



Engineers or Scientists?

Michael Rabinovich • Case Western Reserve University • misha@eecs.cwru.edu

To everyone who's been looking forward to another one of Fred Douglass's brilliant opening columns, I'm the first to admit to being disappointed by its absence. However, as an editor in chief can serve at most four years, it was time for Fred to get his life back. With this column, I'd like to introduce myself as *IEEE Internet Computing's* new EIC. I thank the search committee and *IC's* editorial board and staff for entrusting me with this role. I'm looking forward to a productive tenure. I realize that Fred will be a tough act to follow. During his time at the helm, the magazine has become one of the most impactful IEEE publications. In fact, its impact factor (a metric based on the citations and downloads that IEEE uses to assess its publications) has risen steadily and now places *IC* third out of all IEEE Computer Society magazines and 13th out of IEEE's 125 publications. The only consolation is that Fred will remain on the editorial board, and I know I can always count on his help and candid advice.

Internet computing is a wide and diverse field, and we've all become increasingly specialized in select areas. My areas of interest revolve around Internet performance and infrastructures. With this inaugural column, I'd like to share some thoughts on our field from this perspective.

My university, Case Western Reserve, is fortunate to have Lev Gonick, a prominent visionary, as its CIO. One of Lev's latest initiatives is a test deployment of gigabit Internet connectivity to a number of homes around the university. As we brainstormed research opportunities offered by this deployment, we discussed applications that could possibly utilize this capacity. High-definition TV? That's roughly 8 Mbps, and even the most avid TV viewer would be hard-pressed to consume 125 simultaneous streams. Home surveillance and remote healthcare diagnostics?

Same story. So, we thought it might be handy to be able to download a full-length movie in a few seconds, although this would only consume the full bandwidth for short peak periods. This tremendous upload capacity might also spur further increases in peer-to-peer content delivery. (In fact, we're already aware of one enterprising student who's using this new capacity to redistribute third-party content for a small profit.)

Perhaps homes of the future will have those floor-to-ceiling wall monitors we see in science-fiction movies and some research labs (for example, www.cs.princeton.edu/omni-media); these monitors might use this bandwidth to display life-size high-definition video streams, possibly in 3D. Or immersion rooms will be common, where people can go for a quick virtual trip to a location of their choice. Judging from the past, I have faith that applications will emerge and gobble up whatever network capacity is available. Regardless, what's clear is that this new access capacity will shake up fundamental assumptions about the network we currently take for granted (for example, the last-mile bottleneck and over-provisioned core) and will force rethinking of many mechanisms designed with these assumptions in mind. And this brings up some general questions about our profession.

Internet researchers who are worth their salt have faced the frustrating experience of having to reconcile their wonderful ideas within the confines of the existing technological landscape. I'd venture to say that more brainpower has been spent on tricking the Internet to allow particular techniques than on inventing the new techniques themselves. Consider, for example, the email spam problem. Several protocols have been proposed that would largely eliminate this problem, ranging from technolo-

gies such as Yahoo's DomainKeys, which get rid of email spoofing, to various postage schemes that would make spam economically infeasible. Yet, in the absence of universal adoption of these schemes, researchers are making great efforts to blunt spam's impact through palliative technologies such as filters. But can we view as research the efforts so intimately dependent on particular Internet realities, however inventive these efforts might be? Shouldn't researchers be concerned with fundamental principles and not quirks?

In the past few years, in an effort to sidestep Internet legacy restrictions, government agencies and institutions in the US, Europe, and Japan have been pursuing parallel programs to encourage clean-slate ideas for the Internet of the future (see www.nets-find.net, [\[-internet.eu\]\(http://-internet.eu\), and \[http://nag.nict.go.jp/topics/AKARI_fulltext_e_translated_version_1_1.pdf\]\(http://nag.nict.go.jp/topics/AKARI_fulltext_e_translated_version_1_1.pdf\)\). This has certainly invigorated off-the-wall thinking. However, despite some ideas that expressly attempt to allow more flexibility in the Internet architecture, the Internet is likely to remain dependent on system-wide conventions and standards. Thus, a periodic splash of clean-slate activity won't resolve the general tension between practically impactful, but incremental, innovations that are firmly grounded in the current technological context and abstract ideas that will never leave the pages of the articles that describe them.](http://www.future</p>
</div>
<div data-bbox=)

Yet, I think it should be rather obvious that this tension is artificial as far as research is concerned: both types of innovations have their place in our labs as well as in our publica-

tions, both types of activity require creativity, and both can produce approaches that can be either elegant or unwieldy. One activity type can improve people's lives and the other can enrich their knowledge. This matter wouldn't warrant mentioning if not for two reasons: first, its influence on peer reviewing, in which worthy submissions regularly die after being labeled either as engineering or as divorced from reality; second, the often-confused line between the two activity types. The latter reason is a subtle but important one. Many of us have come across proposals that spend much effort to comply with some legacy constraints of the Internet but not others. Without a clear line between the two activity types, I believe any such work would require a clear justification for the constraints we choose to

IEEE Internet Computing

Editor in Chief

Michael Rabinovich • misha@eecs.cwru.edu

Associate Editors in Chief

M. Brian Blake • mb7@cse.nd.edu
 Siobhán Clarke • siobhan.clarke@cs.tcd.ie
 Maarten van Steen • steen@cs.vu.nl

Editorial Board

Virgilio Almeida • virgilio@dcc.ufmg.br
 Elisa Bertino • bertino@cerias.purdue.edu
 Azer Bestavros • best@cs.bu.edu
 Vinton G. Cerf • vint@google.com
 Fred Douglas* • f.douglas@computer.org
 Schahram Dustdar • dustdar@infosys.tuwien.ac.at
 Stephen Farrell • stephen.farrell@cs.tcd.ie
 Robert E. Filman* • filman@computer.org
 Juliana Freire • juliana@cs.utah.edu
 Carole Goble • cag@cs.man.ac.uk
 Michael N. Huhns • huhns@sc.edu
 Barry Leiba • barryleiba@computer.org
 Samuel Madden • madden@csail.mit.edu
 Cecilia Mascolo • c.mascolo@cs.ucl.ac.uk
 Pankaj Mehra • pankaj.mehra@ieee.org
 Chris Metz • chmetz@cisco.com
 Dejan Milojićić • dejan@hpl.hp.com
 George Pallis • gpallis@cs.ucy.ac.cy
 Charles J. Petrie* • petrie@stanford.edu
 Gustavo Rossi • gustavo@lifa.info.unlp.edu.ar
 Amit Sheth • amit.sheth@wright.edu
 Munindar P. Singh* • singh@ncsu.edu
 Oliver Spatscheck • oliver@spatscheck.com
 Torsten Suel • suel@poly.edu
 Craig W. Thompson • cwt@uark.edu

Shengru Tu • shengru@cs.uno.edu
 Doug Tygar • tygar@cs.berkeley.edu
 Steve Vinoski • vinoski@ieee.org
 * EIC emeritus

CS Magazine Operations Committee

Dorée Duncan Seligmann (chair), Erik Altman, Isabel Beichl, Krish Chakrabarty, Nigel Davies, Simon Liu, Dejan Milojićić, Michael Rabinovich, Forrest Shull, John R. Smith, Gabriel Taubin, Ron Vetter, John Viega, Fei-Yue Wang, Jeffrey R. Yost

CS Publications Board

David A. Grier (chair), Alain April, David Bader, Angela R. Burgess, Jim Cortada, Hakan Erdogmus, Frank E. Ferrante, Jean-Luc Gaudiot, Paolo Montuschi, Dorée Duncan Seligmann, Linda I. Shafer, Steve Tanimoto, George Thiruvathukal

Staff

Lead Editor: Jennifer Gardelle, jgardelle@computer.org
 Content Editor: Rebecca Deuel-Gallegos
 Manager, Editorial Services: Jenny Stout
 Staff Editor: Dennis Taylor
 Publications Coordinator: internet@computer.org
 Production Editor/Webmaster: Monette Velasco
 Contributors: Cheryl Baltes, Greg Goth, Keri Schreiner, Joan Taylor, and Alex Torres

Director, Products & Services: Evan Butterfield
 Senior Manager, Editorial Services: Lars Jentsch
 Manager, New Media & Production: Steve Woods
 Senior Business Development Manager: Sandy Brown
 Membership Development Manager: Cecelia Huffman
 Senior Advertising Supervisor: Marian Anderson, manderson@computer.org

Technical cosponsor:



IEEE Internet Computing
 IEEE Computer Society Publications Office
 10662 Los Vaqueros Circle
 Los Alamitos, CA 90720 USA

Editorial. Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *IEEE Internet Computing* does not necessarily constitute endorsement by IEEE or the IEEE Computer Society. All submissions are subject to editing for style, clarity, and length.

Submissions. For detailed instructions, see the author guidelines (www.computer.org/internet/author.htm) or log onto *IEEE Internet Computing's* author center at ScholarOne (<https://mc.manuscriptcentral.com/cs-ieee>). Articles are peer reviewed for technical merit.

Letters to the Editors. Email content editor Rebecca Deuel-Gallegos, rdeuel-gallegos@computer.org.

On the Web. www.computer.org/internet/

Subscribe. Visit www.computer.org/subscribe/.

Subscription Change of Address. Send requests to address.change@ieee.org.

Missing or Damaged Copies. Contact help@computer.org.

To Order Article Reprints. Email internet@computer.org or fax +1 714 821 4010.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

uphold and those we decide to ignore.

A related issue with a large impact on how we do research is the Internet's fast-moving nature, which gave us the notion of a "hot area." Hot areas change rapidly, and Internet researchers can rarely hope for a straight trajectory in their work. The Internet is full of examples in which a hot area "cools" almost to oblivion, only to come raging back. For example, take the perennial issue of "thick" or "thin" clients. The pendulum swung from mainframes (and "dumb terminals" – the thinnest clients) to PCs to an ill-fated attempt at swinging back with Oracle-led network computers (does anyone remember those?). Now, the pendulum is swinging back vigorously with cloud computing. In fact, some of the arguments for or against each of these paradigms revolve around similar issues, such as ease of administration, user control, economy of scale, and security. Or consider the area of multiprocessor interconnects. This area went out of vogue some time ago but recently came back with a vengeance in the form of datacenter networking. Yet another example is provided by content delivery networks (CDNs). Research interest in CDNs declined

dramatically after the dot-com bust but is flourishing again after the revival of the CDN industry.


These examples notwithstanding, I wouldn't suggest that you should just sit and wait for your favorite area to come back – this would be a high-risk proposition! In fact, most of us who have been at it for a while enjoy the fast pace of our profession and consider it as part of the fun. But, if you're just starting out, it might cause some anxiety: "What if my dissertation area grows stale by the time I graduate? How will I ever get a job?" Well, if worrying is your thing, your advisor can probably point you to some areas with longer horizons. In my field, if there were two areas that could be safely projected to remain important for the foreseeable future, they would be network security and measurements.

The security area will be kept alive by constant competition between defense technologies and ever-more sophisticated attackers. In fact, many believe security problems are going to grow worse before they get better. Increased bandwidth and CPU capacities allow attackers to launch wider attacks with fewer resources. In the short-to-medium term, this will motivate continued improve-

ments in our defenses against classic types of attacks such as spam, various denial-of-service attacks, and phishing. Longer term, this issue will be among the main drivers for re-architecting the Internet.

The evolving networking landscape requires continued measurements because new technologies entail new properties and behaviors that must be measured, characterized, and monitored. Network measurements and monitoring are becoming more challenging as networks increase in size, diversity, and complexity. And, while becoming more difficult to obtain, high-quality measurements will grow in importance because they support other trends such as virtualization and network convergence.

I'd like to conclude these thoughts by a truism that we're all very fortunate to be working in such a rich and still-young field as the Internet. We can choose to float between different areas without the fear of being typecast by our peers, or we can work in areas with longer horizons. We can go back and forth between practical and "out there" research. As I argued earlier, there's no (or shouldn't be) tension between these two types of activities as long as researchers are cognizant of which kind of activities they're pursuing. An important side-effect of the government's focus on the future Internet is the popularization of the term "clean-slate design." By giving this type of research a name, the community has both legitimized it and created a sort of ontology to help researchers better frame their efforts. Now all we have to do is produce good ideas, however they might be labeled, and have some fun while doing it. □

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.



IEEE Intelligent Systems
THE #1 ARTIFICIAL INTELLIGENCE MAGAZINE!



IEEE Intelligent Systems delivers the latest peer-reviewed research on all aspects of artificial intelligence, focusing on practical, fielded applications. Contributors include leading experts in

- Intelligent Agents • The Semantic Web
- Natural Language Processing
- Robotics • Machine Learning

Visit us on the Web at www.computer.org/intelligent