A Research Roadmap
for Context-Awareness-Based
Self-managed Systems

Weishan Zhang[1], Klaus Marius Hansen[2], and Paolo Bellavista[3]

[1] Department of Software Engineering, China University of Petroleum, No. 66
Changjiang West Road, Qingdao, China, 266580
zhangws@upc.edu.cn
[2] Department of Computer Science (DIKU), University of Copenhagen, Njalsgade
128, 2300 Copenhagen S, Denmark
klausmh@diku.dk
[3] DISI, University of Bologna, V.le Risorgimento, 2 - 40136 Bologna, Italy
paolo.bellavista@unibo.it

Abstract. Cloud computing, autonomic computing, pervasive and mobile computing tends to converge, maximizing the benefits from different computing paradigms. This convergence makes emerging applications such as search and rescue applications, smart city and smart planet applications, e.g., to minimize the power consumption of a full city to achieve green computing vision, more promising on the one hand, but more complex to manage on the other hand. Interesting research questions arise due to this convergence. For example, how to efficiently retrieve underlying contexts that are difficult to recognize especially with resource-limited handheld devices, how to make use of these contexts for achieving self-management, and how to process large-scale contexts. These challenges require that researchers from software engineering, artificial intelligence, pattern recognition, high-performance distributed systems, cloud and mobile computing, etc. collaborate in order to make systems work in an efficiently self-managed manner.

1 Introduction

There is a growing trend for the convergence of different computing paradigms, such as cloud computing, pervasive and mobile computing, among others due to the needs of global availability and sharing of information at different scales [5][7]. For example, there are more and more work on mobile cloud computing research [7] where mobile computing node is enhanced by making use of a cloud computing infrastructure [6][4], or creating a light weight cloud infrastructure using the MapReduce architecture [8] [16]. Considering the development of mobile and pervasive computing, the mobile cloud computing is deemed to evolve to a so called pervasive cloud computing, where light weight clouds and classical heavy weight clouds are working harmoniously.

In the converged computing environment, in order to achieve the vision of a smart planet, it is important to obtain meaningful information from various sources at a large scale due to the deployment of the Internet of Things/Services (IoT/IoS). However, it is not always possible to expect stable, reliable, always-available data from surrounding environments, among others due to the uncertainty, fuzziness, unavailability of data from sensors. Most often the processing, retrieval, and reasoning on contexts are to be done (sometimes partially) on small devices, which further calls for a supporting infrastructure, e.g., for dynamic offloading and management of temporarily intermittent connectivity. Therefore, applications and systems tend to become more complex than before to manage and operate in the converged world. This calls for useful self-managing capabilities.

The requirements of self-management capabilities come from different application domains, for example, next-generation e-health applications, smarter cities, up to an envisioned smarter planet where the wide-scale collection of monitoring/sensing data from the whole Earth will enable new forms of mass behavior understanding and powerful predictions (for instance, see the FuturICT Knowledge Accelerator project[1]). But the achievement of self-managing capabilities is hard. Devices/equipment in an application should be able to deal with ever-changing execution environments and information sources while offering their users timely and relevant services. Furthermore, devices should adapt to these changes in a seamless manner by following (often application-specific) quality of service requirements with the support from awareness contexts [22]. Very often the awareness of contexts is hard to achieve also, due to the ambiguity of sensed data, the big scale of the IoT/IoS events, and other factors such as the processing capabilities of handheld devices.

In the following, we will first present the application domains of self-managed pervasive services and discuss of problems that arise in the domain together with some possible solutions and research directions. We then summarize the research challenges of self-management pervasive systems. Finally, we discuss the expected results of our roadmap for near-future research in context-aware self-managed systems.

2   Applications

2.1   E-Health

In many countries, the proportion of people with a high age is increasing, thereby increasing the cost of healthcare. Pervasive health services such as health status monitoring (including for example blood pressure and pulse monitoring) and food recognition and recommendation become essential parts of e-health applications. In these examples, monitoring may be important in treatments for high blood pressure whereas food recognition may be useful for some diseases like Ulcerative Colitis (UC) where food intake may have serious consequences. These

---

[1] http://www.futurict.eu

applications involving small devices require self-managed and optimal running to, e.g., save energy in order to prolong running time.

In these healthcare applications, some contexts are hard to be obtained with resource scarce devices. For example, in a meal, several food items are mixed together maybe with seasoning. In such a case, it is hard to figure out exact food types and calculate the amount of food through hand-held smart phones [10]. Therefore, fast image recognition algorithms running efficiently on mobile phones are needed, as well as a supporting infrastructure like a public computation cloud that can be used for further recognition to improve the results obtained on a mobile phone. For some unknown food that is not in the sample image database, ontological reasoning can be utilized to obtain food type information.

Some interesting research issues that arise are how to make the whole system work in a globally balanced manner to save energy, e.g. through dynamic switching of interaction protocols [23] and which node in the cloud is used to store an ontology knowledge base to get better reasoning performance. Even the mobile phones themselves can be utilized to form a computing cloud in cases the remote cloud is not available, but need processing power in order to handle emergency cases [7]. For such a mobile cloud, self-management to achieve some quality of service requirements is very important, for example having a low response time and saving energy at the same time. Other open issues include how to make decisions on the migration of services/component among cloud nodes at self-managed manner, and which algorithms are used in the migration process.

2.2   Smart Cities

The concept of "smart cities" is emerging as a major innovation driver for better health, environment, life quality and businesses [9]. In a smart city application, the IoT/IoS is typically deployed across a city. For example, traffic jams is a major problem in large cities and here wireless sensor networks can be used to monitor the position and speed of vehicles on road networks in real-time. To handle the incompleteness, complexity and unpredictability of traffic condition information, control strategies like fuzzy sets, case-based reasoning, and genetic algorithms are used to make decisions [19]. Such traffic control systems may also need supporting infrastructures in the form of, e.g., public cloud systems to handle computation intensive tasks like route planning, and store large data sets on vehicles routes. Let us note that also e-health applications, when moving from personal and single-user services towards community-oriented solutions to detect and guide the behaviors of a large set of citizens (e.g., the people under the public health services of an entire region or country), can be seen as a very relevant part of the envisioned smart city applications, sharing with them many technical challenges, as detailed in the following.

Compared to a typical pervasive computing system, a smart city application has some new issues, mainly because the scale of an application is very large, the topology of the networking is unknown due to the uncertainly of the availability of sensors and mobility of some devices, and there is much higher heterogeneity of the inter-connected devices [21]. In such situations, it is hard to realize holistic

and comprehensive context-awareness [5] and thus high-level inference is needed to know the actual contexts and make decisions [21]. This could be using a hybrid of different artificial intelligence approaches as in [19]. For large-scale sensing in IoT applications, collaborative sensing can be used to improve context accuracy and reliability [13].

3   Research Challenges

Considering the above applications, we now summarize the technical challenges for context-awareness and self-management research in order to highlight some directions of potential future work.

Context Modeling, Recognition, and Reasoning. More fundamental research is needed to make the currently available modeling languages more capable of expressing and reasoning over probability, fuzziness, and inaccurate situations. We have seen some work towards this, such as [25], but more efforts are needed, e.g., to make the modeling and reasoning work efficiently on resource-scarce devices. To recognize certain contexts, different techniques are under investigation and not yet ready, at least from the industrial application perspective, e.g., we need faster image recognition algorithms to identify food. Also, as noted in [14], hybrid knowledge representation and reasoning is needed to combine benefits from OWL (Web Ontology Language)[2], RDF (Resource Description Framework), text-based information retrieval/extraction, and machine-learning-based probabilistic handling.

Comprehensive Context-Awareness. Critical and highly dynamic contexts are not completely considered in nowadays real-life applications [22], which partially hinder the realization of comprehensive and holistic context-awareness [5]. Also due to the large scale of the IoT/IoS applications, knowing the global dynamic status of systems, efficiently processing large number of events from different resources, and obtaining meaningful contexts from these events are important and open issues. A cloud computing infrastructure can be used to perform computationally intensive tasks, e.g., to extract contexts, and is helpful to realize comprehensive context-awareness [11], which is worth of further investigations.

Quality of Contexts (QoC). Quality of Contexts (QoC) can help to resolve uncertain and conflicting situations involving context information [1]. For instance, context information may not be accurate, such as for identified food items, as discussed above. Other sources of contexts, such as wireless sensors in the field, may be offline at any time. To handle uncertainty and inaccuracy from various context sources, new supporting ontologies are proposed in [21]. QoC may be used in this process and can be parts of context ontologies [18] used to measure raw, derived,

---

[2] http://www.w3.org/TR/owl2-overview/

and inferred contexts. Due to the large number of events and contexts in IoT/IoS applications, QoC can be used to manage such complexity, e.g., by helping to filter out irrelevant context information and keep only the metadata that are expected to be more useful for a specific user/application/execution situation. How and to what extent an application is tolerant enough for such inaccurate and unreliable contexts needs to be investigated.

Unexpected Adaptations. Domain experts can make decisions in unexpected situations on how to bypass problems [21]. How such kind of totally unexpected situations can be handled to make appropriate decisions computationally and automatically is an unsolved technical challenge nowadays [2]. If completely automatic control is not possible, then how can we predict the situation and time when involvement of human experts is necessary? This challenge, central to self-managed context-aware systems, is particularly important and timely for safety-critical applications.

Hybrid Intelligence. Due to the complexity, scale, uncertainty, and fuzziness nature of real-world self-managed systems, a single intelligence mechanism is hardly sufficient to deal with the 'smartness' requirements in, e.g., smart city applications. The current self-management approaches employing mainly single mechanisms cannot always meet the intelligence requirements that challenging smart applications are calling for [3]. How to use hybrid artificial intelligence techniques to get better reasoning capabilities, as shown in [19] and [24] is an interesting issue. Problems need to be resolved on how to make different mechanisms work harmoniously, how to distribute different reasoning capabilities among different nodes in a (partly mobile/pervasive) cloud, and how to make different mechanisms work efficiently on resource-scarce devices by automatically considering the availability of different resources.

Efficient Management of Pervasive Clouds. Pervasive and mobile cloud become promising choices for enhancing the capabilities of pervasie and mobile computing systems. To support the migration of components and services between different nodes in the named pervasive cloud, a comprehensive self-management context knowledge base is needed for pervasive cloud environments, and correspondingly, a context-management framework for pervasive clouds. Also efficient algorithms and mechanisms are needed to support the migration of services following various QoS requirements [15], which may be conflicting with each other. To achieve this, some hybrid intelligence may be needed to support making decisions at different self-management layers.

Tool Support and Testbeds for Development. Tool support for developing context-awareness and self-management applications are critical for their acceptance and adoption. There are some development environments created (for example in the Hydra (LinkSmart) [22] and MUSIC [20] projects), but tool support for self-managed applications are still in their infancy and cannot easily

provide testbeds for exploring different self-management features or adding new or removing existing self-management features. In the long run, tools should allow end-users to enable programming by themselves [2] and provide easy-to-use self-management capabilities that can be extended with very limited and rapid efforts by developers.

Unified Architecture and Standards. There has been quite some research on middleware for self-managed pervasive systems, e.g. Hydra (LinkSmart) and the self-adaptive architecture of MUSIC. But there is still no unified solution approach that can be considered as a reference architecture for self-managed systems [3]. We are even not closer to the old question of "What does it mean to program hello world for ubicomp" [2]? The three-layered architecture proposed by Kramer and Magee [12] is abstract and extensible enough to incorporate different self-management strategies, and was adopted in Hydra [22]. This work may have the potential to be used as a high-level reference architecture, but needs to be specialized for different domains and applications. Research work also needs to be expanded on autonomic software components, and formally restricting their adaptiveness to avoid any unexpected violations of architecture constraints at runtime. The lack of a unified and widespread recognized architecture of solution exacerbates the interoperability issues of the systems proposed so far; standardization and interoperability are crucial to reduce development costs and to guarantee long-term value of development efforts/investments in the field.

Security and Privacy. As well as for interoperability, so far there is very little work on enforcing conventional and physical boundaries to improve security mechanisms for context sensing [2]. Location of data thus needs to be investigated, especially in the converged computing environments, for example which node is used for storage and processing of private data in a cloud-based solution for self-managed context-aware systems. As it is hard for users to specify their privacy preferences as used in current approaches, there is a strong need to investigate automated practices and techniques for context-based privacy protection [14].

4   Expected Result

Self-management remains as one of the most important research issues for future Internet applications [17]. The realization of self-management needs cross-discipline research, design, and implementation. For example, software engineering research on autonomic software components, architecture driven self-management, dynamic configuration, and adaptation of autonomic components need to be combined.

Software engineering researchers should develop tools and environments to help the creation of self-management applications. Cloud computing researchers should bring ideas on how to make use of the cloud infrastructure to make the self-management and context-awareness working more efficiently in a more scalable way. At the same time, they will benefit from the ideas of context-awareness on

how to optimize the running of cloud systems to provide better services, and expand their services to small, mobile, and embedded devices where to dynamically offload parts of the distributed applications. Artificial intelligence researchers should develop efficient algorithms to run on small devices for conducting efficient reasoning for, e.g., optimizing system configurations and planning how to reach the best configurations. Therefore, collaborations between the different areas of software engineering, context aware computing, system and network engineering, autonomic and self-managed computing, cloud and mobile computing, and artificial intelligence is needed and critical to making the self-managing systems scale in small and resource-scarce devices, with high usability, good performance, and secure services as a result. Given the complexity of the related technical challenges, we claim that cross fertilization among different research areas and strong interdisciplinarity are essential to put together effective solutions for this promising field of context-aware, large-scale, self-managed systems.

References

1. Bu, Y., Gu, T., Tao, X., Li, J., Chen, S., Lu, J.: Managing quality of context in pervasive computing. In: Sixth International Conference on Quality Software, QSIC 2006, pp. 193–200. IEEE (2006)
2. Cáceres, R., Friday, A.: Ubicomp systems at 20: Progress, opportunities, and challenges. IEEE Pervasive Computing 11(1), 14–21 (2012)
3. Cheng, B.H.C., de Lemos, R., Giese, H., Inverardi, P., Magee, J., Andersson, J., Becker, B., Bencomo, N., Brun, Y., Cukic, B., Di Marzo Serugendo, G., Dustdar, S., Finkelstein, A., Gacek, C., Geihs, K., Grassi, V., Karsai, G., Kienle, H.M., Kramer, J., Litoiu, M., Malek, S., Mirandola, R., Müller, H.A., Park, S., Shaw, M., Tichy, M., Tivoli, M., Weyns, D., Whittle, J.: Software Engineering for Self-Adaptive Systems: A Research Roadmap. In: Cheng, B.H.C., de Lemos, R., Giese, H., Inverardi, P., Magee, J. (eds.) Self-Adaptive Systems. LNCS, vol. 5525, pp. 1–26. Springer, Heidelberg (2009)
4. Christensen, J.H.: Using restful web-services and cloud computing to create next generation mobile applications. In: Proceedings of the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages and Applications, pp. 627–634. ACM (2009)
5. Conti, M., Das, S.K., Bisdikian, C., Kumar, M., Ni, L.M., Passarella, A., Roussos, G., Tröster, G., Tsudik, G., Zambonelli, F.: Looking ahead in pervasive computing: challenges and opportunities in the era of cyber-physical convergence. Pervasive and Mobile Computing (2011)
6. Corredor, I., Martínez, J.F., Familiar, M.S.: Bringing pervasive embedded networks to the service cloud: A lightweight middleware approach. Journal of Systems Architecture (2011)
7. Dinh, H.T., Lee, C., Niyato, D., Wang, P.: A survey of mobile cloud computing: architecture, applications, and approaches. In: Wireless Communications and Mobile Computing (2011)

8. Dou, A., Kalogeraki, V., Gunopulos, D., Mielikainen, T., Tuulos, V.H.: Misco: A mapreduce framework for mobile systems. In: Proceedings of the 3rd International Conference on PErvasive Technologies Related to Assistive Environments p. 32. ACM (2010)

9. Komninos, N., Pallot, M., Schaffers, H.: Special issue on smart cities and the future internet in Europe. Journal of the Knowledge Economy, 1–16 (2012)

10. Kong, F., Tan, J.: Dietcam: Automatic dietary assessment with mobile camera phones. Pervasive and Mobile Computing (2011)

11. Kovachev, D., Klamma, R.: Context-aware mobile multimedia services in the cloud. In: Proceedings of the 10th International Workshop of the Multimedia Metadata Community on Semantic Multimedia Database Technologies (2009)

12. Kramer, J., Magee, J.: Self-managed systems: an architectural challenge. In: Future of Software Engineering, FOSE 2007, pp. 259–268. IEEE (2007)

13. Lukowicz, P., Choudhury, T., Gellersen, H.: Beyond context awareness. IEEE Pervasive Computing 10(4), 15–17 (2011)

14. Lukowicz, P., Nanda, S., Narayanan, V., Albelson, H., McGuinness, D.L., Jordan, M.I.: Qualcomm context-awareness symposium sets research agenda for context-aware smartphones. IEEE Pervasive Computing 11(1), 76–79 (2012)

15. Ma, R.K.K., Lam, K.T., Wang, C.L.: excloud: Transparent runtime support for scaling mobile applications in cloud. In: 2011 International Conference on Cloud and Service Computing (CSC), pp. 103–110. IEEE (2011)

16. Marinelli, E.E.: Hyrax: cloud computing on mobile devices using mapreduce. Technical report, DTIC Document (2009)

17. Metzger, A., Cassales Marquezan, C.: Future Internet Apps: The Next Wave of Adaptive Service-Oriented Systems? In: Abramowicz, W., Llorente, I.M., Surridge, M., Zisman, A., Vayssière, J. (eds.) ServiceWave 2011. LNCS, vol. 6994, pp. 230–241. Springer, Heidelberg (2011)

18. Miron, A.D., Satoh, I., Gensel, J., Martin, H., et al.: Modeling and measuring quality of context information in pervasive environments. In: 2010 24th IEEE International Conference on Advanced Information Networking and Applications (AINA), pp. 690–697. IEEE (2010)

19. Nakamiti, G., da Silva, V.E., Ventura, J.H., da Silva, S.A.: Urban Traffic Control and Monitoring – An Approach for the Brazilian Intelligent Cities Project. In: Wang, Y., Li, T. (eds.) Practical Applications of Intelligent Systems. AISC, vol. 124, pp. 543–551. Springer, Heidelberg (2011)

20. Rouvoy, R., Barone, P., Ding, Y., Eliassen, F., Hallsteinsen, S., Lorenzo, J., Mamelli, A., Scholz, U.: MUSIC: Middleware Support for Self-Adaptation in Ubiquitous and Service-Oriented Environments. In: Cheng, B.H.C., de Lemos, R., Giese, H., Inverardi, P., Magee, J. (eds.) Self-Adaptive Systems. LNCS, vol. 5525, pp. 164–182. Springer, Heidelberg (2009)

21. Teixeira, T., Hachem, S., Issarny, V., Georgantas, N.: Service Oriented Middleware for the Internet of Things: A Perspective. In: Abramowicz, W., Llorente, I.M., Surridge, M., Zisman, A., Vayssière, J. (eds.) ServiceWave 2011. LNCS, vol. 6994, pp. 220–229. Springer, Heidelberg (2011)

22. Zhang, W., Hansen, K.M.: Using context awareness for self-management in pervasive service middleware. In: Handbook of Research on Ambient Intelligence and Smart Environments: Trends and Perspectives, vol. 1, p. 248 (2011)

23. Zhang, W., Hansen, K.M., Fernandes, J., Schütte, J., Lardies, F.M.: Qos-aware self-adaptation of communication protocols in a pervasive service middleware. In: 2010 IEEE/ACM Int'l Conference on Green Computing and Communications (Green-Com), & Int'l Conference on Cyber, Physical and Social Computing (CPSCom), pp. 17–26. IEEE (2010)
24. Zhang, W., Hansen, K.M., Kunz, T.: Enhancing intelligence and dependability of a product line enabled pervasive middleware. Pervasive and Mobile Computing 6(2), 198–217 (2010)
25. Zhao, J., Boley, H., Du, W.: A Fuzzy Logic Based Approach to Expressing and Reasoning with Uncertain Knowledge on the Semantic Web. In: Madani, K., Dourado Correia, A., Rosa, A., Filipe, J. (eds.) Computational Intelligence. SCI, vol. 399, pp. 167–181. Springer, Heidelberg (2012)