

Cold Topics in Networking

Jon Crowcroft
The Computer Laboratory, University of Cambridge
Cambridge, UK
jon.crowcroft@cl.cam.ac.uk

ABSTRACT

A major contribution to global warming has been the number of new workshops publishing proceedings with the prefix *hot*. In this article, I propose that we counter this trend in an attempt to remain carbon neutral with a set of “anti-workshops” on *cold* topics.

We suggest a number of heuristics for detecting when a topic has gone cold, and give some examples of the application of these heuristics. Of course, some cold topics warm up again, and so there is a risk of over dampening in our heuristic. Over a long period, dynamic equilibrium should be assured, but from time to time, our scheme may prejudice against surprising results in boring areas of communications research. Nevertheless, it may leave room for more surprising results in interesting areas of research, which cannot be a bad thing.

Categories and Subject Descriptors

C [..]: 2.1 [Packet-switching networks],[Network Communication]

General Terms

General Terms: Algorithms, Performance, Design

1. WHAT MAKES IT CLEAR A TOPIC IS HOT?

Before we can define cold topics, we have to discuss hot topics (after all, in Real Life, cold is an absence of heat, not a presence of some other phenomenon).

The principle technique which we can use to spot a hot topic is to spot small groups of disgruntled researchers gathered in a major conference to discuss why their latest pet papers have not made it into the proceedings. These researchers have had some cool new idea (perhaps because of some disruptive funding, or else a retreat, or some other stealth mode intervention, such as inspiration, imagination etc), and now have to find an outlet for their work. The existing conference PCs (especially in Systems research) are far too critical to let in anything very new. Such new work would certainly have insufficient evaluation, not enough comparisons with other systems, and there might not be a business case for it¹. Thus is born *hot**, for some value of ***, a new playpen for these researchers.

¹I am being ironic here, in case this wasn't obvious

The outcome of this is that the small groups of researchers propose a new workshop (perhaps associated with the major conference) and call it *hot** for some value of ***, perhaps taken from the list {*green, oblivious, byzantine, botnet, automatic, incentive-compatible, strategy-proof, social network, peer-assisted*} .

2. WHAT MAKES IT CLEAR A TOPIC HAS GONE COLD?

One year after a topic was hot, it is mainstream. Sometimes it takes two years. This doesn't mean it has gone cold yet. To really decide if a topic is dead and gone, we need to take a holistic look at several factors:

1. Look at number of papers in 2nd tier conferences.

Simply looking at the number of papers in top level conferences doesn't tell you - maybe a topic is hard, or is rich in many possible approaches. What tells you is when the paper quality slips and lower tier conferences (mentioning no names) start to play “me too”.

So we want to take the rate of decrease in top conferences and the rate of increase in 2nd tier conferences - a gradient, if you like.

2. Look at delta in performance reported in papers. Early papers may shift performance by many times, or by 10s of %age points. Later on, we see schemes improving things by 0.00000001.

Of course, sometimes, later papers are repairing earlier excessive claims, so this on its own is insufficient to indicate a cold topic.

3. Look at for tools that automate the production of writing papers (NS2, PlanetLab, etc).

Computer scientists love to automate things, but of course if you can automate a piece of research, then it isn't research any more.

Once work in more than 50% of papers in an area is more than 50% automatically generated from a combination of pieces from other papers, you can be sure that the area is dead as a Dodo.

It has been remarked that my rules would apply to all of networking, or perhaps to all of Computer Science, or even all research. Also, that much (most?) good work is incremental in nature, and not at the high-temperature end of things. These remarks are assuming I am making a value

judgement about hot topics being necessarily better than warm topics, or that reviving the cold topic is a futile activity. I am not. I am merely observing that there are trends in topics, and that there are characteristics to these trends. You, gentle reader, may wish to be a dedicated follower of fashion, or even a fashion leader, or you may wish to contribute in a more sustained and consistent manner. That is your choice.

3. TOPICS CONSIDERED UNCOOL

On the basis of the above heuristic, and by judicious use of online bibliographic databases, I have concocted the following initial list of topics, once of interest to the ACM Special Interest Group in Communications, that should now be shunned:

DHT and Structured P2P I think we can safely say that we have all seen enough distributed hash tables to last a lifetime of meals of pork luncheon meat and spam. But more seriously, DHTs are now prime time, not only in use for super-peer organisation in operational file-sharing systems, including several commercial P2P-TV systems, but even for controlling really illegal things like the Storm Botnet.

Internet Coordinate Systems We know where the Internet is, and we don't need a compass to find it. More importantly, you cannot create an absolute euclidean coordinate system for the Internet since it isn't embeddable except approximately in nine dimensions². Thus systems like Vivaldi and Meridian are perfectly fine for finding the best (lowest latency) server out of a list, but do not tell you where to aim that cruise missile if you want to take out a particular data center. Triangle Inequality Violations pervade the Internet and no amount of trouble tickets will remove them.

Faster packet classification Too many people are working on making packets go through routers faster - this is a marginal business for sure, since any router algorithm for header processing that doesn't get packets through at line rate is really wasting the capacity (and power - see above) and therefore is going to get rejected for business case reasons.

BGP The Border Gateway Protocol should be shut down by the border police. It was probably cold when it was created. We have several proposals for replacements, and indeed an entire body of work telling us what is wrong. We even have a meta-replacement. Now is the time to deploy.

DoS Denial of service is undefined since IP provides only *best effort*. How can we deny that?

Spam Spam is only something old fogeys get as they are the only people still trying to use e-mail instead of some chat system or messaging on a social network.

Overlays Overlays were only invented as a disruptive idea in the "Looking over the Fence" report to make people think about doing things differently - that didn't mean one should do overlays (or DTN) seriously. No, really.

²I am starting to sound like Spock or Dr Who here

TINA The Intelligent Network Architecture is coming back as we see people struggling to justify the GENI costs, and proposing massively complex in-network smarts. Perhaps one might list the Knowledge Plane here as one of life's less successful *gedanken* experiments.

TCP+AQM Last time I looked there were over 5000 papers on TCP and AQM published. This cannot be sensible. Many of the papers used a version of NS TCP code that was buggy, so it is probable that the authors should all go back and check, and have an entirely new conference and journal on the topic of "Retractions of Erroneous TCPAQM papers".

Multicast No-one uses it much so we should declare it dead. Watch this space, though, as I believe multicast is about to become hot again any decade now, since several large IPTV deployments now use it, and people are starting to compare them with P2P-TV (see above).

Newarch New architectures will never get deployed (for sound business reasons), so we should declare network architecture a cold topic.

Self similarity, long range dependence, large deviations

The papers on this topic are starting to look remarkably like papers a few years ago on this topic, which means the temperature is falling fast, from one of my heuristics. Indeed, self-similarity of papers (and the use of temperature and entropy as metrics) is a sure-fire indication of frost.

MANET Mobile Ad Hoc Networking is a topic with a great future behind it. Now we have Mesh, and Metro WiFi, and 3G cellular in most of the world, why would anyone deploy a network that doesn't work very well? Actually, this reminds me - there may be another heuristic - when papers say how much better they are than previous work that had achieved 45% connectivity, when they achieve 53% connectivity, and fail to look at other systems that achieve 99.999% connectivity, you know you are approaching zero Kelvin in this area. I might want to include handovers here too.

Self Organising Mobile Wireless Sensor Networks Most sensors are not mobile. Most sensors sit on a device and can be configured. Ask an oil man if they would use a biologically inspired algorithm to detect they are about to have a well go haywire.

Small World Networks This is a scam perpetrated by physicists to get publications in Nature on the AS Topology of the Internet, the structure of the web and the nature of human society, as well as being an albatross around the neck of ecologists modelling bird-flight, and bee swarms.

We know there are really only three physicists in the world, so it isn't surprising that the worst case connected network has two of them at most six arms-length apart, unless they get around a table, or all get asked to the same fondue party. It is diametrically opposite random way-point and Poisson point processes in the unit plane with Voronoi diagrams, as a source of endless cute visualisations telling us nothing terribly surprising, and only rarely even useful.

4. REVENGE IS A DISH BEST SAVOURED COLD

It is very likely that I shall now never be asked back on to the SIGCOMM PC. That sure is a relief.

Someone has pointed out that if there is a really, really good New Architecture, then all the topics above become Hot again, which is fine. Others asked me why I haven't included their favourite topic on here, for example, scheduling, mobility models, handovers, privacy, capabilities, congestion control, or cross-layer optimisation. The reason for that is that I have this absolutely stunning idea for a cross-layer optimised privacy-enhancing capability based scheduling algorithm for handovers for users that follow real-world mobility patterns, which I am about to submit to the new Hotcross workshop.

The last word I want to say on the topic is that I do not want to hear that anyone rejected papers from top-tier conferences because they were about topics in my list above, and I do not want to hear what people think I believe to be 2nd tier conferences. After all, when you come up with your PhD thesis topic, it is certainly going to be red hot, but 3 years later when you are finishing your third paper and dissertation, it may not be so any more. Don't let this put you off: This article is (at least partly) for entertainment in the networking fashion industry.

5. ACKNOWLEDGEMENTS

Thanks to Christophe Diot for getting me to put his head above the parapet. Thanks to Richard Mortier for correcting some of the wilder accuracies, and suggesting the ideal subtitle, *futile but amusing*. Thanks to Michalis Faloutsos for words of advice about PhDs. Thanks to Serge Fdida for comments on fairness.