

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 3

March 2012

Contents

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

Sponsored by:



The IEEE STC on Sustainable Computing

Mission

The STC on Sustainable Computing (STC-SC) is one of the initial set of IEEE Special Technical Committees (STC). The STC-SC takes a holistic view of information technology (IT). Any domain that uses or could potentially benefit from using IT is relevant to STC-SC. The two primary goals of STC-SC are: (1) to promote the design and implementation of sustainable computing; and (2) to facilitate computing for sustainability. With respect to these goals, topics relevant to STC-SC include (but are not limited to):

- Energy efficient design and operation of IT equipment (servers, storage and networks).
- Sustainability across the life cycle of IT equipment and processes.
- Measurement and evaluation of the sustainability of existing IT infrastructures.
- Models or methods that facilitate sustainable computing.
- Use of computing to systematically improve the sustainability of non-IT processes.



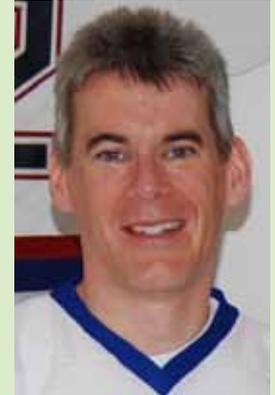
Letter from the Editor

Last month, our membership almost tripled in size! Our STC is becoming a torchbearer for the sustainable computing movement. There are several reasons why. Martin and Ishfaq have put together strong outreach efforts (e.g., this newsletter) and our officers are passionate. Also, it is easy to join (visit the URLs on the front cover). But in this letter, I want to highlight a shared feature between our STC and the sustainable computing movement: They are both grassroots efforts. Members, like yourself, have spoken with your friends in one-on-one settings, "Hey, I am really interested in green computing because sustainability is an important problem to me." You are posting news from this STC to your Facebook page, bringing in new members.

Let's keep it up! Have you mentioned this STC on your home web page? Is it an interest/project on your LinkedIn account? Your efforts can influence others. Thank you for your leadership. Oh, and welcome to the group.

Co-Chair: Ishfaq Ahmad
U. of Texas at Arlington

Co-Chair: Martin Arlitt
HP Labs/U. of Calgary



Officers

Secretary/Treasurer: Giuliano Casale
Imperial College

Conferences Chair: Diwakar Krishnamurthy
University of Calgary

Academic Chair: Niklas Carlsson
Linköping University

Membership Chair: Anirban Mahanti
NICTA

Communications Chair: Abishek 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

Networks & Distributed Systems

by Guillaume Jourjon, NICTA



How many students do we need?

In the context of a series of short courses given at NICTA, I recently presented a series on an introduction to networking. These short courses were organised across five different venues in Australia: two in Sydney, one in Canberra, one in Melbourne and one in Brisbane. Delegates could attend the venues in person or via video conferencing. While it is generally accepted that video conferencing is cheaper, in terms of GHG emission, this post will look at the numbers and try to find a simple relation between the minimum number of students attending remotely and the green effect of setting up a DiffServ network between the aforementioned venues.

The Setup: The course spanned 3 days with 8 hours of presentation every day. Every venue received between 2 and 10 student registrations. Students were to attend in person or as a remote group from their laboratory. Students viewing from a lab streamed the course through an AARNET website. NICTA booked seminar rooms at the Melbourne, Canberra, Brisbane and Sydney laboratories. AARNET deployed a DiffServ service over their network that connected NICTA to the Internet.

We chose this setup because we thought that bringing all the students to a single venue would have cost a lot, in terms of energy. Students outside of major cities would have travelled an average of 1000 kilometres each. However, we ended up having only one student that followed the entire course remotely, while another one was coming in and out. Based on the disappointing outcome of seeing only one student on the remote site (instead of the 20 registered), I decided to see if tele-teaching was really worth it.

The Numbers: The numbers presented hereafter came from four different sources. The first one is “Sustainable Energy — without the hot air”, David JC Mac Kay, a very interesting and comprehensive book available freely on the Internet. The second source is coming from the work of Rod Tucker from the University of Melbourne. I also used the effective electricity consumption of the NICTA building from our facility manager. Finally, the last source I used was the energy consumption for a 737-800 aircraft from Boeing.

If you want to bring students into 1 central venue, it will cost airplane tickets. You also need a conference room and air conditioning unit. For this portion of the analysis, we can ignore the air conditioning unit (you'll see why shortly). We can compute the required consumption for one student for a 1000 km trip:

$$C_{\text{onSite}} = \text{NbLitres} \times \text{PowerOfLitre}$$

Which, based on our sources, gives a consumption per student whatever the length of the course of:

$$C_{\text{onSite}} = 600\text{kWh}$$

This value is obtained as follows; the consumption per passenger of a Boeing 737-800 is around 3L per 100 km. The round trip is about 2000km. Burning 1 L of kerosene produces 10kWh. Hence the consumption depicted above.

Continued on page 6

Minimizing Data Center SLA Violations and Power Consumption via Hybrid Resource Provisioning by Anshul Gandhi, Yuan Chen, Daniel Gmach, Martin Arlitt, and Manish Marwah published in the Proceedings of the Second International Green Computing Conference.

Our first ever Pick of the Month winner received the Best Paper Award at IGCC 2011 and it is easy to see why. It is replete with timely intellectual and practical contributions. On the modelling side, the paper proposes a novel discretization method for performing Fast Fourier Transforms on datacenter traces. On the systems side, there is an adaptive approach to respond to unexpected workload surges. These contributions combine to significantly reduce SLA violations, a worthy goal according to any practitioner. For this month's feature, I was fortunate to interview Anshul Gandhi and Daniel Gmach.

Christopher Stewart: Best paper at IGCC is quite an honor. Given the positive reception from the community and the topic (SLAs and power), I'd imagine that you could have sent this paper to other competing conferences. What led you to submit to IGCC instead of other top forums?

Anshul Gandhi: Surprisingly, while there are a lot of systems conferences, there are very few green computing conferences. IGCC was clearly at the top of the list for green computing conferences.

Daniel Gmach: Thanks, Christopher. I can only speak for myself here but I believe my coauthors thought similar. During my daily work, I came across some IGCC papers from the previous year and was highly impressed by quality and impact. So, I decided that this is a very good conference where I wanted to publish

Christopher Stewart: Let's get into the technical aspects of the paper. In your opinion, what will be the most influential contributions from this work? When did you realize that these contributions could be important?

Anshul Gandhi: The workload discretization, including the DP [dynamic programming] approach, and the combination of predictive and reactive provisioning.... The workload discretization idea was not part of the [original] agenda and was conceived along the way. However, once we implemented our DP based discretization idea...we were confident that this would be an influential contribution.

Christopher Stewart: Section 2 was full of goodies for me. Personally, I believe FFTs are under used in systems research. Who was the FFT expert? Who implemented the workload discretization algorithm?

Daniel Gmach: I agree that FFTs are very useful to predict patterns. They are actually used quite frequently to determine re-occurring patterns in time series. A common alternative are auto-correlation functions, which [can] also work very well. I've used FFTs and autocorrelation functions in the past.... Anshul and Manish have great expertise on FFTs.... The implementation itself was mostly done by Anshul during his internship at HP Labs. He is a very good student who has great theoretical as well as practical background.

Disclaimer: Comments in this article reflect the personal views of the interviewed authors only. These views may not reflect the views of other authors, affiliated institutions, or the publishing organization.

Nominations are Open!

Each issue of the Sustainable Computing Register features a Pick 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.

Selection Process:

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.

Christopher Stewart: Alright, a lighthearted question: How much time was there between when the last graph was produced and when the paper was submitted? Did you sleep that night?

Anshul Gandhi: I hate to disappoint you, but we actually got all our results done a couple of months before we finished our paper. This project was part of my 2010 summer internship, and so we made sure that we finish all the experiments before my internship was done.

Daniel Gmach: In my opinion, good time management and giving yourself enough time to finalize and improve the presentation of the paper is very important. The history of this paper might be a good example for this.

Christopher Stewart: Increasingly, researchers are beginning to look at sustainable systems that use little energy and/or are powered by clean energy. But your work reminds us that these systems better perform well too, as measured by the SLA. What do you see as the key research challenges facing SLA management going forward, especially for sustainable systems?

Anshul Gandhi: I think researchers have realized that focusing on performance SLAs alone is not sustainable. I believe that SLAs need to be restructured, wherever possible, so as to take into account some form of "sustainability". In some scenarios, this is already happening. There are tax credits for businesses that meet certain eco-friendly standards, so there is also a cost motivation for this restructuring. In recent years, researchers have looked at metrics such as "Performance-per-watt" or "queries per joule", or holistic metrics such as PUE.

Thus, the key research challenge here is to come up with SLAs that take sustainability into account, either explicitly (99% of requests should complete in 500ms, using no more than x watts of power per request), or implicitly (a weighted sum of cost for missing performance targets and cost for consuming power).

Christopher Stewart: Thank you both for your time and valuable contribution to the field. Best wishes on your future endeavors. I hope to see you at IGCC 2012.

Resource and Energy Management

by Danilo Ardagna, Politecnico di Milano



Placing datacenters around the world

Popular Internet services use multiple geographically distributed datacenters. For example, Google has 19 locations in the US, 12 in Europe, one in Russia, one in South America, and three in Asia (for more details see <http://www.google.com/about/datacenters/>). The location of the datacenters directly impacts application response times, capital and operational costs, and, from an environmental point of view, carbon dioxide emissions.

Selecting a location involves many important considerations, including its proximity to population centers, power plants, and network backbones; the source of the electricity in the region; the electricity, land, and water prices; the average temperatures at the location (which has a direct impact on the PUE); and finally, also taxes and one-time government incentives. As there can be many potential locations and many issues to consider for each of them, the selection process can be extremely involved and time-consuming. Large Internet companies, such as Microsoft and Google, have their own processes for selecting locations for their datacenters, which they keep confidential.

Iñigo Goiri, Kien Le, Jordi Guitart, Jordi Torres, and Ricardo Bianchini “Intelligent Placement of Datacenters for Internet Services” presented a paper at ICDCS 2011 that proposed a framework for datacenter placement. The authors provide a very detailed cost model, formulate the datacenter location as optimization problems solved by various heuristics, and moreover illustrate the selection tradeoffs quantifying the minimum cost of achieving different response times and availability levels. Finally, the framework can also restrict the selection to locations where green energies are available or chiller-less datacenters can be built.

Networks and Distributed Systems by Guillaume Jourjon

Continued from page 2

On the other hand, you will need to have one A/C unit per site (4.8kW) which in our case is three extra unit compared to the other scenario if we discard the common one to both scenarios. Then you will need to setup the DiffServ network and finally you will consume power for the traffic (21kW for a 10Mbit/s encoding). As the course lasted for 24 hours we have a consumption, regardless of the number of students, of:

$$C_{\text{tele}} = 436.6\text{kWh}$$

Furthermore, we can compute that we need a course of more than 28 hours to make one 1000km trip more efficient in terms of energy than the tele-teaching.

Conclusion: Tele-teaching used less energy than the on-site approach. We can now ask the real question: What is the product of teaching? In a recent blog post, I questioned the current trend in sustainable computing. How do we assess the value of a (so-called) “green solution?” In particular, I wondered what could be considered to be the measurable product of a distributed service. Here, the distributed service is the teaching/learning content, and usually in order to evaluate teaching a post-course questionnaire is given to students in which they assess the course. Unfortunately, we were unable to obtain enough data to reach statistical significance. In my opinion, tele-teaching is the way to go, especially when you have courses over several campuses, but there are still new methods and tools being developed.

STC Updates

by Giuliano Casale, Imperial College



Report from the chair: With our monthly newsletter now up and running, our focus is shifting to other topics. We have followed up with IEEE CS about technically cosponsoring conferences and workshops. Over the next several months we intend to test out the process. At the beginning of this month we wrote a progress report on the activities of our STC thus far and shared it with the chairs of the other STCs. The intent of this report is to share our lessons learned, so that others can leverage what worked well for us, while avoiding things that did not work as anticipated. Lastly, we are exploring options to ramp up our membership growth. We have been challenged to surpass 100 members as quickly as possible. We ask all current members to please encourage their colleagues in the field of sustainable computing to join our STC.

- Martin Arlitt

Membership: 64

Report from Secretary/Treasurer (Giuliano Casale):

- Collected officers' activity reports and prepared monthly STC report.

Report from Conferences Chair (Diwakar Krishnamurthy):

- Set up collaborations with SuperComputing, GCM and CNSM. Circulated newsletter to these communities.
- Approached several other conferences and workshop organizers.

Report from Academic Chair (Niklas Carlsson):

- Sketching on a potential short-feature that may help show case academic and industry groups/people. This effort will be coordinated with Canturk Isci.

Report from Membership Chair (Anirban Mahanti):

- Facebook group has 35 members, Linked-In group has 25, 13 new members from website.
- Circulated invitations to colleagues, this resulted in doubling of the membership count.

Report from Communications Chair (Abhishek Chandra):

- Continued to identify conferences, workshops and journals relevant to sustainable computing.
- Prepared a spreadsheet for inclusion in the monthly newsletter and website.

Report from Policies and Procedures Chair (Stephen Dawson):

- Documenting implemented STC-SC processes and gathering requirements for future processes.

Report from Industry Chair (Canturk Isci):

- Worked on narrowing down the IBM industry feature, identifying the potential focus projects.
- Explored new industry-academic programs related to sustainable computing.

Report from the Newsletter Editor (Christopher Stewart):

- Setting up the voting process for Pick of the Month series
- Looking for short contributions to the newsletter

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	Notification
ICT-GLOW'12	ICT against Global Warming Collaborative and Autonomic Green computing	Vienna, Austria	Sep. 3-7, 2012		Mar. 23, 2012	May 15, 2012
CAGing	Green Communications	Toulouse, France	Jun. 25-27, 2012		Mar. 26, 2012	Apr. 16, 2012
GreenCom	Sustainable Internet	Online	Sep. 25-28, 2012		Apr. 3, 2012	Jun. 15, 2012
SustainIT	Supercomputing	Pisa, Italy	Oct. 4-5, 2012		Apr. 16, 2012	Jun. 30, 2012
SC'12	Energy-aware HPC	Salt Lake City, UT	Nov. 10-16, 2012	Apr. 20, 2012	Apr. 27, 2012	
EnA-HPC 2012	Sustainable computing	Hamburg, Germany	Sep. 12-14, 2012		Apr 29, 2012	May 21, 2012
GreenMetrics 2012	Green Communications and Networking	London, UK	Jun. 11, 2012		May 1, 2012	May 15, 2012
GreeNETS 2012	Computer Systems	Gandia, Spain	Oct. 24-26, 2012		May 1, 2012	Jun. 11, 2012
OSDI'12		Hollywood, CA	Oct. 8-10, 2012		May 3, 2012	Jul. 10, 2012

Journal and Special Issue Call For Papers

Journal	Papers Due	Notification
Sustainable Computing	(Open)	

Conference, Workshop & Symposium Call for Participation

Short Name	Main Topic	Location	Dates
CCSES'12	Green networking and Smart Grids	Orlando, FL	Mar. 30, 2012
EuroSys 2012	Computer Systems	Bern, Switzerland	Apr. 11-13, 2012
WWW 2012	World Wide Web	Lyon, France	Apr. 16-20, 2012
NSDI'12	Networked Systems Design and Imp.	San Jose, CA	Apr. 25-27, 2012
CCGrid 2012	Cluster, Grid and Cloud Computing	Ottawa, Canada	May 13-16, 2012
IPDPS'12	Parallel and Distributed Systems	Shanghai, China	May 21-25, 2012
IGCC'12	Green Computing	San Jose, CA	Jun. 5-8, 2012
Sigmetrics'12	Measurement and Modeling	London, UK	Jun. 11-15, 2012
GCC'12	Green Communications and Networking	Ottawa, Canada	Jun. 15, 2012
HPDC'12	High Performance Distributed Computing	Delft, Netherlands	Jun. 18-22, 2012

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

To advertise a relevant venue, email Abhishek Chandra at 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 computing applications with ecological and societal impacts.

Submit your article at <http://ees.elsevier.com/suscom/>