Technical Report

UCAM-CL-TR-554
ISSN 1476-2986

Number 554

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Towards a field theory for networks

Jon Crowcroft

January 2003

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/TechReports/

Series editor: Markus Kuhn

ISSN 1476-2986

Towards a field theory for networks

Jon Crowcroft

Abstract

It is often claimed that Internet Traffic patterns are interesting because the Internet puts few constraints on sources. This leads to innovation. It also makes the study of Internet traffic, what we might cal the search for the *Internet Erlang*, very difficult. At the same time, traffic control (congestion control) and engineering are both hot topics.

What if "flash crowds" (a.k.a. *slashdot*), cascades, epidemics and so on are the norm? What if the trend continues for network link capacity to become flatter, with more equal capacity in the access and core, or even more capacity in the access than the core (as in the early 1980s with 10Mbps LANs versus Kbps links in the ARPANET)? How could we cope?

This is a paper about the use of field equations (e.g. gravitational, electrical, magnetic, strong and weak atomic and so forth) as a future model for managing network traffic. We believe that in the future, one could move from this model to a very general prescriptive technique for designing network control on different timescales, including traffic engineering and the set of admission and congestion control laws. We also speculate about the use of the same idea in wireless networks.

1 Introduction

There are two dominating reasonable assumptions in much of the work about congestion control and scheduling of traffic in the Internet which we question in this paper:

Traffic Matrix Assumption The first is that traffic is TCP and that applications usage forms a mix of "mice" and elephants. Elephants are long lived, thus research mainly concentrates on steady state behaviour of TCP. Mice are very short lived, but numerous. The behaviour of a large collection of mice is mainly taken as following typical human session arrival statistical behaviour.

Hierarchical Network Assumption The second assumption is that the Internet is hierarchical in the sense that access networks are very much slower than the "core" or cores.

Indeed, some research is directed at "making it so": Recent work on Internet Measurement and provisioning has been motivated by what one might call the search for the Internet Erlang model. As with the PSTN, implicit in thinking is a traffic matrix that has a preference for short lived, and localized calls, over longer, and long-distance - this could be re-enforced by pricing, creating a self-fulfilling prophesy.

End-to-end purists assert (rightly, we believe) that the traffic matrix is the wrong model - to do useful things (like design nets, provision links, design congestion control and admission control schemes and multi-metric multi-path routing schemes etc etc) we want to retain the new model of "networks" - one that captures the dynamics, but is sufficiently simple to be useful. Steady state TCP behaviour, on an hierarchical network provides this.

However, as the growth in the network levels off in the first world, we predict that the "death of distance" will grow in importance as "core" providers lose the ability to deploy capacity faster than fibre access network providers.

We believe that two things will change critically in the network and the source behaviours, which are mutually re-enforcing, and challenge the assumptions above:

The "Slashdot-is-the-norm" Future Rapid changes in both the regions of interest, and the source traffic rates will become both feasible, and common.

The Flat Network Future Capacity will no longer be matched. Access networks will have the ability to dominate massively over transit resources. The network will become a euclidean space - effectively a mesh, where latency is simply distance.

Rather than a disaster, we believe that this scenario fulfils the Internet's original promise - if the network is idle between me and a resource, I should be able to use it immediately, and at no especially increased cost. To retain this capability, we need to extend the operating range, or scaling property of the mechanisms for traffic control over many more orders of magnitude than currently available.

The rest of this paper is about a proposed way that this might be done. The next section outlines the proposal. The subsequent section sketches some implementation ideas. The last section summarises the ideas and suggests some future directions.

2 Proposed Solution

We propose a system of "laws" for traffic management that completely generalise the increase/decrease behaviour of adaption over several scales, but critically incorporate a model that could lead to a systems architecture that is both feasible, and efficient.

The appropriate laws might be chosen from field theory: i.e.

attraction gravity, opposite poles in electro-magnetic fields.

repulsion same pole, and of course momentum! not just

$$1/r^2$$

.

Notice that these laws have two effects, one temporal, the other spatial.

Of course, fields operate in a 3D space, while an arbitrary graph is not planar. However, we believe that the evolution of the network infrastructure will tend towards a Euclidean space- mapping capacity with distance to a plane[13] Then we are more concerned with deviations.

Some laws have already been observed about the recent evolution of networks:

attraction e.g. metcalfe++, idea that the value of a network is

$$n^2$$

for n users. we revise this to say that while that is true for n users in the subset of common interests, for a common service that serves that interest there is

repulsion cost - shadow price (per packet, per session, per aggregate, per service. generally, this is a variable with time depending on congestion on one of these timescales.

2.1 Spatial Reduction of Traffic Pressure

You need a mechanism to manage this - and this is not the odd situation, it is the norm in the Internet. So you want to decrease the aggression of TCP in a region - the interest can be modelled as:

$$GMm/r^2$$

where

$$M = \text{Mass}$$

of large object and m is mass of (typically small) client.

M is distorting the plane, where the plane is a representation of the capacity - i.e. deviations from planar mapping represent the *difference* between expected and actual traffic.

So now we introduce a repulsive "force" (or feedback - see [12][13]) to reduce damage -this force needs to repel *dynamically* - i.e. needs to deal with arrival process of flows not just each flow - hence it needs to be better than

$$1/rtt$$

- so we make it

$$1/(rtt^n)$$

for n as a function of M (not sure what function). This will distribute load *impulses* over the net, and scale things really elegantly, by structuring the cost of hot spots less badly. Future research is needed in how rapidly such a system reaches equilibrium, and the real global optimisation problem it solves.

2.2 Temporal Reduction of Traffic Pressure

There appear to be three timescales (with associated actuators (excl apps)) for traffic management:

- round trip times (0-1s); congestion control, uncontrolled flux
- routing timescales (30mins-4hrs); internal, external
- peering timescales (days-weeks?); economic, pricing, social
- provisioning (weeks-years?); economic, pricing, technology

In the work above, we are looking at the rtt timescale with impacts at longer timescales being predictable from this...

The effect is to provide intermediate aggregation structures (which are always going to be more constrained than the simple utility model). One question is: what on earth do they look like/what characterisation is the most interesting?

The idea is that in contrast to optimization based flow/congestion control of TCP based on individual flows, and parameterized (as in multcp [15]) the idea is a set of users subscribe to a service by downloading a control rule (either then individually running it, or having a proxy run it for their aggregate traffic). The rule now operates like a combination of a packet-timescale admission control, and a shaper. the users express a "willingness-to-pay", and they express a desire for smoothness of adaption or aggressiveness. very slow increase decrease , or very bursty.

The rule is implemented as a control law which has an increase and decrease function,

Instead of controlling a large number of individual flows, and using fine grain, small time scale, service differentiation via parameterizing TCP equation to scale the system over a wider operating range, and to scale the *management traffic* over a wider operating range, we want to aggregate. In some ways, the idea unifies slow start, fast start, and rate control[14].

So what is the mechanism for aggregation?

well instead of distributing a tariff +per source/per net+ in the form of a $/ECN bit/flow, and changing the increase/decrease constants, or even changing the power laws per end user from

$$a * x^0 + b * y^{-1}$$

(classic linear increase, multiplicative decrease) to:

$$a * x^i + b * y^{i-1}$$

(geometric increase/decrease) we could change things to even more generally:

$$a * x^i + b * y^{i-j}$$

for

$$j > 1$$

.

This is sufficient to provide convergence, and some form of stability (according to a naive understanding of control theory[7]).

Now, what we do is apply this law to a bunch of end systems - we can do this by some very simple coding of the law - it represents a power law and a bit of algebra can derive what share of capacity as a function of load and mix and rtt you get at a bottleneck.

The

$$1/RTT$$

, dependence in TCP is another thing we can play with.

This is more promising since it models the *proportional* fair share: the more links/hops you traverse in a wired net (and could be geographic distance in a wireless net) the more resource you use[4]. Respectively the more batter power in mobile wireless devices you consume, either for forwarding in a multi-hop, or for transmission on each hop! If you think of price (== general load/mix) as a repelling force, and increase aggressiveness as an attractive force, as well as route choice . The nearer you are to someone, the more you have in common (see [8][9], and subsequent work on *why* people attach to nodes they are familiar with!).

Then the idea is that we form virtual islands of users based on wealth/market - in a fixed network, we can even use the income to provision islands of providers (virtual providers) - the size and distribution of the islands is a function of the set of initial values and evolutionary trajectories of willingness-to-pay.

We might expect this system to congeal into a set of structures of virtual providers and links ,and price efficiencies. In general, the set of power laws leads to a sort of structural granularity (same way that the field equations lead to the granularity of matter, from gluon, through to molecular, and on to phases of matter).

3 Implementation

Here we sketch some of the components of the proposed implementation:

ECN++ ECN++ is a modified version of ECN, which instead of setting a single bit, sets the parameter to the appropriate degrees of polynomial for the increase/decrease functions, as well as the RTT dependence[1][3][15][4][5].

A sketched of an implementation would be as follows: The TCP (or DCCP equivalent) SYN Ack piggybacks ECN the power scaling function. We could also multicast these values from a region of congestion to a set of routers near the region. Think of ECN as gluon/particle exchange used to carry a

$$signal = force$$

field. One problem is that fields can propagate much faster than material. There are not enough TCP Acks to piggyback our force on on, thus we *also* also need the next piece of the equation:

router participation in choosing (or re-writing) ECN++ parameters and also introducing deliberate loss/delay, plus implementing

$$TTL--$$

operation on ECN++ parameters.

Finally, the algorithm used in the router must use a virtual buffer so that it has enough history of the flows to apply the parameters accurately,., and must also be capable of measuring neighbour delay, as well as logging the aggregate ECN++ rate.

4 Summary, Conclusions and Speculations

We have presented an idea for future network traffic management based on an extension of the congestion control paradigm. Future work on this includes:

Analysis The system of fields should be designed and analysed for convergence and stability.

Speculation on Wireless An ad-hoc wireless network could use much the same system structure: however, the appropriate power laws would need to include cost models of battery, transmit versus receive power, and the relationship of capacity to wireless link distance which is radically different from the capacity in future all optical networks.

Speculation on Multicast How to adapt the field equations to deal with group communication is interesting.

Security (incentives) Given the nature of the approach, how could we avoid denial and theft of service attacks?

5 Acknowledgements

Thanks to Will Ng for discussion, and pointer to curvature model of mutual interest in the web[11]. Thanks to Tom Kelly for corrections on a misunderstanding of TCP, and a conversation with Nina Taft at Sprint ATL[10] which introduced the idea of using fields as an organising principle.

References

[1] "Modelling TCP Throughput: A Simple Model and its Empirical Validation", J. Padhye, V. Firoiu, D. Towsley, J. Kurose, Proceedings of Sigcomm'98, Vancouver, CA, September 1998.

[2] "Weighted Proportional Fairness and Pricing for TCP" Philippe Oechslin, J Crowcroft ACM CCR (see `ftp://cs.ucl.ac.uk/darpa/multcp.ps`).

[3] "Congestion Avoidance and Control", V.Jacobson, *ACM SIGCOMM'88*, August 1988, Stanford, CA, pp.314-329.

[4] "Resource pricing and the evolution of congestion control " R.J. Gibbens and F. P. Kelly. Automatica 35 (1999). `http://www.statslab.cam.ac.uk/ frank/evol.html`

[5] "Admission Control for TCP in the Internet", Richard Mortier and Ian Pratt, Cambridge University Work in Progress

[6] "Making Greed Work in Networks: A Game-Theoretic Analysis of Switch Service Disciplines", Shenker et al in SIGCOMM Symposium on Communications Architectures and Protocols, (London, UK), pp. 47-57, Sept. 1994.

[7] "Modern Control Systems" R Dorf, Addison-Wesley, ISBN 0-201-54279, 1989

[8] "On Power-Law Relationships of the Internet Topology", Michalis Faloutsos, Petros Faloutsos, Christos Faloutsos, ACM SIGCOMM'99, Boston, 1999.

[9] "Linked:- The New Science of Networks", Albert-Lszl Barabsi,Perseus Publishing; ISBN: 0738206679; 1st edition (May 14, 2002)

[10] "Gravity Field model of the Matrix", Nina Taft, Sprint ATL, Personal communication.

[11] Curvature of Co-Links Uncovers Hidden Thematic Layers in the World Wide Web jean-Pierre Eckmann & Elisha Moses Proceedings of the National Academy of Science (US), 99, 5825-5829 (2002)

[12] Paxson, V., & Floyd, S., "Why We Don't Know How To Simulate The Internet", Proceedings of the 1997 Winter Simulation Conference, December 1997. (see `http://ftp.ee.lbl.gov/floyd/papers.html`)

[13] "Models for a self-managed Internet" F. P. Kelly Philosophical Transactions of the Royal Society A358 (2000) 2335-2348.

[14] "Rate control in communication networks: shadow prices, proportional fairness and stability." Kelly, Maulloo, Tan, 1998 Journal of Operational Research Society 49.

[15] "Differentiated End-to-End Internet Services using a Weighted Proportional Fair Sharing TCP" Jon Crowcroft, Philippe Oechslin, ACM CCR, 1998.