

A Parallelism Viewpoint to Analyze Performance Bottlenecks of Parallelism-Intensive Software Systems

Abstract

The use of parallelism enhances the performance of a software system. However, its excessive use can degrade the system performance. In this paper we propose a parallelism viewpoint to optimize the use of parallelism by eliminating unnecessarily used parallelism in legacy systems. The parallelism viewpoint describes parallelism of the system in order to analyze multiple overheads associated with its threads. We use the proposed viewpoint to find parallelism specific performance overheads of an industrial case, a precision critical parallelism-intensive electron microscope software system. Results show that reduction in parallelism requires a profound insight into the thread-model of the system, which can be achieved by using our proposed viewpoint.

Keywords: Architecture Viewpoint; Software Performance; Multithreading; Software Architecture.

1. Introduction

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 its use 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.

Using architecture views and viewpoints to describe an Architecture Description (AD) 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 [5]: “An architecture view is a work product expressing a system’s architecture from the perspective of its concerns.” Whereas, “An architecture viewpoint establishes the conventions for constructing, interpreting and analyzing an architecture view addressing concerns framed by the viewpoint.”

An AD of the system consists of multiple views, which are constructed by using the architectural viewpoints.

In our research work we developed a parallelism viewpoint to frame parallelism specific concerns of the stakeholders and we analyzed it for performance overheads. 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. We extend this support for concurrent systems by describing parallelism related concerns with our viewpoint.

In this paper we describe the proposed parallelism viewpoint. We identify its stakeholders and describe how this viewpoint addresses their concerns by developing parallelism specific models. Subsequently, we analyze these models to find the threads causing performance overheads.

We validate the parallelism viewpoint by using it to describe the parallelism of an industrial case; a large and complex parallelism-intensive software system used for electron microscopes. The viewpoint is generic and can be applied to similar systems.

The remainder of the paper is structured as follows. In section 2 we describe our research problem. In section 3, we outline the building blocks of the proposed parallelism viewpoint and explain our approach to model them with the help of an industrial case. Section 4 contains related work. Finally, in section 5 we draw conclusions and present our future work.

2. Research Problem

A multithreaded environment can perform multiple tasks together quite efficiently. Excessive use of threading can however affect the performance of the system by introducing the following overheads:

Context Switches: In order to execute tasks the operating system scheduler lets threads use the CPU for a specific time, after which the scheduler puts them back in a queue. Allocation and de-allocation of threads to CPU requires some time. As the

number of threads increases context switching increases, which amplifies the context switching overhead.

Read/Write: Distribution of tasks (operations) to multiple threads divides the read/write operations among them. An uneven distribution in which some threads perform too much operations whereas others remain idle, can degrade the system performance

Thread Management: As the number of threads grows their creation, task allocation and monitoring become more complex.

To analyze the performance bottlenecks caused by these overheads we will develop a parallelism viewpoint. The viewpoint contains models that address stakeholder’s concerns related to these overheads. The need for this domain-specific viewpoint arose as the existing concurrency viewpoint only provides support for concerns related to communication and synchronization among concurrent elements [3]. Whereas for the above described overheads, we essentially need support for task and timing related concerns of the system.

By using the parallelism viewpoint we will describe the parallelism of a software system which is used for electron microscopes. It is a client-server distributed system whose design follows a component-based architecture. Figure 1 shows a schematic diagram of the software. An end user of the machine interacts with the microscope hardware through a client and a server application. Because of the heterogeneous nature of the machine, its devices come from multiple domains such as electronics, mechanics and physics. The server software is responsible for data acquisition and control of these devices. Furthermore, it also performs complex computation such as image processing for the end user. It has a large code base with multi-million lines of code and employs several hundred threads to perform various microscopy functions.

3. Parallelism 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 [5]. We adopt a three step approach to outline these fundamental building blocks for our

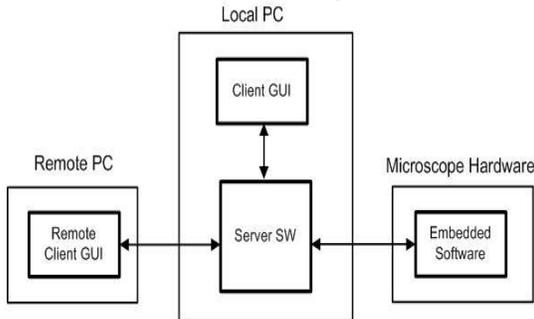


Figure 1 Schematic diagram of the electron microscope software

parallelism viewpoint. First, we identify the concerns related to the aforementioned parallelism specific performance overheads. Subsequently, we catalog the corresponding stakeholders. Finally, we develop parallelism specific model kinds to analyze the overheads. We exemplify these steps with the case mentioned in section 2.

3.1. Concerns and Stakeholders

Table 1 contains parallelism specific concerns and their corresponding stakeholders. We find these concerns necessary to identify the overheads associated with context switches, read/write operations and complex thread management. These concerns were identified with the inputs from the stakeholders of the electron microscope software and researchers from the parallelism domain.

We compile the stakeholders profile by using a template introduced by Koning and van Vliet in [12]. The choice of selecting this template is inspired by the fact that it makes the stakeholder’s position about the viewpoint explicit. The AD of the system contains profiles for all the stakeholders.

3.2. Model Kinds

A model kind is a set of notations and conventions which are used to develop domain-specific models for a viewpoint. One instance of the model kind is an architecture model that shows how concerns of the corresponding viewpoint are being addressed. An architecture description should identify at least one model kind for a viewpoint. The model kind in turn should define notations and conventions for at least one architecture model.

For the parallelism viewpoint we propose five model kinds; Time Distribution, Task Distribution, Task Type, Thread Behaviour and Thread Management. We provide notations and modelling techniques which can be used to frame stakeholders concerns in

Table 1 Parallelism specific concerns and corresponding stakeholders

Concerns	Descriptions	Stakeholders
Time Allocation	Represents the total time used by a thread.	Software architect, Developer, Tester, End User
Task Types	The nature of the tasks performed by threads e.g. file/registry read and write operations	Developer, Tester, System Maintainer
Task Distribution	Number of tasks performed by each thread.	Developer, System Maintainer
Active Time	The time when a thread is using the CPU to perform a task.	Software architect, Tester, System Maintainer
Waiting Time	The time when a thread is waiting for its turn to get the CPU time.	Software architect, Tester, System Maintainer
Execution Elements Management	A way of managing a thread’s life cycle e.g. a thread creation and deletion mechanism	Developer, Testers

the viewpoint. Furthermore, we also explain how model kinds are used to analyze performance overheads discussed in section 2.

Before going into details, we first categorize the runtime information needed for these model kinds and the approach we adopted to harvest that information.

3.2.1. Execution Metamodel

Trosky et al. outline in [6] a metamodel of software execution. The metamodel illustrates in a hierarchical way the mapping of functional components to the execution elements (processes and threads) and describes the activities performed during a typical execution. Trosky used the metamodel to analyze dependencies among executing components. Because the goals of our research work, the analysis of parallelism related performance overheads, differs from the goals set by Trosky, we use a modified form of their metamodel. The metamodel is shown in figure 2. The highlighted parts represent the modifications we made.

The system execution starts with an interaction of the user. A user may perform multiple actions with the system. Since it is not possible to analyze the system execution for all user actions, some representative test scenarios can be used [7]. The selected scenarios should cover all concerns of a viewpoint. A single or a set of representative scenarios can be used for this purpose. For the example case we choose System Startup scenario. Detail of the scenario is given in table 2. We select this as a representative scenario because our example system initializes too many threads at start. This particular scenario will provide insight into unnecessary initialization performed at this phase. We illustrate only one scenario in this paper but the actual AD of the system contains multiple scenarios.

The chosen scenario is divided into a number of tasks which represent a sequence of steps required to fulfill the intended action. Further, these tasks are assigned to software components.

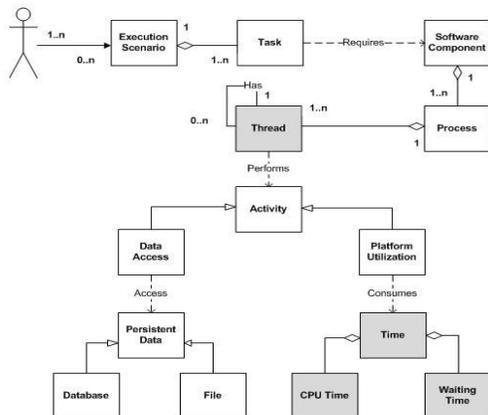


Figure 2 Software system runtime information metamodel

Subsequently, these software components are mapped to execution elements i.e. processes and threads. This mapping is vital, particularly while developing parallelism-intensive systems, as it involves the distribution of workload among execution elements. While analyzing the parallelism of a system, understanding about its workload distribution is necessary as an unbalanced distribution degrades performance of the system. A thread may spawn multiple threads to hand out certain tasks to them. Thread spawning is beneficial to a certain extent that it allows existing threads to create a new thread when needed, however an excessive spawning effects their management. To have a profound understanding about how a system manages its threads information about parent child relationship is necessary. We incorporate this information by modifying the thread element in the metamodel outlined by Trosky.

The lower part of the metamodel shown in figure 2 contains information about the activities performed by a thread for the previously identified tasks. Among activities are platform utilization and read/write operations on a file and/or a database. Understanding about these activities is necessary to frame stakeholders' concerns related to the tasks i.e. task distribution and task types. To address the timing related concerns listed in table 1 we explicitly include active and idle times of a thread to the platform utilization element of the metamodel.

3.2.2. Model Kind Generation

In the previous section we identified the execution elements required to describe parallelism related concerns. Now we describe an approach to harvest that information and develop model kinds from it. Figure 3 shows the overall scheme of our approach which consists of multiple steps. The core purpose of the approach is to harvest information concerning the execution elements, and subsequently use this information to define notation and conventions for the model kinds of our viewpoint.

Information harvesting itself is composed of two subtasks: collecting information and building a repository. In our approach, we use process logs and logs maintained by the software system under investigation. A wide range of commercial and open source monitoring tools are available to populate

Table 2 Representative scenario for the electron microscope software system

Scenario Reference	SC: System Startup
Overview	Starting server application and connecting a GUI application to it.
System Environment	Windows XP
Required System Behaviour	Processes (FeiRBBM, Feibload and FeiBBox processes of the example case) are started on operating system.

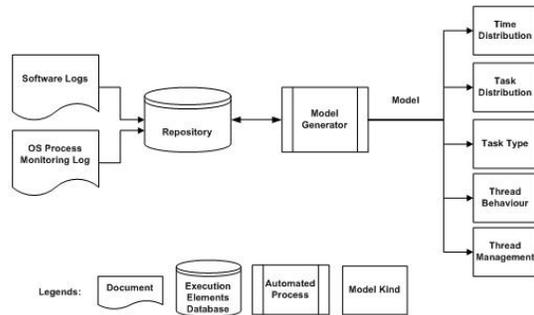


Figure 3 Overview of the proposed approach

various kinds of information about processes. Information from these logs is used to build a repository. The repository transforms this information in a structured form that is suitable for the next step i.e. model generation.

Finally, notations and conventions are built by using the information in the repository. These notations and conventions are domain-specific and to a large extent depend upon the type of the system and the kind of analysis to be performed. Our approach however is quite generic and flexible as we do not limit it to any particular technology or technique.

Now that we have outlined the required execution elements and our approach to extract them we describe the model kinds identified for the parallelism viewpoint. Every model kind will be used to perform an associated analysis. Additionally, they can also be used together to perform other kinds of analysis. For instance, the time distribution and task distribution model kinds can help in finding threads with less CPU time and performing only a very few tasks.

Time Distribution:

An important task of an OS is to allocate computer resources among competing execution entities [7]. CPU time is one of the critical resources a computer system has and that needs to be allocated in an appropriate manner. Threads which are the basic units of execution use their quota of CPU time to perform their tasks [8]. Time distribution is a

model kind that illustrates the total time used by every thread in a system over a period of time. Since the devised approach is scenario based, a single instance of this model kind shows the time distribution across threads for a particular scenario. It mainly addresses the time allocation concern in the parallelism viewpoint.

Time distribution analysis can be performed on this model kind to identify threads consuming no CPU time or a very small amount. Once identified, some strategies can be formulated to tackle with these threads. These threads can be traced in the source code by utilizing debugging techniques. Furthermore, this analysis helps in understanding the activity level of threads such as most active, mediocre and least active. Based on the results a strategy can be outlined to adjust their level.

Figure 4(a) is an instance of the time distribution model kind which we generated for our example case. Along horizontal axis it shows the threads running in the system whereas vertical bar represents the total amount of time consumed by a thread. As shown in the model about half of the threads created at the startup are consuming no CPU time, they were just created for some future purposes. Reason for such initialization could be of two fold. First, it is a design decision to create all threads at startup. But this approach becomes problematic when there are too many threads. Second, these threads are part of some dependent components whose initialization is necessary at starting phase. In this case as well, creating too many idle threads at startup means too many dependencies. In either case, this analysis not only identifies uneven time distribution across threads but also identifies possible improvements in the management of threads. This model kind can be used for multiple scenarios to spot idle threads during the complete execution of the system.

Task Distribution:

A system makes use of multiple threads to distribute its workload. Task distribution model kind portrays this distribution. It shows the total number of tasks performed by every thread of the system.

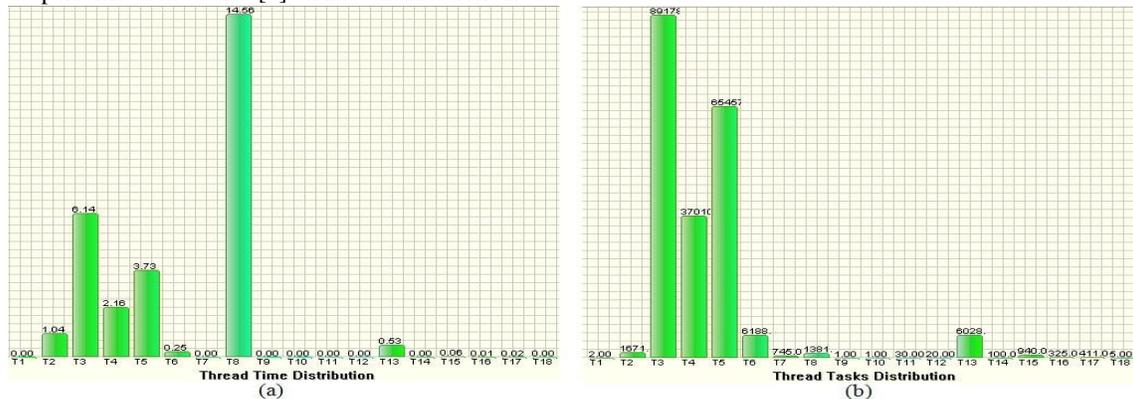


Figure 4 Time and Task distribution model

Similar to the time distribution, this model kind also depicts distribution for a particular scenario. Primarily, it addresses read/write related concerns in the viewpoint.

Task distribution can be analyzed to identify threads performing too many tasks and those with a very small number of tasks. Like the previous model kind it also helps in making a comparative analysis of threads against their number of tasks. Although this distribution depends upon the nature of the system but an uneven distribution highlights the possible improvements in the thread structure. Particularly, when the difference between these numbers is bigger.

Threads with a low number of tasks can share burden from threads performing too many tasks or can be eliminated by assigning their tasks to some other related threads.

Figure 4(b) shows an example model of the task distribution model kind generated for the example software system. In this model kind too, threads are shown along the horizontal axis whereas vertical axis represents the total number of tasks performed by a thread. We can observe an uneven distribution of tasks across threads. Threads on the left side of the model are performing a huge number of tasks whereas threads on the opposite side are calm. The relative difference in these numbers among threads is notable. Some threads are too much active i.e. performing read/write operations whereas a majority is idle.

By using this model together with the time distribution model we can observe that many threads exist in the system but they are useless as they are doing nothing.

Thread Behaviour:

Along with the overall distribution of time across threads it is important to understand the active/idle behaviour of a system. This determines how important a thread is, at least from the timing perspective. Thread behaviour is a model kind that portrays this behaviour by showing activities of a single thread performed during its life cycle. To a large extent associated with the active time and idle time concerns of stakeholders it also deals with the overall time consumption of the thread and context switches. Similar to the previous model kinds it also illustrates thread behaviour based on a particular scenario. An instance of this model kind is shown in figure 5. The model kind shows the sequence of Active (A) and Waiting (W) times of a thread.

Essentially this model kind can be used to analyze threads suitable for thread pooling. Thread pooling is a technique in which a number of threads are created at the start of a software system and are reused to perform a set of tasks that resides in a queue [9]. This technique can improve the performance of a system as it makes available a set of reusable threads. Reusing threads for multiple

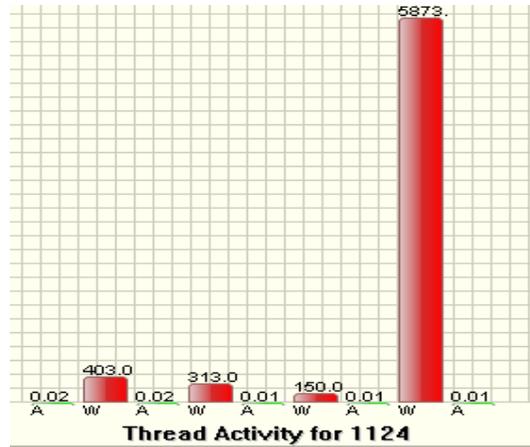


Figure 5 Thread behaviour model

tasks reduces thread creation and deletion overheads. The use of a thread pool is effective when used with a medium size pool and shorter tasks.

Thread behaviour model kind identifies and portrays active/idle behaviour of the threads that can be substituted with a pool of threads. These threads are those which are idle most of the time during their life cycle and performing a very few tasks. The model shown in figure 5 describes the active/idle behaviour of a thread, which is active in total only for a shorter period of 10 ms during its life cycle and is performing no more than 100 tasks. These kinds of threads are perfectly suitable to be replaced with a pool of threads. For the electron microscope software system we find many such threads. We designate threads consuming total time less than 10 ms and performing tasks less than 100 appropriate for thread pooling. These numbers are flexible and can be adjusted if required, provided that the total time consumption and number of tasks remain shorter.

Thread behaviour model kind is also suitable for identifying the number of context switches for every thread as it describes their active and idle positions. An overhead is involved in switching context from one thread to another. For threads with too many context switches a strategy can be planned to eliminate them.

Task Type:

Task type model kind defines notation and conventions to generate models that show the nature of the tasks (operations) performed by threads of a multithreaded system. These types are read and write operations on files and database. Understanding the nature of threads' tasks is important to categorize them and to know what kind of tasks a thread is designated for. This model is helpful in identifying those threads performing related kinds of operations. This information can be used while reducing the number of threads in a system by integrating multiple threads to one. Furthermore, this model kind also gives insight into the frequency and time consumption of each task performed by a thread.

Figure 6(a) outlines a list of tasks performed by each thread running in the electron microscope software system. The left part of the model lists threads running in the systems whereas the opposite side shows task related characteristics of a thread.

Thread Management:

Thread management model kind represents the actions performed during the life cycle of a thread. These actions include thread creation, activities of thread and their deletion. Previous model kinds describe activities of the thread whereas this model kind is mainly associated with the thread creation and deletion.

The example model shown in figure 6(b) describes the thread management structure of the example system. Each cell represents the number of child for a thread, for instance every 4th thread (T4) of the system on average creates 29 child threads. To find an average of the number of child threads multiple samples can be collected for a single scenario.

In the example model we can observe that many threads in the electron microscope system are spawning a large number of threads. However, the distribution of this number is not constant for every execution of the scenario. This highlights the need for improvements in the thread management structure of the system. Excessive spawning can degrade the system performance and makes the thread management structure of the system complex and more error prone. Furthermore, this model kind provides insight into the accumulative overhead associated with the creation and deletion of a thread.

The feedback of the stakeholders of the electron microscope software system was encouraging as they found the parallelism viewpoint quite efficient in identifying the complexities and overheads caused by the excessive use of threads. The results of the analyses identified more improvements opportunities in the system than the stakeholders were anticipating.

The explicit categorization of the parallelism related elements in the system execution made it easy for the stakeholders to retain focus on parallelism related issues. Moreover, with the parallelism specific models the stakeholders found it

easy to capture and understand the parallelism behaviour of the system.

4. Related Work

The use of an architecture viewpoint for modelling stakeholders’ concerns is common practice in both industry and research. Hereunder we present related work from both areas.

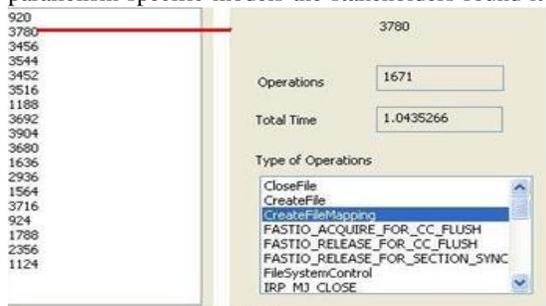
Trosky et al. [4] introduced an approach to construct viewpoints for a large and complex software intensive MRI system from Philips. Similar to our work they also customized the concurrency viewpoint to model domain-specific concerns. However, our work differs in the nature of concerns as we modeled parallelism related concerns compared to control and data flow related concerns.

Flanagan et al. [10] proposed a modular approach called Calvin for analyzing the thread behaviour of multithreaded software systems. They analyzed the system behaviour by performing modular checking of each procedure call made by threads present in the system. In this paper, we also analyzed the runtime behaviour of threads to understand the overheads associated to it. However, unlike Flanagan we use viewpoints for such an analysis.

To achieve similar goals Dean and Shen [11] presented an approach for integrating existing threads in order to reduce the total number of threads. In their work, they improved the performance of the system by overlapping the execution of multiple threads. To improve the performance by reducing the number of threads, our and Dean and Shen’s research work require a change in the thread model of the system.

In [12] Koning and van Vliet propose a four step method for designing viewpoints. Their approach defines explicit relationships between stakeholders’ concerns and viewpoints by developing stakeholders’ profiles. For our work we follow their guidelines while describing the stakeholders of our viewpoint.

Nicholas May [13] performed a survey of architecture viewpoint models. The survey shows that existing viewpoints do not address every concern of a particular domain and need to be



Task Types

(a)

	Thread	T1	T2	T4	T6	T9	T9	T62	T63	T67	T70	T72	T76	TOTAL	Parent Thread
Index															
1	1	4	39	26	11	0	0	0	0	2	0	3	87	7	
2	1	4	22	28	16	0	0	12	0	0	0	3	87	7	
3	1	4	22	31	23	5	1	0	0	0	0	0	88	7	
4	1	4	22	32	23	0	0	0	0	3	0	0	86	6	
5	1	4	41	19	16	0	0	0	5	0	0	0	87	6	
6	1	4	28	25	22	4	0	0	0	0	3	0	88	7	
Avg	1	4	29	27	18.5	1.5	0.1	2	0.8	0.8	0.5	1	87		

Thread Management

(b)

Figure 6 Task type and thread management models

tailored to fulfill this purpose. This strengthens our motivation for constructing a parallelism specific concurrency viewpoint. The survey also finds it necessary to complement viewpoints of different architecture frameworks to broaden the coverage of concerns.

Razavizadeh et al.[14] proposed a framework that generates viewpoint models from the source code of the system. We achieve the same goal by generating models based on the software and operating system logs.

Li and Malony [15] diagnosed the performance bottlenecks of parallel applications with the help of a model-based diagnosis framework called Hercule. In contrast we used an architecture viewpoint to diagnose performance issues.

5. Conclusions and Future Work

In this paper we presented a parallelism viewpoint to analyze parallelism related overheads in existing parallelism-intensive software systems. These overheads include excessive context switches, uneven distribution of Read/Write operations and complex thread management structure. The viewpoint describes parallelism of the system and identifies threads involved in above overheads. Subsequently, these threads can be eliminated to minimize overheads.

For the proposed viewpoint, we identified parallelism specific concerns and introduced a set of models to describe these concerns. The models describe timing and tasks distribution behaviours of the system. We find these behaviours important in understanding the characteristics of the threads running in a system. We also identified parallelism specific execution elements. This explicit identification, we believe is imperative to retain focus on parallelism related issues.

The concurrency viewpoint is used to describe the communication structure of concurrent systems. Adding to this, our viewpoint provides support for describing another important aspect i.e. parallelism of these systems.

By successfully identifying performance bottlenecks of the electron microscope software system with the help of our viewpoint, we showed 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.

We introduced 5 model kinds to describe time and tasks related parallelism specific concerns of the system. The metamodel we discussed for the parallelism viewpoint is extensible and can be used to generate some other models, to address other parallelism related concerns. For instance, models to analyze software and hardware resource utilization of the threads, and their IO operations.

As a part of this research study we are building a flow-latency viewpoint to describe latencies of a flow-intensive system. The parallelism viewpoint and flow-latency viewpoint will be a part of an AD of the electron microscope software. Our future work involves tracing links between the elements of both viewpoints. In particular, we are interested in utilizing the outcome of the viewpoint presented in this paper to understand the effects of the aforementioned overheads on flow latencies of the system.

6. References

- [1] Clements, P., Kazman, R., Klein, M.; *Evaluating Software Architecture*, Addison-Wesley (2002)
- [2] Hofmeister, C., Nord, R., and Soni, D.: *Applied Software Architecture*. Addison-Wesley Longman Publishing Co. Inc.(2000)
- [3] Rozanski, N. and Woods, E. *Software Systems Architecture: Working with Stakeholders Using Viewpoints and Perspectives*. Addison-Wesley Professional (2005)
- [4] Trosky Boris Callo Arias, Pierre America, Paris Avgeriou: *Defining Execution Viewpoints For A Large And Complex Software-Intensive System*. WICSA/ECSA, pp. 1-10. (2009)
- [5] ISO/IEC 42010, "Systems and Software Engineering -- Recommended Practice For Architectural Description Of Software-Intensive Systems" (2007)
- [6] Arias, T. B., Avgeriou, P., and America, P.: *Analyzing the Actual Execution of a Large Software-Intensive System for Determining Dependencies*. In *Proceedings of the 2008 15th Working Conference on Reverse Engineering*, pp. 49-58. WCRE. IEEE Computer Society, Washington, DC(2008)
- [7] Somé, S. S. and Cheng, X.: *An Approach For Supporting System-Level Test Scenarios Generation From Textual Use Cases*. In *Proceedings of the ACM Symposium on Applied Computing*, pp. 724-729. SAC '08. ACM, New York, NY (2008).
- [8] William Stallings: *Operating System: Internals and Design Principles*. Pearson Higher Education (2009)
- [9] Jalal Raissi: *Performance Impact of Thread Pooling in DSOCARE*. In *Proceedings of the IEEE SoutheastCon*, pp. 108 – 113. Memphis, TN (2006).
- [10]. Cormac Flanagan , Stephen N. Freund , Shaz Qadeer , Sanjit A. Seshia,: *Modular Verification of Multithreaded Programs*. *Theoretical Computer Science*, v.338 n.1-3, p.153-183 (June 2005)
- [11]. Dean, A. G. and Shen, J. P.: *Techniques for Software Thread Integration in Real-Time Embedded Systems*. In *Proceedings of the IEEE Real-Time Systems Symposium*, pp. 322. RTSS. IEEE Computer Society, Washington, DC, (1998)
- [12]. Koning, H. and van Vliet, H.: *A Method For Defining IEEE Std 1471 Viewpoints*. *J. Syst. Softw.* 79, 120-131 (2006)
- [13]. May, N.: *A Survey Of Software Architecture Viewpoint Models*. In *Proceedings of the Sixth Australasian Workshop on Software and System Architectures*, pp. 13-24. Melbourne, Australia (2005)
- [14]. Razavizadeh, A., Cimpan, S., Verjus, H., and Ducasse, S.: *Software System Understanding via Architectural Views Extraction According to Multiple Viewpoints*. In: R. Meersman, P. Herrero, and T. Dillon (Eds), *Lecture Notes In Computer Science*, vol. 5872, pp. 433-442, Springer-Verlag, Berlin, Heidelberg (2009)
- [15]. Null Li Li and Allen D. Malony,: *Automatic Performance Diagnosis of Parallel Computations with Compositional Models*. In *IEEE International Parallel and Distributed Processing Symposium*, p. 211(2007)