Open Access

Open Journal of Internet of Things (OJIOT)
Volume 3, Issue 1, 2017
http://www.ronpub.com/ojiot
ISSN 2364-7108

Towards a Model-driven Performance Prediction Approach for Internet of Things Architectures

Johannes Kroß [A], Sebastian Voss [A], Helmut Krcmar [B]

[A] fortiss GmbH, Model-based Systems Engineering, Guerickestr. 25, 80805 Munich, Germany, {kross, voss}@fortiss.org
[B] Technische Universität München, Chair for Information Systems, Boltzmannstr. 3, 85748 Garching, Germany, krcmar@in.tum.de

Abstract

Indisputable, security and interoperability play major concerns in Internet of Things (IoT) architectures and applications. In this paper, however, we emphasize the role and importance of performance and scalability as additional, crucial aspects in planning and building sustainable IoT solutions. IoT architectures are complicated system-of-systems that include different developer roles, development processes, organizational units, and a multilateral governance. Its performance is often neglected during development but becomes a major concern at the end of development and results in supplemental efforts, costs, and refactoring. It should not be relied on linearly scaling for such systems only by using up-to-date technologies that may promote such behavior. Furthermore, different security or interoperability choices also have a considerable impact on performance and may result in unforeseen trade-offs. Therefore, we propose and pursue the vision of a model-driven approach to predict and evaluate the performance of IoT architectures early in the system lifecylce in order to guarantee efficient and scalable systems reaching from sensors to business applications.

Type of Paper and Keywords

Visionary paper: *performance, model, simulation, prediction, evaluation, Internet of Things, IoT, architectures*

1 Introduction

Since several years Internet of Things (IoT) constitutes one of the main future topics for industries [3]. Information and communication technologies for small devices continuously become not only more affordable, but also more powerful regarding processing. This enables these devices to be connected to the Internet. Additionally, big data technologies emerged and enabled organizations to store huge amounts of data and to analyze incoming data streams with sophisticated algorithms in real-time [11]. This has facilitated the evolution of IoT and enables organizations to build solutions for a highly diverse range of use case scenarios in different domains. Therefore, IoT may be considered as an umbrella term for different disciplines that already have longer histories (e.g., industry automation) and, additionally, promotes the integration of these different disciplines, for instance, the automatic combination of

This paper is accepted at the *International Workshop on Very Large Internet of Things (VLIoT 2017)* in conjunction with the VLDB 2017 Conference in Munich, Germany. The proceedings of VLIoT@VLDB 2017 are published in the Open Journal of Internet of Things (OJIOT) as special issue.

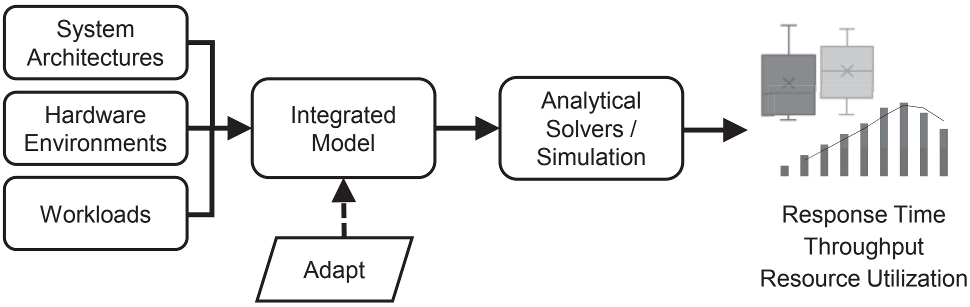

Figure 1: Model-based prediction approach

sensor data with enterprise resource planning data.

Although being promoted very much, only a few IoT use cases are implemented in industry yet. On the contrary, there are already hundreds of IoT platforms and technologies that are waiting to be exploited. In addition, there are several initiatives to define standards [3]; however, their establishment progresses slowly and an oversupply of vendor-specific implementations hamper the development of integrated solutions. For instance, an IoT developer survey with 528 participants conducted by the Eclipse IoT Working Group, IEEE IoT, and AGILE-IoT suggests that security, interoperability, and connectivity represent the three major concerns among all participants for developing IoT solutions [14]. However, for developers and organizations that have already deployed IoT solutions, performance becomes the third concern over connectivity. This reflects our comprehension that performance is not considered sufficiently when building architectures and finalized developments become very costly to counteract on late in the software life cycle.

We emphasize the role and importance of performance in terms of response time, throughput, and resource utilization. It is a vital aspect in planning and building sustainable IoT solutions as they involve multi-domain environments including constrained devices, gateways, and platforms of which the latter combines big data technologies and business applications. All these levels can have a direct impact on the performance of an overall system. Furthermore, evaluating the impact of design choices (e.g., regarding security, interoperability, and platforms) at development time is difficult, especially, for large-scale operations. These are only some of the factors that complicate IoT performance management.

In order to address and solve these issues, we propose the vision of a model-based approach for representing components and performance-influencing factors of IoT architectures and allow for *performance-by-design*. These models shall serve as input for analytical solvers or simulation engines and allow for predicting different performance metrics (Figure 1) [6]. In this way, architectures can proactively be evaluated regarding bottlenecks and scalability. Required resource demands can be planned and the throughput and response time behavior of subsystems can be estimated. The derived performance metrics and predictions shall also contribute to support communication and collaboration between developers (e.g., embedded developers and developers for business applications) and operations.

2 MODEL-DRIVEN PERFORMANCE PREDICTION

Our vision and its realization is driven by the following three research questions, which we use to explain our proposal and intentions:

1. What resource requirements and performance difficulties occur and are relevant on different levels of IoT architectures?
2. Which existing approaches and technologies can be used for implementing the integrated modeling concept?
3. How can existing meta-models and simulation approaches of different levels be integrated and combined?

In order to address the first research question the different levels and developer roles of architectures must be considered. Figure 2 shows a very basic IoT architecture that is reduced to the essential three layers. First, constrained devices and controllers represent the things in IoT. Second, gateways, routers, and smart devices enable fog computing at the edge and may integrate as well as pre-analyze data from devices [13]. Third, platforms process, store, and aggregate data from different sources and enable business applications to analyze and report data to end users. The connectivity and communication among the levels is not limited to one direction. In addition, non-functional requirements

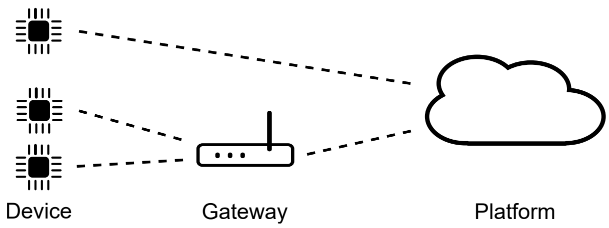

Figure 2: Abstract IoT architecture

such as performance, security, and interoperability are topics that influence all levels. For performance engineers, for instance, the following questions occur:

1. How shall computing resources (e.g., CPU, disk, memory, network) be sized on each level?
2. Shall gateways pre-analyze data and of how many devices per gateway?
3. What is the impact of protocol, security, and interoperability choices on the overall performance?
4. Does the architecture scale linearly with increasing number of devices and gateways?

Since the IoT stack is highly diverse, different developers and engineers are involved in the implementation. Embedded and systems developers are responsible on the device level and also partly on the gateway level. As gateways continuously expand, become smarter, and are able to run sophisticated operating systems, application developers also constitute a part on the gateway level. On the platform level, a mix of data scientists, application developers, and web developers implement the integration and visualization of data. Due to this mix of interests and engineering disciplines, we see the need to investigate the performance requirements on each level and for each role in order to understand influencing factors in a holistic view that need to be included in our model approach.

Similarly, previous and present related modeling approaches consider these disciplines in separated ways. As mentioned, IoT provides and increases the opportunity of combining existing approaches. Use case scenarios arise, for instance, that require capacity planning for devices and gateways based on formal models which are already well understood in the domain of business information systems. Since there is a tremendous number of modeling approaches, the second research question addresses reviewing existing methods and technologies for different levels with regards to our vision. In the following we list one example technology for each level.

For the device level, for instance, AutoFOCUS3[1] represents an integrated model-based tool for the development process of embedded systems [1]. It includes the activities for modeling requirements, software architectures, hardware platforms, and deployments as well as for generating code. The software architecture is built up by different software components that may be connected to each other to allow for interactions and may also be decomposed into multiple hierarchical subcomponents. The hardware architecture includes resources such as processors and memory that can be linked. It also involves platform architectures for execution and runtime environments such as operating systems or Java virtual machines. The integration and combination of these models enables developers to apply different analysis and synthesis methods such as testing, model checking, deployment, and automated scheduling [1].

The Eclipse Framework for Distributed Industrial Automation and Control (4diac)[2][3] is part of the Eclipse IoT ecosystem and represents an instance for modeling the gateway level. It provides an open source infrastructure for distributed industrial process measurement and control systems based on the IEC 61499 standard [17]. In order to model software architectures, 4diac includes an application editor that allows for representing function block networks consisting of one or multiple function blocks and their interaction via events. Similarly, a separate editor is included to model the specification of hardware by modeling devices and resources. By the means of several more editors and an own runtime environment, 4diac supports the development of industrial IoT applications and facilitates portability, interoperability, configurability, and scalability as promoted by IEC 61499 [17].

For the platform level, the performance management work tools (PMWT)[4] provide several integrated approaches to automate, support and integrate performance engineering activities across the software lifecycle [6]. This includes the automatic generation of models for enterprise applications based on performance measurements [5], modeling complex user behavior of applications [15], and simulating the performance of big data applications [9].

In order to address the third research question, we will examine similarities of model-based approaches for the different levels and domains of IoT architectures. For instance, models on architecture-level may often separate their meta-model as illustrated in Figure 1.

[1] http://af3.fortiss.org
[2] https://eclipse.org/4diac/
[3] http://fortiss.org/research/projects/4diac/
[4] https://pmw.fortiss.org/tools/pmwt/

One or several models are used to describe the software and system architecture, its components and its activity flow. Another model is used to describe the hardware and resource environment such as computing nodes with processors, disks, and memory connected via a network. An additional usage model is used to describe the use case scenarios of the software architecture and the workload.

Although implicitly considering performance aspects, present solutions focusing on the device and gateway level usually concentrate on guaranteeing functionality and safety [1]. In contrast, there are a lot of performance models to predict and analyze behavior on the platform level. Existing models on architecture-level that provide the benefits we seek with our vision, however, only focus on classical business applications and involve different requirements. For instance, the workload of business applications is mostly user-driven such as the number of parallel user accesses, whereas IoT applications are mostly data-driven such as the volume, velocity, and variety of incoming data. In addition, massive distributed and parallel computing and resource clusters are properties that are usually not found in business applications. Therefore, we aim at combining existing model approaches and additionally implementing missing functionalities so we are able to model the performance requirements we identified in the first research questions.

In order to be able to predict the performance behavior of the architectures, we will also implement simulators and solvers for deriving different metrics such as response time, throughput, and resource utilization. To evaluate our approach (Figure 1), we plan to model applications from IoT benchmark suites and adapt these models for various scenarios such as increasing number of things and increasing resource capacities. After simulating the models, we will compare simulation results with measurement results of applications from the benchmark that we adapted in the same way.

3 Related Work

As already mentioned, IoT involves and combines a variety of application domains and development stacks. Consequently, there are a lot of related, but also highly diverse approaches of which we try to refer to examples to the best of our knowledge in this Section. There are also several solutions available to model constrained devices, simulate networks within IoT architectures on a very detailed level, and evaluate them for throughput and latency issues. For instance, Wang et al. [16] apply the network simulator OPNET for IoT cloud solutions; Brambilla et al. [4] propose a simulation methodology to test large-scale IoT systems with interconnected devices in urban environments and include several network protocols and different mobility, network, and energy consumption models.

Furthermore, many related approaches specifically analyze and compare the performance of different protocols or technologies on different layers. For scenarios in which devices and gateways do not have a wired connection, for instance, Costantino et al. [7] investigate LTE as a suitable interconnection in terms of its efficiency, bandwidth, and coverage. In contrast, Daud and Suhaili [8] provide a performance evaluation of protocols for the application layer in IoT architectures. Therefore, they compare the hypertext transfer protocol (HTTP) and the constrained application protocol (CoAP) for message formatting, communication, and request handling on different test beds.

There are several developments of performance benchmarks for IoT, however, mostly on the platform level. Arlitt et al. present an analytics benchmark called IoTAbench [2]. It allows for generating, loading, repairing and analyzing synthetic data and was evaluated by the example of a smart metering use case and using a HP Vertica database. Shukla et al. [12] propose another benchmark for distributed stream processing platforms (i.e., Apache Storm) called RIoTBench. They provide different data workloads and generators as well as a set of 27 common IoT tasks for different domains. Furthermore, Medvedev et al. [10] provide an evaluation of different IoT platforms with regards to performance characteristics.

4 Conclusion and Future Work

This paper proposes and pursues the vision of an model-based approach for predicting and evaluating the performance of IoT architectures and systems. It shall support developers and engineers at examining design choices early in the system lifecycle, finding potential bottlenecks, planning and sizing required resources on different levels, and predicting response times from sensors to visual results. We will start our future research and work with combining and integrating modeling approaches for embedded systems with approaches for big data systems as well as for business information systems. Therefore, we are currently developing a first prototype for an integrated modeling environment.

REFERENCES

[1] V. Aravantinos, S. Voss, S. Teufl, F. Hölzl, and B. Schätz, "AutoFOCUS 3: Tooling Concepts for Seamless, Model-based Development of Embedded Systems," in *Proceedings of the 8th International Workshop on Model-based Architecting of Cyber-physical and Embedded Systems*, 2015.

[2] M. Arlitt, M. Marwah, G. Bellala, A. Shah, J. Healey, and B. Vandiver, "IoTAbench: An Internet of Things Analytics Benchmark," in *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering*, 2015, pp. 133–144.

[3] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A Survey," *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[4] G. Brambilla, M. Picone, S. Cirani, M. Amoretti, and F. Zanichelli, "A Simulation Platform for Large-scale Internet of Things Scenarios in Urban Environments," in *Proceedings of the International Conference on IoT in Urban Space*, 2014, pp. 50–55.

[5] A. Brunnert and H. Krcmar, "Continuous Performance Evaluation and Capacity Planning Using Resource Profiles for Enterprise Applications," *Journal of Systems and Software*, vol. 123, pp. 239–262, 2017.

[6] A. Brunnert, C. Vögele, A. Danciu, M. Pfaff, M. Mayer, and H. Krcmar, "Performance Management Work," *Business & Information Systems Engineering*, vol. 6, no. 3, pp. 177–179, 2014.

[7] L. Costantino, N. Buonaccorsi, C. Cicconetti, and R. Mambrini, "Performance Analysis of an LTE Gateway for the IoT," in *Proceedings of the International Symposium on a World of Wireless, Mobile and Multimedia Networks*, June 2012, pp. 1–6.

[8] M. A. Daud and W. S. H. Suhaili, "Internet of Things (IoT) with CoAP and HTTP Protocol: A Study on Which Protocol Suits IoT in Terms of Performance," in *Proceedings of the Computational Intelligence in Information Systems Conference*, S. Phon-Amnuaisuk, T.-W. Au, and S. Omar, Eds., 2017, pp. 165–174.

[9] J. Kroß and H. Krcmar, "Modeling and Simulating Apache Spark Streaming Applications," *Softwaretechnik-Trends*, vol. 36, no. 4, 2016.

[10] A. Medvedev, A. Hassani, A. Zaslavsky, P. Jayaraman, M. Indrawan-Santiago, P. Delir Haghighi, and S. Ling, "Data Ingestion and Storage Performance of IoT Platforms: Study of OpenIoT," in *Preceedings of Second International Workshop on Interoperability and Open-Source Solutions for the Internet of Things, Stuttgart, Germany, November 7*, I. Podnar Žarko, A. Broering, S. Soursos, and M. Serrano, Eds., 2017, pp. 141–157.

[11] M. Schermann, H. Hemsen, C. Buchmüller, T. Bitter, H. Krcmar, V. Markl, and T. Hoeren, "Big Data - An Interdisciplinary Opportunity for Information Systems Research," *Business & Information Systems Engineering*, vol. 6, no. 5, pp. 261–266, 2014.

[12] A. Shukla, S. Chaturvedi, and Y. Simmhan, "RIoTBench: A real-time IoT benchmark for distributed stream processing platforms," Tech. Rep., 2017. [Online]. Available: http://arxiv.org/abs/1701.08530

[13] K. Skala, D. Davidovic, E. Afgan, I. Sovic, and Z. Sojat, "Scalable Distributed Computing Hierarchy: Cloud, Fog and Dew Computing," *Open Journal of Cloud Computing*, vol. 2, no. 1, pp. 16–24, 2015. [Online]. Available: https://www.ronpub.com/ojcc/OJCC_2015v2i1n03_Skala.html

[14] I. Skerrett, "Profile of an IoT Developer: Results of the IoT Developer Survey," 2016. [Online]. Available: https://ianskerrett.wordpress.com/2016/04/14/profile-of-an-iot-developer-results-of-the-iot-developer-survey/

[15] C. Vögele, A. van Hoorn, E. Schulz, W. Hasselbring, and H. Krcmar, "WESSBAS: Extraction of Probabilistic Workload Specifications for Load Testing and Performance Prediction — A Model-Driven Approach for Session-based Application Systems," *Software & Systems Modeling*, pp. 1–35, 2016.

[16] J. Wang, X. Shi, M. Alhussein, L. Peng, and Y. Hu, "SPSIC: Semi-Physical Simulation for IoT Clouds," *Mobile Networks and Applications*, vol. 21, no. 5, pp. 856–864, 2016.

[17] A. Zoitl, T. Strasser, and A. Valentini, "Open Source Initiatives as Basis for the Establishment of New Technologies in Industrial automation: 4DIAC a Case Study," in *Proceedings of the IEEE International Symposium on Industrial Electronics*, 2010, pp. 3817–3819.

AUTHOR BIOGRAPHIES

Johannes Kroß received a B.Sc. degree in Business Information Systems from Middlesex University in London and a M.Sc. degree in Information Systems from Technical University of Munich. At the fortiss institute he is responsible for the Performance Management Group within the Model-based Systems Engineering competence field. The group focuses on all aspects required to ensure that given performance requirements for application systems can be met.

Sebastian Voss has done his Ph.D. in the avionic context at EADS Innovation Works in the department Sensors, Electronics & Systems Integration. Previously, he worked one year for Daimler research & development. At fortiss he is heading the Model-based Systems Engineering competence field. His research interests include techniques and tools for the professional development of software-intensive systems, design space exploration methods and model-based (tool) development in AutoFOCUS3.

Helmut Krcmar studied business management in Saarbrücken and obtained his doctorate in 1983. Afterwards, he worked as a postdoctoral fellow at the IBM Los Angeles Scientific Center and as assistant professor of information systems at the New York University and the City University of New York. From 1987 to 2002, he held the Chair for Information Systems at the University of Hohenheim, Stuttgart. Since 2002 he holds the Chair for Information Systems at the Technical University of Munich (TUM). From 2010 to 2013, he served as Dean of the Faculty of Computer Science.