Journal of Systems Architecture 73 (2017) 17–27

Contents lists available at ScienceDirect

Journal of Systems Architecture

journal homepage: www.elsevier.com/locate/sysarc

Evolutionary planning of virtualized cyber-physical compute and control clusters

Boguslaw Jablkowski*, Ulrich Thomas Gabor, Olaf Spinczyk

Department of Computer Science, TU Dortmund, Germany

ARTICLE INFO

Article history:
Received 7 May 2016
Revised 20 August 2016
Accepted 1 November 2016
Available online 9 November 2016

Keywords:
Cyber-physical systems
Virtual machines
Real-time guarantees
Evolutionary algorithms
Formal performance analysis

ABSTRACT

Virtualization technology has the potential to notably advance the automation process in the domain of cyber-physical systems (CPS). It can improve both dependability and availability as well as significantly reduce the procurement, operation and maintenance costs of such systems. However, in the context of virtualization, research has put the most emphasis on topics of hardware utilization and fault-tolerance. There is little literature on how to model, integrate and consolidate a CPS by means of virtualization. In this paper we present a methodology for planning safe and efficient virtualized cyber-physical compute and control clusters – execution platforms for time-constrained virtual machines (VMs) that encapsulate CPS applications. We discuss the used methods, describe the corresponding models and the required system architecture. In contrast to typical resource allocation problems from other domains (e.g. cloud computing), in this case, the planning process must take real-time requirements of applications into account. In order to achieve this, we combine evolutionary algorithms with formal system performance analysis – in particular algorithms considered in classical scheduling theory. Such an approach allows not only to optimally dimension the compute and control clusters, but also provides strict guarantees regarding the timing predictability of the integrated CPS. Further, the embedment of a formal performance analysis technique notably eases the modeling of a system. As a consequence, the modeling process is fast, flexible and accessible not only to experts but also to system designers as they do not have to struggle with complex and time consuming mathematical formulations. Finally, our approach also provides answers to several practical questions that arise when integrating a CPS by means of virtualization.

1. Introduction

Substantial technological advances in sensing, computing and communication are facilitating the development of cyber-physical systems (CPS). Therefore, it comes as no surprise that the already vast scope of application for CPS is continuously expanding. CPS can be found in various domains such as power systems, robotics, avionic systems, autonomous automotive, intelligent buildings, and many more. They are also believed to be a key enabler for future technologies. However, their flexibility and versatility of use entail some nontrivial challenges. The major ones are their complexity and management. CPS are inherently heterogeneous and distributed. These attributes alone pose a significant challenge in the analysis, development and deployment of CPS. This issue is being further complicated by the fact that recently, in order to use and manage the new opportunities provided by the continuously improving infrastructure, the amount of required software is rising rapidly. As a consequence, modern CPS are evolving into software-intensive systems and it stands to reason that in the future this trend will gather momentum additionally increasing the role of software and therefore the complexity of CPS.

Fortunately, these problems are not new nor are they limited to the field of CPS. There are other domains that have to struggle with similar challenges. A good example are large data centers where comparable issues were successfully tackled by the technique of virtualization. The success of this technique lies in the fact that virtualization allows for a transparent integration and consolidation of system components, thus reducing their complexity and as a consequence eases their management. Furthermore, virtualization is also known for its excellent isolation properties – both in space and time dimension – an essential aspect for providing high degree of fault-containment in safety-critical systems. Moreover, features like live migration, which facilitates proactive maintenance, or high availability solutions that allow to survive hardware failures render this technology so interesting.

* Corresponding author.
E-mail addresses: boguslaw.jablkowski@tu-dortmund.de (B. Jablkowski), ulrich.gabor@tu-dortmund.de (U. Thomas Gabor), olaf.spinczyk@tu-dortmund.de (O. Spinczyk).

http://dx.doi.org/10.1016/j.sysarc.2016.11.001
1383-7621/

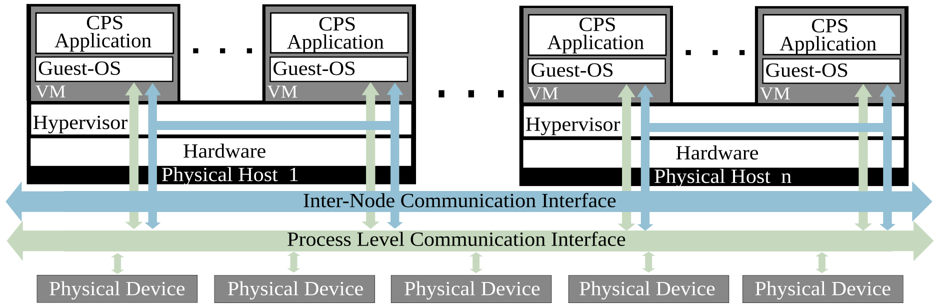

Fig. 1. Architecture overview.

The idea of integrating a CPS into one functional whole on a homogenous platform has the potential to strongly reduce its complexity and therefore the management effort as well as improve dependability and availability of the system. Additionally, it can reduce the procurement, operation and maintenance costs. However, till recently research has put the most emphasis on the topics hardware utilization and fault-tolerance. In the context of virtualization, there is little literature on how to integrate a CPS or how to model and fulfill the strict timing constraints and needed predictability requirements of CPS. This is still a new and open research field.

This paper aims at contributing to this area by addressing the issue of CPS integration by means of virtualization. It proposes a methodology for planning virtualized compute and control clusters that aids at providing an answer to the following question: How to find a suitable mapping between the VMs encapsulating the CPS applications and an appropriately dimensioned hardware and at the same time fulfill the strict timing requirements of CPS? In order to answer this question the following aspects of CPS have to be taken into account:

- Cyber-physical systems are mostly real-time systems. This implies that computational delays of their critical functions are as important as their functional correctness. The violation of a real-time property of a critical task can lead to damage or even a catastrophe. In order to avoid such problems, guarantees regarding timing properties have to be provided.
- Due to the distributed character of CPS, a local – single core or host – timing analysis is not sufficient. A global interference analysis of the system components is required and therefore metrics like end-to-end latencies have to be obtained. This implies that the analysis also has to take both communication as well as data dependency into account.
- CPS control or protection functions have to be characterized by high availability. As a consequence, the analysis has to allow to incorporate high availability solutions (active VM replication) into its model.
- CPS vary significantly in their size. Therefore, the analysis should remain efficient and scale with the increasing size of a CPS.

Finally, the planning process should additionally allow for the optimization of non-functional requirements such as the minimization of latencies, migration overhead, or the amount of needed physical hosts. The first can lead to increased service frequency, the second facilitates dynamic reconfiguration of the system in exchange for higher network bandwidth, the last one reduces costs.

All the constraints and objectives render the discussed problem a multiobjective optimization problem (MOP) from the class of NP-hard problems. The issue is being additionally complicated by the fact that in our case the criteria for providing guarantees regarding the timing properties of a given cluster cannot be represented as a simple function. In order to determine a feasible solution an exhaustive system performance analysis of each postulated VM to hardware mapping has to be conducted. To solve these challenges we propose an approach that combines multiobjective optimization with formal system performance analysis. In particular, we use evolutionary algorithms to tackle the exponential search space of the problem and adopt classical scheduling theory formulas in order to analyze the timing properties of the solution candidate in question and therefore its feasibility.

The contributions of this article can be summarized as follows:

- We propose a holistic approach for the integration of CPS into virtualized cyber-physical compute and control clusters.
- We describe the methodology behind the approach, define the used methods, the corresponding models, as well as the assumed system architecture.
- We combine evolutionary algorithms (EA) with algorithms considered in classical scheduling theory. This allows not only for an efficient search of the solution space but also for a detailed modeling of the problem – including CPU-scheduling, which in turn enables to provide guarantees regarding the timing predictability of the system and – at the same time – to optimize additional non-functional requirements.
- Finally, our approach also provides answers to numerous practical questions that arise when integrating or managing a CPS: How to fulfill the strict requirements without an over-dimensioning of the system? Can additional functions be safely deployed in an already running system? Or, how to efficiently plan maintenance?

The paper is organized as follows. First, a generic architecture for an integrated CPS is being presented in Section 2. Next, in Section 3, we discuss the methodology behind our approach, the applied techniques – including the EA, system performance analysis and integer linear programming – as well as the corresponding models. After presenting the evaluation results, together with solutions to practical questions, in Section 4, we discuss related work in Section 5. Finally, Section 6 summarizes and concludes the article.

2. Virtualized CPS-architecture

The starting point for our approach forms the virtualized CPS architecture. Fig. 1 depicts how the information and communication infrastructure interweaves with the CPS and how we assume the generic architecture of the integrated CPS to look like. Every computational node (physical host) represents a server running a bare-metal hypervisor – a thin software layer running directly on the hardware. The purpose of the hypervisor is the management of the VMs. Each hypervisor is able to concurrently host a specific number of VMs – this number is bounded by the available resources and the applied technology. Today's solutions support up to hundreds of VMs per host. The VMs encapsulate a guest