IJDCST @Oct-Nov-2015, Issue- V-3, I-7, SW-02
ISSN-2320-7884 (Online)
ISSN-2321-0257 (Print)

Integrating Social Network with Cloud for Allocating and Sharing Infrastructure Resources

Shaik Faheemunnisa Begum[1], Sayeed Yasin[2]

[1]M.Tech (CS), Nimra College of Engineering and Technology, A.P., India.

[2]Head of the Department, Dept. of Computer Science & Engineering, Nimra College of Engineering and Technology, A.P., India.

Abstract —Social network platforms have brought a radical change in the way that people communicate and interact. Ever since the 'apps' become more sophisticated, it will become easier for users to share their own services, resources and data via social networks. They have provided an immense opportunity to establish, participate in digital communities as well as the exploring of social relationships. To substantiate this, we present a Social Compute Cloud where the provisioning of Cloud infrastructure occurs through "friend" relationships. In a Social Compute Cloud, resource owners offer virtualized containers on their personal computer(s) or smart device(s) to their social network. However, as users may have complex preference structures concerning with whom they do or do not wish to share their resources, we investigate, via simulation, how resources can be effectively allocated within a social community offering resources on a best effort basis. In the assessment of social resource allocation, we consider welfare, allocation fairness, and algorithmic runtime. The key findings of this work illustrate how social networks can be leveraged in the construction of cloud computing infrastructures and how resources can be allocated in the presence of user sharing preferences.

Keywords — Social Cloud Computing, Social Networks, Cloud Computing, Preference-based Resource Allocation.

I. INTRODUCTION

For the past few years, Cloud computing has gained acclamation for many reasons, most remarkably due to its ability to reduce overheads and costs for consumers by leveraging economies of scale to provide infrastructure, platforms and software as services. Infrastructure providers such as Amazon Elastic Compute Cloud (EC2) rid users of the burdens associated with purchasing and maintaining computer equipment; instead compute resources can be out-sourced to specialists and consumers can obtain access to an "unlimited" supply of resources. Despite its benefits, many businesses and end users are put off by an array of (perceived) uncertainties, as identified in numerous studies (e.g. [1], [2]). Two key issues are the notions of trust and accountability between resource consumers and providers [3]. In this context, trust and accountability encapsulate several different aspects such as security, privacy, ethical practices, transparency, protection of rights, and issues concerning compensation. Addressing these concerns is a significant undertaking, and consequentially, many international research pro- grams have emerged, covering issues such as provider certification and service level agreements.

In this paper we argue an alternative approach to establish trust and accountability in Cloud platforms: a Social Cloud [4]; and advocate a novel preference-based approach to facilitate resource sharing.

A Social Cloud is "a resource and service sharing framework utilizing relationships established between members of a social network." [5]. It is a dynamic environment through which (new) Cloud-like provisioning scenarios can be established based upon the implicit levels of trust that transcend the inter-personal relationships digitally encoded within a social network. Leveraging social network platforms as mediators for the acquisition of a Cloud infrastructure can be motivated through their widespread adoption, their size, and the extent to which they are used in modern society. For example, Facebook surpassed 1 billion users in 2012, and has illustrated that Milgram's 6 degrees of freedom in social networks [6] may in fact be as low as 4 [7]. Users also spend inexorable amounts of time "on" social network platforms – a recent study indicated up to 1 in every 7 minutes of time spent online by all Internet users worldwide [8]. The computational social capital available is also significant: if only 0.5% of Facebook users provided CPU time on their personal computer resources the potential computational power available would be comparable to a www.top500.org supercomputer [9]. Examples of such sharing include: the 25 years of cycle stealing with Condor [10], the 16 years of volunteer computing since the Great Internet Mersenne Prime Search2 and more recently Boinc; which show users are willing to donate personal computer resources to "good" causes.

Our vision of the Social Cloud is motivated by the need of individuals or groups to access resources they are not in possession of, but that could be made available by connected peers. In this paper, we present a Social Compute Cloud: a platform for sharing infrastructure resources within a social network. Using our approach, users can download and install a middleware (an extension to Seattle), leverage their personal social network via a Facebook application, and provide resources to, or consume resources from, their friends through a Social Clearing House. We anticipate that resources in a Social Cloud will be shared because they are underutilized, idle, or made available altruistically.

II. PROBLEM STATEMENT

A Social Compute Cloud is designed to enable access to elastic compute capabilities provided through a cloud fabric constructed over resources contributed by socially connected peers. A Social Cloud is a form of Community Cloud (as defined in NIST's definition of Cloud Computing [3]), as the resources are owned, provided and consumed by members of a social community. Through this cloud infrastructure consumers are able to execute programs on virtualized resources that expose (secure) access to contributed resources, i.e. CPU time, memory and disk/storage. In this model, providers host sandboxed lightweight virtual machines on which consumers can execute applications, potentially in parallel, on their computing resources. While the concept of a Social Compute Cloud can be applied to any type of virtualization environment in this paper we focus on lightweight programming (application level) virtualization as this considerably reduces overhead and the burden on providers; in [11] we explored the use of a more heavyweight virtualization environment based on Xen, however the time to create and contextualize VMs was shown to be considerable.

Architecting a Social Compute Cloud

In line with the challenges outlined above, we identify three areas of functionality needed for the construction of a Social Compute Cloud: A Social Cloud Platform: the technical implementation for the construction and facilitation of the Social Cloud as well as necessary middleware to enable resource sharing between "friends" at the edges of the Internet. A socio-technical adapter: the means to observe and interpret social ties for the elicitation or derivation of sharing preferences. A socioeconomic model: the formulation of a social microeconomic system for the allocation of resources upon the premises of social ties, and preferences with respect to how social ties denote a user-specific willingness to consume and/or provide resources.

A Social Clearing House is an institutionalized microeconomic system that defines how supply is allocated to demand. Smith defined the key components of a microeconomic system for the purposes of exchange. However, this definition is orientated primarily for monetary-based exchanges, which is not the case here. Therefore, a social clearing house captures the following: the protocols used for distributed resource allocation, the rules of exchange, i.e. who can take part, and with whom may they exchange, and the formalization of one or more allocation mechanisms. A social clearing house is therefore the central point in the system where all information concerning users, their sharing preferences and their resource supply and demand is kept. For this reason, the social clearing house requires two databases: to capture the social graph of its users, as well as their sharing preferences, and a resource manager to keep track of resource reservations, availability, and allocations.

A middleware to provide the basic resource fabrics, resource virtualization and sandboxing mechanisms for provisioning and consuming resources. It should also define the protocols needed for users and resources to join and leave the system. For these purposes we selected Seattle as it largely provides the needed functionality. However, Seattle cannot allocate its resources based upon social ties, and was thus extended.

Fig 1 A Social Compute Cloud and its Core Components.

Compute Resources are the technical endowment of users that they provide to and consume from the Social Cloud. Here, resources largely entail personal computers, servers or clusters. However, we note that the latter is unlikely for the average user. We envisage that as the computing industry continues to invest in mobile computing devices that such devices could also be offered within a Social Cloud in the future. Today, however, issues such as network stability and battery life hamper their inclusion. However, despite this researchers are making notable progress in this area.

III. RELATED WORK

With the increasing pervasiveness of social network platforms, adoption of social network structures for different types of collaboration is becoming more common. Key examples are: community and scientific portals like PolarGRID and ASPEN; social science gateways; social storage systems like Friendstore, and omemo.com; network and compute infrastructure sharing web sites such as fon.com; models to share insurance policies amongst social peers (friendsurance. de); and where social networks emerge due via collaboration, e.g. [11].

McMahon and Milenkovic [12] proposed Social Volunteer

Computing, an extension of traditional Volunteer Computing, where consumers of resources have underlying social relationships with providers. This approach is similar to the nature of a Social Compute Cloud, but it does not consider the actual sharing of resources, as there is no notion of bilateral exchange.

Ali et al. present the application of our Social Cloud model to enable users in developing countries to share access to virtual machines through platforms like Amazon EC2. In effect they subdivide existing allocations to amortize instance cost over a wider group of users. Using a cloud bartering model (similar to our previous virtual credit model), the system enables resource sharing using social networks without the exchange of money and relying on a notion of trust to avoid free riding. Like our approach, they use a virtual container (LXC) to provide virtualization within the existing virtual machine instance; however our approach using Seattle's programming level virtualization provides a much more lightweight model at the expense of flexibility.

Mohaisen et al. present an extension to our definition of a Social Cloud. The authors investigate how a Social Compute Cloud could be designed, and propose extensions to several well known scheduling mechanisms for task assignments. Their approach considers resource endowment, and physical network structure as core factors in the allocation problem, which are different considerations for resource allocation. They analyse the potential of a Social Cloud via simulation, using several co-authorship and friendship networks as input. They observe how a Social Cloud performs based upon variations in load, participation and graph structure.

Tan et al. [1] present a similar idea to the basic concept of a Social Cloud. The authors, although not extending beyond a conceptualization, motivate the philosophy of a Social Cloud with the core use case of sharing and exchanging resources within a social network or community to tackle Big Data problems.

Gracia-Tinedo et al. propose a Friend-to- Friend Cloud storage solution, i.e. dropbox via a social network: F2Box. They analyze and discuss how to retain a reliable service whilst using the best effort provisioning of storage resources from friends. They identify that a pure friend-to-friend system cannot compare in terms of quality of service with traditional storage services. Therefore, they propose a hybrid approach where reliability and availability can be improved using services like Amazon's S3. This approach provides a valuable consideration in the realisation of a Social Cloud, but is not necessarily transferable to our setting.

There have been several publications on economic models for a Social Cloud that have developed independently. Zhang et al. [5] and we [6] discuss different types of incentives users face during their participation in a Social Cloud, and describe the challenges of providing the right incentives to motivate participation. While in another study [6], we investigated how the infrastructure of a Social Cloud can be co-operatively provided by the participating members, and present an economic model that takes individual incentives and resource availability into account.

Kuada and Olesen [7] propose opportunistic cloud computing services (OCCS): a social network approach for the provisioning and management of enterprise cloud resources. Their idea is to provide a governing platform for enterprise level social networking platforms consisting of interoperable Cloud management tools for the platform's resources, which are provided by the enterprises themselves. The authors, present the challenges and opportunities of an OCCS platform, but there is no indication that they have yet built an OCCS. Similarly, Diaspora, and My3 [8] apply similar concepts to host the social network on resource provided by their users. Gayathri et al. [9] and Chen and Roscoe [5] discuss the security implications in the construction of a Social Cloud. They respectively pay special attention to, and provide counter measures for, how a Social Cloud can be used to circumvent copyright as well as perform other illicit actions. Whilst the consideration of security implications are critical for the success of a Social Cloud, it is not yet a focal point in our work.

IV. CONCLUSION

In this paper, we have presented a Social Compute Cloud: a platform that enables the sharing of infrastructure resources between friends via digitally encoded social relationships. Using our implementation, users are able to execute programs on virtualized resources provided by their friends. To construct a Social Compute Cloud, we have extended Seattle to access users' social networks, allow users to elicit sharing preferences, and utilize matching algorithms to enable preference-based socially-aware resource allocation. Preference-based resource matching is (in a general setting) an NP-hard problem, makes often unrealistic assumptions about user preferences and most state of the art algorithms run in batch modes. Therefore, we investigated what happens when we apply these algorithms to a Social Compute Cloud under the assumption that resource supply and demand do not fit to a batch allocation model. By applying methods to allocate resources in between Amazon EC2-like periodic allocations, we were able to quickly (in milliseconds) allocate resources temporarily, and then globally optimize resource allocation at the next batch allocation period. Our results are promising and indicate how the allocation of resources could take place in a production Social Compute Cloud.

As future work, we will include additional ways for users to provide their preferences, as well as methods to detect them automatically from their social network. Where examples of the latter include: clustering based on homophily (aspects of similarity), relationship lists and Granovetter-like indicators for relationship strength. This would also enable further and potentially more realistic settings for experimenting with the allocation algorithms. In terms of the Social Cloud platform we will further extend the sandbox to provide additional system calls and social access control so that users can give extended/restricted access rights to groups, for example enabling command line access for family members. These extensions would increase the number of possible applications that could be executed within the Social Cloud and also further extend the social integration of the system. Finally, we aim to investigate how users use and interact with the resources of their friends, and move our implementation towards a production ready system.

V. REFERENCES

[1] N. Cao, C. Wang, M. Li, K. Ren, and W. Lou, "Privacy-Preserving Multi-Keyword Ranked Search over Encrypted Cloud Data," Proc. IEEE INFOCOM, pp. 829-837, Apr, 2011.

[2] L.M. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner, "A Break in the Clouds: Towards a Cloud Definition," ACM SIGCOMM Comput. Commun. Rev., vol. 39, no. 1, pp. 50-55, 2009.

[3] Z. Guo, R. Singh, and M. Pierce, "Building the polargrid portal using web 2.0 and opensocial," in Proceedings of the 5th Grid Computing Environments Workshop (GCE '09). New York, NY, USA: ACM, 2009, pp. 1–8.

[4] R. Curry, C. Kiddle, N. Markatchev, R.Simmonds, T. Tan, M. Ar-litt, and B. Walker, "Facebook Meets the Virtualized Enterprise,"in Proceedings of the 12th

International IEEE Enterprise DistributedObject Computing Conference (EDOC '08). Washington, DC, USA: IEEE Computer Society, 2008, pp. 286–292.

[5] W. Wu, H. Zhang, and Z. Li, "Open social based collaborative science gateways," in Cluster, Cloud and Grid Computing (CCGrid), 2011 11th IEEE/ACM International Symposium on. IEEE, 2011, pp. 554–559.

[6] W. Wu, H. Zhang, Z. Li, and Y. Mao, "Creating a cloud-based life science gateway," in E-Science (e-Science), 2011 IEEE 7th Interna-tional Conference on. IEEE, 2011, pp. 55–61.

[7] D. N. Tran, F. Chiang, and J. Li, "Friendstore: Cooperative online backup using trusted nodes," in Proceedings of the 1st International Workshop on Social Network Systems (SocialNet'08), Glasgow, Scot-land, 2008.

[8] D. D. Roure, C. Goble, and R. Stevens, "The design and realization of the myexperiment virtual research environment for social shar-ing of workflows," Future Generation Computer Systems, vol. 25,no. 5, pp. 561–567, 2009.

[9] G. Klimeck, M. McLennan, S. P. Brophy, G. B. Adams, and M. S.Lundstrom, "nanoHUB.org: Advancing Education and Research in Nanotechnology," Computing in Science and Engineering, vol. 10,pp. 17–23, 2008.

[10] A. McMahon and V. Milenkovic, "Social volunteer computing,"Journal on Systemics, Cybernetics and Informatics (JSCI), vol. Volume 9(4), pp. 34 – 38, 2011.

[11] Z. Ali, R. Rasool, and P. Bloodsworth, "Social networking for sharing cloud resources," in Cloud and Green Computing (CGC), 2012 Second International Conference on, 2012, pp. 160–166.

[12] A. Mohaisen, H. Tran, A. Chandra, and Y. Kim, "Socialcloud: Us-ing social networks for building distributed computing services,"arXiv preprint arXiv:1112.2254, 2011.

Authors:

Student Pic

Guide Pic