Future Generation Computer Systems 29 (2013) 2143–2156

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

eScience in the Social Cloud☆

Kris Bubendorfer [a,*], Kyle Chard [b], Koshy John [a], Ashfag M. Thaufeeg [a]

[a] School of Engineering and Computer Science, Victoria University of Wellington, Wellington, New Zealand
[b] Computation Institute, University of Chicago and Argonne National Laboratory, Chicago, IL, USA

HIGHLIGHTS

- We explore two Social Cloud models designed to support collaborative eScience.
- Social Clouds rely on a collaborative overlay obtained from existing social networks.
- The SoCC supports VM based sharing of computational resources between peers.
- The SoCPR uses social incentives to increase contribution in volunteer computing.
- The prototype and simulations show efficient sharing and increased contribution.

ARTICLE INFO

Article history:
Received 6 March 2012
Received in revised form
29 January 2013
Accepted 6 April 2013
Available online 16 April 2013

Keywords:
Social Cloud
Distributed systems
Cloud computing
Volunteer computing
Collaboration
Resource sharing

ABSTRACT

Social networks offer great potential for fostering collaboration between individuals and amongst groups. This potential collaborative environment is not only applicable for recreation, but can also provide considerable value to diverse research communities. For this reason scientists are increasingly utilizing social networking concepts in projects to form groups, share information, publicize their work and communicate with their peers. This article describes two different approaches to supporting eScience, by providing scientific computing and collaboration within what we term the Social Cloud. In our first approach the social network is used as a collaborative overlay, in combination with the ad hoc creation of infrastructure composed of virtual machine clusters built from resources contributed, by the users, to the Social Cloud. Our second approach is based around the principle of volunteer computing, where the Social Cloud provides researchers with a platform to exploit social networks by reaching out to non technical users who would otherwise be unlikely to donate computational time for scientific and other research. In this article we specifically explore the motivations of users to contribute computational time and examine the various ways these motivations can be catered to through the use of incentives in existing social networks.

1. Introduction

The rapid growth of social networking has created new ways for individuals and companies to communicate and share basic information without geographic barriers. These digital relationships are as important to many individuals as their real world relationships and form a primary part of their daily social interactions. Many applications now use social networks as a platform for authentication, e.g., Facebook Connect and application Portals such as ASPEN [1] and PolarGrid [2]. However, Social Networks offer the potential for much more than simple user authentication or the provision of portals. Indeed, social networks provide a valuable general basis for enabling communication, coordination and discovery for individuals and groups of individuals. Scientists and researchers have been quick to adopt the principles and ideas of social networking, adding them to scientific platforms such as MyExperiment [3] and nanoHub [4] but the scope of these projects is limited by the size and disconnect between their respective communities.

The Social Cloud, first published in [5], takes a different approach by extending collaboration and computation functionality to existing social networks—rather than adding social networking to existing computation and collaboration tools. This has a number of advantages—we can utilize the existing connections between people, exploiting their preexisting and vital communities,

☆ This is a combined and extended version of the papers: K. John, K. Bubendorfer, and K. Chard. A Social Cloud for Public eResearch. In proceedings of the 7th IEEE International Conference on eScience, Stockholm, Sweden, December 2011. And A. Thaufeeg, K. Bubendorfer, and K. Chard. Collaborative eResearch in a Social Cloud. In proceedings of the 7th IEEE International Conference on eScience, Stockholm, Sweden, December 2011.

* Corresponding author. Tel.: +64 4 463 6484; fax: +64 4 463 5045.
E-mail addresses: kris@ecs.vuw.ac.nz (K. Bubendorfer), kyle@ci.uchicago.edu (K. Chard), koshyjohnuk@msn.com (K. John), thaufeashf@myvuw.ac.nz (A.M. Thaufeeg).

0167-739X/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.future.2013.04.003

and we can utilize the social network's familiar and comfortable user interfaces (including the use of groups, incentive mechanisms, feeds and walls) to reduce the learning curve and frustration with yet another system. From a pragmatic technical point of view, we benefit from well developed programming interfaces, reliable architecture, provision of infrastructure, plus authentication and authorization systems. Specifically a Social Cloud is defined in [6] as; *A Social Cloud is a resource and service sharing framework utilizing relationships established between members of a social network*, and relies on the principle that social connections represent contextual trust that can be applied in any scenario where sharing and exchange (essentially collaboration) takes place. A Social Cloud is unique in that it is completely open and ad hoc—it is created, controlled and managed by its users and therefore requires only a lightweight infrastructure external to that provided by the host social network.

Researchers faced with extremely large computations or the requirement of storing vast quantities of data have come to rely on distributed computational models, including clusters, grids and more recently clouds, to provide large scale science capacity. However, distributed computation is perceived as complex, expensive and often has little uptake outside of the eScience community—that is, by the wider group of eResearchers. This article explores two developments within the Social Cloud project in support of eScience and eResearch.

The first project is the Social Collaborative Cloud, outlined in Section 1.1, and the second project in this article is the Social Cloud for Public eScience, which is outlined in Section 1.2.

1.1. Social Collaborative Cloud

For scientists and other eResearchers, multi-institute collaboration is common, however communication is often difficult with face to face meetings occurring only sporadically at conferences and workshops. Social networks can be used to support collaboration by increasing communication channels and also facilitating the discovery of other scientists working on similar projects. Scientific collaborations also often have resources, such as compute, data and meta resources (e.g. workflows), that they wish to share dynamically in groups for the duration of a project. Currently this is a difficult process requiring manual (reciprocal) user account registration, creation, authorization, multiple systems and credentials, multiple user interfaces and so on.

The Social Collaborative Cloud (SoCC), is a framework in which individuals or institutions contribute the capacity of their compute, data and other resources, which are leased through the social network. In the context of computation, members of the Social Cloud can contribute, request, and use Virtual Machines from other members, as well as from Virtual Organizations aligned with project goals by forming Social Network groups. The SoCC enables scientists to share resources for the duration of a collaboration, allows more efficient use of available resources, as scientists with similar projects or interests can naturally consolidate resources towards common goals, and most importantly, it promotes greater uptake beyond the scientific community by utilizing familiar tools, as well as publicity and networking opportunities provided by the host social network. In this way, the social network is not just an add-on, it is instead integral, to the Social Cloud.

1.2. Social Cloud for Public eResearch

Volunteer computing [7] is an alternative means of obtaining large computing resources for eScience and lately eResearch—by getting the public to support specific projects by donating their spare computational and storage resources. The amount of computational time available to researchers is a function of the number of volunteers contributing at any given point of time. While there are a sizable number of volunteers who participate in volunteer computing (e.g. 2.2 million BOINC participants [8]), this is insignificant when compared to the 845[1] million active Facebook users [9] and the potential compute power they could contribute for research—within their own individual limitations.

The Social Cloud for Public eResearch (SoCPR) aims to provide researchers with a platform to exploit social networks to reach out to users who would otherwise be unlikely to donate computational time for scientific and other research oriented projects. We explore the motivations of users to contribute computational time and examine the various ways these motivations can be catered to through established social networks. Specifically, we look at integrating Facebook and BOINC, and discuss the architecture of the functional system and the novel social engineering algorithms that power it.

1.3. The Social Cloud paradigm

What is different about the Social Cloud paradigm? The social network comes first: It is not a cloud or collaboration environment extended with a social network, it is a social network extended with cloud functionality. The people and their networks form the basic infrastructure, the services they access and share are formed around their unique social graph. Individuals contribute their competencies—services that encapsulate; data, storage, computation, algorithms, etc. The users choose how to delegate their contributions between the groups within which they collaborate.

Users fulfill all roles, from provision of resources, management of collaborations through to consumption of services and computation—access to which is performed with familiar social networking tools. In some ways this turns the provision of infrastructure on its head. Rather than fitting the user to the infrastructure, we build the infrastructure around the user.

A Social Cloud is not crowdsourcing as the relationships in the social cloud are essentially[2] symmetric, that is, participants are more-or-less equals who come together under some commonality of interest to benefit from sharing. There is explicitly a preexisting and external underlying set of relationships that connects people within a Social Cloud (no anonymity). Whereas crowdsourcing operates in the master–worker model where workload flows in one direction, while there may be monetary exchange or other compensation this does not in itself constitute sharing.

1.3.1. Advantages in usability

Using a Social Cloud lowers usability barriers: The interface and tools are already familiar to non expert (non computer science) users. Traditionally difficult authorization and access control take place transparently for the owner and users. There is no visible certificate management, the social network provides single sign on and applications can be seamlessly integrated with core functionality. Collaborations are light-weight and dynamic, services and resources can be delegated, removed and accessed using the social network group structure. Users share services for the duration of collaboration. A good role for a Social Cloud is in the early stages of a project, when the costs of dedicated or leased infrastructure would be prohibitive.

1.3.2. Groups in a Social Cloud

Groups in Social Networks, like virtual organizations [10], have an intent, membership and policies that define sharing in

[1] March 2012.

[2] The Social Cloud for Public eResearch is the exception.