which in turn inherits the properties of Service. We intend to extend the supported services with video manipulation and transcoding services in the near future. Users of CDL can also create their own services' descriptions, by simply inheriting from the general Service class. The infrastructure part defines the basic building blocks that reflect the hierarchical structure of the CineGrid Exchange. We define the Element class to describe the common characteristics of elements in the infrastructure. There are two classes which inherit directly from Element, and seven more specific classes that inherit from these two intermediate classes:

Device for single devices.

- **Projector** to represent stand alone video projection devices.
- **Display** to represent video display devices.
- **Host** to represent a single host.

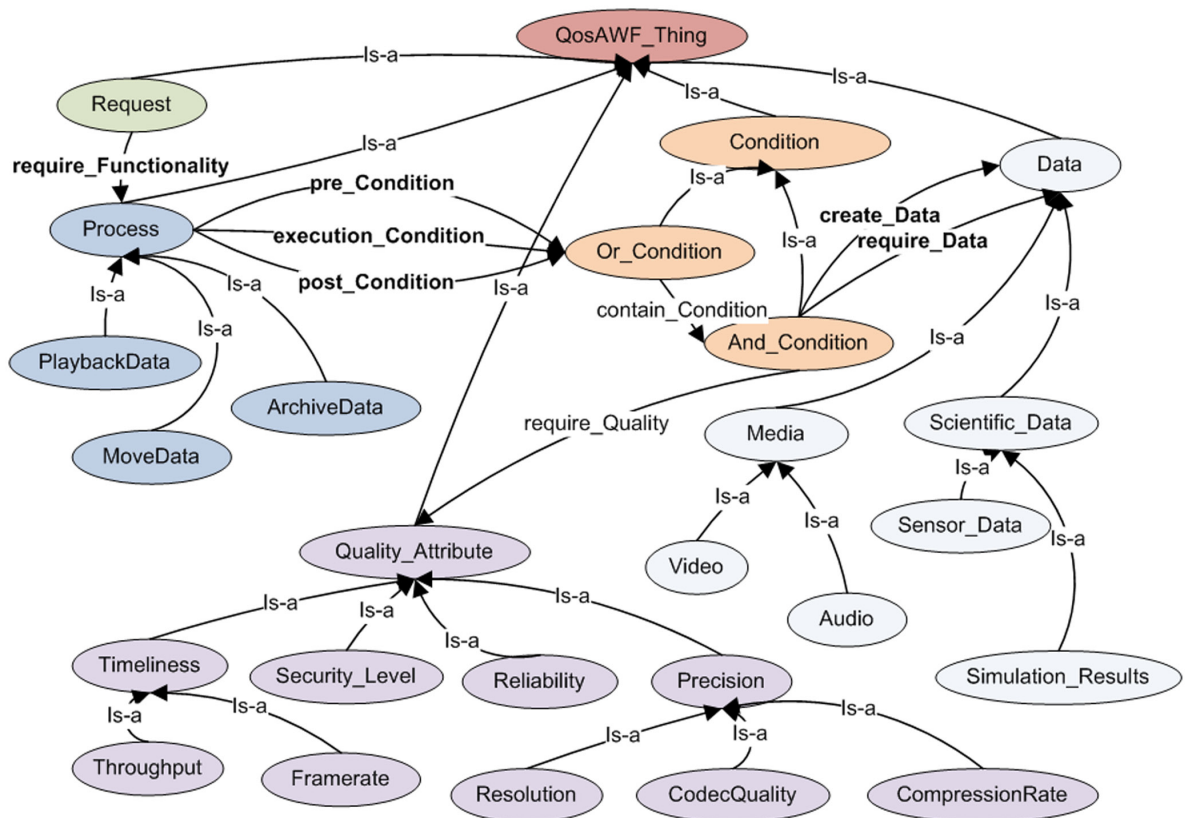
Group to represent element groups.

- **Cluster** to represent computing clusters.
- **Node** to represent a collection of devices working as a single entity.
- **Exchange** to represent an exchange platform (e.g. the CineGrid Exchange).

Elements provide the services to the users. The `cdl:providesService` property links services to elements.

We map NDL to the CDL infrastructure ontology using the `owl:sameAs` property; this allows us to say that a certain object in one namespace is the same as an object in the other namespace. This has a very essential practical implication for the operation of the CineGrid Exchange. Network administrators can describe the network portion in NDL and CineGrid node administrators can link their device objects to the NDL objects, and do reasoning on both the CineGrid Exchange and the supporting network topology.

Figure 7. The abstract workflow process schema



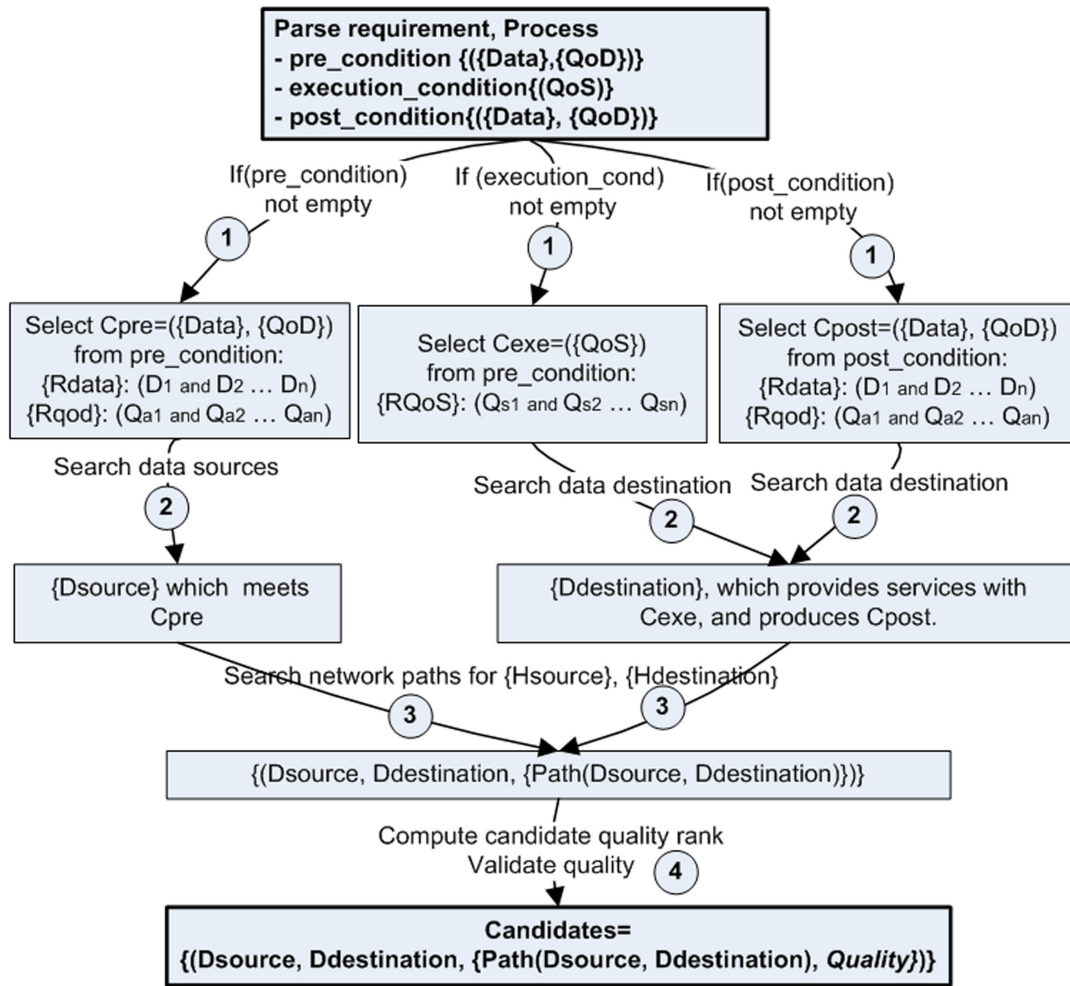
3.4 QoS Abstract Workflow Process Modeling

Based on early work (Caragea & Syeda-Mahmood, 2004; Klusch, Fries, & Sycara, 2006; Bubak, Gubala, Kapalka, Malawski, & Rycerz, 2005), we propose an ontology for describing abstract workflows (*qosawf.owl*). Figure 7 shows the graphical representation.

This ontology defines the basic concepts of workflow processes, pre/post/execution conditions of the process, media data, and quality attributes. A user's request is described as an object of the *Request* class, and a *Request* consists of one or more *Processes* which can be accessed via the *requestFunctionality* property. A *Process* class uses *pre Condition* and *post Condition* to indicate the requirements for *Data* the process requires

and generates, and the quality for the required data. The *Process* class also uses *execution Condition* to indicate the service quality for the process. In the current definition, *Data* contains two specific types: *Media* and *Scientific Data*. And the service quality is modeled as a set of *Quality Attributes*. Based on the QoS taxonomy defined in (Sabata, Chatterjee, Davis, Sydir, & Lawrence, 1997), *Quality attribute* can more specifically be *Precision*, *Timeliness*, *Reliability* and *Security Level*. In our case, where the pre and post conditions consist of requirements for data and the data quality, *and Condition* and *or Condition* are the two most important types. Using the above ontology, a user is able to formulate a request for obtaining and playing back specific video material with a minimal resolution and frame rate.

Figure 8. The network service selection



3.5 Resource Discovery

Figure 8 shows the basic procedures in the network resource selection process. The resource discovery agent 1) parses the input description, 2) searches suitable CineGrid resources which meet the requirements for being the data sources and destination, 3) looks for optimal network paths between them, and 4) computes the quality of resource candidates and proposes solutions, as shown in Figure 8.

1) *Step 1: QoS requirement parsing.* The input of RDA contains functional requirements for data operation (*Process*) and the quality requirement

for both the operation and the data. The current QoS-AWF schema allows one input description to contain only one instance a *Process* concept. The parsing procedure obtains the pre/execution/post condition of the process. The *pre Condition* and *post Condition* of a process contains both requirements for data, such as data type and properties, and for the quality of the data, such as resolution if the data is a video file. The *execution Condition* gives QoS requirements for the process. For instance, the *pre Condition* contains both content and quality requirements for data, as follows:

$$C_{pre} = \{C_{pre}_{data} \text{ and } C_{pre}_{qod}\}$$

$C_{pre}_{data} = \{C_{pre}_{d1} \text{ or } C_{pre}_{d2} \text{ or } \dots C_{pre}_{dn}\}$ in which

$C_{pre}_{di} = \{C_{pred}_{i1} \text{ and } C_{pred}_{i2} \text{ and } \dots C_{pred}_{dim}\};$

$C_{pre}_{qod} = \{C_{pre}_{qod1} \text{ or } C_{pre}_{qod2} \text{ or } \dots C_{pre}_{qodn}\}$ in which

$C_{pre}_{qodi} = \{C_{pre}_{qodi1} \text{ and } C_{pre}_{qodi2} \text{ and } \dots C_{pre}_{qodik}\}.$

The RDA selects an element in the pre/execution/post condition, and uses it as the constraints for the resource search.

2) *Step 2: data and the operation.* From the data requirements derived from the step 1, the hosts that contain the required data, namely *data sources*, and the hosts that will consume or store the data, namely *data destinations*, are identified. From the resource description, the RDA derives the set of storage services that contains the *Data* instance that meets the required type, and quality. In CineGrid, each *Data* instance is associated with a *Meta data* object, which can be accessed via the property *hasMetadata*. Therefore, the sources of data are located by searching instances of *Data* which contain meta data meet the requirements abstracted from the *pre Condition*. Using the property of *cdl:providedby* and *owl:sameAs*, the actual host that stores data can then be derived.

The destination of the data is derived from the process types described in the requirement. As we mentioned above, based on the type of data operations, we abstract three basic process types: *MoveData*, *PlayData*, and *ArchiveData*. For the process of *PlayData*, *post Condition* can be empty, because the process does not generate data. The processes are linked to the actual services of CineGrid via property *implementedBy*. Therefore, the process destination of the data is determined by both the location of the implemented services

and the location of the data required in the *post condition*.

3) *Step 3: network paths.* The next step is to find all network paths between the data sources and destinations. Using NDL, a network path can be found using three properties: *link to*, *connect to* and *switch to*. The *link to* property indicates that two network devices are directly connected via a physical line, while *connect to* refers to a connection which might include unknown devices between the two end points of the path. The *connect to* property is mostly used in the situation where two devices belong to two different domains and the detailed physical connections between them is not clear or not open to public due to administration rules. The *switch to* property is only used in a switch device to indicate the connectivity between different ports in the device. The RDF triples defined in the network topology description give a suitable *graph* representation for finding network paths.

4) *Step 4: quality ranking:* The first three steps return resource candidates, which are represented as (*source, destination, path*). The quality of the resource can be evaluated at multiple levels: 1) the quality of data, 2) the quality of the storage/stream services, 3) quality of the hosts, which provide the services, and 4) the network connection between hosts. From the CDL and NDL ontology, the RDA can abstract the following quality attributes: 1) the quality of data, such as compressed ratio and resolution, from the data catalogue of resources 2) the properties of host, such as its CPU speed, memory size and the available storage space, 3) the network bandwidth of network connections. From the quality attributes and the quality requirements defined for the process, the RDA applies the following rules to filter unqualified candidates from the searched results:

1. The RDA first checks if the data and services meet the quality requirement.
2. Then compute the bandwidth of the candidate network paths, only the candidates that have

bandwidth meet the minimal data transfer rate are kept.

3. The RDA sorts all qualified candidates based on the quality of the hosts that provide data or visualization service, and the bandwidths of the network connections.

We have compared different options to realize the resource search mechanism; we have evaluated several Query languages (RQL, RDQL, N3, Versa, SeRQL, SPARQL) and Rule languages (SWRL, Prolog/RDF lib, JESS etc.). We have finally chosen the RDF library of SWI-Prolog; its triple based manipulation interface is flexible for the high level language we use to implement the agents (Java); it is also easy to access the runtime state of the triples. Finally, the Prolog language provides effective solutions to realize graph path findings. The FIPA (FIPA, 2011) standards provide a suitable architecture to implement distributed agents in our system. The Agent Communication Language (ACL) allows agents to exchange messages using an explicitly defined semantic schema, which allows seamless integration between agents and remote Ontology knowledge bases. In the current prototype, the RDA receives the URI of the user requirements and network resources from the QoS WP. The RDA parses the given abstract workflow and searches the resource description; it returns results in the form of (storage host, visualization host, path, quality rank).

3.6 Network Provisioning

The ultimate goal of our NEWQoS Planner is to automatically find a quality optimal network path for delivering and processing media content. Provisioning a path in the network is an important step to make the network available to workflow. So far we have only integrated the NEWQoS Planner in an ad hoc fashion with our network test bed. The planner can execute some scripts to create network paths in our experimental network. Obviously, such an approach does not scale to

larger networks. There will be problems with authentication and authorization, supporting different kinds of network equipment, compatibility with other systems, et cetera. An easier solution is to integrate with existing network management tools. This will allow the workflows to be used for intra-domain path selection in many more networks, or even inter-domain using the global GLIF network.

There are currently several network provisioning systems that allow integration with other applications. ESnet and Internet2, two large research and education networks in the USA, have developed the On-Demand Secure Circuits and Advance Reservation System (OSCARS) (Gridnets, 2006). This system allows users to create reservations for circuits in the ESnet and Internet2 network. The system can use either MPLS and RSVP to create connections, in the case of the ESnet network, or integrate with Internet2's Dynamic Circuit Network and provision VLANs on their national backbone network.

The OSCARS system allows users to specify different properties that a circuit reservation should fulfill, such as bandwidth, or a specific VLAN number. The OSCARS system also allows applications to use the web service interface for a more direct provisioning service. This kind of integration would be ideal for our NEWQoS Planner.

Another system currently available is the OpenDRAC system (OpenDrac, 2011), originally developed by Nortel Networks. This provisioning system is currently in use on the SURFnet network in the Netherlands. The management system provides the network operator with the tools to manage and monitor the network, but also has an interface for users to request lightpaths. Depending on the access rights of the user he can request lightpaths from several locations with different capacities. OpenDRAC allows users to specify other attributes of the circuit as well, such as bandwidth, VLAN ID, etc., depending on the capabilities of the underlying network. The OpenDRAC system also features a web service

interface, which allows for simple integration with other applications.

Currently, the different management systems such as OpenDRAC and OSCARS are not directly compatible, meaning that it is not possible to create a reservation that goes from a domain managed by OpenDRAC to an OSCARS managed domain or vice-versa. There currently is a demonstration project going, called Fenius, to implement a simple inter-domain interface between these provisioning systems to allow for the automatic set up of inter-domain circuits. This has been demonstrated successfully at the SuperComputing 2010 conference (Fenius, 2011). In the future this will converge to a standard currently in development in the Open Grid Forum, called the Network Service Interface (Nsi-Wg, 2011). This standard will allow provisioning systems to interact with each other to automatically create inter-domain circuits for the users and their applications.

4. A USE CASE

The system presented in section 3 was originally developed in the context of CineGrid. An important mission of the CineGrid project is to provide a dedicated network environment to connect distributed parties from different domains to share large quantities of very high quality digital media, such as the high definition video material used in the movie industry. The results reach beyond the workflow field, and they can be beneficial to understand how advanced network connections enhance the digital media delivery in the academic and education context. In this section, we will demonstrate how the designed architecture works in the following use case, and discuss the technical considerations to prototype the system.

4.1 Basic Scenario

We are focusing on a *digital media delivery on demand* use case: the goal is to retrieve media

material from the infrastructure, and request quality guaranteed connections to deliver the data to qualified nodes for further processing, such as playback or visualization. Using the proposed agent framework, the use case will be prototyped as follows:

1. The user uses the schema provided by the system to describe the name and properties of the media, and to specify the quality requirements for visualizing the data.
2. The QoS WP parses the user input and creates queries for the RDA to look for data sources of the media.
3. Based on the input requirements, the RDA looks for the data repositories that contain the required media, and the visualization devices that meet the required playback quality. Then the RDA looks for all possible network paths between the sources and the visualization devices.
4. The RDA returns a list of candidates in the form of (source, destination, path) triplets, and the candidates are ordered based on the quality they provide. The QoS WP selects the best candidate from the list and sends it to the RPP and WCA to make a resource provisioning plan, and to create a workflow that can deliver the media from the source to the visualization device, and to play it back in the visualization device.
5. To help RPP and WCA make the provision plan and the workflow compliant to a specific workflow engine, the QoS WP also explicitly tells the RPP and WCA what language of the third party engine will use.
6. After receiving the scripts generated by the RPP and WCA, the QoS WP sends them to the third party engine to execute the provisioning plan and the delivery plan.

Figure 9. The test bed of the use case and the screen snapshot of the user portal

The figure consists of two parts. On the left is a screenshot of a web browser showing the CineGrid distribution center Amsterdam portal. The browser address bar shows the URL `http://cinegrid.uvalight.nl/portal/contentbrowser/`. The page header includes the CineGrid logo and the text "distribution center Amsterdam". Below the header is a navigation bar with links: "ome | About | Browse Content | CDL Demo | cinegrid.org | cinegrid.nl". A search bar contains the text `C:\Users\zhiming\Projects\NewQoSPlanner-0.1\ontology\dewaag.r`. Below the search bar, it shows "1 matching items for: De Waag". There are several filter dropdown menus: "All nodes", "All sizes", "All formats", "All durations", and "Sort by name". A thumbnail image of a building is shown next to the title "De Waag". To the right of the thumbnail is a description: "A shot of De Waag on the Nieuwmarkt in Amsterdam." Below the description are the following details: "Available formats: 4096x2160 dxt (2.0 GB), 1920x1080 dxt (1.6 GB)", "Created: 2 years, 8 months ago", "Author: Cinegrid.", and "Categories:".

On the right is a network topology diagram. It shows a central "SURFnet" cloud connected to three external nodes: "National Film&TV Association", "SARA", and "De Waag Society". Below this, a "University of Amsterdam" cloud is connected to "FNWI" and "Internet". The "University of Amsterdam" cloud contains several nodes: "UvA Starlight", "Force10", "Houdini", "Cinegrid", "Cinevideo", "Cineedit", and "Cinestream". The "Internet" cloud is connected to "Cinegrid", "Cinevideo", "Cineedit", and "Cinestream". The diagram shows various network links with speeds like "1Gb" and "10Gb".

4.2 Prototype

We are prototyping our ideas using a small portion of the CineGrid infrastructure as a test bed. Four locations in Amsterdam host CineGrid resources and are connected via dedicated and configurable circuits provided by SURFnet. Using the description languages presented in section 3.3, we have described the use case test bed. Four locations (UvA, SARA, De Waag and the Dutch Film and TV institute) in Amsterdam are connected with up to two dedicated switchable 1Gbit/s links, which can be dynamically changed between locations using the openDRAC network provisioning software used by SURFnet. In the use case, a portal is developed as the user front end to the NEWQoSPlanner. The portal allows a user to search a movie in two ways: browsing the entire media base by category, or searching a specific movie by providing relevant meta information. The user can also specify the quality requirements for playing back the movie. The Portal encodes all

the user input using the schema discussed in 3.4 and sends it to the QoSPlanner. The QoSPlanner returns the results as a list of ordered candidates based on their quality. The user can then select an optimal option and play it back. The use case has also been partially demonstrated in the Super Computing Conference 2010 (Zhao, Koning, Grosso, & de Laat, 2010). Figure 9 shows the screen snapshot and the test bed topology.

4.3 Discussion

The media delivery use case presented in the previous section demonstrates the feasibility of including network QoS in the resource planning of the high level data intensive applications. The planner extends the control loop of high level applications to tune the behavior of the network infrastructure. In our research we initially chose workflow as an application execution manager; however, our solution does not preclude its use

in other systems that have control on the network behavior, *e.g.* clouds' front ends.

The planner currently assumes that all the infrastructure descriptions share NDL and CDL as the schema. Descriptions of a large scale environment are often composed and maintained by the different owners. Therefore, these descriptions do not always share the same level of details, and do not provide the same level of information to external parties due to different administrative policies. Often network administrators prefer not to publish the entire network topology description to minimize security risks. In such cases, the planner has to obtain the descriptions from all involved parties in order to do a resource query on the entire environment. Even then, the application might not have access to all the information needed, *e.g.* the quality attributes that applications require may not be explicitly stated in the collection of descriptions.

There is clearly a big gap between the requirements from high level applications and the availability of the semantic information provided by the distributed environments. On the one side, information from different infrastructure domains may be partially accessible and may have overlaps or conflicts on certain infrastructure due to evolution and or the delay of maintenance of the descriptions. On the other side, applications require different types of information from the descriptions to find suitable resources or to make decisions on resource allocations. To apply the planner in a large scale infrastructure, we plan to develop an information preprocessing framework that will ensure that the infrastructure descriptions meet the requirements of the planner.

5. FUTURE RESEARCH DIRECTIONS

In this chapter, we discussed how network QoS and workflow systems are used to optimize the data movement processes in the media delivery. The use of dedicated network paths, such as the

ones described previously, certainly increases the QoS that can be offered to streaming applications. We believe this trend toward pay-per-use advanced network services will increase; as service provider will realize the earning potential of offering this guaranteed quality to specific applications. It is foreseeable that the same model will slowly trickle down to the single individual customers. Still we envision two main evolutions in the offering of these services: one related to increases in network bandwidths and one to a change in network provisioning model.

First, higher and higher bandwidths are in fact available to send data between end points and can certainly be used for streaming video; 40Gbit/s channels are now a reality (Dumitru et al., 2010) and they will increase the potential for remote collaborations on digital material editing and production. While certain quality of service aspects will automatically improve by the larger network pipes available, we believe the problem of selecting, configuring and matching resources at the edge of the networks will still require the same semantic based approach to workflow delivery we have started to develop.

Second, we are moving already towards the use of Next Generation Ethernet as the supporting technology for education and research networks; here individual applications use a QoS-enabled vLAN in the core infrastructure, created with the use of Provider Backbone Transport and Provider Link State Bridging technologies. This will require the proper evaluation of the semantic information that needs to be included in the network ontologies describing these networks, as proper path and resource selection will rely on smart and usable categorizations. It is now possible to combine semantic descriptions of the network, and the available media content on the network to provide a QoS aware workflow for media delivery. This is only the first step into making available all the relevant resources for generic workflows.

Third, workflow systems emerge as a key service to glue different levels of technologies

and hide the underlying details from the high level applications. However, it also introduces new challenges in application development and the validation of workflow results. An important issue in our research agenda is to develop a suitable semantic model is needed for logging and querying workflow processes with the network QoS information. With this model, namely provenance model, the runtime information of the workflow and the network events will be recorded for the further querying for reproducing execution scenarios of the workflow. Furthermore, we also plan to extend the ontologies to other kinds of resources. This will make it possible to define any kind of computing workflow that can involve searching for content, computation, data transport and visualization.

6. CONCLUSION

Quality control at the network level is crucial for workflow applications in which large data movement is the performance bottleneck. Advanced network infrastructures provide guaranteed services for high level data intensive applications. To bridge workflow requirements and the services provided by the network we propose to use the semantic web technology. We developed the QoSAWF ontology to provide lightweight solution to describing QoS requirements for data operation related workflow processes.

Our network resource discovery agent provides a necessary service for tuning data transfer processes from the application level. The NEWQoSPlanner is the first step towards the direction of network quality adaptive workflow planner, and it can play a role in the development of CDNs based on the latest hybrid network architectures in place of a traditional internet fabric.

REFERENCES

- W3C. (2010). *Resource description framework*. Retrieved from <http://www.w3.org/RDF/>
- Alaya, I., Solnon, C., & Ghdira, K. (2007). Ant colony optimization for multi-objective optimization problems. In *Proceedings of IEEE International Conference on Tools with Artificial Intelligence*, (pp. 450–457).
- Avanes, A., & Freytag, J. (2008). Adaptive workflow scheduling under resource allocation constraints and network dynamics. In *Proceedings of VLDB Endowment*, 1(2), 1631–1637.
- Bellifemine, F., Poggi, A., & Rimassa, G. (2001). JADE: A FIPA2000 compliant agent development environment. In *Proceedings of the Fifth International Conference on Autonomous Agents*, (pp. 216–217). ACM Press.
- Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The Semantic Web. *Scientific American*, 284, 34–43. doi:10.1038/scientificamerican0501-34
- Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., & Weiss, W. (1998). *An architecture for differentiated service*. Request for Comments 2475, IETF. Retrieved June 10, 2011, from <http://www.ietf.org/rfc/rfc2475.txt>
- Braden, R., Clark, D., & Shenker, S. (1994). *Integrated services in the internet architecture: An overview*. Request for Comments 1633, IETF. Retrieved June 10, 2011, from <http://www.ietf.org/rfc/rfc1633.txt>
- Bubak, M., Gubala, T., Kapalka, M., Malawski, M., & Rycerz, K. (2005). Workflow composer and service registry for grid applications. *Future Generation Computer Systems*, 21(1), 79–86. doi:10.1016/j.future.2004.09.021

- Caragea, D., & Syeda-Mahmood, T. (2004). Semantic API matching for automatic service composition. In *WWW Alt. '04: Proceedings of the 13th international World Wide Web Conference on Alternate Track Papers & Posters*, (pp. 436–437). New York, NY, USA.
- Case, J., Mundy, R., Partain, D., & Stewart, B. (2002). *Introduction and applicability statements for internet-standard management framework. RFC 3410*. Informational.
- De Laat, C., & Blom, J. (2000). User-level performance monitoring program. In *Proceedings of TERENA Network Conference 2000*, Lisbon, Portugal.
- De Laat, C., Radius, E., & Wallace, S. (2003). The rationale of the current optical networking initiatives. *3rd Biennial International Grid Applications-Driven Testbed Event. Future Generation Computer Systems*, 19(6), 999–1008. doi:10.1016/S0167-739X(03)00077-3
- Deelman, E., Gannon, D., Shields, M., & Taylor, I. (2009). Workflows and e-Science: An overview of workflow system features and capabilities. *Future Generation Computer Systems*, 25(5), 528–540. doi:10.1016/j.future.2008.06.012
- DMTF. (2011). *Common information model (CIM)*. Retrieved from <http://www.dmtf.org/standards/cim/>
- Dumitru, C., Koning, R., & de Laat, C. (2010). *ClearStream: End-to-end ultra fast transmission over a wide area 40Gbit/s Lambda*. Demo Supercomputing 2010.
- Fang, Q., Peng, X., Liu, Q., & Hu, Y. (2009). *A global QoS optimizing Web services selection algorithm based on moaco for dynamic web service composition* (pp. 37–42). International Forum on Information Technology and Applications.
- Farrel, A., & Bryskin, I. (2006). *GMPLS: Architecture and applications* (1st ed.). Morgan Kaufmann.
- Fenius. (2011). Retrieved June 17, 2011, from <http://code.google.com/p/fenius/>
- FIPA. (2011). *The Foundation for Intelligent Physical Agents*. Retrieved from www.fipa.org
- Fortino, G., Russo, W., Mastroianni, C., Palau, C. E., & Esteve, M. (2007). CDN-supported collaborative media streaming control. *IEEE Multimedia Magazine*, 14(2), 60–71. doi:10.1109/MMUL.2007.29
- Grasa, E., Figuerola, S., Recio, J., Lopez, A., Palol, M., & Ribes, L. (2006). Video transcoding in a Grid network with user controlled LightPaths. *Future Generation Computer Systems*, 22(8), 920–928. doi:10.1016/j.future.2006.03.003
- Gridnets. (2006). Retrieved June 17, 2011 from <https://oscars.es.net/OSCARS/docs/papers/gridnets.pdf>
- Grossman, R., Gu, Y., Hanley, D., Sabala, M., Mambretti, J., & Szalay, A. (2006). Data mining middleware for wide-area high-performance networks. *Future Generation Computer Systems*, 22(8), 940–948. doi:10.1016/j.future.2006.03.024
- Ham, J., Dijkstra, F., Grosso, P., Pol, P., Toonk, A., & de Laat, C. (2008). A distributed topology information system for optical networks based on the semantic web. *Optical Switching and Networking*, 5(2–3), 85–93.
- Ham, J., Dijkstra, F., Travostino, F., Andree, H., & de Laat, C. (2005). *Using RDF to describe networks*. Future Generation Computer Systems, Feature topic iGrid.
- Harada, F., Ushio, T., & Nakamoto, Y. (2007). Adaptive resource allocation control for fair QoS management. *IEEE Transactions on Computers*, 1(56), 344–357. doi:10.1109/TC.2007.39

- Holub, P., Matyska, L., Liska, M., Hejtmanek, L., Denemark, J., & Rebok, T. (2006). High-definition multimedia for multiparty low-latency interactive communication. *Future Generation Computer Systems*, 22(8), 856–861. doi:10.1016/j.future.2006.03.014
- IETF. (2011). *Netconf working group*. Retrieved June 17, 2011, from <http://www.ops.ietf.org/netconf/>
- International Telecommunications Union. (ITU) (March 2000). *Generic functional architecture for transport networks. Recommendation ITU-T G.805*. Retrieved June 17, 2011 from <http://www.itu.int/rec/T-REC-G.805/>
- Jo, J., Hong, W., Lee, S., Kim, D., Kim, J., & Byeon, O. (2006). Interactive 3D HD video transport for e-science collaboration over UCLP-enabled GLORIAD lightpath. *Future Generation Computer Systems*, 22(8), 884–891. doi:10.1016/j.future.2006.03.006
- Juve, G., & Deelman, E. (2008). Resource provisioning options for large-scale scientific workflows. In *Proceedings of ESCIENCE '08: The 2008 Fourth IEEE International Conference on eScience*, (pp 608– 613). Washington, DC: IEEE Computer Society.
- Kay, A. (1984). Computer software. *Scientific American*, 251(3), 53–59. doi:10.1038/scientificamerican0984-52
- Klusch, M., Fries, B., & Sycara, K. (2006). Automated semantic web service discovery with OWLS-MX. In *AAMAS '06: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, (pp. 915–922). New York, NY, USA.
- Lecue, F., & Mehandjiev, N. (2009). Towards scalability of quality driven semantic web service composition. In *Proceedings of IEEE International Conference on Web Services*, (pp. 469–476).
- Li, Y., Chen, M., Wen, T., & Sun, L. (2008). Quality driven web services composition based on an extended layered graph. In *Proceedings of International Conference on Computer Science and Software Engineering*, (pp. 53– 156).
- Li, Y., & Xu, Z. (2003). An ant colony optimization heuristic for solving maximum independent set problems. In *Proceedings of International Conference on Computational Intelligence and Multimedia Applications*.
- Maassen, J., Verstoep, K., Bal, H. E., Grosso, P., & de Laat, C. (2009). Assessing the impact of future reconfigurable optical networks on application performance. In *Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Processing*, (pp. 1-8).
- Massonet, P., Deville, Y., & Neve, C. (2002). From AOSE methodology to agent implementation. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi Agent Systems*, (pp. 27–34). ACM Press.
- Michlmayr, A., Rosenberg, F., Leitner, P., & Dustdar, S. (2009). Service provenance in QoS-aware web service runtimes. In *Proceedings of IEEE International Conference on Web Services*, (pp. 115–122).
- Moreau, L., Freire, J., Futrelle, J., & Robert, E. McGrath, Myers, J., & Paulson, P. (2008). The open provenance model: An overview. *Provenance and Annotation of Data and Processes*, (pp. 323–326). Berlin, Germany: Springer-Verlag.
- Nan, K., Ma, Y., Zhang, H., & Chen, G. (2006). Transfer, processing and distribution of cosmic ray data from Tibet. *Future Generation Computer Systems*, 22(8), 852–855. doi:10.1016/j.future.2006.03.015
- Nichols, K., Blake, S., Baker, F., & Black, D. (1998). *Definition of the differentiated services field (DS Field) in the IPv4 and IPv6 headers*. Request for Comments 2474. Retrieved June 17, 2011, from <http://www.ietf.org/rfc/rfc2474.txt>

Nsi-Wg. (2011). Retrieved June 17, 2011, from <http://forge.ogf.org/sf/projects/nsi-wg/>

OGF. (2011). *Open Grid Forum homepage*. Retrieved June 17, 2011, from www.ogf.org

OpenDrac. (2011). *The open dynamic resource allocation controller*. Retrieved June 17, 2011, <http://www.opendrac.org>

Raicu, I., Zhao, Y., Dumitrescu, C., Foster, I., & Wilde, M. (2007). Falcon: A fast and light-weight task execution framework. In *SC '07: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*, (pp. 1–12). New York: ACM.

Rajbhandari, S., Contes, A. F., Rana, O., Deora, V., & Wootten, I. (2006). Trust assessment using provenance in service oriented applications. In *Proceedings of International Conference on Enterprise Distributed Object Computing Workshops*.

Rosenberg, F., Leitner, P., Michlmayr, A., Celikovic, P., & Dustdar, S. (2009). Towards composition as a service - a quality of service driven approach. In *Proceedings of International Conference on Data Engineering*, (pp. 1733–1740).

Sabata, B., Chatterjee, S., Davis, M., Sydir, J., & Lawrence, T. F. (1997). Taxonomy of QoS specifications. In *Proceedings of IEEE International Workshop on Object-Oriented Real-Time Dependable Systems*, (pp. 0-100). IEEE Computer Society.

Sage. (2010). *Scalable adaptive graphics environment*. Retrieved June 17, 2011, <http://www.evl.uic.edu/cavern/sage/>

Shimizu, T., Shirai, D., Takahashi, H., Murooka, T., Obana, K., & Tonomura, Y. (2006). International real-time streaming of 4K digital cinema. *Future Generation Computer Systems*, 22(8), 929–939. doi:10.1016/j.future.2006.04.001

Smarr, L., Brown, M., de Fanti, T., & de Laat, C. (2006). Special Issue on iGrid2005. *Future Generation Computer Systems*, 22(8).

Sobieski, J., Lehman, T., Jabbari, B., Ruszczczyk, C., Summerhill, R., & Whitney, A. (2006). Dynamic provisioning of LightPath services for radio astronomy applications. *Future Generation Computer Systems*, 22(8), 984–992. doi:10.1016/j.future.2006.03.012

The SURFNet. (2002). *The surfnet homepage*. Retrieved June 17, 2011, <http://www.surfnet.nl/>

Travostino, F., Daspit, P., Gommans, L., Jog, C., de Laat, C., & Mambretti, J. (2006). Seamless live migration of virtual machines over the MAN/WAN. *Future Generation Computer Systems*, 22(8), 901–907. doi:10.1016/j.future.2006.03.007

Venugopal, S., Buyya, R., & Ramamohanarao, K. (2006). A taxonomy of data grids for distributed data sharing, management and processing. *ACM Computing Surveys*, 38(1), 1-53. ISSN 0360-0300

Wooldridge, M., & Jennings, N. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152. doi:10.1017/S0269888900008122

Yu, J., Kirley, M., & Buyya, R. (2007). Multi-objective planning for workflow execution on grids, In *Proceedings of IEEE/ACM International Workshop on Grid Computing*, (pp. 10–17).

Zhao, Z., Belloum, A., & Bubak, M. (2009). Editorial: Special section on workflow systems and applications in e-Science. *Future Generation Computer Systems*, 25(5), 525–527. doi:10.1016/j.future.2008.10.011

Zhao, Z., Koning, R., Grosso, P., & de Laat, C. (2010). *Quality guaranteed media delivery on advanced network*. Demo in Supercomputing 2010.

ADDITIONAL READING

Braun, T., Diaz, M., Gabeiras, J., & Staub, T. (2008). *End to end quality of service over heterogeneous networks*. Springer Verlag.

KEY TERMS AND DEFINITIONS

Advanced Network: Networks that provide enhanced services to end users, in particular with focus on greater Quality of Service. Advanced networks often make use of the latest technologies, such as fully photonic or all optical devices to create circuits for transport of applications data.

Multi Agent System: A set of software agents that work together in a system. The agents may cooperate, compete, or both via some common infrastructure. A multi agent system is not simply a collection of disjoint set of autonomous agents.

Provisioning: The procedure of configuring network elements according to the user service requirements and make them ready for the customer to actually use the service.

Quality of Service (QoS): A measure of the ability that a service can provide to its consumers. In the network context, QoS refers to the performance attributes such as delay variation, bandwidth, and packet loss rate, which are also called QoS metric.

Scientific Workflows: The workflows used in scientific experiments.

Workflows: Sequences of steps or tasks defined in the business processes or scientific experiments. A workflow management system provides modeling mechanisms for describing workflow logics, automates the execution of workflow steps, and provides necessary support at different levels to allow users to interact with the execution.

ENDNOTE

- ¹ Acknowledgements - The Global Lambda Integrated Facility (GLIF) Map 2011 visualization was created by Robert Patterson of the Advanced Visualization Laboratory (AVL) at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign (UIUC), using an Earth image provided by NASA with texture retouching by Jeff Carpenter, NCSA. Data was compiled by Maxine D. Brown of the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC). Support was provided by GLIF, NCSA/UIUC, the State of Illinois, and US National Science Foundation grants # OCI-0962997 to EVL/UIC. For more information on GLIF, see <http://www.glif.is/>. The GLIF map does not represent all the world's Research and Education optical networks, and does not show international capacity that is dedicated to production usage. The GLIF map only illustrates excess capacity that its participants are willing to share with international research teams for applications-driven and computer-system experiments, in full or in part, all or some of the time. GLIF does not provide any network services itself, and researchers should approach individual GLIF network resource providers to obtain lightpath services.