

Sustainable C[★]omputing

REGISTER

Monthly bulletin 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 2 Issue 2

Table of Contents

- 2 Letter From the Editor
- 2 Call for Pick of the Month Nominations
- 3 Pick of the Month: Greening Geographical Load Balancing
- 5 Column: Resource and Energy Management
- 6 STC Updates
- 7 Upcoming Events

Sponsored by:



From the Editors Desk

by Christopher Stewart, The Ohio State University



In 1968 and 1969, the Boston Celtics were back to back champions of the National Basketball Association (NBA). Perhaps surprisingly, their coach, Bill Russell, was in only his second season as a coach in 1968. Even more shocking, coach Russell also played center. He is the only player-coach in NBA history to win a championship.

Well, I am now a second-year editor for the Sustainable Computing Register. I want to do more than edit. Starting in the next issue, I will contribute a regular column on economic and policy incentives for sustainable computing. While my business cards say distributed systems, economics and policy have always interested me. I doubt that I will be as prolific of a writer as Bill Russell was a center, but I hope you all will take a look at my columns from time to time. Sustainable systems, like no other technology-related field, are heavily influenced by national and global policies.

Also, I want to remind all members to get involved with Pick of the Month voting. Our process changed this year to increase participation. Anyone reading this letter can visit <http://www.cse.ohio-state.edu/~cstewart/potm> and vote. The process is set up to encourage self-interested behavior. Specifically, I want to highlight the following incentives: First, you should nominate your own papers. Second, you should vote for your own papers (no rule against that). Third, you should cajole your colleagues to vote for your papers (stuff the ballot box). Fourth, you should cajole other voters to vote for your papers (explicitly encouraged via a comments section). Finally, if your paper wins, you should vote other really good papers (keeping the prestige of the award high).

The Pick of the Month highlights great work in the field of sustainable computing because we publish in a diverse range of conferences. We need your participation and passion.

- Chris

Call for Nominations: Pick of the Month

The Pick of the Month series highlights top research papers and projects within our community. It is community-driven award with an open nomination process and selection.

Anyone can nominate a paper by simply sending an email to either Prof. Christopher Stewart or Prof. David Chiu. Then vote and converse about your nomination online by:

1. Visiting the URL below and clicking through to the Pick of the Month survey
2. Posting (anonymous) comments describing why the paper deserves pick of the month recognition. You can even comment about why one nominee is better than another.
3. Every month, nominees with the most votes are eligible for Pick of the Month. SC STC Information Officers will decide among eligible candidates.

Try it out now. <http://www.cse.ohio-state.edu/~cstewart/potm>

PICK OF THE MONTH April 2013

Greening Geographical Load Balancing by Zhenhua Liu, Minghong Lin, Adam Wierman, Steven Low, and Lachlan Andrew published in the Association for Computing Machinery's International Conference on Measurements and Modeling of Computer Systems (SIGMETRICS)

We resume our Pick of the Month series with Greening Geographical Load Balancing, a SIGMETRICS 2011 paper that has been incredibly impactful. The paper builds on a now widely accepted trend: When datacenters scale out, the spread out across geographically diverse locations. Amazon, Google, and Microsoft have datacenters in almost every corner of the globe. Even large enterprise companies in banking and insurance are operating datacenters multiple locations.

Geographically distributed locations offer the benefits of diversity. One location is often unlike another. At the time that our pick of the month was published, datacenter managers knew that they could exploit different geographic and temporal pricing differences. The Pick of the Month showed that 1) naively routing work for low cost could increase energy and carbon footprints and 2) wisely routing work for low cost and low energy keeps costs low in a sustainable fashion. Note, the word "showed" understates the paper's rigor. In this issue, we sit down with the lead authors, Zhenhua Liu and Minghong Lin.

Christopher Stewart: This work is becoming one of the most cited works in the field of sustainable computing. When I looked, Google Scholar estimated over 70 citations in just 2 years. What are your thoughts about the community's response to this work? Why has it been received so well?

Zhenhua and Minghong: We are really happy to see that the community appreciates our efforts. The idea of using geographical load balancing to extract benefits from heterogeneity is widely used in subsequent papers, e.g., to reduce power cost and/or GHG emissions, to ease the incorporation of renewable energy. Some following work also extended our model, which is very exciting. The reason why it has been received so well is, in our mind, that we tried our best to provide a solid model together with practical distributed algorithms. Together with the GreenMetrics paper "Geographical load balancing with renewables" (best student paper), we provided numerical results to quantify the benefits and identify the potential of geographical load balancing in demand response, as verified by subsequent works.

Christopher Stewart: OK. We'll publish the plug to your GreenMetrics paper. (laughing) Mainly because it is good work as well and is also a pick of the month nominee. Let me ask about the process of innovation for you guys. As the title suggests, this paper builds on prior work on geographical load balancing by adding a "green" element. The key observation is that load balancing for low energy costs actually wastes energy. To me, this is a really clever observation. How did you uncover this problem?

Continued on page 4

Pick of the Month continued from page 3

Zhenhua and Minghong: This was first uncovered from our model. By constructing a reasonable model and characterizing the problem, we found this seemingly surprising result. Interestingly, we noticed that similar phenomenon has been revealed in other area, e.g., Jensen’s paradox, which states the demand, is expected to increase if the cost decreases. We then conducted numerical simulations to validate our observation. Therefore, we think mathematical models are very useful to provide insight into system design.

Christopher Stewart: Students in my special topics courses often miss the importance of distributed solutions. Can you explain the virtues of distributed load balancing compared to a centralized approach?

Zhenhua and Minghong: Generally, distributed load balancing algorithms can provide better robustness, faster response, lighter communication and synchronization overhead, and other benefits. The benefits are more significant for a system with a large number of points, e.g., Akamai. Even for Google, from our collaborations with them, distributed solutions are usually preferred if possible and they actually make a huge effort on this and sometimes they are pushed by the system scale to transit to distributed solution. From the theory side, our distributed algorithm extended existing gradient-based method and guaranteed convergence even when the objective function is not smooth enough.

Christopher Stewart: Your paper includes an entire subsection labeled “The Importance of Dynamic Pricing.” Your paper helped me understand the importance of dynamic pricing in sustainable computing systems. Can you talk a little about this mechanism and why it is useful?

Zhenhua and Minghong: Before we have some emission regulation or carbon tax, it is really hard, if not impossible, for data center operators to focus on reducing emissions other than electricity bill, although some big companies such like Apple and Google did this due to increasing public pressure and future energy security consideration. Among different options, e.g., coincident peak charging, or ancillary market participation, one possible approach is real-time dynamic pricing. In our paper, we illustrated by numerical simulations that the non-renewable energy usage could be significantly reduced if real-time electricity pricing was used proportionally to the instantaneous non-renewable energy generation. We also pointed out if the pricing was not set in this way, the reduction can be limited, which was later supported by the Sigcomm 2012 paper “It’s not easy being green”. Therefore, we really hope the transition to smart grid can provide more incentives for customers such as data centers to utilize their demand flexibilities to act as “virtual storages” through various demand response programs.

Christopher Stewart: Thank you both for the excellent paper and for this insightful interview. I fully expect to interview you again since this paper is part of bigger body of excellent work. Congratulations.

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.

Resource and Energy Management

by Danilo Ardagna, Politecnico di Milano



This month's feature is an overview of the PREFigure (Performance, Power, and Reliability Framework for Disk Drives) framework for managing disk reliability and power.



Contributing authors from left to right: Feng Yan, Xenia Mountrouidou, Alma Riska, and Evgenia Smirni

PREFigure: As data centers grow bigger and central to many service providers and enterprises, storage systems have drawn attention. Despite the emerging new storage technologies such as solid state drives (SSDs), hard disk drives (HDDs) still host the overwhelming majority of enterprise and consumer data. Specifically, hard disk drives continue to store most of the aging data, which are expected to grow in size over the years. While data center power consumption amounts to about 20% of all IT power consumption, as much as 20% to 40% of data center power consumption is attributed to disk drives. Because not all data in a vast data center is accessed simultaneously, a compelling approach for reducing power consumption in data centers is to spin down idle hard disk drives. This approach is routinely deployed in storage systems that serve as archival or backup systems and is being exploited even in high-end computing environments.

Spinning down disk drives to save energy is a challenging open problem because it can impact user performance and disk reliability. If requests arrive while the disk is still in a power saving mode, they can be delayed by the disk reactivation process. In addition, even though the disk drives are under utilized, the idle times can be highly fragmented. Idle periods that are long enough for efficient power savings are limited. Finally, every power up/down wears out the disk drive. As a result, disk drives should be put into power saving modes only for a limited number of times to prevent premature wear out.

Common practice methods try to address these challenges by idle waiting for a fixed amount of time or guide scheduling decisions based on the past utilization. However, none of these methods provides performance guarantees or takes into consideration disk reliability. PREFigure overcomes these short-comings. In its core, there is a framework that uses as input user- or system-level constraints such as the number of allowable power ups/downs of a disk within a time period (strict constraint) and the user acceptable potential performance degradation of future IOs (soft constraint), and estimate the projected power savings as well as provide a strategy on when and for how long these power savings should be scheduled. The histogram of past idle times plays a key role in PREFigure as it projects future power savings based on statistical information that is monitored or extracted from this histogram. Probabilistic interpretation of all of the above information leads PREFigure to define robust schedules for power saving modes. PREFigure captures the workload changes in the system by updating the histogram of idle times as well as information about the sequence of idle times. Such updates enable the adjustment of the schedules of power saving activation to the workload dynamics.

We have evaluated PREFigure on enterprise disk drive traces with a wide range of idleness characteristics. We have seen excellent agreement between the results from PREFigure's analytic estimations and the trace driven simulations, suggesting it has practical benefit.

STC Updates

By Danilo Ardagna, Politecnico di Milano



Membership: 430

Officers reports from April 1, 2013 to April 29, 2013

Report from Secretary/Treasurer (Danilo Ardagna):

- Collected officers' activity reports and prepared monthly STC report

Report from Conferences Chair (Diwakar Krishnamurthy):

- Solicited collaborations from organizers of DCPERF, USENIX ATC, HPDC, ICS, SIGMETRICS, ISCA, WETICE, IGCC, GCN, and IWQoS
- Continuing to update list of conferences related to STC

Report from Academic Chair (Niklas Carlsson):

- Identified and contacted candidates for future community highlight features

Report from Membership Chair and vice-Chair (Sergey Blagodurov, Matthew Forshaw):

- Counted the number of members every week for the past month.
- Sending invitations to potential STC-SC members and continuing working on the new invitation list

Report from Communications Chair (Abhishek Chandra):

- Prepared the upcoming events for announcement with the communications vice chair

Report from Policies and Procedures Chair (Stephen Dawson):

- Contacting colleagues to find a vice-chair

Report from Industry Chair (Canturk Isci):

- Working with the academic chair to identify the next set of community highlights
- Starting work on a next industry highlights feature

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

- Contributed material for newsletter and blogs

Report from the Newsletter Editor (Christopher Stewart):

- Produced April newsletter
- Set up public website for Pick of the Month

Report from the Web Master (Yan Shvartzshnaider):

- Added a resource page including links to relevant websites and interesting articles

Upcoming Events

By Abhishek Chandra, University of Minnesota
and Bhuvan Urgaonkar, Penn State University

Conference, Workshop & Symposium Call For Papers

Short Name	Main Topic	Location	Dates	Papers Due	Notification
OnlineGreenComm	Green solutions for and based on communication technologies		Oct. 29-31, 2013	May 10, 2013	Jul. 26, 2013
CGC	Cloud and Green Computing	Karlsruhe, Germany	Sep. 30-Oct. 2, 2013	May 12, 2013	Jul. 10, 2013
IEEE BigData	Big data	Santa Clara, CA, USA	Oct. 6-9, 2013	Jun 2, 2013	Aug. 10, 2013
INFOCOM	Computer Networking	Toronto, Canada	Apr. 27-May 2, 2014	Jul. 21, 2013	Nov. 25, 2013

Journal and Special Issue Call For Papers

Journal		Papers Due	Notification
Sustainable Computing			
Elsevier COMCOM	Special issue associated with SUSTAINIT 2013	(Open)	
Elsevier PMC	Special issue associated with SUSTAINIT 2013		

Conference, Workshop & Symposium Call for Participation

Short Name	Main Topic	Location	Dates
HotOS	Operating Systems	San Ana Pueblo, NM, USA	May 13-15, 2013
E2DC	Energy-efficient Data centers	Berkeley, CA, USA	May 20, 2013
e-Energy	Computing and Communication for Smart Grid, Energy-efficient Computing/Comm.	Berkeley, CA, USA	May 22-24, 2013
HPDC	High performance distributed computing	New York City, NY, USA	Jun 17-21, 2013
Sigmetrics	Performance Modeling	Pittsburgh, PA, USA	Jun 17-21, 2013

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

STC-SC Officers

Chair: Anirban Hahanti, NICTA	Policies and Procedures: Stephen Dawson, SAP
Secretary/Treasurer: Danilo Ardagna, Politecnico di Milano	Industry Chair: Canturk Isci, IBM
Conferences: Diwakar Krishnamurthy, University of Calgary	Editor: Christopher Stewart, Ohio State University
Conferences: Amarjeet Singh, IIT-Delhi	Editor: David Chiu, Washington State University Vancouver
Academics: Niklas Carlsson, Linköping University	Information: David Carrera, UPC BarcelonaTech
Membership: Sergey Blagodurov, Simon Fraser University	Information: Danilo Ardagna, Politecnico di Milano
Membership: Matthew Forshaw, Newcastle University	Information: Fan Dongrui, Inst. of Computing Technology
Communication: Abhishek Chandra, University of Minnesota	Information: Guillaume Jourjon, NICTA
Communication: Bhuvan Urgaonkar, Penn State	Web Master: Yan Shvartzshneider, University of Sydney