

# Developing highly complex distributed systems: A software engineering perspective

Marco Autili · Paola Inverardi · Patrizio Pelliccione · Massimo Tivoli

Received: date / Accepted: date

**Abstract** What is a highly complex distributed system in the Future era? What are the needs that may drive the development of such systems? And what is their life cycle? Is there any new challenge for Software Engineering (SE)? In this paper we try to provide a partial answer to the above questions by characterizing few application domains that we consider of raising interest in the next years. Our thesis is that there is a need to rethink the whole software process for such systems. The traditional boundaries between static and dynamic activities disappear and development support mingles with run time support thus invading the middleware territory.

**Keywords** Future Internet · Service-oriented Computing · Service-oriented Middleware · Cloud Computing · Middleware-based Software Process.

## 1 Introduction

The near future envisions a pervasive heterogeneous computing infrastructure that enables end-user, with

---

This work is supported by the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreements: number 257178 (project CHOReOS - Large Scale Choreographies for the Future Internet - [www.choreos.eu](http://www.choreos.eu)), and number 231167 (project CONNECT - Emergent Connectors for Eternal Software Intensive Networked Systems - <http://connect-forever.eu/>).

---

Marco Autili, Paola Inverardi,  
Patrizio Pelliccione, Massimo Tivoli  
Università degli Studi di L'Aquila, Italy  
E-mail: [marco.autili@di.univaq.it](mailto:marco.autili@di.univaq.it),  
[paola.inverardi@di.univaq.it](mailto:paola.inverardi@di.univaq.it),  
[patrizio.pelliccione@di.univaq.it](mailto:patrizio.pelliccione@di.univaq.it),  
[massimo.tivoli@di.univaq.it](mailto:massimo.tivoli@di.univaq.it)

different needs and expectations, to provide and access software services on a variety of devices. This vision leads to the development of software systems that are more and more complex and large scale, namely Highly Complex Distributed Systems (HCDSs). To ensure that users experience the best Quality of Service (QoS) according to their needs and specific contexts of use, HCDSs need to be context-aware, adaptable, and dependable. Heterogeneity of the underlying communication infrastructure, mobility inducing changes to the availability of resources and continuously evolving requirements also call for adaptiveness and context-awareness.

The development and the execution of HCDSs is a big challenge and it is far to be solved. Indeed, HCDSs cannot rely on the classical desktop-centric assumption that the system execution environment and the networking environment are known a priori at design-time. Rather, HCDSs will need to cope with variability, as software gets deployed on an increasingly large diversity of computing platforms and operates in different networking environments that cannot be statically anticipated. Therefore, context-awareness and adaptiveness become two key aspects to consider while developing and running HCDSs. While providing/consuming services, HCDSs need to be aware of and adaptive to their *context*, that is, the combination of user-centric data (e.g., information of interest for users according to their current circumstance, like requested QoS) and resource/computer-centric data (e.g., status of devices and network, like availability of resources).

Engineering the development and supporting the execution of HCDSs raises numerous challenges that involve languages, methods and tools, spanning from design- to run-time, from validation to evolution. Although different software engineering techniques have

been studied to deal with distributed, adaptive, and context-aware systems, these initiatives are mainly confined to specific research areas: software architectures [4, 7, 19], middleware [29, 32], component-based development [30], service-oriented systems [11], etc. There is consensus, however, that the results obtained so far are inadequate to deal with the challenges that HCDSs will face in the future. HCDSs require a coordinated strategy that covers the whole life cycle within one large ambitious action. As firstly stated in [23] and then mentioned in [24, 5, 3], the software development process life cycle needs to be rethought by breaking the traditional division among development phases. This is achieved by moving some development activities from design-time to deployment- and run-time, hence asking for new and more efficient techniques to support run-time activities. Hence, for HCDSs, the way software will be produced and used radically changes. From an engineering perspective, enabling user-centrism, context-awareness, and adaptiveness requires the adoption of a dynamic development process that never stabilizes, rather it is permanently under maintenance.

Software engineering best practices suggest the exploitation of a middleware that, through the provision of proper features, supports distributed applications by masking the distribution and heterogeneity of the execution and networking environment [25]. In this respect, our thesis is that, for HCDSs, the frontier between application and middleware cannot be a priori fixed once the application domain has been identified, rather it varies as the application concerns vary. This moves forward the view of proposed middleware-based development processes, where an application-domain-specific middleware is selected and employed once and for all, as proposed in [25]. The envisioned dynamic development process includes a perpetual phase in which the run-time-support features, offered by the inherently dynamic middleware-layer, are selected and customized, assembled and evolved according to the characteristics of: the context, the execution environment, and the user needs.

In this paper we propose a perpetual development process model for HCDSs (Section 2). We then consider different application domains together with their supporting middleware to show how for HCDSs the frontier between application- and middleware-layer is unlikely to be tangible (Section 3).

## 2 A Perpetual Engineering Process Model for Highly Complex Distributed Systems

Our vision is that the engineering of HCDSs needs a dynamic software process where development-time ac-

tivities mingle with run-time support, hence invading the middleware territory. This means that the dynamic software process we envisage asks for an explicit characterization of the properties of the software artifacts and of their assumptions, by taking into account perpetual evolution of both the application- and middleware-layer.

We recall that HCDSs need to be always in a up-to-date state, that is, they must be able to dynamically evolve in order to interact with a continuously changing environment that might affect the system behavior. The changes might be of different nature, e.g.: changes of network topology (e.g., due to mobility), where new nodes appear and existing nodes vanish; changes of the set of available components, that is new components may become available, existing components may disappear, and new connections may be dynamically established; changes in the system requirements that can happen while the system is providing services to its users.

In general, the process view focuses on the set of activities, which characterize the production and the operation of a software system, and the set of artifacts elaborated/produced by such activities. These are traditionally divided into activities related to the actual production of the software system (i.e., before deployment) and activities that are performed when the system can be executed and goes into operation (i.e., after deployment). Specification, Design, Validation, and Evolution activities vary depending on the organization and the type of system being developed. Each activity works on suitable abstractions (models) of the system and requires its language, methods and tools. Therefore, in the process life cycle, systems are represented (through models) at very different levels of abstraction, from requirements specification to code.

So far, software complexity has been addressed by exacerbating the dichotomy development-/static-/compile-time versus execution-/dynamic-/interpret-time concentrating the management of and reasoning on models, as much as possible, at development-time. The mobile, ubiquitous and heterogeneous scenario in which HCDSs operate poses new requirements on this standard process. The evolutionary nature (adaptiveness, context-awareness, user-centrism) of HCDSs makes unfeasible a static-oriented development process since it would require, before the system is in execution, to predict the system behavior with respect to, virtually, any possible change, from application- down to middleware-layer. In the scenario in which HCDSs operate, possibly evolving user requirements and needs have to be perpetually guaranteed through adaptation. Whatever the change nature is, when it occurs, the cur-

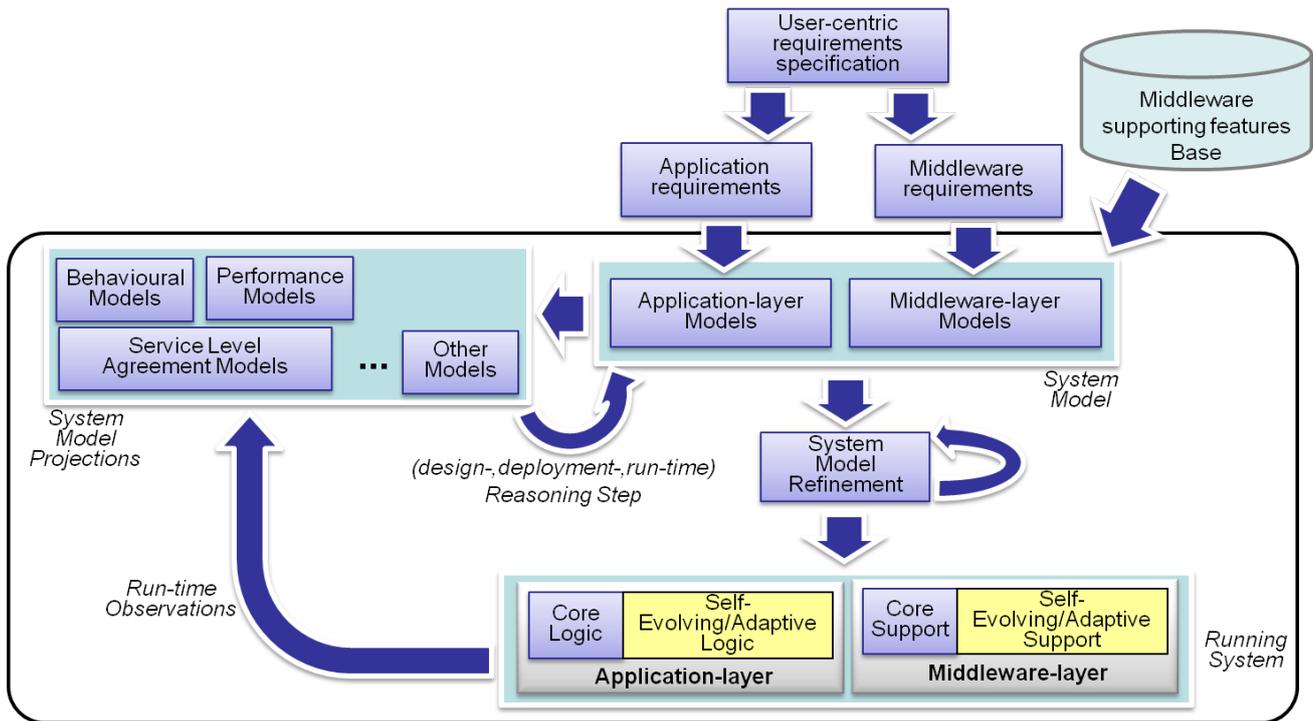


Fig. 1 The perpetual engineering process model

rent state of the system, including its current context and execution environment, is observed in order to suitably react to the change. This means that the needed models must be available at run-time so that reasoning steps on them become part of the adaptation activity carried on while the system is executing, i.e., at run-time. Such a *perpetual engineering process model* therefore has to explicitly account for complex reasoning steps at run-time when all the necessary pieces of information of the system are available.

The envisioned process model sketched in Figure 1 intends to support the systematic development of HCDSs from their design to their execution. The process model is based on model transformation techniques that serve the role of (i) refining user-centric requirement specifications to produce system code by accounting for middleware concerns, (ii) producing reasoning-technique-specific models (i.e., system model projections) from system models, embedding reasoning feedbacks into system models, and keeping them synchronized, (iii) propagating the interpretation of run-time observations of the running system to reasoning-technique-specific models, which in turn will enact possible changes to system models.

Referring to the point (i) above, as shown in Figure 1, system models are obtained by refining user-centric requirement specifications into application and middleware requirements, and then into system mod-

els. System models are the starting point for a refinement chain that leads to the code. This approach promotes system’s code synthesis that starts from the (a) application-layer models of the system and from (b) the middleware-layer models. Input (a) accounts for functional and non-functional concerns. Input (b) comes from the exploitation of a *middleware-layer supporting features* base. Indeed, our process model envisions a scenario in which middleware-layer supporting features are publicly available (e.g., into a registry) and can be discovered and selected to satisfy the specified middleware-layer requirements. Then, model-to-code transformations are used to build both the core code and the “generic code” of the system, from the application- down to the middleware-layer. The core code is the frozen unchanging portion of the system’s code. The generic code is an adaptive code that embodies a certain degree of variability (i.e., different ways of implementing a sub-system) making it capable to evolve. This code portion is evolving in the sense that, based on contextual information and possible changes of the user needs, the variability can be solved by considering alternative code.

As discussed in point (ii), model-driven techniques support the integration of different reasoning techniques, by propagating reasoning results into the various models, and by keeping them synchronized. Indeed, changes occurring in a model may have a strong impact

on other interoperating models. Therefore, through changes propagation, models are kept in a consistent state. This issue becomes a complex task when dealing with multiple model notations. Such a task is inevitable and requires to be managed by a dedicated approach. Recent development of model-driven techniques to bridge among different notations [17] makes this vision feasible.

Finally, referring to the point (*iii*), the interpretation of run-time observations needs to be propagated back to the system models level. These observations are used to feed the reasoning process that can induce changes to the system model projections, as well as to the system models [16, 8].

### 3 Breaking the Frontier between Application- and Middleware-layer

Middleware abstract the networking and computing environment by providing well known reusable solutions to recurrent problems like heterogeneity, interoperability, security, dependability. In this respect, since all these dimensions match distributed application requirements, middleware facilitates the development of distributed applications.

Depending on the application domain of interest, different supporting middleware have been developed. As far as evolution of HCDSs is concerned, in the following, we consider some application domains together with their computing paradigms. We then discuss how from software engineering perspective the frontier between application- and middleware-layer is unlikely to be distinguishable, once and for all, prior deployment.

#### 3.1 Context-aware Adaptable Systems

During the last decade, context-awareness and adaptation have been receiving significant attention in many research areas [2, 21, 22, 12]. The need for adapting software applications becomes obvious for new development paradigms, such as mobile and pervasive computing, when dealing with context variations. The main drivers of the pervasive computing paradigm are then context-aware and adaptable software applications for resource-constrained mobile devices. These are characterized by their heterogeneity (e.g., in terms of hardware platform, operating system, and programming language) and limitedness (e.g., limited processor speed, limited battery lifetime, and slow unreliable and intermittent connection). Context awareness refers to the capability of observing pieces of information of the

surrounding environment of the system that may influence the system behavior, but are out of the system control (e.g., networks, user needs, resources, device status). Adaptability refers to systems whose behavior can be adjusted in response to the context variations in order to keep requirements satisfied.

The development and execution of context-aware adaptable applications is a big challenge for the research community and it is far to be solved. The main difficulty is to provide (*i*) an easy-to-use and powerful programming technique for developers to actually program adaptable applications, and (*ii*) a context-aware run-time support to properly handle contextual situations. In the literature many valuable approaches have been proposed to this purpose, e.g., [18, 37, 15, 21, 28, 27, 20, 34, 35, 36]. In particular, the work in [21] points out that current mainstream programming languages and runtime environments provide little help for specifying adaptation and for supporting context awareness. This leads to a system design more complex and convoluted than needed, thus motivating the authors to discuss the need for a context-oriented programming approach that supports context-dependent variations.

In general, three different development approaches towards adaptable applications can be distinguished: (*i*) self-contained applications that embed the adaptation logic as a part of the application itself; therefore, they are a-priori instructed on how to handle dynamic changes in the environment; (*ii*) tailored applications that are the result of an adaptation process which has been applied on a generic version of the application, at deployment time at latest; (*iii*) middleware-based adaptable applications in which the middleware embeds the adaptation logic in the form of meta-level information on how the applications can be adapted.

Self-contained adaptable applications are inherently dynamic but suffer the pay-off of the inevitable overhead imposed by the adaptation logic. On the contrary, tailored adapted applications are suitable also for limited devices, but are dynamic only with respect to the environment at deployment time, while remaining static with respect to the actual execution, i.e., they cannot natively self-adapt to runtime changes in the execution environment. Middleware-based adaptable applications emerge as the result of a line of research, started in the early 00', looking for general purposes solutions to externalize adaptation mechanisms from the application logic implementation [13]. To enable such an externalization, many middleware-based solutions have been proposed, e.g., [6, 10]. While externalizing adaptation mechanisms, the middleware not only virtualizes communications but also offers proper supporting features to access and manage the context.

However, when tightly coupled to mobile devices and ubiquitous computing, context-aware adaptable applications become distributed in nature and intrinsically complex to be engineered and implemented. In fact, engineering middleware-based context-aware adaptable applications for mobile devices, is not easy due the limitedness and heterogeneity of the targeted execution environment. Indeed, it is unfeasible a “one-for-all” and “all-in-one” middleware solution that embeds the adaptation logic together with all the needed meta-level information, while fitting the resource constraints imposed by the different mobile devices.

A right trade-off can be obtained by a hybrid approach that brings together the advantages of middleware-based solutions, tailored and self-contained solutions. The approach would allow for externalizing adaptation by customizing/selecting not only the application logic but also the middleware features in support of adaptation mechanisms. Self-containment would bring the possibility of embedding those portions of the adaptation logic that are application specific and, as such, cannot be offered by the middleware as a general reusable solutions to frequently encountered problems. Clearly, such a mixed approach requires a future engineering process where the core logic of the application and the core support of the middleware are both subject to evolution. That is, after a certain tailoring step (that takes into account the resources status of the devices, its context, and user needs in a given instant), the core logic and the core support may evolve in response to context changes, e.g., some resources degrade or disappear.

Next section discusses the domain of service-oriented systems and further elaborates on context-aware adaptable systems within this domain.

### 3.2 Service-oriented Systems

The notion of Service-Oriented Architecture (SOA) has been receiving significant attention within the software engineering community. SOA has been largely accepted as a well founded reference architectural style for a wide class of distributed software systems. SOA relies on the Service-Oriented Computing (SOC) paradigm that considers services as building blocks for developing distributed applications. Networked devices and their hosted applications are then abstracted as autonomous loosely coupled services that are amenable to be integrated into a network of services to create flexible and dynamic processes, thanks to the service-oriented interaction pattern (depicted in Figure 2). With network connectivity being embedded in most computing devices, any networked device may seamlessly consume

but also provide software applications over the network. SOC then introduces natural design abstractions to deal with ubiquitous networking environments [9].

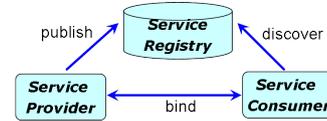


Fig. 2 Service oriented interaction pattern

Service-Oriented Middleware (SOM) supports the service-oriented interaction pattern through the provision of proper features for deploying, publishing, discovering and binding services at run time, as well as, features for supporting extra-functional crosscutting concerns such as service semantics and QoS. Services play a central role in this vision as effective means to achieve interoperability between heterogeneous parties of a business process and independence from the underlying infrastructure. At the same time they offer an open platform to build new value added service-based systems as a collaboration of available services. To support this view, the notion of *choreography* is recently receiving growing interest as a key concept in forming HCDSs. Available services are discovered in the network and choreographed to fit users’ needs.

By following the SOA loose-coupling principle, services come from different sources, and are developed in different ways and, in general, without any coordination among developers. In this scenario, for a given system, uniformity of the adopted SOC technologies and middleware cannot be assumed in general. Therefore, an integration solution at the middleware level is strongly required. Targeting the enterprise domain, Enterprise Application Integration (EAI) and Enterprise Service Bus (ESB) emerged as integration paradigms. In particular, an ESB can play the role of an intermediary among heterogeneous middlewares and offered interaction paradigms through the integration of a set of reusable bridges. Another approach [25] to protocol interoperability relies on reflection, by dynamically plugging-in the most appropriate communication protocols according to the protocols sensed in the environment.

However, when engineering choreography-based service-oriented systems, where a large number of heterogeneous services are called to collaborate, it is unfeasible to assume a full knowledge of all the heterogeneity facets, e.g., all the interaction protocols run by these services. In fact, a choreography specification predicates on a set of abstract roles that, in the general case, are fulfilled only at run-time. In other words,

the binding phase can take place only after concrete services have been discovered as suitable participants able to fulfill the required roles. Moreover, considering the Future Internet<sup>1</sup> as the evolution of the current Internet, the service world is evolving at a very fast pace. For instance, up to now, a large amount of RESTful and/or WS\* services can be discovered all over the Web (see ServiceFinder<sup>2</sup>, WebServiceList<sup>3</sup>, RemoteMethods<sup>4</sup>, WSCE<sup>5</sup>, ProgrammableWeb<sup>6</sup>).

In this scenario it is likely to have several services providing similar functionalities through different interfaces and/or interaction protocols. This implies that unforeseen communication protocols will emerge, newly coming services will appear, and running choreographies are required to evolve. This prevents the selection of a specific middleware that, in the best case, can support all the bridges that solve interoperability among all the kind of services that are known at selection time. This means that the frontier between a choreography-based service-oriented application and its SOM cannot be a priori established once and for all, rather it varies as the choreography evolves.

Considering mobile devices with limited capabilities that access and provide services, makes matters worse. Indeed, mobile computing environments are much more heterogeneous than conventional ones (see Section 3.1). This further restricts the possibility of selecting a middleware that provides seamless integration and interoperability through a general purpose service-oriented interaction patterns (Figure 2). The development of such a general purpose middleware is per se challenging since it should have built-in support for as many mobile applications, old and new, as possible. On the other hand, mobile SOA applications need to rely on a run-time support for service access and provision, and will do so for long time in the future.

Therefore, a future engineering process is required to enable the development of adaptable and lightweight SOA-based application together with SOM solutions capable to simultaneously evolve. When saying “capable to simultaneously evolve” we mean that both the application logic and the middleware support should be “upgraded or downgraded” to obtain the right balancing that fits the restrictions imposed by different mobile devices and permits to keep the evolution pace of the service world.

### 3.3 Cloud Systems

The general promise of Cloud computing [38], as a more recent paradigm with respect to Grid computing, is to offer distributed computing infrastructure together with a set of technologies for elaborating and storing data. CPU power and ready-to-use software are provisioned by sharing and coordinating resources within a virtual organization. The technology behind cloud computing promises cheaper and flexible IT components, that may be rented instead of acquired upfront, and quickly scaled up and down according to the end-user’s needs.

Mature virtualization technology is currently allowing the generalization of modern, more flexible service offerings ranging from virtualized hardware (Infrastructure as a Service - IaaS) such as Amazon EC2<sup>7</sup>, to application platforms (Platform as a Service - PaaS) such as Salesforce and Google Apps Engine, to complete applications such as Google Apps (Software as a Service - SaaS) [31]. More specifically, IaaS refers to very basic computing capability machines with operating systems and storage. PaaS is between IaaS and SaaS and refers to an environment where to build and run an application platform in the cloud using whatever pre-built components, and interfaces are provided by that particular PaaS platform. SaaS refers to a software delivery model in which ready-to-use applications and related data are hosted in the cloud and made available over a network.

As it is clear from the discussion above, cloud systems constitute the most representative example of HCDSs. Promoting the philosophy of “Everything as a Service”, cloud computing offers the modular and distributed virtualization of a complete system at different levels, with the (promised) high added value of paying only for what is needed when it is needed, by making the used resources vary in accordance with user needs. Whether a user decides to make use of SaaS, PaaS, or IaaS or any combination of them for a particular solution depends on a number of factors, including the availability of a SaaS application that directly suits the needs, the level of expertise in developing custom applications over PaaS, and the amount of required flexibility in many resources notably IaaS ones.

Since many years, SOC middleware are considered to be good candidates to support Grid and Cloud computing [26]. On the one hand, many Grid-oriented middleware have been proposed, e.g., OurGrid [14], InteGrade [33], to enable the execution of computationally-intensive applications on sets of highly distributed clusters of machines. On the other hand, Cloud comput-

<sup>1</sup> <http://services.future-internet.eu/>

<sup>2</sup> <http://www.service-finder.eu/>

<sup>3</sup> <http://www.webservicelist.com/>

<sup>4</sup> <http://www.remotemethods.com/>

<sup>5</sup> <http://www.uoguelph.ca/~qmahmoud/WSCE/index.html>

<sup>6</sup> <http://www.programmableweb.com>

<sup>7</sup> <http://aws.amazon.com/ec2>

ing as a more recent paradigm providing virtualization mechanisms specifically calls for a Cloud-oriented middleware for supporting more elastic and on-demand provision of remote networked resources at different levels, i.e., at SaaS, PaaS, or IaaS level. Even though at these three levels we can have a uniform representation of services, each level requires specialized means for managing services at run time, as well as, specialized features for supporting extra-functional crosscutting concerns.

In the following we mention some middleware that support cloud computing at the IaaS, PaaS, and SaaS levels.

**Cloud middleware at the IaaS level:** OpenNebula<sup>8</sup> (from the FP7 Reservoir project<sup>9</sup>) is an open-source platform for the management of virtualized data centers to build any type of IaaS cloud: private, public, virtual private and hybrid. OpenNebula aims at providing an open, flexible, extensible, and comprehensive management layer to automate and orchestrate the operation of virtualized data centers by leveraging and integrating existing solutions. Eucalyptus<sup>10</sup> (which is used by Ubuntu for their cloud offering) enables the creation of IaaS private clouds, with no requirements for re-tooling the organization's existing IT infrastructure or need to introduce specialized hardware. Eucalyptus promotes a modern infrastructure virtualization software to create elastic pools that can be dynamically scaled up or down depending on applications workloads. Cloud.com<sup>11</sup> is an open source cloud computing platform for building and managing private and public cloud. It enables simple and cost effective deployment, management, and configuration of cloud computing environments, regardless of where they are deployed.

**Cloud middleware at the PaaS level:** WSO2 Stratos<sup>12</sup> is an enterprise-grade and open PaaS, which provides the core cloud services and essential building blocks required for developing SaaS and cloud applications. Oracle Fusion Middleware [31] identifies in the platform level of PaaS the right balance between flexibility and ease of use for the cloud customers. It proposes a foundation for realizing private cloud by delivering effective dynamic resourcing. Dynamic resourcing is complemented by user interaction technologies together with modularity, sharability, and composability, thus providing a powerful self-service platform of reusable components.

**Cloud middleware at the SaaS level:** In [1] it is presented an architecture for achieving multi-tenancy

at the SaaS level. Multi-tenancy allows a single application to emulate multiple application instances. This enables users to run services in a multi-tenant SOA framework and to build multi-tenant applications. This middleware is implemented on top of the WSO2 Carbon platform, already mentioned above when presenting cloud middleware at the PaaS level. CloudPointe<sup>13</sup> is a mobile SaaS middleware to enable mobile access to digital assets and shared documents on the cloud. CloudPointe solves data duplication, privacy, integrity, and security issues.

In line with the cloud computing philosophy, which distinguishes the three computing levels mentioned above, all the efforts towards the realization of cloud middleware adopt a separation-of-concerns approach. To this end each middleware provides run-time support to the specific level it was born for. The perpetual engineering process model in Figure 1 can be then instantiated at any level spanning from SaaS to IaaS. This calls for a new dimension of flexibility to be taken into account when engineering cloud-based HCDSs; this flexibility somehow subverts the three-levels vision of cloud computing, where middleware at higher levels mask details of lower levels. Thus, as for context-aware adaptable systems and SOA systems, this means that the frontier between the application and middleware cannot be defined once and for all. Even though the physical division imposed by the layered architecture of cloud computing is still valid, the future development process should provide a logical vision that crosscuts the three levels of middleware. This allows us to reach the best tradeoff by suitably selecting the most appropriated supporting features of the different levels in the cloud.

## 4 Conclusion and Final Remarks

In this paper we try to provide a partial answer to some questions concerning the development of highly complex distributed systems (HCDSs). From a SE perspective we identify what are the needs that may drive the development of such systems and what is their life cycle, hence discussing new challenges for SE?

We consider different kind of HCDSs within three application domains, namely, context-aware adaptable systems, service-oriented systems, and cloud systems. Our thesis is that these systems require a futuristic dynamic development process that breaks the traditional boundaries between static and dynamic activities as well as, the frontier between the application and middleware layer.

<sup>8</sup> <http://opennebula.org/>

<sup>9</sup> <http://www.reservoir-fp7.eu/>

<sup>10</sup> <http://www.eucalyptus.com/>

<sup>11</sup> <http://www.cloud.com/>

<sup>12</sup> <http://wso2.com/cloud/stratos/>

<sup>13</sup> <http://www.cloudxy.com>

This means that, for HCDSs a standard development process, where an application-domain-specific middleware is selected once and for all and employed [25], is unfeasible. The envisioned dynamic development process includes a perpetual phase in which the run-time-support features offered by the middleware-layer are selected and customized, assembled and evolved together with the application logic, the characteristics of the execution context, as well as, user needs. In other words, both the application logic and the middleware support should be capable to simultaneously evolve, meaning that they should be “upgraded or downgraded” to obtain the right balancing that fits the application requirements, the computing and the networking environment, while satisfying the user expectations.

## References

1. Azeez, A., Perera, S., Gamage, D., Linton, R., Siriwardana, P., Leelaratne, D., Weerawarana, S., Fremantle, P.: Multi-tenant soa middleware for cloud computing. In: Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on, pp. 458–465 (2010). DOI 10.1109/CLOUD.2010.50
2. Baldauf, M., Dustdar, S., Rosenberg, F.: A survey on context-aware systems. *Int. J. Ad Hoc Ubiquitous Comput.* **2**(4), 263–277 (2007)
3. Baresi, L., Ghezzi, C.: The disappearing boundary between development-time and run-time. In: Proceedings of the FSE/SDP workshop on Future of software engineering research, FoSER '10, pp. 17–22. ACM (2010)
4. Bencomo, N., Blair, G.S.: Using architecture models to support the generation and operation of component-based adaptive systems. In: Software Engineering for Self-Adaptive Systems, pp. 183–200 (2009)
5. Blair, G., Bencomo, N., France, R.B.: Models run.time. *Computer* **42**, 22–27 (2009). DOI <http://doi.ieeecomputersociety.org/10.1109/MC.2009.326>
6. Blair, G.S., Coulson, G., Andersen, A., Blair, L., Clarke, M., Costa, F., Duran-Limon, H., Fitzpatrick, T., Johnston, L., Moreira, R., Parlavantzas, N., Saikoski, K.: The design and implementation of open orb 2. *IEEE DISTRIBUTED SYSTEMS ONLINE* **2**, 2001 (2001)
7. Bucchiarone, A., Pelliccione, P., Vattani, C., Runge, O.: Self-Repairing systems modeling and verification using AGG. In: WICSA/ECSA'09, pp. 181–190 (2009)
8. Caporuscio, M., Marco, A.D., Inverardi, P.: Model-based system reconfiguration for dynamic performance management. *Journal of Systems and Software* **80**(4) (2007)
9. Caporuscio, M., Raverdy, P.G., Issarny, V.: UbiSOAP: A Service Oriented Middleware for Ubiquitous Networking. *IEEE Transactions on Services Computing* **99** (2011)
10. Capra, L., Emmerich, W., Mascolo, C.: Carisma: Context-aware reflective middleware system for mobile applications. *IEEE Transactions on Software Engineering* **29**, 929–945 (2003)
11. Cardellini, V., Casalicchio, E., Grassi, V., Lo Presti, F., Mirandola, R.: Qos-driven runtime adaptation of service oriented architectures. In: ESEC/FSE '09, pp. 131–140. ACM (2009)
12. Cheng, B.H.C., de Lemos, R., Giese, H., Inverardi, P., Magee, J. (eds.): Software Engineering for Self-Adaptive Systems [outcome of a Dagstuhl Seminar], *Lecture Notes in Computer Science*, vol. 5525. Springer (2009)
13. wen Cheng, S., cheng Huang, A., Garlan, D., Schmerl, B., Steenkiste, P.: Rainbow: Architecture-based self-adaptation with reusable infrastructure. *IEEE Computer* **37**, 46–54 (2004)
14. Cirne, W., Brasileiro, F., Andrade, N., Costa, L., Andrade, A., Novaes, R., Mowbray, M.: Labs of the world, unite!!! *Journal of Grid Computing* **4** (2006)
15. Cowan, C., Black, A., Krasic, C., Pu, C., Walpole, J., Consel, C., Volanschi, E.: Specialization classes: An object framework for specialization. In: Fifth IEEE International Workshop on Object-Orientation in Operating Systems. Seattle, Washington (1996)
16. Epifani, I., Ghezzi, C., Mirandola, R., Tamburrelli, G.: Model evolution by run-time parameter adaptation. In: Proceedings of the 31st ICSE, pp. 111–121 (2009)
17. Eramo, R., Malavolta, I., Muccini, H., Pelliccione, P., Pierantonio, A.: A model-driven approach to automate the propagation of changes among architecture description languages. *Software and Systems Modeling*
18. Gamma, E., Helm, R., Johnson, R., Vlissides, J.M.: Design Patterns: Elements of Reusable Object-Oriented Software. A.-Wesley Professional Computing Series. A.-Wesley (1995). GAM e 95:1 1.Ex
19. Garlan, D., Cheng, S.W., Huang, A.C., Schmerl, B., Steenkiste, P.: Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer* **37**(10), 46–54 (2004)
20. Henriksen, K., Indulska, J.: Developing context-aware pervasive computing applications: Models and approach. *Pervasive and Mobile Computing*, In **2**, 2005 (2005)
21. Hirschfeld, R., Costanza, P., Nierstrasz, O.: Context-oriented programming. *Journal of Object Technology* **7**(3), 125–151 (2008)
22. yi Hong, J., ho Suh, E., Kim, S.J.: Context-aware systems: A literature review and classification. *Expert Systems with Applications* **36**(4), 8509–8522 (2009)
23. Inverardi, P.: Software of the future is the future of software? In: U. Montanari, D. Sannella, R. Bruni (eds.) *Trustworthy Global Computing, Lecture Notes in Computer Science*, vol. 4661, pp. 69–85. Springer Berlin / Heidelberg (2007)
24. Inverardi, P., Tivoli, M.: Software engineering. chap. The Future of Software: Adaptation and Dependability, pp. 1–31. Springer-Verlag, Berlin, Heidelberg (2009)
25. Issarny, V., Caporuscio, M., Georgantas, N.: A perspective on the future of middleware-based software engineering. In: L. Briand, A. Wolf (eds.) *Future of Software Engineering*. IEEE-CS Press (2007)
26. Issarny, V., Georgantas, N., Hachem, S., Zarras, A., Vassiliadis, P., Autili, M., Gerosa, M., Hamida, A.: Service-oriented middleware for the future internet: state of the art and research directions. *JISA* **2**, 23–45 (2011)
27. Keays, R., Rakotonirainy, A.: Context-oriented programming. In: *MobiDe '03: Proceedings of the 3rd ACM international workshop on Data engineering for wireless and mobile access*, pp. 9–16. ACM (2003)
28. Kochan, S.: *Programming in Objective-C*. Sams (2003)
29. Liu, H., Parashar, M., Member, S.: Accord: A programming framework for autonomic applications. *IEEE Transactions on Systems, Man and Cybernetics* **36**, 341–352 (2005)
30. Peper, C., Schneider, D.: Component engineering for adaptive ad-hoc systems. In: *SEAMS '08*, pp. 49–56. ACM (2008)

31. Piec, M.: Platform-as-a-service private cloud with oracle fusion middleware. <http://www.oracle.com/us/technologies/cloud/036500.pdf> - last acces Oct. 28th, 2011 (2009)
32. Rouvoy, R., Barone, P., Ding, Y., Eliassen, F., Hallsteinen, S.O., Lorenzo, J., Mamelli, A., Scholz, U.: Music: Middleware support for self-adaptation in ubiquitous and service-oriented environments. In: *Software Engineering for Self-Adaptive Systems*, pp. 164–182 (2009)
33. da Silva e Silva, F.J., Kon, F., Goldman, A., Finger, M., de Camargo, R.Y., Filho, F.C., Costa, F.M.: Application execution management on the InteGrade opportunistic grid middleware. *Journal of Parallel and Distributed Computing* **70** (2010)
34. Tanter, É., Gybels, K., Denker, M., Bergel, A.: Context-aware aspects. In: *Software Composition, LNCS*, vol. 4089, pp. 227–242 (2006)
35. Villazón, A., Binder, W., Ansaloni, D., Moret, P.: Advanced runtime adaptation for java. In: *Proceedings of GPCE'09*, pp. 85–94. ACM (2009)
36. Villazón, A., Binder, W., Ansaloni, D., Moret, P.: Hot-wave: creating adaptive tools with dynamic aspect-oriented programming in java. In: *Proceedings of GPCE'09*, pp. 95–98. ACM (2009)
37. Volanschi, E.N., Consel, C., Muller, G., Cowan, C.: Declarative specialization of object-oriented programs. In: *OOPSLA*, pp. 286–300 (1997)
38. Zhang, Q., Cheng, L., Boutaba, R.: Cloud computing: state of the art and research challenges. *Journal of Internet Services and Applications* **1**(1) (2010)