CineGrid Amsterdam: Research Challenges and Solutions

Cosmin Dumitru, Paola Grosso, Cees de Laat

Abstract

CineGrid Amsterdam has stimulated both creation of high quality media content and advanced research in computer network and systems. In this article we focus on the research outcome of the project. We highlight three main results : the CineGrid Description Language, the CineGrid Amsterdam Portal and the Vampires framework. We also present the way these relate to the CineGrid community. We believe that given the current diversity in the computing resources landscape, user-centric systems require more attention. In this article we present some of our current work that addresses this issue with the CineGrid community as a use-case.

1 CineGrid Amsterdam

The CineGrid Amsterdam(CineGrid-Amsterdam 2014) program ran from 2010 to 2013, putting together a large number of Dutch organizations from different backgrounds. In this article we provide an overview of the activities and results of the project. We want in particular focus on the research challenges that arose during the project, and even more importantly on the solutions that were implemented to address them. In addition, we describe our current work related to the CineGrid research direction. We conclude by highlighting the still open issues and indicate possible directions for future work for the international CineGrid community.

1.1 The consortium

In the period around 2007/2008 most of the activities in the area of cinematic 4K productions were taking place in the United States and Japan. Several Dutch groups that were already involved in the international CineGrid(CineGrid 2014) collaboration came to the conclusion that it was necessary to join forces, procure funding and try to create a focal point for this type of research in Amsterdam. Amsterdam is a city with a large concentration of creative industry, scientific research centers and commercial parties; its central position in Europe make it an ideal location for leading this type of research.

The University of Amsterdam (UvA), the Waag Society, SURFsara and SURFnet shaped the project, which was ultimately funded by the European

Union and by the city of Amsterdam. The intended goal for CineGrid Amsterdam was to create a strong synergy between the international CineGrid collaboration and the Dutch community. Furthermore, the partners intended to create a facility for the production and display of 4K movies. They also focused on research solutions addressing the unique challenges of data distribution, data processing and data storage. All of these intended results were ultimately accomplished. The project has in fact created the CineGrid Amsterdam facility located at de Waag, it has promoted and hosted several international events in Amsterdam for the whole community and it has produced a number of movies and short documentaries; last but not least it has produced scientific methods and software that have been tested by the whole CineGrid community to operate its content exchanges.

1.2 Architecture

The CineGrid collaboration relies on the existence of distributed and individually operated resources interconnected by dedicated high speed networks. The resources are grouped in "CineGrid Exchanges" (Liu, Schulze, Herr, Weekley, Zhu, Osdol, Plepys, and Wan 2011). Each exchange has a different resource pool but ultimately the same purpose, to support *digital media asset management, distribution and preservation.*

In Amsterdam the *Exchange* is operated by the University of Amsterdam. It includes a number of storage servers and a dedicated cluster, DAS4(DAS-4 2014), used for media workflows. The network is provided by SURFnet, the Dutch National Research and Education Network (NREN). Orchestration is performed with the help of the CineGrid Portal (see 2.2).

2 Research results

The University of Amsterdam was the main academic partner in the CineGrid Amsterdam collaboration. The UvA scientific focus had one underlying theme: providing end users with seamless methods to retrieve, store, distribute and process content in the CineGrid Exchanges. To achieve this goal, UvA identified three missing elements:

1. a common vocabulary to describe the resources and the services within CineGrid;
2. an easy to use and intuitive web based resource orchestration tool for the CineGrid community;
3. a system that, while transparent to the user, would enable the use of all infrastructure and network capabilities present in the Exchanges.

The answers to the above were ultimately:

1. the CineGrid Description Language (see Sec.2.1);

2. the CineGrid Amsterdam portal (see Sec.2.2);

3. the Execution Engine (see Sec.2.3).

2.1 CineGrid ontology

The CineGrid Description Language (CDL) is a OWL ontology describing the services and infrastructure resources present in CineGrid (Koning, Grosso, and de Laat 2011). CDL is integrated with the more general Infrastructure and Description Language(INDL) ontology (Ghijsen, van der Ham, Grosso, Dumitru, Zhu, Zhao, and de Laat 2013), an ontology that provides technology independent descriptions of computing infrastructures. When possible CDL re-uses concepts from INDL, and it augments them and expands them for all the CineGrid specific characteristics. The CDL ontology has two distinct parts:

- *a resource ontology, which describes all the elements that are part of CineGrid;*

- *a service ontology, which describes the tasks a device can perform for the users of CineGrid. CineGrid devices can potentially perform multiple types of tasks, possibly at the same time. These tasks maps into services; and the user of the ontology deals directly with services.*

In figure 1 we present a CDL example description of the CineGrid Amsterdam Exchange. It includes 3 main layers: the service layer - transcoding and storage services, the computing infrastructure layer - physical hosts, and the networking layer - network switches and ports.

In addition UvA developed a metadata format based on similar semantic web concepts to describe the characteristics of the content available in the exchanges. For example, aspects like resolution, video codec, file properties are all included in the ontology.

CDL provides the language to create CineGrid services across domains. The metadata facilitates content curation, archival, retrieval and exchange. Both are tightly integrated in the CineGrid portal.

2.2 CineGrid Portal

The CineGrid Amsterdam portal was developed to serve three types of users:

- *content creators, by providing them with an easy to use tool to upload their content to the common storage facilities, while annotating them for further retrieval;*

- *infrastructure administrators, by providing them with the interfaces to add and remove network and computing resources present in the CineGrid Amsterdam infrastructure;*

- *content users, by allowing them to search among the available material and to schedule streaming and processing to the desired locations.*

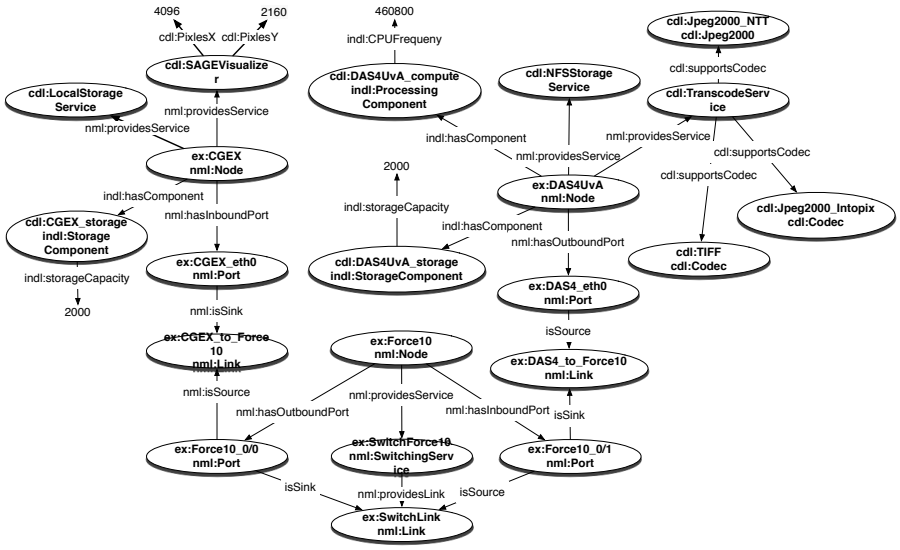

Figure 1: CDL Description of CineGrid Amsterdam Exchange

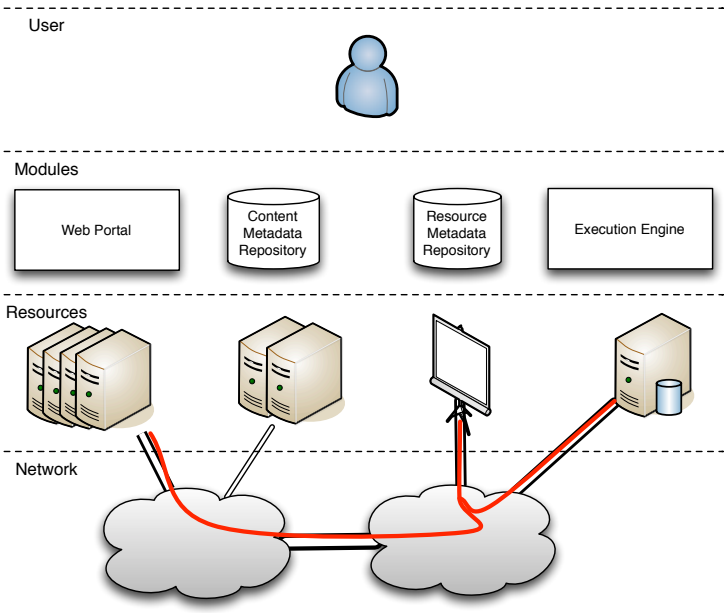

Figure 2: General Architecture of the CineGrid Amsterdam Exchange

The web Portal relies on two distinct repositories to perform its tasks: a content metadata repository and a resource metadata repository.

The content metadata repository splits the content metadata in two parts: one concerning the intrinsic properties of the assets, like authorship/copyright information, creation date and the other one related to the image properties, such as file format, resolution, compression options, bit depth, etc.

The resource metadata repository contains information about resources and their properties. As we said CDL distinguishes between a resource and a service ontology; this in turn reflects in CineGrid resources implementing one or more CDL service. The services then consume data stored in the exchanges. This decoupling between the functionality of resources and their physical representation allows dynamically coupling services at runtime. In addition, it hides the infrastructure complexity from the users, allowing them to focus on the media workflow and not on the resource orchestration. Figure 2 presents the layered system architecture of the CineGrid Portal.

2.2.1 Uploading content

Users of the Amsterdam exchange can upload their content to the system by providing the required metadata information; the curator of the exchange will determine if this new content will be accepted and published in the repository. Content might differ in terms of allowed usage, from being publicly accessible to more restrictive cases where there are specific rules and access rights to be respected. Uploading can be performed either through the public Internet or via dedicated high speed links offered by SURFNet.

2.2.2 Searching and selecting content

A user of the system will use the web portal to select the content of interest. The portal offers different search criteria, such as name based or location based searches. Once the pieces of content that match the search query have been identified, the user is presented with the corresponding choice of clips. Each result contains a web-friendly version of the high-quality media to allow fast screening of the content, to ensure that the selected item is the one desired.

2.2.3 Streaming and processing content

Finally the user can select the modality to stream the content to its desired final destination with an extra processing step, if needed. The execution engine is in charge of the actual setup of the resources and the actual content streaming or processing.

2.3 The Execution Engine

The execution engine operating in the CineGrid Amsterdam Exchange is in charge of the resource selection and the content streaming and processing. These operations require the engine to identify for each content type and for each operation on the content the best match of computing and networking resources. Furthermore, the engine needs to setup the network paths between the sources

and destination nodes whenever the traffic will travel across dedicated and dynamically provisioned paths.

When focusing on the computing we need to remember that by construction all CineGrid exchanges operate heterogeneous resources; in addition, it is possible to acquire additional computing capacity from commercial parties or externally linked cloud computing providers. The engine needs to find the best balance between two parameters: 1. the performance of the selected nodes, *i.e.* the time they require to complete a task; and 2. the total costs when purchasing the resources externally to the exchange.

When focusing on the network the engine needs to interface with the underlying provisioning systems, such as the Network Service Interface(NSI) (Roberts, Kudoh, Monga, Sobieski, MacAuley, and Guok 2014) framework. NSI creates the proper network connections between the resources and allow the access to high bandwidth dedicated network paths.

All of this aspects have been research and implemented in the Vampires framework, that is ultimately usable outside of the CineGrid context.

3 Current work

Our current work has a user-centric focus on computing and networking resource selection and use. Users in the CineGrid community are presented with a complex set choices due to the high heterogeneity and size of the collaboration. This is not specific just to CineGrid but is a more general aspect in the cloud computing landscape. This makes our work reusable outside the collaboration.

In the following sections we present two developments related to this challenge; the Vampires framework - a resource scheduler for cloud environments and the Open Cloud Exchange - a concept for multi-cloud usage.

3.1 Vampires

The processing component of the Execution Engine introduced in section 2.2 is implemented using the *Vampires framework*(Dumitru, Oprescu, Živković, van der Mei, Grosso, and de Laat 2014). The framework allows runtime optimization for workloads of two objectives - execution time (performance) and cost. The workload that Vampires deals with is modeled as a Bag-of-Tasks, the dominant workload in the CineGrid exchange. The cost aspect is very important if the resources are rented from cloud service providers, while the performance is relevant in the case when users need to execute workloads under time constraints. The two objectives are contradictory with regard to the type and amount of computing and networking resources needed.

Execution involves a three step process in which candidate resources are first sampled in order to extract performance metrics related to the workload. Using the metrics and a formal model of the system, together with a specific resource configuration, Vampires can predict the expected performance without actually executing it. Of particular importance is the behavior of the system in the

presence of I/O bottlenecks (limited network bandwidth) as they are usually the limiting factor when it comes to scaling resources. In (Dumitru, Oprescu, Živković, van der Mei, Grosso, and de Laat 2014) we presented a detailed model and evaluation of this execution prediction method. Vampires presents the user with relevant resource configurations. The set of relevant configurations represents the Pareto front of the workload. A resource configuration from the Pareto front is not dominated by any other configuration neither in performance nor in cost. The user will select the one which better matches the original requirements.

3.2 Open Cloud Exchange (OCX)

The rise in popularity of commercial cloud computing providers has brought along new challenges and opportunities. One of them is the combined use of multiple cloud computing resources. By having available multiple offers users can select those which suit their workload better, even if they're not all from the same provider. Many factors can influence the choice of providers: the cost and performance aspects, geographic location, availability of "exotic" hardware (e.g. GPUs, FPGAs , many core or very high memory systems, low latency internal networking etc), law jurisdiction, possibility of having dedicated network connectivity to the user's site and so on.

The Open Cloud Exchange (OCX) (Demchenko, Ham, Ngo, Matselyukh, Filiposka, Laat, and Escalona 2013) is a concept which describes a dedicated place for inter-connection and peering between cloud providers and customers. The OCX permits users to create dedicated network connections between two or more different cloud providers, in an application transparent way.

A challenging aspect of having such a system in place is how would users make the most of the offers available. We have extended the Vampires framework to allow interaction between resources of cloud providers which participate to the OCX. Currently we are exploring the performance of the framework by using CineGrid workloads, like image transcoding. An interesting aspect is the ability to create network links which are neither over provisioned or under provisioned with regard to the application requirements.

4 Conclusions and future work

The CineGrid Amsterdam project has developed novel methods that enable content providers to exploit the most advanced computing and networking infrastructures. We presented here three major research results that are now available to the whole CineGrid community: the CineGrid Description Language, a ontology for resources and services for high quality media; the interactive portal for content and resource selection, streaming and processing; and the Vampires framework, a cloud scheduler and execution engine. The portal and the cloud scheduler both address the distributed nature of the CineGrid resources, where

computing and storage facilities needed for the processing and storage of content may be located in different countries and/or continents.

The portal allows for the selection of the appropriate resources and if needed triggers creation of the network paths across sites. Vampires includes a queuingnetwork inspired model in its performance prediction, allowing it to deliver in realistic estimates of the content processing times.

Our current focus is on helping end-users to fully utilize CineGrid-like environments, where heterogeneity and geographic distribution increase the complexity of operating them. We envision that solutions like the Open Cloud Exchange will make inter-provider collaboration seamless from the perspective of the user, ultimately reducing cost and improving performance.

References

CineGrid (2014). Cinegrid. `http://www.CineGrid.org/`. Accessed: 2014-09-01.

CineGrid-Amsterdam (2014). Cinegrid amsterdam. `http://CineGrid.nl/`. Accessed: 2014-09-01.

DAS-4 (2014). Das-4 (the distributed asci supercomputer 4. `http://www.cs.vu.nl/das4/`. Accessed: 2014-09-01.

Demchenko, Y., J. v. d. Ham, C. Ngo, T. Matselyukh, S. Filiposka, C. d. Laat, and E. Escalona (2013). Open cloud exchange (ocx): Architecture and functional components. In *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on*, Volume 2, pp. 81–87. IEEE.

Dumitru, C., A.-M. Oprescu, M. Živković, R. van der Mei, P. Grosso, and C. de Laat (2014). A queueing theory approach to pareto optimal bags-of-tasks scheduling on clouds. In *Euro-Par 2014 Parallel Processing*, pp. 162–173. Springer.

Ghijsen, M., J. van der Ham, P. Grosso, C. Dumitru, H. Zhu, Z. Zhao, and C. de Laat (2013). A semantic-web approach for modeling computing infrastructures. *Computers & Electrical Engineering 39*(8), 2553 – 2565.

Koning, R., P. Grosso, and C. T. A. M. de Laat (2011). Using ontologies for resource description in the CineGrid Exchange. *Future Generation Computer Systems 27*(7), 960–965.

Liu, S., J. P. Schulze, L. Herr, J. D. Weekley, B. Zhu, N. V. Osdol, D. Plepys, and M. Wan (2011). Cinegrid exchange: A workflow-based peta-scale distributed storage platform on a high-speed network. *Future Generation Computer Systems 27*(7), 966 – 976.

Roberts, G., T. Kudoh, I. Monga, J. Sobieski, J. MacAuley, and C. Guok (2014). GFD.212: NSI Connection Service v2.0. Technical report.