

Doctoral Research Abstract

PhD Title: Model-Based Enhancement of Software Performance for Precision Critical Systems

Name: Naeem Muhammad, KULeuven, Belgium

Promoter: Prof. Dr. ir. Yolande Berbers, KULeuven, Belgium

Academic Supervisor: Nelis Boucké, KULeuven, Belgium

Industrial Supervisor: Sjir van Loo, Embedded System Institute, The Netherlands

Problem Statement

Quality attributes are characteristics of a system which are used to evaluate the quality of the system. Because quality attributes are often linked with more than one functional component of the system, modelling and analyzing quality attributes is a difficult task. The difficulty increases further with system complexity, which is the case for our problem domain: electron microscope systems. Electron microscopes are highly complex devices whose design combines quite a number of disciplines: (high-voltage) electron physics, electromagnetism and electron optics, image sensors and image processing, mechatronics, instrument conditioning (vacuum, heat), electronics (amplifiers, sensors, embedded control) and software engineering. Quality attributes for the electron microscope precision critical systems are mainly performance, evolvability, reliability and accuracy. Modelling, prediction and analysis of these quality attributes at architectural level is a difficult task but very cost effective.

Performance is the most important quality attribute for such systems as multiple aspects are involved in it. Among these aspects are throughput, efficiency, response time, concurrency and flow latencies. Although various techniques for analyzing the performance at the architectural level are already available (e.g. ATAM, SAAM, CBAM) [1] a more precise and realistic analysis can be achieved by utilizing the system implementation, in the case of legacy systems.

Aims and Objectives

The primary aim of this doctoral study is to provide modelling and analysis support for parallelism and flow latencies aspects of the performance quality attribute, by utilizing the implementation of the system. To achieve this aim the following objectives have been set:

1. Selection of suitable modelling languages
2. Development of conceptual architectural models of the system
3. Modelling the parallelism behaviour of the system

This work has been carried out as a part of the Condor project (<http://www.esi.nl/Projects->Condor>) at FEI company under the responsibilities of the Embedded Systems Institute (ESI). This project is partially supported by the Dutch Ministry of Economic Affairs under the BSIK program.

4. Modelling system flows and their latencies
5. Provide analysis support for analyzing performance bottlenecks associated with parallelism and system flows
6. Developing tool support for modelling and analyses activities
7. Application of the developed techniques and methods on an industrial case: the FEI electron microscope software

Approach

Using architecture views and viewpoints to describe an architecture description of a system is a common approach [1][2]. They describe the architecture of the system for some related concerns of a set of stakeholders. According to ISO 42010 [3]:

“An architecture view is a work product expressing a system’s architecture from the perspective of its concerns.”

“An architecture viewpoint establishes the conventions for constructing, interpreting and analyzing an architecture view addressing concerns framed by the viewpoint.”

A viewpoint essentially includes modelling formalisms, modelling and analysis techniques, and design rules. In this PhD, to provide support for analyzing the issues associated with the parallelism and flow latencies of the system, we develop the Parallelism and Flow Latency architecture viewpoints. Furthermore, we also study the impact of the parallelism overheads on the latencies of the system. The highlighted part in figure 1 shows the contribution and the overall approach of the PhD. The Parallelism and Flow Latency viewpoints are part of the execution view of an architecture description.

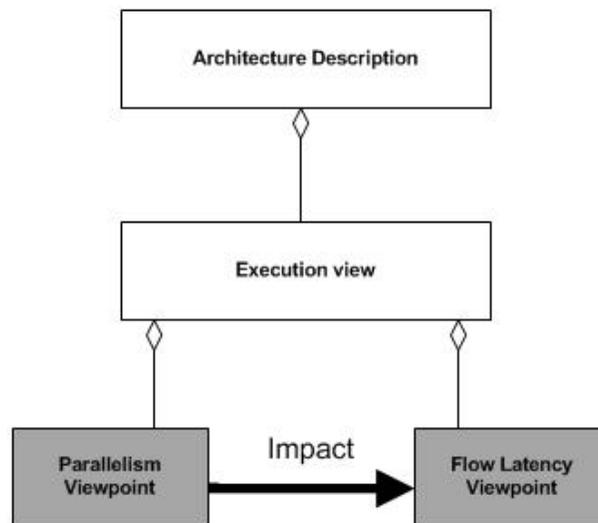


Figure 1: Research contribution and approach

Parallelism Viewpoint

The use of parallelism (multithreading) for software systems is becoming increasingly more important because of the potential benefits it provides. Multithreaded applications are considered to be more efficient because of their better software and hardware resource utilization caused by the parallel execution of tasks. Nowadays most of the widely used hardware and operating systems support multithreaded applications, therefore the use of parallelism has increased substantially. However, overheads and difficulties associated to parallelism amplify dramatically as the number of threads increase. Among them are context switches overhead, incorrect distribution of Read/Write operations and a complex thread management structure. Especially in legacy systems, where design decisions are usually not available, excessive use of multithreading proves to be problematic. Identification and mitigation of these issues in such cases is extremely hard.

One of the goals of this PhD is to provide support for modelling parallelism behaviour of the system to analyze these issues. In this perspective we have proposed an architecture viewpoint called Parallelism Viewpoint.

The Parallelism Viewpoint is a domain-specific form of the concurrency viewpoint which is used to describe the concurrent structure of a system. The concurrency viewpoint mainly provides support for describing concerns related to the communication and synchronization mechanisms of concurrent systems [4]. We extend this support for concurrent systems by describing parallelism related concerns with our viewpoint.

Essentially, a viewpoint must explicitly describe the concerns of a particular domain, identify the stakeholders of these concerns and specify a set of model kinds. For the proposed viewpoint we identified parallelism specific concerns, associated stakeholders and introduced a set of 5 models (Time distribution, Task distribution, Task types, Thread behaviour and Thread management) to describe these concerns. In this research work we also explicitly identified parallelism specific execution elements to retain focus on parallelism related issues. Figure 2 shows the methodology for developing models of the Parallelism Viewpoint.

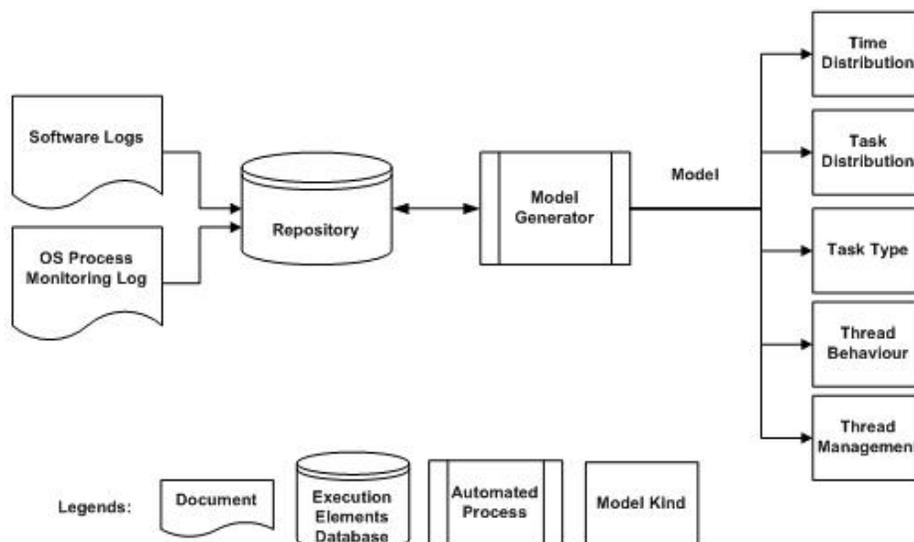


Figure 2: Parallelism Viewpoint model kinds generation

In our approach, we use process logs and logs maintained by the software system under investigation to harvest information and subsequently store it in an architectural repository. By using the information in this repository a model generator generates parallelism-specific model kinds. A model kind is a set of notations and conventions which are used to develop domain-specific models for a viewpoint.

By successfully identifying performance bottlenecks of the electron microscope software system with the help of our viewpoint, we have shown that our approach is generic and that it can be used for other parallelism-intensive software systems too. This argument is further strengthened by the fact that the description of our viewpoint is not bound to any particular technology.

Flow Latency Viewpoint

As a part of this PhD research study we are developing a flow latency viewpoint to describe latencies of a flow-intensive system. It includes identifying stakeholders of the system, their flows related concerns and developing a set of models to model and analyze issues with the flow latencies.

Currently, in our research work we have devised an approach for identifying various levels of abstractions in flows by creating connections among them [6]. A flow-intensive system consists of multiple flows. Each flow has a single starting and ending point but may internally consist of many distinct flows. Together these distinct flows compose a single flow at higher abstraction, which is called a composite flow. Such composite flows are most interesting from an architecture abstraction point of view since they often represent more abstract flows of information through the system.

We selected the Architecture Analysis and Design Language (AADL) as a modelling formalism for developing flow models of the system. AADL is one of the best known and most actively used architecture description language for embedded systems [5]. Though AADL is suitable for modelling flows we find its support insufficient for establishing connections among them, which is required in this case to connect distinct flows. We extended AADL to provide such support and provided tool support to analyze latencies of the composite flows of the system. In our solution we connected adjacent flows by introducing a new AADL property called `Link_Flow`. The new property holds a string value representing the identifier of the flow that is to be linked. Subsequently, this value is manipulated by the tool we developed, to analyze the latencies of the composite flows.

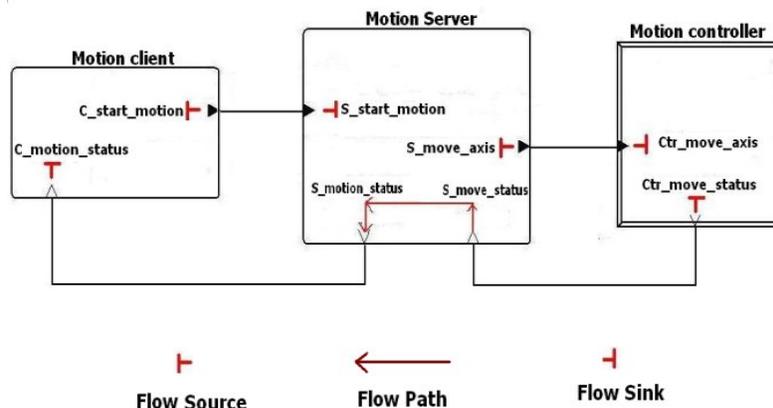


Figure 3 Flows in motion-subsystem

We successfully applied the proposed solution to an industrial case of a motion-subsystem, a subsystem of the electron microscope software. Figure 3 shows three components of the motion-subsystem and three distinct flows flowing in them, starting with a flow source and ending with a flow sink. These distinct flows contribute to a single composite flow. In the case of our example system these flows contribute to a single flow of a command send by the client to move the specimen in the electron microscope. Our proposed solution will enable system architects to model and analyse system flows at different abstraction levels. We will extend this research work to develop the Flow Latency viewpoint.

Summary

The goal of this PhD is to develop state of the art methods and techniques for modelling and analyzing the performance of the system at architectural level by utilizing its implementation. The focus of this research is on parallelism and flow latencies aspects of the performance.

To provide support for modelling parallelism behaviour of the system to analyze associated overheads we have proposed an architecture viewpoint called Parallelism Viewpoint. Results show that the Parallelism Viewpoint provides a profound insight into the parallelism behaviour of the system, required to identify and mitigate the associated performance overheads.

Currently, we are developing the Flow Latency Viewpoint to analyze issues with the latencies of the system flows. The viewpoint will explicitly identify the flows related concerns, associated stakeholders and a set of models to model these concerns.

The Parallelism Viewpoint and Flow Latency Viewpoint will be a part of an architecture description of the electron microscope software. Another goal of this PhD is to trace links between the elements of both viewpoints. In particular, we are interested in utilizing the outcome of the Parallelism Viewpoint to understand the effects of the parallelism overheads on flow latencies of the system.

References:

- [1] Clements, P., Kazman, R., Clein, M.; *Evaluating Software Architecture: methods and case studies*, Addison-Wesley (2002)
- [2] Hofmeister, C., Nord, R., and Soni, D.; *Applied Software Architecture*. Addison-Wesley Longman Publishing Co. Inc.(2000)
- [3] ISO/IEC 42010, “Systems and Software Engineering -- Recommended Practice For Architectural Description Of Software-Intensive Systems” (2007)
- [4] Rozanski, N. and Woods, E.; *Software Systems Architecture: Working with Stakeholders Using Viewpoints and Perspectives*. Addison-Wesley Professional (2005)
- [5] Feiler, P.H.; Lewis, B.A.; Vestal, S.; *The SAE Architecture Analysis & Design Language (AADL) A Standard for Engineering Performance Critical Systems*. In *IEEE International Symposium on Computer-Aided Control Systems Design*, pp.1206-1211 (2006)
- [6] Naeem Muhammad, Yves Vandewoude, Yolande Berbers, Sjir van Loo,; *Modelling composite end-to-end flows with AADL*, In *workshop on the definition, evaluation, and exploitation of modelling and computing standards for Real-Time Embedded Systems*. Dublin, Ireland (2009)