Cloud Computing and Software Agents: Towards Cloud Intelligent Services

Domenico Talia
ICAR-CNR & University of Calabria
Rende, Italy
talia@deis.unical.it

Abstract — **Cloud computing systems provide large-scale infrastructures for high-performance computing that are "elastic" since they are able to adapt to user and application needs. Clouds are used through a service-oriented interface that implements the *-as-a-service paradigm to offer Cloud services on demand. This paper discusses Cloud computing models and architectures, their use in parallel and distributed applications, and examines analogies, differences and potential synergies between Cloud computing and multi-agent systems. This analysis is lead having in mind the goal of implementing high-performance complex systems and intelligent applications by using of Cloud systems and software agents. The convergence of interests between multi-agent systems that need reliable distributed infrastructures and Cloud computing systems that need intelligent software with dynamic, flexible, and autonomous behavior can result in new systems and applications.**

Keywords - Cloud computing; multi-agent systems; service oriented computing;

I. INTRODUCTION

Cloud computing provide elastic services, high performance and scalable data storage to a large and everyday increasing number of users [1]. Cloud computing enlarged the arena of distributed computing systems by providing advanced Internet services that complement and complete functionalities of distributed computing provided by the Web, Grid computing and peer-to-peer networks. In fact, Cloud computing systems provide large-scale infrastructures for high-performance computing that are dynamically adapt to user and application needs.

Today Clouds are mainly used for handling highly intensive computing workloads and for providing very large data storage facilities. Both these goals are combined with the third goal of potentially reducing management and use costs. At the same time, multi-agent systems (MAS) represent another distributed computing paradigm based on multiple interacting agents that are capable of intelligent behavior. Multi-agent systems are often used to solve problems by using a decentralized approach where several agents contribute to the solution by cooperating one each other. One key feature of software agents is the intelligence that can be embodied into them according to some collective artificial intelligence approach that needs cooperation among several agents that can run on a parallel or distributed computer to achieve the needed high performance for solving large complex problems keeping execution time low.

Although several differences exist between Cloud computing and multi-agent systems, they are two distributed computing models, therefore several common problems can be identified and several benefits can be obtained by the integrated use of Cloud computing systems and multi-agents. The research activities in the area of Cloud computing are mainly focused on the efficient use of the computing infrastructure, service delivery, data storage, scalable virtualization techniques, and energy efficiency. In summary, we can say that in Cloud computing the main focus of research is on the efficient use of the infrastructure at reduced costs. On the contrary, research activities in the area of agents are more focused on the intelligent aspects of agents and on their use for developing complex applications. Here the main problems are related to issues such as complex system simulation, adaptive systems, software-intensive applications, distributed computational intelligence, and collective learning.

Despite these differences, Cloud computing and multi-agent systems share several common issues and research topics in both areas have several overlaps that need to be investigated. In particular, Cloud computing can offer a very powerful, reliable, predictable and scalable computing infrastructure for the execution of multi-agent systems implementing complex agent-based applications such when modeling and simulation of complex systems must be provided. On the other side, software agents can be used as basic components for implementing intelligence in Cloud computing systems making them more adaptive, flexible, and autonomic in resource management, service provisioning and in running large-scale applications.

For these reasons and for others that we discuss later, this paper investigates research work in the two areas and point out potential synergies that deserve to be analyzed. The paper discusses Cloud computing models and architectures, their use in parallel and distributed applications, and examines analogies, differences and potential synergies between Cloud computing and multi-agent systems. Analysis is led having in mind the goal of implementing high-performance complex systems and intelligent applications by using both Cloud computing systems and software agents. Section II introduces Cloud computing concepts and reviews some research activities. Section III discusses multi-agent systems and some research topics that are related to Cloud computing. Section IV presents some ideas on using intelligent software agents to improve the performance and functionality of Clouds, whereas Section V discusses how Cloud computing platforms can be used for the efficient execution of MAS. Section VI concludes the paper.

II. CLOUD COMPUTING

A. A definitition

Since the Cloud computing paradigm has been conceived several definitions have been given. Some of them focus on on-demand dynamic provisioning of processing and storage resources, others emphasize the service-oriented interface and the exploitation of virtualization techniques. The National Institute of Standards and Technology (NIST) have given a complete reference definition [2]. NIST defined Clouds as follows: "*Cloud computing is a pay-per-use model for enabling available, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*" Moreover, according to NIST: "*Cloud model promotes availability and is comprised of five key characteristics, three delivery models, and four deployment models*."

The key characteristics of Clouds are: On-demand self-service, ubiquitous network access, location independent resource pooling, rapid elasticity, and pay per use. Fig. 1 summarizes the main aspects of Cloud computing system both from the technical side and the business side [3].

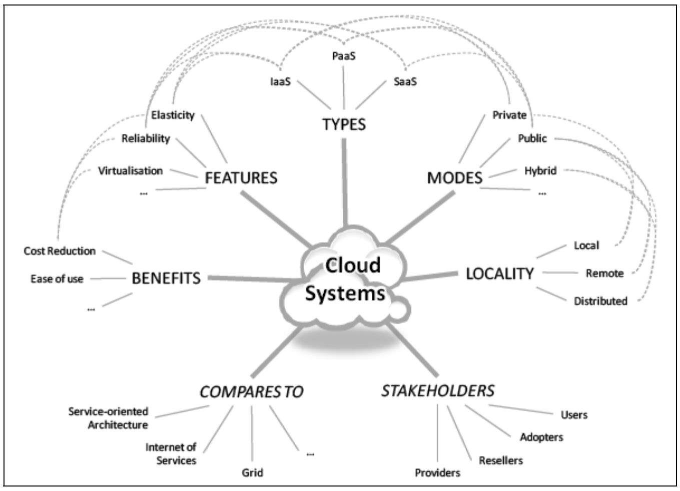

Figure 1. Main aspetcs of Clouds (from [3]).

The delivery models of Clouds are very important because they define three different types of Cloud computing systems:

- *Infrastructure as a Service* (*IaaS*). The capability provided to the user is to rent computing, storage, networks, and other computing resources where the user is able to deploy and run software, which can include operating systems and/or applications. The user does not manage or control the hardware Cloud infrastructure but has control over operating environments, storage, deployed applications, and possibly select networking components. Examples for commercial Cloud infrastructures are Amazon EC2 and Rackspace.

- *Platform as a Service* (*PaaS*). The functionality provided to the user is to deploy onto the Cloud infrastructure consumer-created applications using programming languages, compilers and toolkits supported by the provider (e.g., Java, .Net). The consumer does not manage or control the underlying cloud infrastructure, network, servers, operating systems, or storage, but the consumer can control the deployed applications and possibly the application hosting environment configurations.

- *Software as a Service* (*SaaS*). The capability provided to the consumer is to use the provider's applications running on a Cloud infrastructure and accessible from various client devices through a thin client interface such as a Web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure, network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings

B. Cloud Deployment models

About five years ago, when the first Cloud infrastructure has been deployed by Amazon, the online bookseller company that took the decision to start a new business selling computing resources to companies and private users, the only deployment model was the *Public Cloud* one. It is a pay-per-use IaaS Cloud infrastructure that is owned by an organization selling Cloud services to the general public or to enterprises. Thus, it is public because it can be rent by anyone for developing and/or running any kind of applications. To use Amazon services, users must provide a credit card account and can spend from few cents to thousands or millions of dollars depending on the number of used resources and the usage time.

After this early Cloud version, other deployment models different from *Public Clouds* have been designed and implemented (see Fig. 2):

- *Private Cloud.* The Cloud infrastructure is owned or leased by a single organization and is operated only for that organization. No public access to it is permitted. This model can be used in case of strict data privacy and/or security requirements.

- *Community Cloud.* The Cloud infrastructure is shared by a limited number of organizations and supports a specific community that has shared concerns (e.g., goals, security requirements, policy, and compliance issues).

- *Hybrid Cloud.* This fourth class of Cloud infrastructure is a composition of two or more Clouds (private, community, or public) that although they are unique entities, are combined together by standardized or proprietary technology that enables data and application portability (e.g., Cloud federation).

- Private clouds
  - ✓ enterprise owned or leased
- Community clouds
  - ✓ shared infrastructure for specific community
- Public clouds
  - ✓ Sold to the public, mega-scale infrastructure
- Hybrid clouds
  - ✓ composition of two or more cloud

Figure 2. Deployment models for Clouds.

Cloud computing is the most recent results of the advancement of several computer technologies both from the hardware side, such as virtualization and multi-core architectures, and from the software side like cluster computing, Grid computing, Web services, service-oriented architectures, autonomic computing, and large-scale data storage. In particular, virtualization in Cloud computing is the key element that separates system functionality and implementation from physical resources. By exploiting virtualization techniques, a Cloud infrastructure can be partitioned in several parallel virtual machines, dynamically configured according to the user requirements and devoted to run independent applications concurrently. Virtualization separates applications from hardware and users from other users giving them the feeling that a large-scale computing infrastructure is devoted to their applications by meeting a given quality of service (QoS). Virtualization is also used to isolate applications avoiding that if one fails that other can fail too. Finally, virtualization is a way to improve security and privacy of concurrent applications running on the same Cloud.

C. Example of Cloud systems

As we can infer from the previous descriptions, the Cloud computing paradigm represents an advancement of the existing computing services available over the Internet. In particular, Cloud infrastructures adopted the Web services paradigm for delivering new capabilities beyond the Web capability.

Several companies set up large Cloud facilities and built programming environments where developers can program applications as Cloud software services. Just to mention some example, Amazon on his EC2 and S3 Cloud platforms implemented Elastic BeanStalk, Microsoft implemented .Net technology on Azure, Google provides the AppEngine, and VMware has Cloud Foundry.

On the other side, the research community developed open source software that can be deployed and configured on servers, computer farms or data centers for implementing private, public, community or hybrid Cloud infrastructures or for inter-Cloud computing facilities. Examples of these systems are OpenNebula, Eucaliptus, OpenQRM, Puppet, and OpenStack. These open source software projects are also working to systems and services that allow Cloud-to-Cloud interoperability and federation.

III. AGENT COMPUTING

An agent is a computational entity that acts on behalf of another entity (or entities) to perform a task or achieve a given goal. Agent systems are self-contained software programs embodying domain knowledge and having ability to behave with a specific degree of independence to carry out actions needed to achieve specified goals. They are designed to operate in a dynamically changing environment.

Agents typically include a set of features. The main features of agents include the following:

- *Autonomy*: the capacity to act autonomously to some degree on behalf of users or other programs also by modifying the way in which they achieve their objectives.
- *Pro-activity:* the capacity to pursue their own individual set goals, including by making decisions as result of internal decisions.
- *Re-activity:* the capacity to react to external events and stimuli and consequently adapt their behavior and make decisions to carry out their tasks.
- *Communication and Cooperation*: the capacity to interact and communicate with other agents (in multiple agent systems), to exchange information, receive instructions and give responses and cooperate to fulfill their own goals.
- *Negotiation*: the capability to carry out organized conversations to achieve a degree of cooperation with other agents.
- *Learning*: the ability to improve performance and decision making over time when interacting with the external environment.

Although a single agent can act and run to perform a given task, the agent paradigm was conceived as a distributed computing model where a set of agents interact one another by exchanging information and cooperating to perform complex tasks where interaction, intelligence, adaptation and dynamicity are key issues to be handled.

This means that even if we can define an agent in isolation, the agent paradigm can find its complete exploitation if we consider agents as entities acting in a collection of agents, therefore implementing the so called multi-agent system paradigm. In fact, it is rather difficult to imagine that an agent will exist and operate only as a stand-alone entity and will never interact with other agents (real or artificial) in its environment. Also information agents, or personal agents, which are mainly supposed to work as stand alone entities in solving problems, will certainly improve their behavior and achieved results if cooperate with other agents to receive information, to delegate task execution or to exchange knowledge that improve the agent role and contribution. According to these considerations, the social dimension of agents is one of its essential features.

As stated by Sycara [4], the characteristics of MASs are that

- each agent has incomplete information or capabilities for solving a problem and, thus, has a limited viewpoint on the global task to be done;

- there is no system global control;
- data are decentralized; and
- computation is asynchronous.

These features are typical of decentralized computing paradigms. In fact, being a distributed computing paradigm, multi-agent systems share several characteristics with other distributed paradigms like actors (which we can consider as their progenitors), concurrent objects, peer-to-peer networks, Grid computing, sensor networks, autonomic computing, and Cloud computing. At the same time, it is worth to notice that agents possess some properties, as discussed before, that differentiate them from other distributed computing models.

Commonalities and differences among these distributed computing models can be exploited for the integrated use of some of the technologies that are based on them. For example, decentralized applications based on multi-agent systems can be developed on Grid systems or on peer-to-peer networks. At the same time, applications based on sensor networks can use use distributed intelligence techniques by means of a multi-agent system with learning and pro-activity features.

In the past years several agent programming environments supporting specific agent architectures and providing libraries of interaction protocols like Jason, 3APL, JACK, Claim, SyMPA, JADE, Cougaar, Jadex, and ZEUS have been developed. Moreover, software engineering methodologies like Gaia, Tropos, and AUML have been designed to analyze and design agent-based systems. Efforts have been done to standardize some features or facilities of agent systems, such as has been done with FIPA and KQML for inter-agent communication. These environments, toolkits and methodologies are enabling technologies for implementing MAS applications on traditional computing systems. However they can be more interesting if they will be available of distributed computing infrastructures like Grid, Cloud or P2P networks for supporting the development of large-scale MAS applications achieving high performance and scalability.

However, despite the potential common space where agent technology and Cloud computing infrastructures can be effectively used to produce innovative models, techniques, systems and applications, till today only a few research activities that make use of both these technologies are performed. In the literature a very limited number of papers can be found on agents and Cloud integration [5, 6, 7, 8].

In the next two sections, we discuss two main approaches for the integrated use of agents and Cloud systems. The first one is based on the principle that agent flexibility, intelligence, pro-activity, and autonomy can be used in Cloud computing platforms to produce new advanced Cloud solutions and services that offer new functionalities and intelligent services that todays are not yet available in current Cloud computing infrastructures.

The second one is centered on the idea that Cloud infrastructures can offer an ideal platform where run MAS-based systems, simulations and applications because of its large amount of processing and storage resources that can be dynamically configured to run large-scale MAS-based software at unprecedented scale.

IV. Clouds Using Agents

Cloud computing is a novel technology that has been designed and implemented in the past five years, mainly due to industry that was looking to a large-scale scalable computing infrastructure for implementing and selling service-oriented commercial solutions. Whereas much of the current effort on Cloud computing was devoted to the production of Cloud infrastructures and technologies for supporting virtualization and data centers, little attention has been devoted to introduce innovative methods for users and developers to discover, request, assemble and use Cloud computing resources. Autonomous and flexible agents and MASs are suitable tools for negotiating user access, automating the resource and service discovery, and composition, trading, and harnessing of Cloud resources.

A new discipline, called *agent-based Cloud computing* must be set for providing agent-based solutions founded on the design and development of software agents for improving Cloud resources and service management and discovery, SLA negotiation, and service composition. Autonomous agents can make Clouds smarter in the interaction with users and more efficient in allocating processing and storage to applications.

In large-scale data centers, agents can search, filter, query and update the massive volumes of data that are stored. We can envision a scenario where *Cloud agents* working on our and operating systems behalf, to provide intelligent data access services, monitoring services, processor-to-application assignment strategies, and energy-efficient use of Cloud computing infrastructures.

Research activities must be carried out to implement effective agent-based solutions for Cloud computing. This work should be done towards the three different *-as-a-Service delivery classes. In IaaS infrastructures, agents can be used to help the intelligent provisioning of basic resources to user applications, whereas in Paas infrastructures, agent can play a role in the efficient deployment and execution of programming environments that developers use for application implementation. Finally, in SaaS Cloud infrastructures, agents can be programmed to optimize the use of applications provided as services and the management of the underlying hardware/software infrastructure taking care of its efficient utilization and, at the same time, for maintaining the declared QoS.

In Clouds, there also is the need to design and implement techniques and methodologies that adapt to the dynamic behaviors of Cloud computing environments. Autonomic techniques may help providers and users to reach this goal. Multi-agent systems that are able to handle with changing configurations, heterogeneity, and volatility, can provide a promising approach for addressing this requirement. Last but not least, security and trust are two very critical issues in Cloud computing as data and software are stored, accessed and run on machines that are not owned or directly managed by owners of data and software. Agent-based models and algorithms for trust and security in Cloud infrastructures could be very useful.

In summary, if agent-based solutions will be introduced in the software infrastructure of Clouds we will have:

- Intelligent and flexible Cloud services,
- Autonomous and pro-active services,
- Autonomic Clouds.

V. Agents Using Clouds

Complex agent-based applications or large-scale simulations based on MASs often require high-performance computing systems and large data storage devices. Therefore, Cloud infrastructures can offer an ideal platform where to run MAS-based systems simulations and applications because of its large amount of processing and memory resources that can be dynamically configured for executing large agent-based software at unprecedented scale.

Agent-based applications can rely on Cloud computing infrastructures to access and use vast amounts of processors and data. So this approach would allow to offload the compute-intensive agents to the appropriate subsets of processes and storage elements in a Cloud. The entire MAS application can run on a Cloud infrastructure or only the most compute-intensive part of it can be hosted in the Cloud, whereas the light part can run on a local server or simply on the client PC. In this way agents can become more efficient and, at the same time, lighter and smarter. This can be obtained because, by using powerful Cloud facilities, agents can improve their intelligence and accurateness by running more sophisticated algorithms. In fact, the amount of storage and processing power of a Cloud-enabled MAS is larger than in other computing environments, making it more powerful.

Cloud-enabled agents can coupling agents and large-scale dynamic distributed computing platforms bringing big new opportunities to the agent computing area and expanding agent's knowledge beyond the possibilities offered by traditional computing platforms.

Virtualization mechanisms offered by Cloud computing can be exploited for efficient composition of parallel machines where to execute large scale concurrent agents with real-time constraints or needing high performance for achieving results in reasonable time.

Agents implemented in Cloud systems can adapt to available virtual machines by using the basic properties of agents such as autonomy, pro-activity, negotiation and learning. Since Clouds are elastic, they can expand and shrink based on demand of users or applications. This property is very useful for the scalable execution of MAS applications and simulation that are able to adapt to the available resources.

In summary, agent can find in Cloud computing infrastructures the appropriate platform where to run and access large data. This opportunity must be exploited for implementing efficient MASs and, from a more general point of view, for advancing the way to design and implement a new generation of large-scale software agents.

VI. Conclusions

The marriage between Clouds and Agents can be convenient for both parties. We discussed in the paper how this can be done and which scientific areas and issues must be involved to carry out research work that will produce intelligent Cloud services and high-performance multi-agent systems on Clouds. The convergence of interests between multi-agent systems that need reliable distributed infrastructures and Cloud computing systems that need intelligent software with dynamic, flexible, and autonomous behavior will result in new systems and applications. Both research communities must be aware of this opportunity and should put in place the joint research activities needed to reach that goal.

References

[1] M. Armbrust, et al., "A view of cloud computing," Communications of the ACM, vol. 53, no. 4, pp. 50-58, April 2010.

[2] www.nist.gov/itl/cloud, June 2011.

[3] Cloud Computing Expert Group, "The Future of Cloud Computing," Report from European Commission, January 2010.

[4] K. P. Sycara, "Multiagent systems," AI Magazine, vol. 19, no. 2, pp.79-92, 1998.

[5] B. K.-M. Sim, "Towards complex negotiation for cloud economy," Proc. GPC 2010, Springer, LNCS 6104, pp. 395-406, 2010

[6] R. Aversa, B. Di Martino M. Rak S. Venticinque, "Cloud agency: A mobile agent based cloud system," Proc. Int. Conference on Complex, Intelligent and Software Intensive Systems, Krakow, Poland, pp. 132-137, February 2010

[7] B.-Q. Cao, B. Li, Q.-M. Xia, "A service-oriented Qos-assured and multi-agent cloud computing architecture," Proc. CloudCom 2009, LNCS 5931, pp. 644-649, 2009.

[8] I. Lopez-Rodriquez, M. Hernandez-Tejera, "Software *Agents* as *Cloud Computing* Services," Proc. 9th Int. Conference on Practical Applications of Agents and Multiagent Systems - PAAMS 2011, Springer, Salamanca, Spain, pp. 271-276, 2011.