Semantic-Based Workflow Composition for Video Processing in the Grid

Gayathri Nadarajan[1] Yun-Heh Chen-Burger[2] James Malone[3]
Artificial Intelligence Applications Institute (AIAI)
Center for Intelligent Systems and their Applications (CISA)
School of Informatics, University of Edinburgh, UK
[1]G.Nadarajan@sms.ed.ac.uk, [2]jessicac@inf.ed.ac.uk, [3]j.malone@ed.ac.uk

Abstract

We outline the problem of automatic video processing for the EcoGrid [1]. This poses many challenges as there is a vast amount of raw data that need to be analysed effectively and efficiently. Furthermore, ecological data are subject to environmental changes and are exception-prone, hence their qualities vary. As manual processing by humans can be time and labour intensive, video and image processing tools can go some way to addressing such problems since they are computationally fast. However, most video analyses that utilise a combination of these tools are still done manually. We propose a semantic-based hybrid workflow composition method that strives to provide automation to speed up this process. The requirements for such a system are presented, whereby we aim for a solution that best satisfies these requirements and that overcomes the limitations of existing Grid workflow composition systems.

1. Introduction

1.1. Background

The EcoGrid project [1] is aimed at utilising state-of-the-art Grid technologies to establish a cyberinfrastructure for ecological research. This includes the integration of geographically distributed sensors, computing power and storage resources into a uniform and secure platform. Scientists can conduct data acquisition, data analysis and data sharing on this platform. The infrastructure is divided into four components; network, data streaming, data management and workflow enactment. The National Center for High-performance Computing (NCHC), Taiwan and Artificial Intelligence Applications Institute (AIAI), the University of Edinburgh have forged a research collaboration on workflow enactment in the EcoGrid.

Recognising the needs for real-time automatic and continuous information gathering through EcoGrid that offers a unique and immense opportunity for long term ecological monitoring and planning, NCHC has installed and managed Wireless Sensor Nets in several national parks in Taiwan. The information collected is stored in and made available through EcoGrid for access. The information collected includes surveillance videos in the Fu-Shan National Park covering the entire area for observing natural lives and protecting them from potential poachers, audio recording of frogs of rare species, under-sea coral reef and marine life observation stations and more. Due to the continuous and non-intrusive methods deployed, such monitoring and recording efforts have already made ecological discoveries of significant importance that traditional methods otherwise could not have made.

1.2. Ecological Motivation and Challenges

Continuous data collection in the EcoGrid, however, poses a great challenge as how this data may be transformed into useable information for the ecologists and in a timely fashion. For instance, one minute of video clip typically takes 1829 frames and is stored in 3.72 Mbytes. That translates into 223.2 MB per minute, 5356.8MB per day and 1.86 Terabytes per year for one operational camera, and due to the unpredictability of nature, one may not easily skip frames as they may contain vital information. Based on our own experience, one minute's clip will on average cost manual processing time of 15 minutes. This means that one year's recording of a camera would cost human experts 15 years' effort just to perform basic analysing and classifying tasks. Currently there are three under water cameras in operation and this will cost a human expert 45 years just to do basic processing task. This is clearly an impractical situation and more appropriate automation methods must be deployed. An introduction to the problem is provided in [2].

The Grid [4] is an infrastructure for next-generation e-Science applications aimed at enabling resource sharing and coordinated problem-solving between computers and people in a distributed and heterogeneous manner. The Semantic Grid [11] is an extension of the current Grid in which information and services are given well-defined and explicitly represented meaning, better enabling this cooperation. This requires means for composing and executing complex workflows. Considerable research and development efforts have been made towards the development of workflow management systems for the Grid. Apart from the ecological challenge presented above, we are also investigating means to improve the technology of Grid workflow management systems by providing semantic capabilities to the workflow composition engine. By semantic capabilities we mean machine-processable knowledge and information, provided by enriched representations such as ontologies and Semantic Web-based languages [8, 6]. This would allow for a high degree of easy-to-use and seamless automation to facilitate flexible collaborations between researchers and virtual organisations, as imposed by the Semantic Grid community. Employing the power of distributed processing within the Grid will also help immensely with overcoming some of the ecological obstacles mentioned above.

The rest of the paper is organised as follows. The requirements for a suitable system are detailed in Section 2, using an example video processing task. Section 3 presents several existing Grid workflow composition systems, their characteristics and application domains, while Section 4 compares them and highlights their limitations and suitability for our needs. In Section 5, we propose a hybrid method within a three-layered framework as a solution. This proposed framework aims to fill in some of the gaps that current workflow composition systems have yet to achieve. Section 6 concludes and discusses future work to be done.

2. Requirements

Given the ecological problem at hand, we outline a set of requirements that will correspond to the functionalities of a suitable system in our framework.

Process Automation. A mechanism for automatically processing large amounts of data is required. As we are performing video analysis on real-time data, this task can be broken down into a set of processes applied in sequence to the raw data. For example, the task of video annotation could be described in the following high-level steps that will need to be run automatically:

Keyframe detection → Classification by characteristics → Annotation

Performance-Based Selection. For each of these processes, one or more software tools can be used to achieve the particular function. Depending on the quality of the data, the combination of tools with the *best performing capability* for the overall stated goal should be selected. Thus the mechanism should be able to match the goal of the process with the capabilities based on a performance measure.

Iterative Processing. Each process in turn, could also be further decomposed into sub-processes. For instance, the 'Classification by characteristics' process could be made up of a feature extraction sub-process followed by a recognition sub-process that could be performed in a loop. Thus, a mechanism to support iterative processing is required.

Adaptive, Flexible and Generic Architecture. Following from its definition, the Grid is a dynamic environment where availability of resources and their load can change notably from one time point to the next. An ideal system would be one that is extremely sensitive and adaptive to changing environments. It would also possess a generic capability, whereby it should be able to perform a variety of video analysis tasks provided as input. Furthermore, it should be able to generate new sequences of solutions. For this it should include techniques that allow it to incorporate its knowledge of previous experiences in order to generate new solutions as part of its learning. Thus an adaptive, generic and flexible approach is required.

Semantic-Based Compatibility. As well as traditional Grid and video processing requirements, the approach that we take should also be semantic-based, that is to integrate ontologies and Semantic Web-based languages, such as Web Ontology Language (OWL) [8] and more recently, Temporal Resource Description Framework (Temporal RDF) [6] to allow for effective reasoning in compliance with Semantic Grid requirements. Ontologies help with the sharing and reuse of common vocabularies between domains, as well as help to perform reasoning and inferencing about them. Semantic Web-based languages provide a means for representing machine-processable information.

A workflow composition mechanism would be ideal for automating repetitive tasks such as keyframe detection. However, it remains a big challenge to provide a framework that encompasses all the requirements stated above. We now turn to existing workflow composition systems, and investigate how far they go in fulfilling these requirements.

3. Related Work

In this section we provide a brief overview of several influential Grid workflow composition systems and discuss their capabilities, advantages and application areas. Up-to-date surveys [14, 10] provide a more comprehensive study on workflow systems for the Grid.

Pegasus. Pegasus [3] aims to support large-scale data management in particle physics experiments by mapping abstract workflows, which are user-defined or automatically generated, to their concrete forms that are executable in the Grid. This is done by taking an abstract description of a workflow and finding the appropriate data, software and Grid resources to execute the workflow. The concrete workflow is produced with a set of submit files necessary for its execution through a scheduler. Pegasus utilises deferred planning to generate partial executable workflows based on already executed tasks and the currently available resources by a partitioner. This allows for dynamic scheduling. However, Pegasus does not support looping.

Triana. Triana [13] is an open-source environment that allows users to construct workflows in a graphical manner. It has been used for text, speech and image processing tasks. A user creates a workflow by dragging the desired units from a toolbox and dropping them onto the workspace. Units are interconnected by dragging a cable between them. The resulting composite graph can be executed or saved. Workflows defined in these formats can also be read and handled by Triana. The GUI allows users to make changes to the workflow by adding, deleting or changing the sequence of execution by drag-and-drop. Additionally, Triana supports looping constructs, which is desirable from a process modelling point of view.

Taverna. Taverna [9] is a collaboration between several European academe and industries under the myGrid project [12] aimed at supporting biologists and bioinformaticians. It is used for executing scientific workflows in the Grid by utilising the Freefluo workflow enactment engine (Available at http://www/freefluo.sourceforge.net). It provides graphical interfaces that allows workflow manipulation and workflow progress invocation easily. It also provides an implicit iteration mechanism and possesses fault tolerance capability. Taverna is particularly suitable for tasks that can handle simultaneous processing as it supports concurrency.

Kepler. Kepler [7] is a visual, community-driven project with an extendable open source platform. It allows scientists from several different domains to design and execute scientific workflows. The Kepler system models a workflow as a composition of independent components that communicate through well-defined interfaces. It is based on a modular design where different execution models can be easily plugged into the workflows without changing any other components within the workflows. Kepler supports looping and also has good reliability as it is able to produce partial results even when an entire workflow fails.

4. Limitations of Current Solutions and Motivation for an Improved Approach

The systems mentioned above possess some similarities and differences that are worth investigating in order to assess their suitability and limitations for our framework.

In terms of composition itself, Pegasus differs from Triana, Taverna and Kepler because instead of its abilities to compose and execute the workflows directly, it maps abstract workflows to their concrete forms, which are then executed by a scheduler. It also provides adaptivity through a partitioner that uses planning to produce partial executable workflows. The role of planning is important for our system as it allows for dynamic process selection and composition based on a given set of goals.

Triana, Taverna and Kepler provide a basis for users to create and run scientific workflows. All three systems contain similar elements; Triana's tasks are conceptually the same as Taverna's processes and Kepler's actors. The approach in Kepler is very similar to Triana in that the workflow is visually constructed from actors (java components), which can either be local processes or can invoke remote services such as Web services. In addition, our framework will incorporate elements to distinguish goals from the capabilities associated with the processes.

In terms of applicability, Pegasus would best suit a domain with well-defined requirements and where the overall goal could be determined from a given set of rules and constraints. Triana is well-suited for composing complex workflows for Web services and Peer to Peer services. Taverna is also suitable to be used in Web and Grid services contexts, but its use may be limited to composing simple workflows, whereas Kepler works very well for composing workflows for complex tasks but it has yet to reach its potential as a fully Grid-enhanced system since the system it is built upon is primarily aimed at modelling concurrent systems.

Incorporation of Semantic Web technologies within present systems are still limited. The use of such technologies should not be exclusively independent, rather they should be fully intergrated into the system. Existing systems do not provide ontological handling nor integration, instead they make use of separate ontology tools to define and manipulate ontologies. We wish to integrate the capability to read OWL files and provide results and information in OWL and Temporal RDF within our system.

It is apparent that not a single one of the systems above is able to fulfill all the requirements and would therefore not fit our needs. However, in terms of performing complex tasks such as video processing, Kepler would stand out as a very promising system. Thus, we would consider using it as an underlying workflow composition mechanism within our framework. In addition, we wish to take advantage of the most prominent and advantageous features exhibited by these systems, as well as incorporate new features that will allow us to fill the remaining gaps in the EcoGrid challenge.

5. Proposed Framework - Hybrid Method

Based on the analysis of existing systems and motivation provided in the previous section, we propose a semantic-based hybrid workflow composition method within a three-layered framework. Figure 1 illustrates the framework using video annotation as an example.

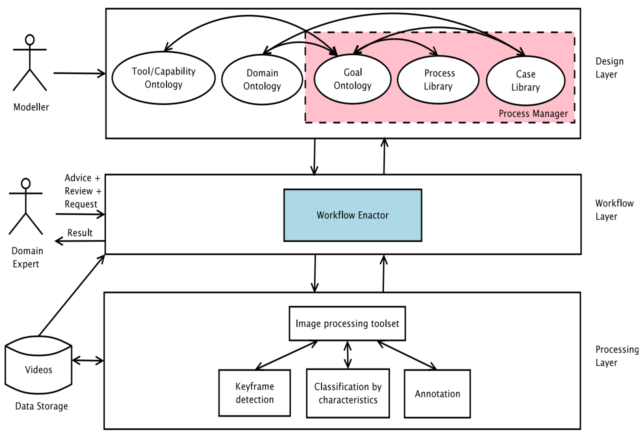

Figure 1. Overview of Hybrid Workflow Composition Framework for Video Annotation.

5.1. Design Layer

The design layer contains components that describe the domain, goals, capabilities and processes to be carried out in the system. These are represented using ontologies, elaborated in section 5.1, and two libraries. A modeller is able to manipulate the components of the design layer, for example populate the libraries and modify the ontologies.

Process Manager. The process manager is responsible for selecting and composing the sequence of processes to be executed in the workflow based on the tools available to perform the task. It has a Planning component, which comprises the goal ontology and the process library, and a Case-Based Reasoning (CBR) component, which is the case library. The process library holds instances of processes that could be executed to perform the tasks of the workflow. Once a set of tools have been identified for achieving the goal, the process library composes the process sequence for execution. The goal ontology and process library, together, constitute the Planning component of the process manager.

However, there could be more than one possible solution or no solution for satisfying a particular task. While the Planning component is responsible for matching the goal with the capabilities (tools) and selecting the procedural steps to be taken, the CBR component is responsible for finding the closest solution from past scenarios for cases where Planning alone doesn't work. The case library keeps track of previous viable solutions and finds a *similar* solution to match the current problem. The heuristic for the closest solution is based on a similarity measure [5].

Ontologies. We have opted to incorporate three ontologies to keep the goals separate from the capabilities and to provide meaning for the process within a semantically integrated system. Each ontology holds a vocabulary of classes of things that it represents; the goal ontology contains the classes tasks (e.g. "annotation") while the capability ontology contains the classes of video and image processing tools. The domain ontology provides meaning for annotation, for example concepts such as "blur", "clear", "bright" are included. The use of ontologies is beneficial because they provide a formal and explicit means to represent concepts, relationships and properties in a domain. They play an important role in fulfilling semantic interoperability, as highlighted in section 4. A workflow system with full ontological integration has several advantages. It allows for cross-checking between ontologies, addition of new concepts into the workflow system and discovery of new knowledge within the system.

5.2. Workflow Layer

This layer acts as the main interface between the design and processing layers. It ensures the smooth interaction between the components, access to and from various resources such as raw data, video and image processing toolset, as well as the provision of the final output to the user. A workflow enactor acts as the interpreter of the events that occur within the system.

Scenario. In the annotation example provided above, the flow of processing is as follows. The user inputs a request into the system. The workflow enactor fetches data in the form of a video clip from the data storage and delegates the user request to the design layer. Based on the Planning and CBR approaches, a workflow for annotation will be established and passed to the workflow enactor. This is then fed to the processing layer for further action. The processing layer will communicate the annotated video back to the workflow enactor, which will pass this result to a domain expert who will provide feedback. This could be used as a basis for improving the performance of the system. Finally, the case library is updated with this solution.

5.3. Processing Layer

The processing layer consists of a set of image and video processing tools that will act on the data. The functions of these tools are represented in the capability ontology in the design layer. Once a workflow has been established, these tools may work on the videos directly. It should be noted that for each capability, there could be more than one tool available. Depending on the quality of the video and the task, each tool may perform with a different level of accuracy. Thus, having domain experts provide feedback on the performance of a particular combination of tools would be beneficial to the system. We also anticipate the use of machine learning techniques to assist with performance measure predictions for the image processing tools.

6. Discussion and Outlook

We have proposed a generic framework that enables an adaptive workflow enactment for video processing in the EcoGrid. This framework is based on a self-learned workflow composition method that utilises a hybrid approach of AI Planning and Case-Based Reasoning. We believe that full automation will be successful if the workflow tool has a full understanding of video processing tasks, but this is impossible for machine-implemented systems. Humans could provide feedback on some tasks to help improve this situation and enhance the performance of the system.

Implementation issues will need to be addressed for the proposed framework. One suggestion would be to develop a layer on top of an existing Grid workflow system. This would provide for a system with full Grid capabilities, which is highly desirable. However, such a system could be too difficult to implement. Another suggested method is to deploy a workflow enactor on an existing process modelling tool with semantic enhancements. This method could be easier to implement, but the scope of the resulting system may be limited and not fully Grid-enabled.

In discussing the approach taken by our framework, it's appropriate to consider process-oriented and data-oriented paradigms for video processing. The former treats processing (how things are done) as the first class primitive over data while the latter treats data (what is being manipulated) as the first class primitive over processing. Our framework, by utilising planning, supports the process-oriented paradigm, as opposed to most existing workflow composition systems. Prioritising processing over data, however, does not imply that the data is not of concern. By focusing on the techniques applied to the data, one is also implicitly giving importance to the data because the methods manipulate the characteristics of the videos and extract useful information from them. Thus the process-oriented approach complements the data-oriented approach.

A similar approach to the process-oriented paradigm is the state-based paradigm, where the next step to be executed is given by a set of allowed states from the current state based on predefined conditions. Although this method is useful for distributed environments, such as agent-based systems, its capability is limited to performing simple tasks, as defining states for complex tasks such as video processing would lead to a big state space.

We expect our work to contribute to strengthening video processing for workflows and vice-versa. Active and notable efforts in the development of Grid workflow systems could benefit from the semantic- and performance-based emphases that our framework proposes.

References

[1] EcoGrid National Center for High Performance Computing, Taiwan, 2006. http://ecogrid.nchc.org.tw/.

[2] Y.-H. Chen-Burger and F.-P. Lin. A Semantic-based Workflow Choreography for Integrated Sensing and Processing. In *Proceedings of the 9th IEEE International Workshop on Cellular Neural Networks and their Applications*, 2005.

[3] E. Deelman, J. Blythe, Y. Gil, C. Kesselman, G.Mehta, S. Patil, M. Su, K. Vahi, and M.Livny. Pegasus: Mapping Scientific Workflows onto the Grid. In *Across Grids Conference 2004, Nicosia, Cyprus*, 2004.

[4] I. Foster and C. Kesselman, editors. *The Grid: Blueprint for a Future Computing Infrastructure*. Morgan Kaufmann, 2nd edition, 2003.

[5] K. Gupta and A. Montezemi. Empirical Evaluation of Retrieval in Case-Based Reasoning Systems using Modified Cosine Matching Function. In *IEEE Trans. Systems, Man and Cybernetics*, 1997. Vol. 27(5), pp 601-612.

[6] C. Gutierrez, C. Hurtado, and A. Vaisman. Temporal RDF. In *Proceedings of European Conference on the Semantic Web (ECSW05)*, Jan 2005.

[7] B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. Lee, J. Tao, and Y. Zhao. Scientific Workflow Management and the Kepler System. *Concurrency and Computation: Practice & Experience*, 2005.

[8] D. McGuinness and F. van Harmelen. *OWL Web Ontology Language*. World Wide Web Consortium (W3C), 2004. http://www.w3.org/TR/owl-features/.

[9] T. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, M. Greenwood, T. Carver, K. Glover, M. Pocock, A. Wipat, and P.Li. Taverna: A Tool for the Composition and Enactment of Bioinformatics Workflows. *Bioinformatics*, 2004. Vol. 20, No. 17, pp 3045-3054.

[10] J. Rao and X. Su. A Survery of Automated Web Service Composition Methods. In *Proceedings of First International Workshop on Semantic Web Services and Web Process Composition (SWSWPC 2004), San Diego, CA, USA*, 2004.

[11] D. D. Roure, N. R. Jennings, and N. R. Shadbolt. The Semantic Grid: Past, Present, and Future. In *Proceedings of the IEEE Volume 93, Issue 3, pp 669-681*, March 2005.

[12] R. Stevens, A. Robinson, and C. Goble. MyGrid: Personalized Bioinformatics on the Information Grid. *Bioinformatics*, 2003. Vol. 19, Suppl. 1, pp i302-i304.

[13] I. Taylor, B. Schutz, B. S. Sathyaprakash, M. Shields, I. Wang, A. Harrison, C. Robinson, E. Al-Shakarchi, I. Kelley, and R. Zhu. The Triana Project, Cardiff University, 2003. http://www/trianacode.org.

[14] J. Yu and R. Buyya. A Taxonomy of Workflow Systems for Grid Computing. *Journal of Grid Computing*, Jan 2006. Volume 3, Numbers 3-4, pp 171-200.