

EUPaaS

Elastic Ubiquitous Platform as a Service for Large-scale Ubiquitous Applications

Fei Li¹, Schahram Dustdar¹, Jakob Bardram², Martin Serrano³, Manfred Hauswirth³,
Vasilios Andrikopoulos⁴ and Frank Leymann⁴

¹*Distributed Systems Group, Vienna University of Technology, Vienna, Austria*

²*Cetrea A/S, Aarhus, Denmark*

³*Digital Enterprise Research Institute, National University of Ireland Galway, Galway, Ireland*

⁴*Institute of Architecture of Application Systems, University of Stuttgart, Stuttgart, Germany*

{li, dustdar}@dsg.tuwien.ac.at, bardram@c3a.dk, {Martin.Serrano, Manfred.Hauswirth}@deri.org,
{Vasilios.Andrikopoulos, Frank.Leymann}@iaas.uni-stuttgart.de

Keywords: Platform as a Service, Cloud Computing, Ubiquitous Computing, Healthcare, Hospital.

Abstract: Large-scale ubiquitous computing applications are rapidly emerging in the fields of pervasive healthcare, smart cities and so on. They present unprecedented challenges to state-of-the-art ubiquitous systems in the respects of accommodating fluctuating user demands, handling volatile data quality and adaptation to complex system and user contexts. Driven by a motivating scenario in future mega-hospital environment, we propose to exploit the potential of cloud computing in supporting large-scale ubiquitous computing applications. This position paper will present the novel concept of EUPaaS (Elastic Ubiquitous Platform as a Service), outline the key research topics, and propose a cloud-based ubiquitous application platform.

1 INTRODUCTION

In recent years, new business and research opportunities have been increasingly emerging from the need for large-scale ubiquitous computing systems (Weiser, 1991), including for example pervasive healthcare, smart cities, and so on. The development and provisioning of these systems face a range of significant challenges. From a system's point of view, they are characterized by employing large numbers of mobile devices and specialized sensors connected in a volatile network setup. They generate and process massive amounts of real-time data while coping with fluctuation of user demands. From the user's point of view, these systems need to ensure responsiveness for ad-hoc usages, support user mobility, adapt to user context, and provide highly personalized usage experience. Both viewpoints need to be coherently supported in large-scale ubiquitous computing systems.

Future mega-hospital¹ is a typical application environment of such systems. The scale of a future mega-hospital will be in the range of 10,000 clinicians; 100,000 hospitalizations and 900,000 out-

patient treatments per year; covering 380,000 m^2 ; 35,000 ambulances and cars arriving every day; and having 1,300 beds. Although ubiquitous computing technologies have been successfully applied in hospital environments (Bardram, 2009), providing context-awareness to such amount of users, with many of them in time- and life-critical procedures, poses unprecedented challenges, including elasticity, large-scale context-awareness, reliable data processing, application adaptation and so on.

This paper proposes to address the challenges in mega-hospital and enable large-scale ubiquitous computing systems by consolidating the core concepts of two computing paradigms—Cloud computing and Ubiquitous computing. Conceptually, ubiquitous computing implies using massive numbers of small computational entities to reach all aspects of human life. Intuitively, the virtually infinite computing power of the cloud can help resource-constrained ubiquitous systems to reach large-scale. However, the state-of-the-art ubiquitous computing systems have mostly been applied on a small, personal scale and focused on personal mobility and context-awareness. The ongoing work on combining ubiquitous computing and cloud computing still focuses on leveraging the computing and storage resources of the cloud

¹Plan of the New University Hospital at Aarhus: <http://www.dnu.rm.dk/english>

(Satyanarayanan, 2010; Chun et al., 2011). To this end, we present *Elastic Ubiquitous Platform as a Service (EUPaaS)*, which is focused on exploiting the potential of clouds in supporting large-scale ubiquitous systems. Specifically, this position paper will outline a set of novel concepts in our ongoing research on the convergence of the two complementary computing paradigms, and propose a cloud-based ubiquitous application platform. Overall the research presented in this position paper will achieve the following objectives:

- To genuinely transfer the key characteristics of cloud computing, especially elasticity and service orientation, into large-scale ubiquitous systems;
- To support the development and execution of mobile, ubiquitous, context-aware applications in PaaS paradigm;
- To realize cloud-native ubiquitous applications that possess the characters of both cloud and ubiquitous applications;
- To provide real-time data services to large amounts of ubiquitous applications by quality-aware, real-time data processing on clouds;
- To compose ubiquitous device services and real-time data into adaptive applications through service-oriented computing.

The paper is organized as follows: Section 2 presents a use case and corresponding challenges in mega hospital environments; In Section 3 we present the novel concepts of the work; then Section 4 proposes the EUPaaS architecture under development; the paper is concluded in Section 5.

2 UBIQUITOUS COMPUTING IN FUTURE MEGA-HOSPITAL

Ubiquitous computing technologies have proven particularly applicable in a hospital setup (Bardram, 2009). Sensors are deployed hospital-wide for location tracking of staff, patients, and equipment; video surveillance; bed monitoring; and biomedical sensing. All of these sensor inputs help build a real-time overview of the clinical logistics at the hospital and are used for context-aware information visualization on large displays and on mobile devices. Moreover, context-aware notifications are forwarded to relevant clinical personnel on a timely fashion using the most appropriate devices. However, as the size of a hospital is increasing, the scalability of these types of ubiquitous computing systems becomes significantly challenging. This is especially true in the case of future

mega-hospitals—even the most simple context-aware applications based on a location tracking infrastructure will need to handle location updates from 10,000-20,000 entities in a large area with a sampling rate of e.g. 1 minute. More importantly, a large amount of time- and life-critical hospital procedures need to be supported at the same time in the future mega-hospital. The following use case illustrates one of these procedures.

2.1 Clinical Logistics for an Emergency Case

Monday morning, Mr. Hansen is not feeling well and his left arm is hurting. His wife takes him to the Emergency and Accident Department (E&AD) at the local hospital for examination. At the E&AD, the Emergency Clinical Logistic System (ECLS) helps the head nurse to quickly assign a doctor, a nurse, and an examination room. In the examination room, his medical information and real-time data on ECG, pulse, and O₂ is displayed on the interactive ECLS display. When examining Mr. Hansen, the doctor realizes that this is a critical and acute case that needs immediate transfer to the University Hospital (DNU) for acute specialized surgery coronary angioplasty. Together with the head nurse, the doctor uses ECLS to initiate a transfer request from this local E&AD to the Department of Cardiac and Thorax Surgery (Dept. T) at DNU. The head nurse at Dept. T is notified about the incoming patient via the ECLS system, and accepts the transfer.

Upon acceptance, the medical data of Mr. Hansen is relayed through the regional healthcare backbone infrastructure. At Dept. T, the ECLS interactive displays at the nursing station start to show real-time information about Mr. Hansen. Porters and an ambulance are automatically scheduled and are notified to pick up the patient in the local hospital and drive him to DNU. The ECLS displays show the location of the ambulance as it is approaching. Once in the ambulance, medical information on Mr. Hansen and his triage is entered in the ambulance medical record, and real-time data on ECG, pulse, and O₂ is obtained through sensors in the ambulance. All of this data is continuously displayed at the nursing station of DNU.

In the meanwhile, based on real-time knowledge on the location, workload, medical profile, and work schedule of the doctors and nurses in Dept. T, the ECLS system starts suggesting a surgeon and a team of operating nurses to the incoming patient case. The final assignment is approved and done by the nursing station, and the involved staff is notified through their ECLS mobile phones. Upon confirmation of

availability, the real-time information about the patient is available from the mobile phones of the assigned medical staff. At the same time, an available operation room and all the things needed for a surgery are booked by the ECLS.

Mr. Hansen arrives and the porters can use their mobile ECLS system to see where to take him. In the operating room everybody is ready for surgery. The surgery is successful and Mr. Hansen is moved to the intensive care unit (ICU) for recovery. Once recovered, he is moved to the patient ward where his wife is waiting for him. She has been able to follow his treatment along the way using the patient/relative ECLS mobile phone app.

2.2 Challenges

The use case alone can be realized with the technologies developed in the past research on context-aware systems. However, supporting such procedure at a very large scale presents unprecedented challenges to both hospital IT platform and ubiquitous applications.

- *Elasticity.* The work load of a hospital is very uneven across each day and week. Relatively few activities are taking place on a Sunday afternoon and night, mostly emergencies, whereas there is a steep increase in activity (and hence infrastructure demand) starting Monday morning at 7 o'clock. Moreover, since a hospital to a large degree is designed to handle critical situations like emergencies, acute patients being transferred from other smaller hospitals, and epidemic situations, its peak load cannot be predicted. There is therefore a fundamental need for elasticity in a hospital's IT systems.
- *Context-awareness in Large-scale.* For a hospital at the scale of the DNU, a large number of emergency processes, like in the use case, are performed in parallel day and night. Besides emergencies, many more routine and less-critical processes of the hospital are also being carried out all the time. Context-awareness is needed prevalently in these processes. There have been many solutions to provide contextual information to applications through dedicated sensors. However, it requires new approaches to acquiring and sharing context for large amounts of applications and users on a large-scale, heterogeneous sensory infrastructure.
- *Data Quality and Data Processing.* In a life-and-time-critical environment like a hospital, data quality has significant impact on users. For example, in the use case, if the location information

of a doctor is not updated in 3 minutes by RFID-based tracking (this could happen frequently because of signal interference), the ECLS needs to alternatively use WiFi location tracking based on the doctor's wireless connection on mobile phone. Furthermore, the system needs to accommodate differentiated data quality requirements for different applications. In the use case, timeliness and accuracy of biological signs are essential for the treatment of Mr. Hansen. In future mega-hospitals, these kinds of data will need to be delivered from large amount of devices to many doctors reliably, timely, and accurately even when the hospital is at peak load, e.g. at 7 o'clock Monday morning.

- *Application Environment.* In mega-hospitals like DNU, many applications like ECLS and other clinical systems are used by a large number of users at the same time. For different users, an application may employ different sensors, require different data processing schemes, and provide information on different user interfaces. To ensure the required service quality for a user, an application needs to adapt to many factors, including availability of devices, fluctuation of data quality, changes of user context, etc. Furthermore, a mega-hospital is not only a centralized healthcare facility, but also a center of regional healthcare ecosystem. Many healthcare services need to perform beyond the boundary of a single hospital and collaborate with other facilities. As suggested by the use case, the system should allow applications to seamlessly interact with authorized outside entities like ambulances, devices and local hospitals.

3 ELASTIC UBIQUITOUS PLATFORM AS A SERVICE

To address the aforementioned challenges and enable large-scale ubiquitous computing systems, this paper proposes Elastic Ubiquitous Platform as a Service (EUPaaS), which constitutes a set of novel concepts to be investigated in the intersection of cloud computing and ubiquitous computing.

• Elastic Ubiquitous Computing

Elasticity is the infrastructural capability that computing resources allocated to an application can be scaled up and down according to the workload of the application while maintaining a certain service level. Elasticity has not been considered so far in ubiquitous applications because of their limited scales and closed system environments.

However, with the emergence of large-scale ubiquitous systems like future mega-hospital, it is important to recognize the need for elasticity in ubiquitous computing environments and realize it on the cloud. Elastic ubiquitous computing means that computing resources allocated to an application can be scaled up and down depending on the number and types of devices the application employs, on the needs to acquire and process the data provided by the devices, and on the quality level needed by the application for a certain situation.

- **Platform as a Service for Ubiquitous Applications**

Platform as a Service (PaaS) (Liu et al., 2011) is defined as "the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations." The PaaS model allows application providers to focus on their application logic by reusing software services on cloud and transparently leveraging the elasticity of cloud platform. For web applications, the tools provided by PaaS usually include identity management, data as a service, messaging infrastructure and so on. For ubiquitous applications, it is necessary to extend these basic platform services with the key facilities of ubiquitous computing, including virtualized device access, device life-cycle management, real-time data services, and adaptive application runtime environment. By offering these facilities as services on the cloud, the development of ubiquitous applications boils down to a composition of these services.

- **Cloud-native Ubiquitous Applications**

The EUPaaS will enable cloud-native ubiquitous applications that inherently possess the following essential cloud and ubiquitous computing characteristics:

- Elasticity—Applications will be consistently responsive to users in need regardless of the load on underlying resources.
- Multi-tenancy—Applications can provide each user with a highly personalized and contextualized usage experience by virtually partitioning data and devices.
- Virtualization—Applications can dynamically bind and use heterogeneous ubiquitous devices

that are virtualized as services on the cloud.

- Context-awareness—Applications will be responsive to users surroundings, activities, physical conditions and any perceivable information that may affect users needs.

- **Real-time Quality-aware Data Service**

Producing, processing, and consuming real-time data is a fundamental aspect of ubiquitous applications. Thus, it is desirable to extend cloud data services with the capability of processing large-amount of real-time data. Furthermore, in order to accommodate the volatile data quality in ubiquitous environments while delivering reliable services, quality of real-time data should be evaluated and assured for applications. The ongoing research on data quality has been focused on database applications (Batini and Scanapieco, 2006), which generally deal with persistent and relatively stable data. And in ubiquitous computing field, the research results on Quality of Context (QoC) are not yet feasible enough to be adopted in large-scale context-aware applications (Bellavista et al., 2013). It is necessary to extend and adapt conventional data quality metrics to suit the needs of real-time data in ubiquitous environments. Quality assurance mechanisms should be developed based on these metrics and provided as quality-aware data services for applications. Furthermore, the capacity of cloud can be exploited to support real-time data quality-awareness in large-scale ubiquitous computing environments.

- **Adaptive Context- and Data Quality-aware Applications**

The applications in ubiquitous environments need to be adaptive to various situations while keeping them functioning and responsive. As demonstrated in the mega-hospital environments, it is important to enable applications to adapt to complex context changes (location, activity, users concerned, and resources available), and to data quality fluctuations (accuracy, freshness, completeness). To cope with such multiple distinct volatilities of the environment, past research on self-adaptive systems and context-aware computing should be put into perspective of cloud-based applications: in offering ubiquitous applications with a PaaS environment, such adaptiveness should be provided as part of the platform capabilities and leveraged by applications through their cloud-based run-time environment.

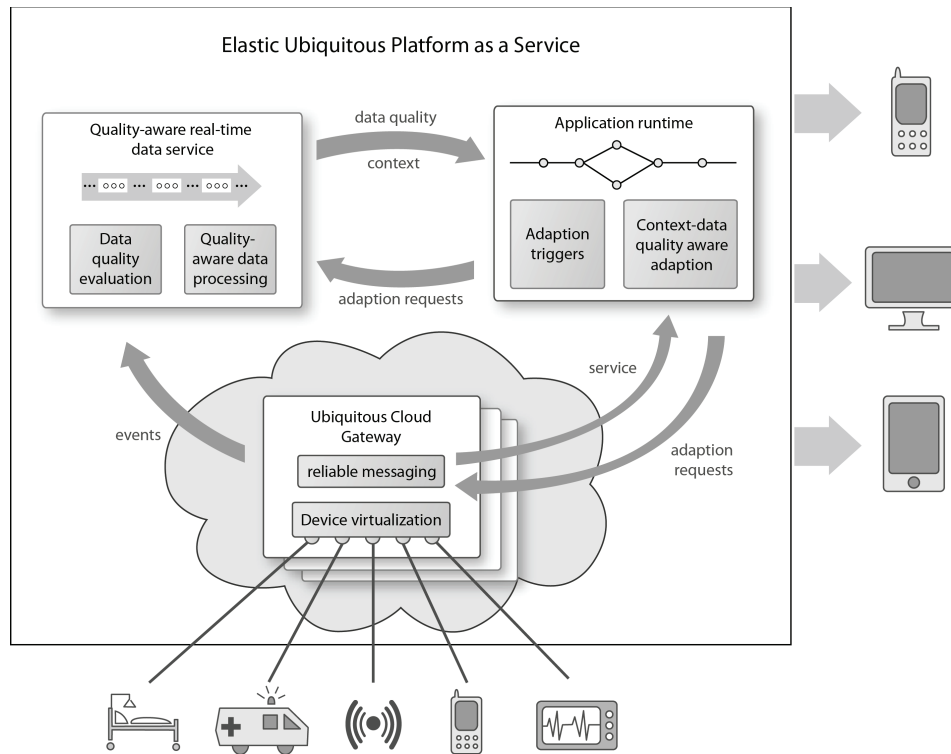


Figure 1: The EUPaaS architecture.

4 THE EUPaaS ARCHITECTURE

This section presents the EUPaaS architecture that realizes the aforementioned core concepts.

EUPaaS follows a federated service-oriented middleware architecture. The foundation for this architecture is the UCG (Ubiquitous Cloud Gateway), which is dedicated to integrating large amounts of sensors and devices within a cloud infrastructure. It provides a virtualization layer for heterogeneous entities through existing service standards. Given the ad-hoc and dynamic nature of ubiquitous environments, a life-cycle model of virtualized services is being developed in order to monitor these entities. At the core of UCG is a messaging infrastructure to manage the reliability of communication and interoperability of devices. The messaging infrastructure takes into consideration both the diverse requirements of ubiquitous applications, as well as the nature of mobile, volatile, light-weight devices. The UCG is designed as a cloud-native component with support for multi-tenancy (Strauch et al., 2012) and elasticity. Overall, the federation of UCGs is the key to the success of the EUPaaS architecture. The federation eventually allows EUPaaS platform deployments to share load between UCG instances, integrate with geographi-

cally distributed devices, and to connect different autonomous domains in mega-hospital environments.

In the EUPaaS architecture, quality-ensured, real-time data are provided to applications as a service. In previous work, we have developed a set of feasible metrics and applied it to real-world datasets (Li et al., 2012). Based on these metrics, formal semantic annotations is being developed for annotating data streams with quality information. The evaluation framework particularly considers a) the data quality change during data processing, e.g. aggregation of data streams with different accuracies, and b) the data quality perceived by applications, e.g. the completeness of multiple streams required in a certain situation. The data quality assurance approach under development is based on statistical methods, which can dynamically compensate for quality fluctuation according to application requirements, available resources, and the characters of the stream (e.g. aggregative, single, etc.). The data quality model, and evaluation and assurance mechanisms will be implemented into a large-scale real-time data processing system on cloud². Eventually, the quality-aware data provisioning capability is open to ubiquitous appli-

²<http://www.gigaspace.com/datagrid>

cations through service interfaces. Applications can then specify their data quality requirements while acquiring real-time information of the application environments.

EUPaaS applications carry out the adaptive execution of complex ubiquitous applications. An EUPaaS application model based on service composition is being developed to assess the impact of context and data quality changes to the application execution. Since the adaptation of applications can be triggered by various factors, the architecture is not focused on supporting isolated adaptation mechanisms for different situations. Instead, we aim at supporting multiple adaptation mechanisms from a wide range of options provided by the system, which can include replacement of services, changing of data sources, and modification of the application logic. Adaptation triggers are defined based on real-time quality information from the data services, a system wide view of resources, and the application context.

5 CONCLUSIONS

This position paper presented our vision and ongoing work of EUPaaS (Elastic Ubiquitous Platform as a Service). The work is motivated by the need for large-scale ubiquitous systems in future mega-hospitals. The challenges we have faced mainly include elasticity, large-scale context-awareness, reliable data processing and application adaptation. To address these challenges, we propose to enable the seamless convergence of cloud computing and ubiquitous computing by researching and realizing a series of novel concepts, including elastic ubiquitous platform as a service, cloud-native ubiquitous applications, real-time quality-aware data service and context- and data quality-aware applications. The research and development work is in progress and will be integrated into the EUPaaS software platform.

It is worth noting that the challenges addresses in this work are not limited to hospital environments. The need for large-scale ubiquitous systems, propelled by the emergence of Internet of Things, is evident in applications in smart cities, smart healthcare, smart transportation and so on. The proposed cloud-based approach will not only enable context-awareness in large-scale, but also extend the application domains of cloud computing to better support the applications based on the growingly ubiquitous sensors and smart devices.

ACKNOWLEDGEMENTS

This work is sponsored by Pacific Controls Cloud Computing Lab (PC³L), a joint lab between Pacific Controls L.L.C., Dubai, United Arab Emirates and the Distributed Systems Group of the Vienna University of Technology.

REFERENCES

- Bardram, J. E. (2009). Activity-based computing for medical work in hospitals. *ACM Transactions on Computer-Human Interaction*, 16(2):1–36.
- Batini, C. and Scannapieco, M. (2006). *Data Quality-Concepts, Methodologies and Techniques*. Springer Berlin Heidelberg New York.
- Bellavista, P., Antonio, C., MARIO, F., and FOSCHINI, L. (2013). A Survey of Context Data Distribution for Mobile Ubiquitous Systems. *ACM Computing Surveys*, 45(1).
- Chun, B.-G., Ihm, S., Maniatis, P., Naik, M., and Patti, A. (2011). CloneCloud-Elastic Execution between Mobile Device and Cloud. In *Proceedings of the sixth conference on Computer systems - EuroSys '11*, page 301, New York, New York, USA. ACM Press.
- Li, F., Nastic, S., and Dustdar, S. (2012). Data Quality Observation in Pervasive Environments. In *The 10th IEEE/IFIP International Conference on Embedded and Ubiquitous Computing (EUC 2012)*, Paphos, Cyprus.
- Liu, F., Tong, J., Mao, J., Bohn, R., Messina, J., Badger, L., and Leaf, D. (2011). Nist cloud computing reference architecture. *NIST Special Publication*, 500:292.
- Satyanarayanan, M. (2010). Mobile computing: The next decade. In *Proceedings of the 1st ACM Workshop on Mobile Cloud Computing & Services Social Networks and Beyond - MCS '10*, pages 1–6, New York, New York, USA. ACM Press.
- Strauch, S., Andrikopoulos, V., Sáez, S. G., Leymann, F., and Muhler, D. (2012). Enabling Tenant-Aware Administration and Management for JBI Environments. In *(to appear in) Proceedings of the 5th International Conference on Service-Oriented Computing and Applications (SOCA'12)*. IEEE Computer Society Conference Publishing Services.
- Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3):94–104.