Reheating Cold Topics(in Networking)

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

ABSTRACT

It has been proposed that research in certain areas is to be avoided when those areas have *gone cold*. While previous work concentrated on *detecting* the temperature of a research topic, this work addresses the question of *changing* the temperature of said topics. We make suggestions for a set of techniques to *re-heat* a topic that has gone cold. In contrast to other researchers who propose uncertain approaches involving *creativity*, l͡ateral thinking and *imagination*, we concern ourselves with deterministic approaches that are guaranteed to yield results[1].

Categories and Subject Descriptors

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

General Terms

General Terms: Algorithms, Performance, Design

Keywords

Communications Systems Research.

1. WHAT MAKES IT CLEAR A TOPIC IS HOT?

In a previous article (`http://ccr.sigcomm.org/online/?q=node/305`) I proposed that it is possible to identify research topics that are *cold*, and gave several heuristics for making this determination. Here, I propose a family of techniques to reverse the process that caused a topic to go cold.

2. HOW TO REHEAT A COLD TOPIC?

When a topic is almost dead, it may not be completely dead and gone, and (as in the Princess Bride) it may be possible to bring it back from the brink. Here we look at systematic methods that can be used to re-vivify a topic, or at least to warm it up a bit.

Extrapolate or Invert Aficionados of Science Fiction literature (especially the classic works of John Wyndham and J.G. Ballard) will be familiar with the idea of taking some factor in life, and extrapolating it to an extreme - for example, what if one day it starts raining, and doesn't stop for a thousand years? What if there was only one gender, or eleven? What if it was not possible to fly?

In computing terms, we can ask the same questions. For example, what would happen if we had absolutely no caches anywhere? What if we had thin servers (like thin client idea but at the other end)?

Multi-disciplinize A second technique for adding spice to yesterday's research is to mix in some research from a different subject domain. What can we learn from botany, for example, when thinking about congestion control?

Back to basics From time to time, in its own clam pace, advances are made in Theory. it is well worth looking at these, as sometimes these advances, when understood and applied, rewrite *accepted wisdom*.

New technology New technology arrives, and then we figure out how to fit it in to the old paradigm. More rewarding is to see how to use the new technology in the most natural way leading to a paradigm shift. The replacement of WIMPs (Windows, Icons, Mice, Pointers) with new interaction modes (multi-touch screens, accelerometers, camera/gesture recognition) is revolutionising HCI right now. How could new technologies re-vitalise our view of communications research topics? I'm tempted to point people at the ever-inspiring XKCD (`http://xkcd.com/190/`) for one noteworthy example.

It has been remarked that these techniques would apply to all of networking, or perhaps to all of Computer Science, or even all research. This is clearly true. They would probably also work to automate much of Hollywood scriptwriting.

3. RE-HEATING THOSE UNCOOL TOPICS

Now we are in a position to demonstrate the techniques above, by applying them to the list of topics from the previous article:

DHT and Structured P2P Use of multi-disciplinarity might help here (technique 2, above): for example applying DHTs to placing passengers and bags on a plane, and

[1]Good research must be feasible, but novel. The tricks here are finding out what may have changed to make something novel, also now feasible, or vice versa.

being able to find and reconcile them again; or a structured approach to share trading, carbon trading, air traffic control, might be worth exploring.

Internet Coordinate Systems Technique 4 (new technology) is a clear winner here - a coordinate system for mobile devices would be really useful - being able to find all those remote controls lost around the house would be handy, and knowing where all the cars on the road really are is clearly a high value proposition. A specific instance of this would be RFID tag location.

Faster packet classification As networks get faster, we're starting (finally) to see proposals to deploy fiber-to-the-home, in Europe and North America. However, those agencies protecting us from the bad guys still need to look at all our packets. Hence we need all optical packet classifiers. This might be a little bit easier than all optical packet switching, and might be a useful step on the way to that laudable goal. Hence technique 4 (new technology - optical logic) applies.

BGP I actually cannot think of a way to save BGP as a topic. Likewise, ATM[2]. More seriously, there are several proposals afoot to use BGP to carry location/identity mapping information for IPv6 and LISP. one could also use BGP's filtering mechanisms to carry DDoS and Spam filter rules. All of these would stress BGP further and require fundamental research into replacing it (meta-routing?).

DoS Technique 1 applies nicely here - DoS is thought of as bad - Storm (a botnet often used to generate Denial of service attacks) is thought of as bad. But Storm is a very nice distributed computation system (cloud or grid computing)

Spam With the advent of youtube, we now get sent zillions of humorous clips - I predict that video spam is definitely the next big thing (technique 4). The use of rural BPO (villagers in India manually classifying content) to identify spam is one possibility. This would entail social science and economics research.

Overlays We can apply technique 1 and say, what if every packet spawned a new overlay - a meta-lay?

TINA This stood for The Intelligent Network Architecture. It was the antithesis of IP (dumb network, smart end systems), with all the intelligence for applications in the center. We could combine programming language research with network management, and so a multidisciplinary approach (technique 2) applies. Alternatively, we could try to re-implement a fully decentralised TINA using BAR-T (Byzantine Altruistic Rational Tolerant) techniques (back to basics).

TCP+AQM Technique 3 comes to the rescue - how would congestion control work securely in conjunction with AQM for multicast mobile nodes?

Multicast Technique 4 is applicable - how could we multicast Google advertisements, or Google's filter rules?

Newarch In the original article, I argued that Newarch would never get deployed. Since most research is never deployed, this is a poor argument against researching this topic. Indeed, almost any of the techniques for extending the lifetime and scope of an old topic amount to new architecture.

Self similarity, long range dependence, large deviations Traffic keeps changing. How self similar is Facebook traffic. What are the power laws in Twitter? Technique 4 (new technology) will always rescue this topic.

MANET MANET was a hot topic for a while mainly because of US military needs. However, MANET is already recognized as a viable way to organise file sharing between cars - the VANET is a good solid, and clueful, purposeful research. Indeed, being inter-disciplinary, applying social networking graph knowledge to the formation of MANETs (or, similarly, for VANETs, collecting data on how often you see other cars in real life using number plate recognition) would be cool.

Self Organising Mobile Wireless Sensor Networks This is an example of a topic that was simply too too kitchen-sink - all we have to do is remove any one buzzword, and we will have a respectable research area again.

Small World Networks Understanding how Facebook actually works as a system is a nice challenge for systems people. Hence, technique 4 is just fine here. Also, determining how to create or modify preferential attachment rules to cause a system to form as a small world would be interesting (technique 1, inversion).

4. THE PROOF OF THE PUDDING IS IN THE EATING

So there you have it: a surefire way to find a new research topic in the dying embers of an old topic. Of course, this does mean that we expect to see a whole new era of research open up, and that the standards for papers to get accepted for SIGCOMM will be higher than ever. Good. On the other hand, there is nothing so satisfying as being the first person in a brand new topic, and taking the risk that no-one will accept or even understand why you are working on it. Take risks.

5. ACKNOWLEDGEMENTS

Thanks to David Cottingham and Christophe Diot for feedback on some of the calories and joules expended here. Thanks to Keshav for editorial guidance.

[2]short of relabelling it as MPLS