

Elastic Process Optimization and Scheduling in the Cloud

Philipp Waibel¹

Abstract: Business processes in today's industry are getting more complex and resource-intensive from day to day. Cloud computing offers methods and technologies that can help managing this steadily increase of complexity in the business process executions. A new class of Business Process Management Systems (BPMS), called elastic BPMS (eBPMS), is using Cloud computing resources for the enactment of so-called Elastic Processes (EP). Those eBPMS are combining traditional BPMS with Cloud controllers and are thus able to use Cloud resources for the EP enactment in an ad hoc manner. This paper focuses on our current work on a research eBPMS called ViePEP and discusses open research questions in the field of eBPMS.

Keywords: Elastic Processes, Elastic BPMS, ViePEP, Business Process, Cloud Computing

1 Introduction

The focus of Business Process Management (BPM) is the management and optimization of companies' business processes. The daily work with business processes is supported by so-called Business Process Management Systems (BPMS). Such a BPMS supports the customer during the complete BPM lifecycle, including the definition of a business process and its enactment [We12]. A business process is a composition of several human tasks as well as software tasks that have to be executed on computational resources. Furthermore, the design of a business process can range from simple sequential executed process steps to more complex processes with loops, splits or choices [va03]. A special case of business processes, called workflows, are processes that only consist of software services [Ho15].

Workflow execution requests may occur in a regular interval, as well as in an ad hoc manner. Especially in the latter case, the rapidly changing resource requirements can lead to a suboptimal usage of the available computational resources (e.g., CPU or RAM). This may lead to over-provisioning, where more resources than needed are available, or to under-provisioning, where less resources than needed are available and the system may not be able to fulfill the request in time [Sc15].

Cloud computing offers a promising solution for the optimal usage of computational resources. By using well-known basic principles of Cloud computing, like leasing and releasing computational resources on-demand, BPMS are able to elastically react to rapidly changing resource requirements. A BPMS that is using Cloud resources to execute workflows is called an *elastic BPMS* (eBPMS) and the corresponding business processes are called *elastic processes* (EPs) [Du11]. Besides the management of the process lifecycle, an eBPMS is also a Cloud controller that leases and release Cloud resources, deploys services on them, and starts the services execution according to a schedule [Ho16].

¹ TU Wien, Distributed Systems Group, Argentinierstrasse 8, 1040 Wien, p.waibel@infosys.tuwien.ac.at

This work-in-progress paper aims at presenting our former work on an eBPMS called ViePEP and discusses further open research questions in the field of eBPMS.

2 Vienna Platform for Elastic Processes

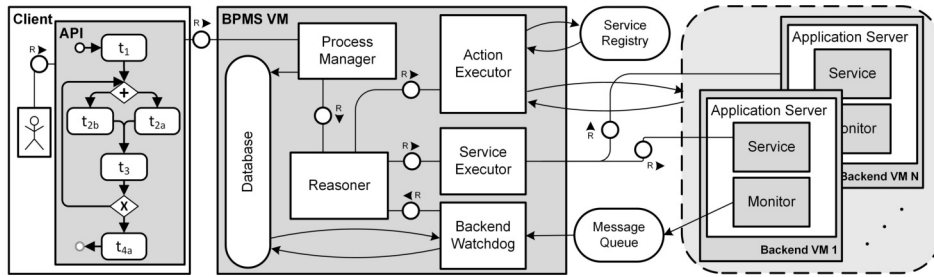


Fig. 1: ViePEP - Architecture

ViePEP (Vienna Platform for Elastic Processes) is a state-of-the-art eBPMS solution that combines the functionalities of a traditional BPMS with a Cloud controller that is able to enact software services on Cloud resources in a cost-efficient way [Sc13a]. This is done by continuous monitoring, optimization and scheduling of the available Cloud resources, in form of Virtual Machines (VMs), and service enactment while considering predefined Service Level Agreements (SLAs) to guarantee a certain level of Quality of Service (QoS).

As can be seen in the Figure 1, ViePEP consists of five top level entities. The first entity, the *Client*, is responsible for the modeling of the business process, the definition of optional SLAs and the issuing of business process enactment requests. The second entity, called *BPMS VM*, is the core of ViePEP. It is responsible for the enactment of the process and the optimization of the Cloud resources by solving an optimization problem in continuous intervals. The BPMS VM is also responsible for the leasing and releasing of VMs on Cloud resources and the execution of *Backend VMs* on them. Those Backend VMs are then responsible for the execution of services that represent the process steps. Finally, the *Message Queue* is used by the Backend VMs to send monitoring information to the BPMS VM. The *Service Repository* hosts the services that can be deployed on Backend VMs.

Due to space constraints we can not discuss each entity in greater detail, but the interested reader can find a more detailed description in [Sc13a].

3 Research Questions

Our former work on ViePEP raises several research questions, which are discussed in the following.

In the recent years the concept of Docker's significantly gained popularity. Compared to a traditional VM, Docker is a more light-weight system, which can be used to execute services that need less additional computational resources by surpassing the need of an

own operating system and is thus cheaper to execute than a VM. At the current status, ViePEP executes services on VMs. By updating the used optimization problem to allow Docker-based service enactment, the overall execution costs can be reduced even further.

To be able to optimize the usage of the Cloud resources as cost-efficient as possible while considering predefined SLAs, ViePEP analyzes the workflow including the process steps that are not yet executed. During this analysis, ViePEP tries to predict and schedule the execution of the upcoming process steps. This analysis raises two new research questions: The first question concerns the execution duration of a process step. Currently ViePEP uses Ordinary Least Squares Linear Regression, a 2-dimensional regression, to predict the execution duration of a process step. However, for many real-world services, this prediction mechanism is not sufficient and a multi-dimensional resource prediction is needed. The second question concerns the execution of a workflow. For a sequential workflow, the scheduling is straightforward. However, real-world workflows can be more complex and include XOR-gates and loops. In the first case, an analyzing software cannot predict which route the actual execution will take. Here, ViePEP considers the route with the longest execution duration and schedules according to these times, i.e., ViePEP applies a worst-case approach. Furthermore, if the process model contains loops, ViePEP cannot predict how often this loop will be taken. At the current state, ViePEP considers a predefined amount of loop rounds. By using stochastic analyzes both solutions can be enhanced. This means that a “smart” BPMS can learn from previous process executions and with this information, predict how often the loop will be executed or which branch of the XOR-gate will be taken.

Finally, the architecture of ViePEP can be extended by changing the central component *BPMS VM* (see Section 2) into a decentralized component. This adds the possibility to distribute the load, which is currently processed at one instance of the BPMS VM, onto several BPMS VM instances, so the BPMS VM does not become a bottleneck.

4 Related Work

To the best of our knowledge, relatively little research has been done in the field of eBPMS, as well as scheduling and resource usage of EPs. Besides our own work [Ho16, Sc13a, Ho13, Sc13b, HSD13], the work of Cai et al. [CLG13], Wei and Blake [WB14] and Bessai et al. [Be13] are worth mentioning. For a full discussion of the related work, we refer to [Sc15, Ho16].

Similar to ViePEP, in the work of Cai et al. [CLG13] the authors define an optimization model that provides a cost-efficient Cloud resource usage for EPs. However, their work only considers single processes. Therefore, their approach is not able to share resources among different processes like it is done by ViePEP. Bessai et al. also provide algorithms for cost- and time-efficient scheduling of business processes in Cloud computing environments [Be13]. In the work of Wei and Blake [WB14], the authors also provide algorithms for using Cloud resources for the execution of business processes. In comparison to the work of Bessai et al., they also allow QoS constraints. However, in comparison to ViePEP, they do not consider deadline constraints.

Acknowledgment

This work is partially supported by the Commission of the European Union within the CREMA H2020-RIA project (Grant agreement no. 637066).

References

- [Be13] Bessai, Kahina; Youcef, Samir; Oulamara, Ammar; Godart, Claude: Bi-criteria strategies for business processes scheduling in cloud environments with fairness metrics. In: 7th Int. Conf. on Research Challenges in Information Science. IEEE, pp. 1–10, 2013.
- [CLG13] Cai, Zhicheng; Li, Xiaoping; Gupta, Jatinder N.D.: Critical Path-Based Iterative Heuristic for Workflow Scheduling in Utility and Cloud Computing. In: 11th Int. Conf. on Service Oriented Computing. volume 8274 of LNCS. Springer, pp. 207–221, 2013.
- [Du11] Dustdar, Schahram; Guo, Yike; Satzger, Benjamin; Truong, Hong Linh: Principles of Elastic Processes. IEEE Internet Computing, 15(5):66–71, 2011.
- [Ho13] Hoenisch, Philipp; Schulte, Stefan; Dustdar, Schahram; Venugopal, Srikumar: Self-Adaptive Resource Allocation for Elastic Process Execution. In: 6th Int. Conf. on Cloud Computing. IEEE, pp. 220–227, 2013.
- [Ho15] Hoenisch, Philipp; Hochreiner, Christoph; Schuller, Dieter; Schulte, Stefan; Mendling, Jan; Dustdar, Schahram: Cost-efficient scheduling of elastic processes in hybrid clouds. In: 8th Int. Conf. on Cloud Computing. IEEE, pp. 17–24, 2015.
- [Ho16] Hoenisch, Philipp; Schuller, Dieter; Schulte, Stefan; Hochreiner, Christoph; Dustdar, Schahram: Optimization of Complex Elastic Processes (accepted for publication). IEEE Transactions on Services Computing, NN:NN–NN, 2016.
- [HSD13] Hoenisch, Philipp; Schulte, Stefan; Dustdar, Schahram: Workflow Scheduling and Resource Allocation for Cloud-based Execution of Elastic Processes. In: 6th Int. Conf. on Service Oriented Computing and Applications. IEEE, pp. 1–8, 2013.
- [Sc13a] Schulte, Stefan; Hoenisch, Philipp; Venugopal, Srikumar; Dustdar, Schahram: Introducing the Vienna Platform for Elastic Processes. In: Performance Assessment and Auditing in Service Computing Workshop at 10th Int. Conf. on Service Oriented Computing. volume 7759 of LNCS. Springer, pp. 179–190, 2013.
- [Sc13b] Schulte, Stefan; Schuller, Dieter; Hoenisch, Philipp; Lampe, Ulrich; Dustdar, Schahram; Steinmetz, Ralf: Cost-Driven Optimization of Cloud Resource Allocation for Elastic Processes. Int. Journal of Cloud Computing, 1(2):1–14, 2013.
- [Sc15] Schulte, Stefan; Janiesch, Christian; Venugopal, Srikumar; Weber, Ingo; Hoenisch, Philipp: Elastic Business Process Management: State of the Art and Open Challenges for BPM in the Cloud. Future Generation Computer Systems, 46:36–50, 2015.
- [va03] van der Aalst, Wil M. P.; ter Hofstede, Arthur H. M.; Kiepuszewski, Bartek; Barros, Alis-tair P.: Workflow Patterns. Distributed and Parallel Databases, 14(1):5–51, 2003.
- [WB14] Wei, Yi; Blake, M. Brian: Proactive virtualized resource management for service workflows in the cloud. Computing, 96(7):1–16, 2014.
- [We12] Weske, Mathias: Business Process Management: Concepts, Languages, Architectures. Springer, 2nd edition, 2012.