Transport Area Working Group B. Briscoe
Internet-Draft BT & UCL
Expires: April 18, 2007 October 15, 2006

Flow Rate Fairness: Dismantling a Religion
draft-briscoe-tsvarea-fair-00.pdf

Status of this Memo

Copyright Notice

Abstract

We were moved to write this memo because the applied research and standards communities in networking are using completely unrealistic and impractical fairness criteria. The issue is not whether they should use this or that allocation scheme; they don't even allocate the right thing and they don't allocate it between the right entities. We explain as bluntly as we can that sharing out flow rates (as TCP and many other popular fairness mechanisms do) has no intellectual heritage from any concept of fairness in philosophy or social science, or indeed real life. Comparing and controlling flow

rates alone will never achieve fairness and should never again be claimed as a fairness mechanism for production networks. Instead, a realistic fairness mechanism must share out the `cost' of each users actions on others.

Status

This memo is posted as an Internet-Draft with an intent to eventually progress to informational status.

Table of Contents

1. Introduction

"But he has nothing on at all."

The Emperor's New Clothes, Hans Christian Anderson

With a few notable exceptions, Internet applied research and standards seem to be afflicted by a collective delusion that fairness between traffic competing for network resources can be achieved by controlling relative flow rates alone. It cannot. To be absolutely clear, this accusation covers widely deployed algorithms such as TCP congestion control and all the variants of fair queuing.

Controlling relative flow rates _alone_ is a completely impractical way of going about the problem. To be realistic for large-scale Internet deployment, relative flow rates should be the _outcome_ of another fairness mechanism, not the mechanism itself. That other mechanism should share out the `cost' of one user's actions on others---how much each user's transfers restrict other transfers, given capacity constraints. Then flow rates will depend on a deeper level of fairness that has so far remained unnamed in the literature, but is best termed `cost fairness'.

The metric required to arbitrate cost fairness is simply volume of congestion, that is congestion times the bit rate of each user causing it, taken over time. In engineering terms, for each user it can be measured very easily as the amount of data sent that the system fails to serve. Or with explicit congestion notification (ECN [RFC3168]) the amount of each user's data to have been congestion marked. Importantly, unlike flow rates, this metric integrates correctly across different flows on different paths and across time.

What we call cost fairness has been in the literature for nearly a decade, but it hasn't been put so bluntly before. We were moved to spell it out unambiguously, because this isn't just some dry academic fairness debate that might change allocation percentages somewhere in the third decimal place. The outcomes due to flow rate fairness that we see on the Internet today are hopelessly unlike the outcomes that would result from cost fairness.

Not that the outcomes we see are the deliberate intent of flow rate fairness. They are the random result of an absence of fairness control, because flow rate fairness isn't even capable of reasoning about questions like, "How many flows is it fair to start between two endpoints? or over different routes?" or, "What rate is fair for a flow that has been running longer than another?".

Resource allocation and accountability are two issues that reappear on every list of requirements for a new Internet architecture [NewArchReq]. We could have started filling this architectural vacuum a decade ago, but architecture not only requires foundational ideas, it also requires consensus. In 1997, the basis of the dominant consensus was completely undermined, but it didn't even notice.

While everyone prevaricates, novel p2p applications have started to thoroughly exploit this vacuum with no guilt or shame, by just running more flows for longer (after all, they are using TCP, which is fair isn't it?). In response, ISPs are introducing kludges like volume caps or throttling specific applications using deep packet inspection. Innocent experimental probing has turned into an arms race. The p2p community's early concern for the good of the Internet is being set aside, aided and abetted by large commercial concerns, in pursuit of a more pressing battle against the ISPs that are fighting back. The rest of the Internet is suffering heavy collateral damage. This trend has spread beyond the p2p community. There is now no shame in opening multiple TCP connections, or offering VoIP or video streaming software without any congestion control.

Whether the prevailing notion of flow rate fairness has been the root cause or not, there will certainly be no solution until the networking community gets its head out of the sand and understands how unrealistic its view is. Only then will there be a climate in which solutions can be adopted.

But isn't it a basic article of faith that multiple views of fairness should be able to co-exist, the choice depending on policy? Absolutely correct---and we shall return to how this can be done later. But that doesn't mean we have to give the time of day to any random idea of fairness.

Fair allocation of rates between flows isn't based on any respected definition of fairness from philosophy or the social sciences. It has just gradually become the way things are done in networking. But it's actually self-referential dogma. Or put more bluntly, bonkers.

We expect to be fair to people, groups of people, institutions, companies---things the security community would call `principals'. But a flow is merely an information transfer between two applications. Where does the argument come from that information transfers should have equal rights with each other? It's equivalent to claiming food rations are fair because the boxes are all the same size, irrespective of how many boxes each person gets or how often they get them.

Because flows don't deserve rights in real life, it is not surprising that trying to allocate rate fairly to flows has two inherent loopholes the size of barn doors when it is attempted in a non-co-operative environment. If at every instant a resource is shared among the flows competing for a share, any real-world entity can gain by i) creating more flows than anyone else, and ii) keeping them going longer than anyone else.

We appeal to the networking community to quietly set aside fairness between flow rates. It had its time, but now it has been shown to be unfounded, unrealistic and impractical. Papers and standards proposals using it should be rejected. And no-one should ever have to cater for it in future Internet protocols---it places stupid arbitrary requirements on the system that are impossible to meet and wouldn't achieve any meaningful sort of fairness even if they could be met.

Alternatively, someone should write a defence of flow rate fairness. Continuing to use flow rate fairness as the dominant ideology, without rebutting Kelly's seminal 1997 paper that undermined it, just leaves the Internet community divided into religious sects, making a mockery of the scientific process towards consensus.

2. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Fair Allocation of What Among What?

The issue with flow rate fairness is far more basic than whether allocations should be max-min, proportional or whatever. Flow rate fairness doesn't even allocate the correct thing. And it doesn't allocate it among the correct entities either. At this most basic level we will contrast the two main contending views:

- o Allocate rate among flows (flow rate fairness)
- o Allocate congestion cost among the bits sent by economic entities (cost fairness)

When cost fairness was proposed, it stated its case in terms of the dominant belief system---flow rate fairness. Unfortunately, this meant that the dominant belief system didn't notice it had been struck an intellectual death blow. Its believers carried on regardless and it remains dominant today.

As a result, one sees talk of weighted proportional fairness in the same context as proportional, max-min (or min-max) fairness as if they are all members of the same set. They are not. Weighted proportional fairness has an extra weight parameter w that all the others lack. With weighted proportional fairness, the interesting bit is what regulates users in their choice of w. Otherwise, it would hardly be a useful definition of fairness to say it is fair for flow A to go w times as fast as flow B, if the user behind flow A has free choice of w.

In fact, it turns out that the interesting bit is nothing to do with flows, or their weights. For internetworking the _only_ interesting definition of fairness depends on the allocation of cost among the bits sent by economic entities, regardless of which flows the bits are in. A user's choice of w then depends on that.

3.1. Structure of Memo

The body of this memo is structured around our question, "Fair allocation of what among what?":

- Section 4 answers the "...of what...?" question, explaining why fair allocation of costs is a sufficient and realistic form of fairness, and allocation of rate is not. A sub-section also explains why TCP fair rate control (TFRC) causes greater congestion costs than TCP, because it wrongly tries to achieve equality of flow rate.

- Section 5 answers the "...among what?" question, explaining why fairness among flows can only be myopic whereas fairness among economic entities naturally brings history into the reckoning. Also fairness among flows is shown to be hard, if not impossible to enforce, while enforcing fairness among economic entities is practical.

Having debunked the dominant ideology of flow rate fairness, and replaced it with cost fairness, in Section 6 we discuss how other forms of fairness can be asserted locally. Then, before we draw conclusions, Section 7 maps the progression of seminal ideas in the literature on which this memo is based. A FAQ Web page [FairFAQ] is planned to answer some frequently asked questions that didn't fit easily into the main flow of the memo.

4. Cost, not Benefit

The issues of fair allocation of resources comes under the domain of political economy (or philosophy). In Section 6 we will discuss how different fairness policies can co-exist. But to answer our question, "Fair allocation of what?" we start from the premise used in economics (and life) that fairness concerns comparing benefits, costs or both.

The benefit of a data transfer can be assumed to increase with flow rate, but the shape and size of the function relating the two (the utility function) is unknown, subjective and private to each user. Flow rate itself is an extremely inadequate measure for comparing benefits: user benefit per bit rate might be ten orders of magnitude different for different types of flow (e.g. SMS and video). So different applications might derive completely different benefits from equal flow rates and equal benefits might be derived from very different flow rates.

Turning to the cost of a data transfer across a network, flow rate alone is not the measure of that either. Cost is also dependent on the level of congestion on the path. This is counter-intuitive for some people so we shall explain a little further. Once a network has been provisioned at a certain size, it doesn't cost a network operator any more whether a user sends more data or not. But if the network becomes congested, each user restricts every other user, which can be interpreted as a cost _to all_---an externality in economic terms. For any level of congestion, Kelly showed [wPropFair] that the system is optimal if the blame for congestion is attributed among all the users causing it, in proportion to their bit rates. That's exactly what routers are designed to do anyway. During congestion, a queue randomly distributes the losses so all flows see about the same loss rate (or ECN marking rate); if a flow has twice the bit rate of another it should see twice the losses. In this respect random early detection (RED [RFC2309]) is slightly fairer than drop tail, but to a first order approximation they both meet this criterion.

So in networking, the cost of one flow's behaviour depends on the congestion volume it causes which is the product of its instantaneous flow rate and congestion on its path, integrated over time. For instance, if two users are sending at 200kbps and 300kbps into a 450kbps line for 0.5s, congestion is (200+300-450)/(200+300) = 10% so the congestion volume each causes is 200kx 10%x 0.5 = 10kb and 15kb respectively.

So cost depends not only on flow rate, but on congestion as well. Typically congestion might be in the fractions of a percent but it varies from zero to tens of percent. So, flow rate can never alone serve as a measure of cost.

To summarise so far, flow rate is a hopelessly incorrect proxy both for benefit and for cost. Even if the intent was to equalise benefits, equalising flow-rates wouldn't achieve it. Even if the intent was to equalise costs, equalising flow-rates wouldn't achieve it.

But actually a realistic resource allocation mechanism only needs to concern itself with costs, because normally we use the market economy to handle the benefits side, as we shall now explain.

In life, as long as people cover the cost of their actions, it is generally considered fair enough. If one person enjoys a hot shower more than their neighbour enjoys the toast they made with equal units of electricity, no-one expects the one who enjoyed the shower to have to pay more. If someone makes more of their lot in life than another, some complain it's not fair, but most call this envy, not unfairness.

Market economics works on the same premise (unsurprisingly given life and market economics are closely related). In fact, a market adjusts supply to meet demand so that benefits are as fairly distributed as is consistent with the pre-existing inequalities in life. But this fairness between benefits is an _outcome_ and it is only met as long as there is a market mechanism to make people accountable for the costs of their actions (as long as various market failures are avoided).

We deliberately say `make people accountable' to avoid the phrase `make people pay', because users tend to prefer to pay a flat rate subscription for Internet access in which case their ISP is likely to limit the congestion they are able to cause to what they have paid for (Section 5.2.2).

If we do make users truly accountable for the cost of the congestion they cause, a form of fairness between flow rates emerges automatically. As everyone increases the rate of each of their flows, congestion rises. As congestion rises, everyone pays due regard to the share of the cost attributed to them. So, each individual will want their congestion control algorithm to continuously adjust its rate to maximise their net utility---benefit minus cost. Kelly [wPropFair] shows that even if each user keeps their utility function private but we _model_ all the different users by an arbitrary weight that scales their utility function relative to the others, users will allocate themselves flow rates in proportion to the share of the cost that they cause---weighted proportional fairness.

But such a flow rate allocation is not the measure of fairness, it is merely a possible _outcome_ caused by cost fairness, given some assumptions about how to model the shape of users' private utility functions. Enforcing underlying cost fairness is in itself a sufficient form of fairness. We repeat: _the resulting relative flow rates are not the measure of fairness_.

Most importantly, but perhaps tangentially to our focus on fairness, Kelly proved cost fairness would maximise the aggregate of everyone's utility across the whole Internet. This is why cost fairness is so important, as other forms of fairness cannot be better, unless some major flaw is found in the assumptions. Kelly _et al_ also proved that, even though relative flow rates would likely be very different from those seen today, the Internet would remain stable given reasonable constraints and assumptions [wPropStab].

While on the subject of assumptions, we should add that the benefit of a real-time application depends on jitter, not just transfer rate. But simple scaling arguments show that it will be possible for network operators to minimise congestion delay as networks increase in capacity ([SelfMan] S.2), an argument supported by recent research showing that router buffers are often significantly oversized [BufSizUp].

Proponents of flow-rate fairness might be forgiven for aiming for an `unrealistic' form of fairness if a `realistic' form was difficult to implement in practice. However this is not the case at all, because congestion costs are already used by Internet congestion control---only minor changes are needed. In fact, it is flow rate fairness that is completely impractical to enforce---see Section 5.1.

But how would users "allocate themselves flow rates in proportion to the share of the cost that they cause"? If they were made accountable for congestion, they would install a version of TCP with a weight parameter (for example MulTCP [MulTCP]), at least for TCP-based applications. Of course, most users wouldn't want the fuss of weighting each individual flow. But if they chose to set policies on average for large classes of flows (or to accept the defaults set by application developers), the resulting suboptimal outcome for themselves would be their own private choice to trade optimality against hassle. The underlying fairness criterion would still be met: that people should be accountable for the costs they cause to others.

In contrast, with flow-rate fairness, two flows may cause orders of magnitude different costs to others (for instance if one has been running orders of magnitude longer) by running at equal rates. Nowhere do we find any justification for the dogma that flow rates must be equal to be fair. Nowhere do we find any rebuttal of Kelly's destruction of flow rate fairness, even after nine years.

4.1. TCP-Fairness Equalises Flow Rates Not Costs

An algorithm that controls flow rate in response to congestion is defined as TCP-compatible if it complies with the specification for TCP congestion control [RFC2581]. An algorithm that converges on the same flow rate as TCP at equilibrium but with different dynamics is called TCP-friendly, but not TCP-compatible. Certain streaming applications won't work unless they are allowed a more sluggish response to congestion than TCP's, so researchers invented the concept of `TCP-fairness' to define fair use of the network in competition with TCP-compatible flows.

`TCP-fair' congestion control currently has proposed standard status in the IETF [RFC3448], and it is incorporated into one of the congestion control profiles of the new datagram congestion control protocol (DCCP [RFC4342]) that is also a proposed standard. Unfortunately `TCP-fairness' was defined as equality of flow rates without regard to costs. Consequently, as we shall explain, `TCP-fair' flows with a sluggish response to congestion cause more congestion than TCP-compatible flows on the same path.

To be clear, the extra congestion caused by `TCP-fair' flows, relative to TCP-compatible ones, is not our major concern. We merely use this case to illustrate the broken logic of flow rate fairness. The fairness problems outlined in the next section (`Economic Entities not Flows') afflict _both_ TCP-compatibility and `TCP-fairness' and have far greater impact on Internet users than this minor extra problem with TCP-fair flows.

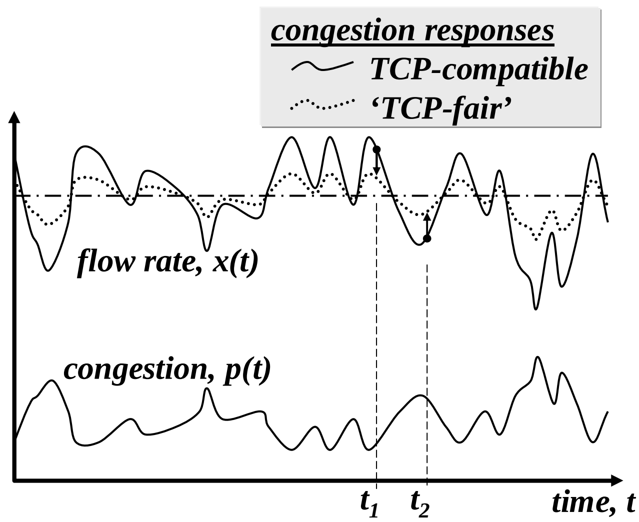

Figure 1: Schematic showing `TCP-fair' flows cause more congestion than TCP. A TCP-fair flow is smoother than a TCP-compatible one but with the same mean rate if measured over long enough time. Therefore at times of high congestion (t_2) it uses more bandwidth than TCP while at times of low congestion (t_1) it uses less.

In order to fairly allocate network resources, all TCP-compatible and TCP-fair congestion control algorithms communicate with each other through the congestion signals coming from network resources. So they actually all have the information necessary for cost-based fairness. But the goal of TCP-fairness was chosen to be equalisation of flow-rate not of flow-rate x congestion. If a flow needs a sluggish response to congestion, TCP-fair rate control keeps it to the same equilibrium rate that a TCP-compatible flow would achieve across the same path. This complex reverse engineering results in a flow that causes more volume of congestion than TCP would in similar circumstances. In terms of rate, it _seems_ fair, but in terms of cost it is not. If a flow with a more sluggish rate response is to cause an equal volume of congestion relative to a TCP flow on the same path, on average it will have to go slower.

To explain, we need to remember that both congestion and flow rate vary over time. A more nimble congestion response like TCP's can mirror changing congestion fairly faithfully. It reduces its rate quickly during periods of higher congestion and increase again more quickly whenever congestion falls. In Figure 1 the resulting schematic plots of congestion and flow rate are shown as mirror images of each other. A more sluggish rate response is not as good at tracking the fast-changing congestion process. So the sluggish flow more often uses higher bandwidth when congestion is high, and more often uses lower bandwidth when congestion is low, causing more volume of congestion on average. Giving more during times of plenty doesn't compensate for taking it back during times of scarcity.

Incidentally, during the standardisation of TCP-fairness, all sorts of other issues arose when trying to equate a real-time flow to a TCP flow. Unlike typical TCP streams, some real-time applications have variable packet sizes and many have a maximum _required_ rate, sometimes also varying rapidly. In contrast, TCP streams have no maximum desired rate. Cost fairness is capable of encompassing all these issues, whereas flow-rate fairness required (and still requires) continual patching with new arbitrary ideas about fairness for each new circumstance.

We are not saying it is easy for a sluggish flow to infer what congestion volume a TCP flow would have experienced. However, as long as congestion costs are accounted for, we don't have to equalise costs _per flow_ anyway---as we are about to explain.

5. Economic Entities not Flows

5.1. Something to Integrate the Allocations

Imagine loaves of bread are regularly delivered to a famine-struck refugee camp. Each time a loaf is brought out, a queue forms and the loaf is divided equally among those in the queue. If the individuals who appear in each queue are always different, except for one who always appears in every queue, would it still be fair to share each loaf equally among those in each queue?

Of course not---commercially realistic fairness policies must depend on an individual's history. But if that isn't a convincing argument, it doesn't have to be. We don't have to show that fairness policies _should_ depend on history, only that realistic ones _probably will_. So a fairness mechanism that claims to support commercially realistic fairness policies needs to be structured so that individual history can be added without destroying scalability. And here, `individual' means some real-world entity with an economic existence, not a flow.

Router-based flow rate fairness mechanisms tend to have to be myopic. To be otherwise would seem to require holding the history of most Internet connected individuals on most routers. Because, at most routers, a flow from nearly any individual in the world might appear. So instead, router-based schemes tend to share out flow rate at each instant without regard to individual history---and hence without regard to commercial reality.

One reason for our frustration with the networking community's focus on flow rate fairness is that the TCP/IP-based architecture of the Internet already has a structure very close to that required to arbitrate fairness based on the costs that individuals cause, rather than on flow rates. Instead of arbitrating fairness on routers, fairness is already arbitrated scalably at the endpoints where the congestion costs of each individual are already collected together.

Congested routers generate cost signals (losses or ECN marks) that are carried to the transport causing the congestion, piggy-backed in the packet stream either as gaps in the transport stream or as ECN marks. These congestion signals are already fed back to the sending transport by nearly all transport protocols. And congestion control algorithms like TCP already adapt their flow rates in response to congestion. So all we would need to change would be to use a weighted TCP algorithm [MulTCP] (or equivalent for inelastic applications) that could weight itself under the control of a process overarching all the flows of one user, which would take into account the user's cost history across all flows.

Of course, there is no incentive for anyone to voluntarily subject themselves to such fairness (nonetheless, they already subject themselves to TCP which voluntarily halves its rate whenever it senses congestion). But as we shall see in Section 5.2.1, policing fairness between individuals (and between networks) at their point of attachment to the Internet has already been solved, whereas getting every router to police fairness between every individual connected to the Internet is a pipedream, because it would be really complicated for routers to have to know about individuals.

So, how come one attachment point can arbitrate fairness between everyone on the Internet when it only knows about locally attached individuals? Do we have to add some fully connected mesh of co-ordination messages between every end-point in the world? The answer is no, because, in a very subtle sense, we already have such a mesh. The thing that keeps fairness at one end-point in line with all the others is the commonly aligned understanding of _cost_ throughout the globe. Cost in any part of the world has an exchange value with cost in any other part, because, wherever there's an Internet attachment, there's a connection with the global economy. Even if some localised authority asserts a non-economic variant of fairness between some sub-set of users (e.g. in a university or corporation), the authority as a whole will still align its understanding of cost with that of the global economy (see Section 6).

So far we have talked of volume of congestion as a cost to other users without calibrating it---specifying how it relates to monetary cost. In fact, in a competitive market, the monetary cost assigned to congestion volume turns out to be the same as the marginal cost of the capacity needed to alleviate the congestion [PrCong] (see FAQ [FairFAQ] for details). The user would be unlikely to see this as a direct charge, instead a flat subscription fee would be considered to include it, in which case the ability to cause congestion would have to be limited by a policer.

Once we have a connection between Internet fairness and economic fairness, the problem of myopia just melts away. The cost that one individual causes is integrated over time for all her flows, by that individual herself. The more cost she causes to others over time and over flows, the more cost she is made to suffer herself. No system has to guess how quickly different people discount benefits and costs experienced in the past, because both are discounted privately by the user just like all the other benefits and costs everyone assimilates in their daily lives.

To manage a system the size of the Internet as a whole, flow rate fairness comes nowhere near being up to the job. It just isn't realistic to create a system the size of the Internet and define fairness within the system without reference to fairness outside the system---in the real world where everyone grudgingly accepts that fairness usually means "you get what you pay for".

Of course, there will be no need to be too precise about that rule. Perhaps some people will get more than they pay for and others less. Perhaps some people will pay for what they get (pay as you go), while others will prefer to be limited to getting what they have paid for (fixed contract). Perhaps some people will be prepared to pay for what others get, and so on. And, as we shall see (Section 6), pockets that may be the size of whole countries can define fairness their own way, within the constraint that the whole pocket pays for what it gets.

But whatever `business model' is used, if individuals run up massive volumes of congestion in small increments over a long time, the balance can be stored in that ingenious invention, the customer account, because we have a technical metric that can be equated to a financial metric: cost. And that other ingenious invention, the networking business is well versed in the art of taking deposits, limiting spending within a credit limit, managing the rate at which credit can be built up and so on. And, of course, the concept of a customer account also naturally ensures that a user cannot escape accountability merely by roaming or mobility. Finally, note well that this `business' and `account' terminology doesn't preclude peer-to-peer creations that arbitrate the resources of a self-provided community network [ArchP2pEcon].

Of course, the details of all this dirty commercial reality don't have to concern the research or the networking standards communities. It is sufficient to design protocols so that congestion costs _can_ be integrated together at some higher layer across different flows and across time, so that senders _can_ be made accountable for the congestion they cause. Systems and protocols intended for Internet deployment do not have to _always_ realise the sort of fairness over time that we find around us in the real world, but they must _be able_ to.

5.2. Enforcement of Fairness

This section drives the final nail into the coffin of flow rate fairness, exposing flaws that even those within the box have to turn a blind eye to, in order to convince themselves that the world within the box is perfectly consistent.

5.2.1. Cheating with Whitewashed or Split Flow Identities

In the real world of deployed networks, if it is easy to cheat the fairness mechanism to get an unfair allocation, it's hardly a useful fairness mechanism. All known flow rate fairness mechanisms are wide open to cheating.

For instance, if I am the customer of a system giving max-min flow rate allocations, it is in my interest to split the identities of my flows into lots of little flows until they are all less than the minimum allocation. Then the system will dance to my tune and reduce the allocations of everyone else in order to increase all the little allocations of my flows. The more I split my traffic down across more and more identifiers, the larger share of the resource all my flows taken together will get.

If a history-based fairness mechanism (Section 5.1) believes it should allocate fewer resource to one flow identifier that it considers has already been given enough, it is trivially easy for the source behind that identifier to create a new identifier for its traffic with a whitewashed reputation.

And it's no good imagining that a router will be able to tell which flow IDs are actually all from the same entity (either in the security sense or the economic sense), because routers have to arbitrate between flows emanating from networks many domains away. They cannot be expected to know which sets of flow identifiers should be treated as a single entity. Flows between a pair of IP addresses may even be attributable to more than one entity, for instance, an IP address may be shared by many hundreds of people on a Web or e-mail hosting site.

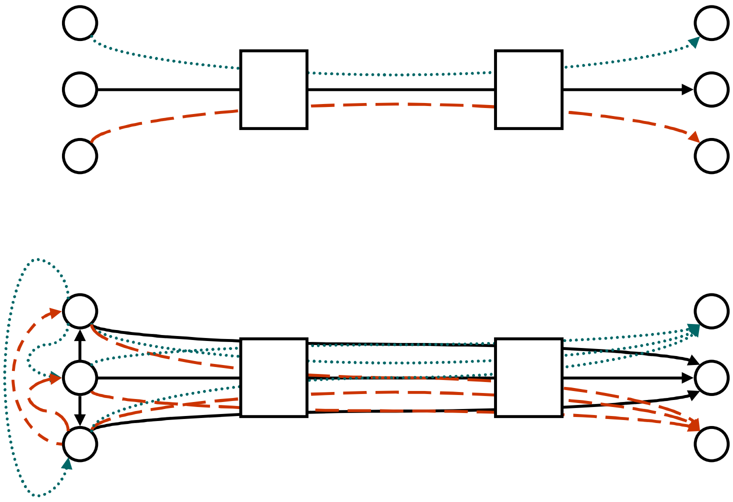

Figure 2: Splitting flow identifiers to cheat against flow rate fairness.

Bottleneck policers [pBox],[XCHOKe],[AFD], suffer from the same inherent problem. They look for a flow ID at a bottleneck that is consuming much more bit rate than other flows in order to police use of TCP. But anyone can cheat by simply running multiple TCP flows. If the policer looks for cheating pairs of source-destination IP addresses, without regard to port numbers, a pair of corresponding nodes can still cheat by creating extra flows from spoofed source addresses after telling each other out of band where to send acknowledgements (or just not using acks). Alternatively, pairs of corresponding nodes can collude to share a portion of each other's flows.

For instance, if the three pairs of nodes in Figure 2 are trying to communicate, the senders can act as stepping stones for each other so that their three (n) flows appear as nine (n^2) across the bottleneck link in the middle. In effect, they have created a routing overlay, much like BitTorrent file-sharing software does. If one pair of naive nodes competes for this bottleneck against n pairs of nodes adopting this strategy, it will get about n times smaller share than each of the other pairs, assuming n is large.

Given identifiers can generally be freely created in cyberspace, it is well-known that they shouldn't be relied on for resource allocation (or more generally for negative reputation) [FrRideP2p],[CheapPseud]. Kelly [wPropFair] chose cost-based fairness (his term was `pricing per unit share') because it was immune to this problem---it allocates cost to bits not to flows and hence doesn't rely on any cyber-identifiers.

In summary, once one accepts that fairness should be based on concepts from social science, fairness can only be meaningful between entities with real-world identities---humans, organisations, institutions, businesses. Otherwise two entities can claim to have arbitrarily many flows between them, making fairness between flows completely meaningless.

5.2.2. Enforcing Cost Fairness

If enforcing flow rate fairness is impractical, is enforcing cost fairness any more achievable? Happily, the Internet's architecture is already suited to carrying the right cost information for cost fairness mechanisms to be enforced in a non-co-operative environment.

Kelly's stated motivation for his focus on pricing was so that the system would be applicable to a non-co-operative environment. In 1999, Gibbens and Kelly went further, pointing out [Evol_cc] that ECN [RFC3168] provided an ideal basis on which to base cost fairness. The idea was simply for network operators to ECN mark traffic at congested routers without regard to flows, then to apply a price to the volume of traffic carrying ECN marks, which would make the transport endpoints accountable for the congestion they caused.

However, understandably, the idea of Internet retailers charging their end-customers directly for congestion met strong resistance. Customers are known to be highly averse to unpredictable charges for services ([PMP] S.5) so duration charging for each Internet flow was unlikely to replace flat monthly charging.

Many threw out the baby with the bath water, associating Kelly's theoretical work solely with its suggested implementation. But over the ensuing years, an active research community has sought to keep the underlying theory but wrapped around with a more realistic incarnation.

Indeed the recent proposal called re-ECN [Re-TCP] does just that. It requires no change to typical flat rate Internet contracts, but it enables addition of a per-source policer that can limit the volume of congestion a customer causes over, say, a month, thus enforcing cost fairness. Although Gibbens & Kelly rightly identified that standard ECN reveals the necessary information for cost-based fairness, it doesn't reveal it in the right place for network layer policing---where the _sender_ attaches to the network. In the current TCP/IP architecture, congestion information emerges from the end of a forward data path, which is the last point in the feedback loop that a network operator can reliably intercept it---the wrong end for policing the sender.

Re-ECN is based on a pattern of feedback called re-feedback [Re-fb], which overloads the standard structure of congestion signalling in IP datagrams, by forcing the sender to honestly declare path congestion in packets it sends on the forward data path (hoping to use the last undefined bit of the IPv4 header). Then it is possible to enforce cost fairness in a per-user policer. Of course, the policer would act as a deterrent, encouraging the end-user to use a weight-based congestion control such as MulTCP [MulTCP], as already described.

Re-ECN congestion information aggregates naturally, giving downstream networks the necessary bulk information to enforce cost-based fairness at their borders with their neighbouring upstream networks. Whether a network asserts cost-based fairness or some other fairness policy between its own upstream users is its own choice---indeed it can choose not to intervene at all. But the re-ECN information it has to forward through its border router allows its downstream neighbour to penalise the whole upstream network for the costs it has allowed its users to cause to downstream users. This enables a variety of fairness policies to co-exist (including absence of policy, which ensures incremental deployment). But each network has an incentive to limit the costs that its users cause to others on the Internet. So, on the grander scale, networks have to be fair to each other, which is the subject of the next section.

6. Fairness between Fairnesses

A social anthropologist would be able to give numerous examples of tribes and societies holding differing opinions on fairness. But, just as gravity pre-dated Newton, the invisible hand of the (maturing) market had been allocating resources in society long before Adam Smith noticed, particularly where the larger picture of trade between societies was concerned. However, monarchs, governments, charities and so on have also been stamping their own view of fairness on this backdrop, sometimes less equal sometimes more.

But, we must also recognise that society's view of fairness is heavily influenced by the fairness that a market would produce [SovJstce]. In terms of alpha-fairness [aFair], flow rate equality (alpha=infinity) is defined as extremely fair. But in life few expect to get an equal share of the cake for nothing. As a society, we accept that a reasonably competitive market mechanism does produce a `realistic' form of fairness; a form of fairness that people grudgingly accept they have to live with, where the buyer gets no more than she pays for, at a competitive price that reflects the effort expended by the seller.

Even if different allocation schemes are chosen locally, perhaps taking account of social inequality, on a global scale arbitration between local views on fairness is largely through market economics---we are not asking anyone to judge whether this is good or bad, it just is. This doesn't imply we believe that economic forces are somehow above policy control. Rather, we observe that market forces (aside from wars) have been the default _global_ resource allocation mechanism over many centuries. In the Greco-Roman civilisations, in the Buddhist and Confucian worlds, and later in the Islamic world, trade was a necessary, but not the primary, aspect of life. Over the last two decades, Western civilisations have been going through a phase of `economics imperialism', where attempting to exert policy control over economics has been viewed as counter-productive. However, we must not assume the current globalisation trend [Saul05] heralds the end of history.

The Internet should be able to reflect this pattern as societal forces shift and different local fairness regimes come and go---`design for tussle' [Tussle]. On the whole, interworking between most parts of the Internet must _be able_ to be based on market economics, while other fairness criteria can be applied locally. For instance, a University might choose to allocate resources to each student equally rather than by how much their parents can afford. But the resources one whole University gets relative to another institution depend on how much each pays their service provider. Whole countries might arrange to subsidise a minimum universal service obligation for Internet _usage_, but still, the country as a whole would be expected to pay its way in the world. On the other hand, in market-led countries, commercial ISPs might solely allocate resources proportionate to customer subscriptions. Local pockets of heterogeneity will exist, from computer clubs to NATO, but the overall fabric of resource allocation that glues all these pockets together at the (inter)network layer is likely to be based on market economics.

This is what we mean by `realistic'---fitting the commercial reality of a global market economy. We are fully aware that the power of market economics can be stretched too far; controlling aspects of society where economic assumptions break down (prompting Samuelson to describe Friedman as "...somebody who had learned how to spell banana but didn't know where to stop" [Swed90]). But we are not advocating that one religion should replace another---market economics replacing flow rate fairness. However, in the case of Internet resource allocation, it must at least _be possible_ to use market economics, despite its known failings, given it is currently the most appropriate tool for managing the large-scale interactions.

A market is meant to optimise allocations in the face of conflicts of self-interest. If we want to assert other fairness regimes, we must recognise this acts against how a market would work. If we don't understand how to overcome self-interest, its invisible hand will force its will on us some other way, distorting our attempts to work against it. This is why the loop holes in flow rate fairness are being so thoroughly exploited.

And this is our point. A market _mechanism_ has to be _designed_. A weak design will be exploited mercilessly. The designs behind flow rate fairness are worse than weak. They are not even aware that, as resource allocation mechanisms, they _should_ be able to meet the stringent requirements of a good market mechanism, such as forgery-resistant `currency', information symmetry, internalisation of externalities and so forth.

If we did wish to promote the cause of equality, equalising flow rates would in no way achieve our ends. In fact, it would only promote the cause of selfishness and malice, because flows don't equate to people, so its broken logic can be thoroughly exploited. Only by providing a bullet-proof mechanism for the market to express itself, can we then move on to allocate resources locally in other ways.

7. The Seminal Literature

For a rigourous tutorial on the various form of fairness, the reader is referred to Le Boudec [ccFairTut].

Max-min flow rate fairness has a long history in networking, with research to find distributed (router-based) max-min algorithms starting in 1980 [DeMaxMin] and Nagle proposing a novel approach in 1985 [RFC0970]. All these early `fair queuing' algorithms gave equal rights to each source. In 1989, to solve the problem of some sources deserving more rate than others, the authors of `weighted fair queuing' (WFQ) proposed that per-source destination pair would be a better model of the size of different sources. It was admitted that a source could deny service to other sources by faking transfers with numerous destinations, but a reasonable tradeoff between efficiency and security was required [WFQ]. Recently, an approach called Justice [Jstce] has proposed a return to (weighted) per source fair queuing, but with configurable link weights throughout the network. However, all these `fair queuing' approaches allocate bit rate as their measure of fairness.

TCP congestion control was also introduced in the late 1980s [TCPcc], based on the assumption that it would be fair if flow rates through a single bottleneck converged on equality.

In 1991, Mazumdar _et al_ [UtilFair] pointed out that there was nothing special about max-min fair rate allocation, and that other _ad hoc_ definitions of fairness perhaps based on ratios of individual demands would be no less valid. Instead Mazumdar _et al_ advocated that it would be precise to base a definition of fairness on game theory, specifically the Nash bargaining solution. This resulted in proportional fairness, but still using the rate allocated to flows as the measure of fairness.

In 1997, Kelly considered that Mazumdar's use of co-operative game theory was unlikely to be relevant to public networks where fairness would have to be enforced. Instead he introduced _weighted_ proportional fairness [wPropFair], which finally broke the link between fairness and flow rates. However, the break in tradition wasn't obvious because the new form of fairness could easily be expressed in terms of flow rates, essentially using the weight of a flow as a `fiddle-factor'.

Kelly showed that all a network had to do to achieve fairness in its economic sense (cost fairness) was to share the cost of congestion among bits (not flows). Then, as long as the network made users experience the cost of their bits, users could choose any size flows they wished. But their choice would be regulated by their own trade off between how much they valued bit rate and the charge for congestion.

Kelly's fairness with respect to bit rate per unit charge could also be (and was) framed in terms of fairness between flows by allowing the user an arbitrary choice of weight per flow. But Kelly pointed out that a flow could be divided into sub-flows without changing the overall rate allocation to all the sub-flows taken together; the user merely had to imagine that the weight she assigned to one flow could be subdivided proportionately into its sub-flows.

Kelly's work built on MacKie-Mason & Varian's seminal paper on the economics of networks from 1995, "Pricing Congestible Network Resources" [PrCong]. This work explained the dual role of congestion costs in controlling demand and regulating supply, in welfare maximising, competitive and monopoly markets.

In his 1997 paper, Kelly framed cost fairness in terms of weighted proportional fairness of flow rates in order to relate to an ATM technology context. With ATM's flow-based user-network interface, users had to declare the weight they chose for their flows to the network. But by 1998 Kelly _et al_ applied this work [wPropStab] to an Internet setting where flows were not part of the user's interface with the network, so flow weights could become a purely private device, internal to the user's rate control algorithm. Nonetheless, the _outcome_ at the flow level was still weighted proportional fairness, and the underlying fairness that produced this outcome was still based solely on sharing the cost of congestion among bits.

Back in 1995, Shenker had identified two main types of network traffic: elastic and inelastic, distinguished respectively by their concave and sigmoid utility functions [FundUtil]. Whatever the utility function, Kelly teaches us that covering congestion costs is sufficient to achieve fairness. But then the outcome (in terms of flow rates) depends on the type of utility function:

- Weighted proportionally fair flow rates will be the outcome for elastic traffic streaming;

- Inelastic traffic flows hit a discontinuity once congestion rises beyond a certain level, at which point each is better off with zero rate, leading to a need for some form of admission control, whether self-admission control or arbitrated by the network [DCAC].

- Key & Massoulie identified a third major class of network traffic where utility is derived solely from the duration required to complete transfer of a fixed volume of data [UtilFile]. They showed that, if cost fairness applied, self-interested congestion control would toggle between full line rate and zero (with occasional probes). Such behaviour alone destabilises the network, but it can be stabilised by mixing with streaming traffic [FairIntgr]. Research on the second order incentives necessary to encourage stability continues.

Since these seminal papers in the late 1990s, theoretical refinement has continued, but the main thrust of research has been to find more realistic and practical ways of applying the insights, a process which is now bearing fruit (see Section 5.2.2).

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

The whole of Section 5.2 discusses how there are no known ways of enforcing flow rate fairness securely in a non-co-operative environment like the current Internet, whereas practical, secure solutions have been proposed for enforcing cost-fairness.

10. Concluding Remarks

10.1. A Cautionary Note

In 1997, Kelly showed [wPropFair] that, assuming everyone covered their costs, max-min flow rate fairness could be contrived by supposing users all valued bit rate with an unrealistically extreme set of utility functions that were all identical and that all valued low bit rate infinitesimally less than high bit rate. But, despite this damning evidence, we have continued to see schemes with routers doing max-min fair rate allocation.

To spell Kelly's result out even more bluntly, max-min fair rate allocation satisfies the policy goal, "To all equally without regard to their wants or needs, from all [others] without regard to cost." Such a view might be the motto of the revolutionary monster raving wacko party, but it should have no place in large-scale networking. Knowing this, what reason would anyone have to take max-min flow rate fairness seriously, ever again?

Further, once the idea of fairness based on integrating costs over time is understood, what reason would anyone have to take any form of instantaneous per-flow rate fairness (whether max-min or TCP) seriously, ever again?

Even if a system is being designed somehow isolated from the economy, where costs will never have to relate to real economic costs, what possible reason could there be for adopting these forms of fairness that so badly relate to real life fairness?

10.2. Conclusions

In much of the networking community you have to put fairness in terms of flow rates, otherwise your work is `obviously' irrelevant. At minimum, you are an outcast, if not a heretic. But actually it is flow rate fairness itself that has no basis in philosophy or science, let alone `commercial reality'. It is a classic case of a hegemony where those living within the box don't recognise the existence of the box, let alone that there is a world outside the box.

Outside the box, cost fairness was derived from economic concepts of fairness back in 1997. When flow rate fairness is seen through the wider eyes of this economic analysis it is clearly broken, even on its own terms. The criticism is far more damning than merely whether allocations are fair. Both the thing being allocated (rate) and what it is allocated among (flows) are completely daft---both unrealistic and impractical. However, the Internet community continues to judge fairness using flow rates, apparently unaware that this approach has been shown to have no intellectual basis.

And to be clear, this accusation applies to the so-called `fairness' that emerges from the TCP algorithm and the various fair queuing algorithms used in production networks. In fact, these flow rate fairness algorithms are myopic in both space and time---they are completely unable to control fairness at all, because they don't adjust depending on how many flows users create and for how long.

In real life, fairness generally concerns costs or benefits. Flow rate doesn't come anywhere near being a good model of either. User benefit per bit rate might be ten orders of magnitude different for different types of flow. And cost depends on the product of bit rate with congestion, which is very variable and nothing like bit rate alone.

But even worse, there is no evidence whatsoever that fairness between flows relates in any way to fairness between any real-world entities that one would expect to treat fairly, such as people or organisations. If fairness is defined between flows, users can just create more flows to get a larger allocation. Worse still, fairness between flows is only defined instantaneously, which bears no relation to real-world fairness over time. In contrast, cost fairness has realistic answers to all these questions.

Further, cost fairness is practical to enforce, unlike flow rate fairness, which seems inherently broken in this respect. We believe cost fairness is a coherent way forward with all the technical barriers overcome, or close to being overcome.

The only outstanding barrier is a religious one. This memo has been written from frustration that no-one in the box believes that the voices that seem to be coming from outside the box should be listened to. It seemed the only way forward was to force the issue, by making the box look ridiculous in its own terms. The applied networking community must justify its preposterous position on fairness with reference to some previously respected notions in philosophy or social science. In this memo, we have shown how the whole house of cards is unlikely to be up to such rigour.

10.3. Further Work

This memo has focused on the fairness ideas we see in the production networks around us today. However, our real concern is that these broken ideas also pervade the community working on replacing the Internet architecture. It is well-known that TCP congestion control is running out of dynamic range and many proposals for hi-speed replacements have been put forward. But these replacements must also simultaneously meet the much harder requirement of encompassing a range of realistic fairness criteria, as outlined in Section 6.

XCP is a router-based hi-speed congestion control mechanism that claims to allow different fairness criteria to be configured. However, XCP fairness is based on the myopic flow-rate-based view that we have so roundly criticised in this memo. For instance, XCP claims to be able to achieve a weighted proportional fair rate allocation ([XCP] S.6), but it glosses over how it regulates each user's choice of the weight---there is no direct congestion information in the XCP protocol that could be used to make each user accountable for their choice. Further research is needed to establish whether a combination of XCP's protocol fields could yield this information, and if so, whether such an approach would be immune to cheating.

We also believe it will be necessary to be able to apply different fairness criteria to different subsets of users of a network, and subsets across an internetwork. We cannot immediately see how this would be feasible with router-based approaches like XCP, but it would be straightforward with end-to-end approaches like re-feedback [Re-fb]. Therefore we plan to focus on achieving hi-speed congestion control in an edge-policed end-to-end control architecture.

11. Acknowledgements

Thanks are due to Scott Shenker for persuading me to write this, Louise Burness for insight into why people think the way they do, Ben Strulo for giving a better way of expressing it and Marc Wennink for challenging the ideas. Also thanks to Arnaud Jacquet, Jon Crowcroft, Frank Kelly and Peter Key for their useful reviews. However, the author alone shoulders the blame for any offence caused by the bluntness of style.

12. Comments Solicited

Comments and questions are encouraged and very welcome. They can be addressed to the IETF Transport Area mailing list <tsv-area@ietf.org>, and/or to the authors.

13. References

13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, April 1998.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.

[Re-TCP] Briscoe, B., Jacquet, A., Salvatori, A., and M. Koyabi, "Re-ECN: Adding Accountability for Causing Congestion to TCP/IP", draft-briscoe-tsvwg-re-ecn-tcp-02 (work in progress), June 2006.

13.2. Informative References

[AFD] Pan, R., Breslau, L., Prabhaker, B., and S. Shenker, "Approximate Fairness Through Differential Dropping", CCR 33(2)23--40, April 2003.

[ArchP2pEcon] Strulo, B., Smith, A., and J. Farr, "An Architecture for Peer-to-Peer Economies", Proc. 3rd Int'l Conf. On Peer-to-Peer Computing (P2P 2003) pp208--209, 2003, <http://csdl.computer.org/comp/proceedings/p2p/2003/2023/00/20230208.pdf>.

[BufSizUp] Ganjali, Y. and N. McKeown, "Update on Buffer Sizing in Internet Routers", ACM SIGCOMM CCR 36, October 2006, <http://www.sigcomm.org/ccr/drupal/?q=node/72>.

[CheapPseud] Friedman, E. and P. Resnick, "The Social Cost of Cheap Pseudonyms", Journal of Economics and Management

Strategy 10(2)173--199, 1998.

[DCAC] Gibbens, R. and F. Kelly, "Distributed connection acceptance control for a connectionless network", Proc. International Teletraffic Congress (ITC16) pp941--952, 1999, <http://www.statslab.cam.ac.uk/~frank/dcac.html>.

[DeMaxMin] Jaffe, J., "A Decentralized, "Optimal", Multiple-User, Flow Control Algorithm", Proc. Fifth Int'l. Conf. On Computer Communications pp839--844, October 1980.

[Evol_cc] Gibbens, R. and F. Kelly, "Resource pricing and the evolution of congestion control", Automatica 35(12)1969--1985, December 1999, <http://www.statslab.cam.ac.uk/~frank/evol.html>.

[FairFAQ] Briscoe, B., "Cost Fairness FAQ", Web page , October 2006, <http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/2020comms/fairfaq.html>.

[FairIntgr] Key, P., Massoulie, L., Bain, A., and F. Kelly, "Fair Internet traffic integration: network flow models and analysis", Annales des Telecommunications 59 pp1338--1352, 2004, <http://citeseer.ist.psu.edu/641158.html>.

[FrRideP2p] Feldman, M., Papadimitriou, C., Chuang, J., and I. Stoica, "FreeRiding and Whitewashing in Peer-to-Peer Systems", Proc. Workshop on Practice and Theory of Incentives in Networked Systems (PINS'04) pp228--236, 2004, <http://doi.acm.org/10.1145/1016527.1016539>.

[FundUtil] Shenker, S., "Fundamental Design Issues for the Future Internet", IEEE Journal on Selected Areas in Communications 13(7)1176--1188, 1995, <http://citeseer.ist.psu.edu/shenker95fundamental.html>.

[Jstce] Eriksson, J., Faloutsos, M., and S. Krishnamurthy, "Justice: Flexible and Enforceable Per-Source Bandwidth Allocation", Proc. Networking pp1206--1218, 2005, <http://www.cs.ucr.edu/~krish/jakob_networking05.pdf>.

[MulTCP] Crowcroft, J. and Ph. Oechslin, "Differentiated End to End Internet Services using a Weighted Proportional Fair Sharing TCP", CCR 28(3) 53--69, July 1998, <http://

www.cs.ucl.ac.uk/staff/J.Crowcroft/hipparch/pricing.html>.

[NewArchReq] Braden, R., Clark, D., Shenker, S., and J. Wroclawski, "Developing a Next-Generation Internet Architecture", DARPA white paper , July 2000, <http://www.isi.edu/newarch/DOCUMENTS/WhitePaper.pdf>.

[PMP] Odlyzko, A., "A modest proposal for preventing Internet congestion", AT&T technical report TR 97.35.1, September 1997, <http://www.dtc.umn.edu/~odlyzko/doc/modest.proposal.pdf>.

[PrCong] MacKie-Mason, J. and H. Varian, "Pricing Congestible Network Resources", IEEE Journal on Selected Areas in Communications, `Advances in the Fundamentals of Networking' 13(7)1141--1149, 1995, <http:// www.sims.berkeley.edu/~hal/Papers/ pricing-congestible.pdf>.

[RFC0970] Nagle, J., "On packet switches with infinite storage", RFC 970, December 1985.

[RFC2581] Allman, M., Paxson, V., and W. Stevens, "TCP Congestion Control", RFC 2581, April 1999.

[RFC3448] Handley, M., Floyd, S., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 3448, January 2003.

[RFC4342] Floyd, S., Kohler, E., and J. Padhye, "Profile for Datagram Congestion Control Protocol (DCCP) Congestion Control ID 3: TCP-Friendly Rate Control (TFRC)", RFC 4342, March 2006.

[Re-fb] Briscoe, B., Jacquet, A., Di Cairano-Gilfedder, C., Salvatori, A., Soppera, A., and M. Koyabe, "Policing Congestion Response in an Internetwork Using Re-Feedback", ACM SIGCOMM CCR 35(4)277--288, August 2005, <http:// www.acm.org/sigs/sigcomm/sigcomm2005/ techprog.html#session8>.

[Saul05] Saul, J., "The Collapse of Globalism and the Reinvention of the World", Pub: Viking, Canada ISBN: 0-670-06367-3, 2005.

[SelfMan] Kelly, F., "Models for a Self-Managed Internet", Philosophical Transactions of the Royal

Society 358(1773)2335--2348, August 2000, <http://www.statslab.cam.ac.uk/~frank/smi.html>.

[SovJstce] Siderenko, S., "Characteristics of Perceptions of Social Justice in the Contemporary USSR", Chapter in: Transitional Agendas: Working Papers from the Summer School for Soviet Sociologists Ch3 pp41--45, 1991, <http://lucy.ukc.ac.uk/csacpub/russian/siderenko.html>.

[Swed90] Swedberg, R., "Economics and Sociology. Redefining their Boundaries: Conversations with Economists and Sociologists", Pub: Princeton University Press , 1990.

[TCPcc] Jacobson, V., "Congestion Avoidance and Control", Proc. ACM SIGCOMM'88 Symposium, Computer Communication Review 18(4)314--329, August 1988, <ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z>.

[Tussle] Clark, D., Sollins, K., Wroclawski, J., and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", ACM SIGCOMM CCR 32(4)347--356, October 2002, <http://www.acm.org/sigcomm/sigcomm2002/papers/tussle.pdf>.

[UtilFair] Mazumdar, R., Mason, L., and C. Douligeris, "Fairness in Network Optimal Flow Control: Optimality of Product Forms", IEEE Transactions on Communications 39(5)775--782, May 1991, <http://dionysos.cs.unipi.gr/~cdoulig/Fairness%20in%20network%20optimal%20flow%20control%20optimality%20of%20product%20forms.pdf>.

[UtilFile] Key, P. and L. Massoulie, "User policies in a network implementing congestion pricing", Proc. Workshop on Internet Service Quality and Economics , December 1999, <http://www.marengoresearch.com/isqe/index.htm>.

[WFQ] Demers, A., Keshav, S., and S. Shenker, "Analysis and Simulation of a Fair-Queueing Algorithm", ACM SIGCOMM CCR 19(4)1--12, September 1989, <http://portal.acm.org/citation.cfm?id=75248>.

[XCHOKe] Chhabra, P., Chuig, S., Goel, A., John, A., Kumar, A., Saran, H., and R. Shorey, "XCHOKe: Malicious Source Control for Congestion Avoidance at Internet Gateways", Proceedings of IEEE International Conference on Network

Protocols (ICNP-02) , November 2002,
<http://www.cc.gatech.edu/~akumar/xchoke.pdf>.

[XCP] Katabi, D., Handley, M., and C. Rohrs, "Congestion Control for High Bandwidth-Delay Product Networks", ACM SIGCOMM CCR 32(4)89--102, October 2002, <http://www.ana.lcs.mit.edu/dina/XCP/>.

[aFair] Tang, A., Wang, J., and S. Low, "Is Fair Allocation Always Inefficient", Proc. IEEE Conference on Computer Communications (Infocom'04) , March 2004, <http://netlab.caltech.edu/pub/papers/fairness-infocom2004.pdf>.

[ccFairTut] Le Boudec, J-Y., "Rate Adaptation, Congestion Control and Fairness: A Tutorial", Web page , November 2005, <http://ica1www.epfl.ch/PS_files/LEB3132.pdf>.

[pBox] Floyd, S. and K. Fall, "Promoting the Use of End-to-End Congestion Control in the Internet", IEEE/ACM Transactions on Networking 7(4) 458--472, August 1999, <http://www.aciri.org/floyd/end2end-paper.html>.

[wPropFair] Kelly, F., "Charging and Rate Control for Elastic Traffic", European Transactions on Telecommunications 8 pp33--37, 1997, <http://www.statslab.cam.ac.uk/~frank/elastic.html>.

[wPropStab] Kelly, F., Maulloo, A., and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability", Journal of the Operational Research Society 49(3) 237--252, 1998, <http://www.statslab.cam.ac.uk/~frank/rate.html>.

Author's Address

Bob Briscoe
BT & UCL
B54/77, Adastral Park
Martlesham Heath
Ipswich IP5 3RE
UK

Phone: +44 1473 645196
Email: bob.briscoe@bt.com
URI: http://www.cs.ucl.ac.uk/staff/B.Briscoe/

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.