

# Sustainable C\*omputing

## REGISTER

Monthly newsletter of the IEEE Computer Society Special Technical Community on Sustainable Computing  
*Providing quick access to timely information on sustainable computing.*

<http://www.linkedin.com/groups/IEEE-STC-on-Sustainable-Computing-4092681>

<http://www.facebook.com/groups/STC.Sustainable.Computing/>

<http://stc-sustainable-computing.ieee.net/members>

Volume 1 Issue 10

October 2012

## Contents

- 2 Letter from the Editor
- 3 Column: Resource and Energy Management
- 5 Column: Networks and Distributed Systems
- 6 Pick of the Month Winners
- 7 STC Updates
- 8 Upcoming Events

Sponsored by:



## Letter from the Editor



Looking back, 2012 has been a great year for our STC. This newsletter, which was first published in January, now has 8 regular contributors. The Pick of the Month series, started in March, is becoming a prestigious and popular award. Anirban and Sergey have led a successful membership drive. We now have more than 300 members. To borrow a phrase from American politics, "The state of our STC is strong." Looking forward, 2013 promises more growth. Anirban will take on the role of STC chair. He has a lot of ideas on how to make our STC even more engaging and relevant to our members. Our newsletter will expand its web presence. Today, I am writing to ask *you* to get involved!

We need more Officers to keep up with our STC's growth. In 2013, each one of our current Officers has been tasked with finding 1 or 2 Vice Officers. Vice Officer is like an apprenticeship. You will share the same responsibilities as an Officer but at lighter load. Ideally, after an Officer has served for a few years, a Vice Officer will step up to fill the position. Anirban and Sergey paved the way for this model as Membership Officers.

How does it help you in your career? Vice Officer is a great way to gain visibility in the community. Did I mention that we have over 300 members and growing? There is not a stronger network of researchers interested in green computing, smart grids, smart buildings, and energy efficiency in the world.

Each Vice Officer position will be different. However, I will conclude my letter by describing the Vice Editor position. In this role, you will publish the newsletter 4 times a year. You will coordinate with the STC Chairs, Information Officers, and Communication Officers to do so. You will also oversee the selection of 4 Picks of the Month. I think the time commitment would likely be less than 12 8-hour days over the calendar year. And we can work together to schedule those days. I look forward to working with you.

## Officers

Co-Chair: Martin Arlitt, University of Calgary/HP Labs  
Co-Chair: Ishfaq Ahmad, Univ. of Texas at Arlington  
Secretary/Treasurer: Giuliano Casale, Imperial College  
Conferences Chair:  
Diwakar Krishnamurthy, University of Calgary  
Academic Chair:  
Niklas Carlsson, Linköping University  
Membership Chair: Anirban Mahanti, NICTA  
Vice Membership Chair:  
Sergey Blagodurov, Simon Fraser University  
Communications Chair:  
Abhishek Chandra, University of Minnesota  
Policies and Procedures: Stephen Dawson, SAP  
Industry Chair: Canturk Isci, IBM  
Newsletter Editor:  
Christopher Stewart, The Ohio State University  
Webmaster:  
Raoufhsadat Hashemian, University of Calgary  
Information Officer:  
David Carrera, Universitat Politècnica de Catalunya  
Information Officer:  
Danilo Ardagna, Politecnico di Milano  
Information Officer:  
Fan Dongrui, Institute of Computing Technology  
Information Officer: Guillaume Jourjon, NICTA

# Resource and Energy Management

by Danilo Ardagna, Politecnico di Milano



This month's column was written by Qi Zhang and Raouf Boutaba of the David R. Cheriton School of Computer Science at the University of Waterloo, Canada. The content extends their paper, “Dynamic Service Placement in Geographically Distributed Clouds” published in the proceedings of the 32nd IEEE International Conference on Distributed Computing Systems (ICDCS). The paper proposes a solution for the run-time management of Cloud infrastructures based on Model Predictive Control. The analysis is comprehensive, considering even the problem of resource contention among multiple selfish users. Here, a user is a service providers hosting their application in the cloud data centers. Competing service providers are modelled as a Nash Game. The results are very interesting and highlight the need of coordination among the cloud provider and the service providers to effectively manage resources under heavy workloads. I am happy to introduce Qi and Raouf's featured article this month.

## **Featured Article: Service Placement across Geographically Distributed Data Centers By Qi Zhang and Raouf Boutaba**

Cloud computing is a form of Internet computing, where online service applications are hosted in data centers and delivered to end-users over the Internet. To provide efficient support for such a computing model, cloud providers today build data centers in geographically distributed locations, with the goal of finding a balance between service response time, capital investment, operational cost as well as environmental-friendliness (the March 2012 issue of the STC Sustainable magazine reported a detailed study on selecting locations for building data centers). Currently, Google has data centers in over 19 locations around the globe, whereas Microsoft runs Windows Azure over 8 data centers across Europe, Asia and the United States ([link](#)). As cloud companies continue to expand their business in cloud computing, these numbers are expected to increase in the near future.

Given the variety of data center locations, cloud customers (mainly service providers) are facing the problem of determining in which data center(s) their service applications should be placed. Typically, the objective of a service provider is to choose appropriate sites for service hosting to fulfill appropriate performance requirements as specified in Service Level Agreements (SLAs), while minimizing total resource rental cost. However, in the context of cloud computing, there are several factors that make the problem challenging:

- Service demand originates from multiple sources (e.g., access networks), and is subject to change over time. For many day-to-day services such as online transactions, the service demand from each geographical region often exhibits a periodic (day-and-night) pattern, where demand is higher during day time than during night time. However, large demand variations such as spikes can also occur occasionally. In these situations, provisioned servers may not have sufficient capacity to handle all service demands in a timely manner. It is therefore necessary to dynamically provision service capacities by adding and removing active servers according to service demand.

Continued on page 4

## Featured Article: Service Placement across Geographically Distributed Data Centers

- A primary concern of the service provider is to achieve the required performance specified in SLAs. SLAs are often expressed in terms of request response time, which is largely affected by the latency between end-users and the data center hosting the service application. As a result, the location of the data center and the distance to end-users are crucial factors that have to be considered when service providers make decisions about the placement of their services.
- Resource prices offered by cloud providers can also fluctuate over time. Indeed, many parts of the world are adopting dynamic energy pricing schemes, and this, in turn, results in fluctuating resource prices. Obviously, adjusting resource allocation according to resource price can lead to significant savings for the service provider. So when it comes to choose the appropriate data center for hosting services, it is necessary to consider not only the diversity of prices in the marketplace but also the resource price fluctuation over time in each of the locations.
- Lastly, there is often a cost associated with reconfiguration (i.e., the cost of adding and removing servers). The consideration of reconfiguration cost is important for ensuring system stability and minimum management overhead. For instance, adjusting the number of servers typically incurs switching costs for service setup (e.g., VM image distribution) and tear-down (e.g., saving data and states of the server to storage devices). Here, stability is perhaps a more important objective. Greedily adapting resource allocation to price and demand fluctuation can lead to frequent oscillation of service placement decisions, which can incur high cost in terms of replication and migration.

The ICDCS paper studied the service placement problem in dynamic cloud computing environments. Specifically, the authors have considered all the aforementioned factors and proposed a control theoretic framework based on Model Predictive Control (MPC) for making online service placement decisions. MPC is a well-known technique for solving dynamic optimization problems in online settings, and has been successfully applied to many industries. As shown in the paper, in the cloud context, MPC is also able to provide a simple yet effective solution for the dynamic service placement problem.

However, even though solving the service placement problem for a single service provider seems important enough, there is an even bigger issue when considering a more realistic scenario with multiple service providers placing services at the same time. The challenge in this setting is that each data center has limited capacity. When a large number of service providers try to scale-up their service infrastructure simultaneously, it is possible for the resource demand to exceed the data center capacity. In this case, service providers have to compete for resources in each data center. This competition among service providers can be modeled as a dynamic game where each service provider is a player that makes selfish decisions. The outcome of the game can be described by Nash Equilibria, which are stable configurations where no individual service provider can further reduce its cost by changing its service placement decisions. The quality of a Nash Equilibrium (NE) is measured by the total social welfare, which in this case, is the sum of the total costs of all service providers. Through the analysis provided in the paper, it turns out that the best NE of the game is the one that achieves the optimal social welfare. However, the worst NE can be arbitrarily worse than the optimal social welfare, if the cloud provider rejects all remaining requests once the data center capacity is fully used.

This result is quite interesting from a theoretical standpoint, but also has strong real-world implications.

Continued on page 6

# Networks & Distributed Systems

by Guillaume Jourjon, NICTA



## Is DIY sustainable?

We can achieve more energy efficiency in networking thanks to two main operations; 1) by optimizing the energy consumption for a specific networking service and 2) using current and future networking technology to improve all services. This article focuses on the latter; the use of technology to improve the productivity of communities or working group. In particular, the article focuses on Do It Yourself (DIY) technologies. This well-known approach is usually taken by the gentlemen farmer, Sunday renovators, or tech moms who want to improved their house during the weekend. Recently more and more technical frameworks are made available in order to provide a cheap DIY opportunities.

DIY initiatives in ICT have recently gained a lot of momentum and media coverage with the success of the fab labs from MIT and in particular their 3D printers. In addition to very interesting tools, the fab labs offer a set of tools and workshops in order to build some cheap and potentially sustainable tools. However, when I read about all these beautiful, interesting and well designed initiatives, I sometimes find it difficult to clearly understand the sustainability of these projects. In particular, coming from an engineering perspective I most of the time assume (certainly in a naïve manner) that if the production process is well done and engineered, we should have less waste than if everybody built their own object. In the remainder of this article, I would like to focus on two recent DIY projects.

FabFi (shown to the right):

The first DIY project is called the FabFi project. FabFi aims at providing both long range point-to-point and mesh hybrid-wireless broadband transmission infrastructure. This project gives you some guidance to build the parabolic reflectors as shown below as well as all the configuration of the wireless network. This project is very impressive and already operational in some places (India, Kenya and Afghanistan).



DIY Video Conferencing:

The July issue of IEEE spectrum (pp 18-19) included a great story about setting up a low cost a video conference service based on Skype and some pretty specific, brand new hardware. The author does not claim any sustainability for this solution, but sometimes in the media the DIY term is associated with a green and community driven aspect. These two projects while part of the same category do not provide the same ecological footprint. In particular, the first one offers a specification based on very generic hardware while the second requires new hardware (higher footprint) with specific requirements.

This kind of community projects brings a lot of advantage but maybe the introduction of some kind of labels or ratings should be introduced in order to better identify if their result is sustainable or not.



## Featured Article: Service Placement across Geographically Distributed Data Centers

Specifically, the reason for the bad Nash Equilibria is that certain service providers will not get their preferred locations due to late arrival of their requests, and as a result, suffer a heavy penalty in terms of performance and cost. To alleviate this situation, the simplest solution is the over provisioning of data center capacity through continuous capacity upgrades. However, this is not practically feasible or always possible in a dynamic environment with occasional, unforeseen demand spikes. Another, more promising solution is to change either the resource allocation scheme or the pricing scheme used by the cloud providers. In this case, the cloud provider acts as a coordinator that adjusts resource allocation according to the supply and demand, with the goal of achieving the optimal social welfare. For instance, the cloud provider can request each service provider to submit its future requests ahead of time, and gradually adjust either the price or resource allocation until they converge to a stable and near optimal outcome. From a service provider perspective, this is a reasonable model as social welfare can be perceived as a measure of customer satisfaction.

As it has been pointed out in the paper, there is much more to be done in the area of resource management in dynamic, geographically distributed cloud computing environments, where resources are shared among service providers that make selfish decisions. The work reported in the paper provides a foundation for future investigations in this area.

### Recent Pick of the Month Winners

March	Minimizing Data Center SLA Violations and Power Consumption via Hybrid Resource Provisioning	IGCC 2012 (best paper)
April	Capping the Brown Energy Consumption of Internet Services at Low Cost	IGCC 2011 (best paper)
May	Power Budgeting for Virtualized Data Centers	Usenix ATC 2011
June	Blink: Managing Server Clusters on Intermittent Power	ASPLOS 2011
Sept.	Leveraging Stored Energy for Handling Power Emergencies in Aggressively Provisioned Datacenters	ASPLOS 2011

### Call for Nominations

The Sustainable Computing Register features Picks of the Month, a research publication or industry project that has significantly advanced the field of sustainable computing. The goal is to increase awareness within our community about high-impact, transformative research.

1. Members can submit worthy papers and industry projects by emailing me.
2. Submissions endorsed by 2 STC-SC officers will advance to public vote on Facebook.
3. By visiting our Facebook page, all members can vote for their favorite paper. At the end of each month, the paper with the most votes will become a Pick of the Month (provided the authors agree to be interviewed). Papers nominated but not selected for more than four (4) months will be removed.

### Requirements for nominees:

- The paper must have been published in a peer-reviewed, research forum.
- The paper must be related to sustainable computing, e.g., energy efficiency, renewable-powered computing, smart grid, life cycle of ICT, smart buildings, etc.
- The paper must have been published in the last 2 years.
- Industry projects must have shown significant practical impact or intellectual contribution.

# STC Updates



By Giuliano Casale, Imperial College

Membership: 303

Report from Secretary/Treasurer (Giuliano Casale):

- Collected officers' activity reports and prepared monthly STC report
- Investigating into financial regulations
- Danilo Ardagna accepted to serve as vice Secretary/Treasurer
- Organizing upcoming newsletter contribution from Prof. Erol Gelenbe, Imperial College

Report from Conferences Chair (Diwakar Krishnamurthy):

- Contacted organizers of CGC 2012 and UCC 2012 to publicize our STC
- Coordinated with LCN 2012 organizers to display a short presentation and poster at the event
- Professor Amarjeet Singh of IIIT, Delhi has kindly consented to be the vice-chair of conferences
- Working on revamping the conference calendar

Report from Academic Chair (Niklas Carlsson):

- Working with the industry chair on a student award proposal

Report from Membership Chairs (Anirban Mahanti and Sergey Blagodurov):

- Created contact list for E-Energy and CloudCom

Report from Communications Chair (Abhishek Chandra):

- Continued to identify conferences, workshops and journals relevant to sustainable computing
- Prepared a spreadsheet with information about upcoming call for papers and call for participation
- Started looking for a vice officer

Report from Policies and Procedures Chair (Stephen Dawson):

- Documenting STC-SC processes and gathering requirements for future processes
- In the process of identifying a successor to bring forward the policy document drafting
- Started looking for a vice officer

Report from Industry Chair (Canturk Isci):

- Working on the STC Community Highlights feature with a potential research group
- Outlining the next Industry Highlight feature for STC

Report from Information Officers (Danilo Ardagna, Fan Dongrui, Guillaume Jourjon, David Carrera):

- Contributed material for newsletter and blogs
- DA: David Carrera has been involved in the STC as new Information Officer and will provide contributions to the new "Big Data and Data-Center Computing" blog and to the registry

Report from the Newsletter Editor (Christopher Stewart):

- Looking into publicity avenues for the newsletter
- Opened the call for a vice-editor position

# Upcoming Events

By Abhishek Chandra, University of Minnesota



The following venues are all requesting submissions on subtopics related to sustainable computing or IT for sustainability.

## Conference, Workshop & Symposium Call For Papers

Short Name	Main Topic	Location	Dates	Abstracts Due	Papers Due
<i>SIGMETRICS</i>	<i>Performance Evaluation</i>	<i>Pittsburgh, PA, USA</i>	<i>Jun. 17-21, 2013</i>	<i>Nov. 2, 2012</i>	<i>Nov. 9, 2012</i>
CCGrid	Cluster, Cloud and Grid Computing	Delft, Netherlands	May 13-16, 2013		Nov. 12, 2012
ICDCS	Distributed Computing Systems	Philadelphia, USA	Jul. 8-11, 2013		Nov. 12, 2012
ENERGY 2013	Smart Grids, Green Communications	Lisbon, Portugal	Mar. 24-29, 2013		Nov. 12, 2012
ISCA	Computer Architecture	Tel Aviv, Israel	Jun. 23-27, 2013	Nov. 14, 2012	Nov. 21, 2012
SEIT 2013	Sustainable Energy IT	Halifax, Canada	Jun. 25-28, 2013	Dec. 20, 2012	
HPDC	High Performance Distributed Computing	New York, USA	Jun. 17-21, 2013	Jan. 14, 2013	Jan. 21, 2013
USENIX ATC	Computer Systems	San Jose, CA, USA	Jun. 26-28, 2013	Jan. 23, 2013	Jan. 30, 2013
GreenMetrics	Sustainable Computing	Pittsburgh, PA, USA	Jun. 17, 2013		Apr. 15, 2013

## Journal and Special Issue Call For Papers

Sustainable Computing

## Conference, Workshop & Symposium Call for Participation

Short Name	Main Topic	Location	Dates
<i>CGC 2012</i>	<i>Cloud and Green Computing</i>	<i>Xiangtan, China</i>	<i>Nov. 1-3, 2012</i>
<i>UCC 2012</i>	<i>Utility and Cloud Computing</i>	<i>Chicago, IL, USA</i>	<i>Nov. 5-8, 2012</i>
<i>GreeNETS 2012</i>	<i>Green Communications and Networking</i>	<i>Gandia, Spain</i>	<i>Oct. 24-26, 2012</i>
<i>SC'12</i>	<i>Supercomputing</i>	<i>Salt Lake City, USA</i>	<i>Nov. 10-16, 2012</i>
GreenCom	Green Computing and Communications	Besancon, France	Nov. 20-23, 2012

Visit <http://stc-sustainable-computing.ieee.net/venues> for more information.

To advertise a relevant venue, email Abhishek Chandra at [chandra@cs.umn.edu](mailto:chandra@cs.umn.edu).

To purchase a printed copy of the Register visit <http://www.magcloud.com/>

## **Sustainable Computing: Informatics and Systems**

The journal for sustainable computing research

Sustainable computing research spans computer science, electrical engineering, sustainability science, and many other engineering disciplines. SUSCOM publishes research findings related to energy-aware and thermal-aware management of computing resources, as well as research on the ecological and societal impacts of computing.

**Now accepting submissions.**