

Using Self-interest to Prevent Malice

Fixing the Denial of Service Flaw of the Internet

Bob Briscoe

<bob.briscoe@bt.com> BT Research and UCL,
B54/77, Aadastral Park, Martlesham Heath, Ipswich, IP5 3RE, UK

October 22, 2006

Abstract

This paper describes the economic intent of a proposed change to the Internet protocol. Denial of service is the extreme of a spectrum of anti-social behaviour problems it aims to solve, but without unduly restricting unexpected new uses of the Internet. By internalising externalities and removing information asymmetries it should trigger evolutionary deployment of protections for Internet users. To be worthwhile architectural change must solve the last stages of the arms race, not just the next. So we work through the competitive process to show the solution will eventually block attacks that other researchers consider unsolvable, and that it creates the right incentives to drive its own deployment, from bootstrap through to completion. It also encourages deployment of complementary solutions, not just our own. Interestingly, small incentives in the lower layer infrastructure market amplify to ensure operators block attacks worth huge sums on the black market in the upper layers.

1 Introduction

Infrastructure must serve a very large population as faithfully as possible even during flash crowds of demand [JKR02]. All infrastructure therefore faces a dilemma that denial of service (DoS) attacks exploit. It must be able to distinguish a flood of bogus requests from a flash crowd of genuine demand, even when both happen together. Indeed attacking during a flash crowd is the most cost-effective strategy an attacker can adopt—at the time when the infrastructure is most valuable to most genuine users, and when it takes least extra malicious effort to tip it into overload.

Most researchers believe that making an attack imitate a genuine flash crowd is the limit of what will be possible [HR06]. By using the economics of the system as a whole, we believe the solution described in this paper

has the potential to suppress an attack during a genuine flash crowd. It also introduces strong incentives for network operators to deploy other, complementary measures to suppress DoS.

The Internet infrastructure efficiently delivers data packets¹ to their destination address. But it also efficiently delivers floods of packets to any targeted victim. Anyone with a grievance can recruit a ‘zombie’ army of other people’s computers to flood out the service of their chosen victim by filling the link to their computer with data packets—a distributed denial of service (DDoS) attack.

In 2005, the CSI/FBI ranked denial of service as the fourth most costly type of computer crime in the US; behind viruses, unauthorised access and info-theft [CSI05].² But, unlike the top three crimes, a potential victim is currently powerless to defend itself against a well-crafted DDoS attack. Defences require concerted action for the common good across the whole Internet. Currently, the best defence is to avoid being conspicuous.

The early ‘script kiddie’ attacks in the late 1990s tended to hit trophy targets. But over the intervening years, organised crime moved in, tending instead to hit businesses that would quietly give in to extortion demands.

In 2004, one large ISP saw on the order of 20 DDoS attacks per day, with about 1 in 3 affecting customer business, another reported 6 or 7 attacks ongoing at any one time, while yet another had never experienced an attack [Han05]. In the first half of 2005 Symantec reported a worrying 680% growth in these attacks (Fig 1).²

Robot armies or ‘botnets’ of tens of thousands of computers are routinely amassed by infecting them with ‘bot’ software. Bot armies are openly sold for a fee of about

¹A packet is a container for a small amount of data, with a ‘header’ at the front containing minimal information sufficient only for its delivery, much like a postal envelope carries its payload inside and the address on the front.

²One must treat any data on the size of security problems with care, given it generally emanates from those with an incentive to inflate it, it is more costly to collect representative data from small companies than large and there is rarely open disclosure of the data collection methodology.

Presented at The Workshop on the Economics of Securing the Information Infrastructure, October 23-24, 2006, Washington DC, USA. <<http://wesii.econinfosec.org/>>

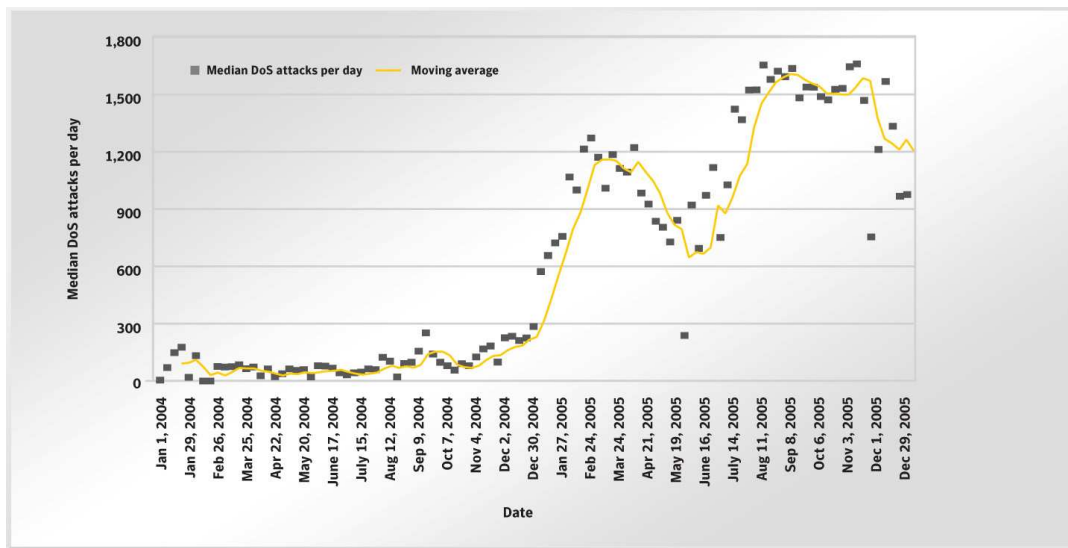


Figure 1: DoS attacks worldwide per day (2004-5). Source: Symantec [Sym06]

5 cents per bot-week [Han05], or the equivalent in kind (avoiding traceable financial transactions). The total number of bot-infected computers observed participating in attacks around the world on any one day stabilised at about 10,000 by the end of 2005 [Sym06]. However, they are not the same 10,000 each day, and larger armies can be marshalled if required—attack forces of hundreds of thousands are not uncommon, which can saturate even the 10Gbit per sec high speed links within the core of a network.

We believe DDoS is best prevented by treating it for what it is—extremely high congestion. Genuine uses of the Internet automatically respond to congestion, typically sending more slowly the more congestion they detect. But this polite response is entirely voluntary. Despite the Internet having become the information infrastructure on which large parts of the global economy depend, the stability of the whole infrastructure depends critically on this polite response—an inheritance from the spirit of mutual trust that prevailed in the Internet’s formative years. This is “The denial of service flaw of the Internet”, but the lack of any network self-defence against traffic attacks is of even wider concern.

Our wider goal has been to turn this unfettered freedom into ‘freedom with responsibility’, but without restricting the openness of the Internet to innovative new applications. In one sense, mitigating denial of service has merely been a pleasant consequence of our wider work.

Our approach is fundamentally economic, treating congestion as a negative externality—a detrimental side effect of users’ actions on others. But we *don’t* believe the answer is to *directly* pass on the cost of congestion to end-users, in some vain hope that this will encourage every end-user in the world to always be more careful with their virus checker. However, we *do* want to pass on the

cost of congestion to the networks closest to the computers causing congestion. Our intent is to create extremely strong incentives for networks to deploy traffic policers to prevent their own customers causing denial of service attacks.

Most approaches to mitigating DoS make binary good/bad decisions about flows, and use filters to take binary allow/block actions. We take a continuous rather than discrete approach (again, reflecting economic intuition), where the more any source causes sustained congestion through the network, the more a traffic policer throttles it. There is no conception of crime or punishment, only increasingly bad behaviour and increasingly starved privileges, with extreme starvation as the natural response to extremely bad behaviour, such as DDoS bots—persistent sources of extreme congestion. We prefer this *proportionate* approach rather than attack detection and filtering, which will remain an arms race with the associated risk of false negatives.

Still borrowing from economics, our solution fixes the information asymmetry that has made it impossible to solve the problem of congestion externalities in the Internet. It ensures that an estimate of the total amount of congestion about to be encountered downstream³ must be written into each packet in order to get through the congestion. As packets pass through the network, requisite congestion information is subtracted from the packet stream as congestion is encountered. The aim is that, if the congestion information becomes persistently negative before reaching the destination, packets will be discarded so that only non-negative parts of the stream continue onward.

³‘Downstream’ means further through the network in the direction of data flow (in this paper we do *not* use ‘downstream’ for its other meaning, as in the direction of service supply from lower layer infrastructure to higher).




Figure 2: Incentives and Desired Outcome: For most network providers to deploy per-user ingress policers (or, at least, other ingress DDoS defences). Note: Variable charges complement fixed interconnect charges (not shown).

This sounds similar to an old idea, but previous work by Kelly [GK99] was placed in the context of Internet technology emerging at the time. Therefore it was based on information being written into the packet as it experienced congestion. Our solution, which we call re-feedback, effectively reverses the information flow without having to change routers. So packets carry a prediction of the congestion they are about to encounter, rather than a record of what they have just encountered. Previously, all that could be achieved was to pass the cost of congestion to the receiver, which added insult to injury, forcing it to pay to be the victim of an attack. Re-feedback claims to solve the whole range of Internet resource sharing problems, but this paper focuses on just DDoS, which is the most extreme and therefore challenging case.

The idea of re-feedback is very simple. But it adds to the art of controlling Internet congestion, which is not renowned as an accessible subject for outsiders. In this paper, we have taken great pains to demystify the engineering behind the idea of re-feedback, aiming for an audience of specialists in economics, business or public policy. We seek review from these disciplines, posing the questions, “Have we misunderstood or misapplied the economics? Have we got the balance right?”

However, the contribution of this paper is not the technical idea of re-feedback itself⁴. Our more ambitious aim is to build an economic argument around re-feedback to show that it provides strong incentives to bootstrap its own deployment *and* to push deployment onward towards

⁴Those with a technical background may prefer to read extensive descriptions of the proposal either in slightly outdated overview [BJCG+05] or in depth [BJSK06]. Those attempting to break the scheme should read the latter, *not* the present paper, which takes liberties for brevity. [BR05] puts the idea in its wider commercial and technical context.

completion.

Ideally, complete deployment would consist of an unbroken ring of traffic policers around the whole Internet—policers that could even starve out attack traffic during a flash crowd. But, in fact, re-feedback should provide strong incentives to deploy any selection of measures against network DDoS until, between them, they prove effective. Our eventual aim is to create an analytical argument, but this initial paper merely gives the intuition.

The outcome we desire and expect (due to competitive pressure) is illustrated in Fig 2. Consider sending customer S_1 who pays a *fixed* fee to network N_A . In turn, network N_A contracts with neighbouring network N_B to pay for any congestion caused in or beyond N_B .⁵ N_B contracts similarly with N_C . Now N_A finds its revenue is fixed, but its costs vary depending on how much congestion it allows its customer to cause in other networks. If S_1 is taken over by bot-software that causes congestion in N_C as shown, then the rate that N_A has to pay N_B (which it pays to N_C) for congestion suddenly goes through the roof. This gives N_A a strong incentive to deploy the policer shown, so that it can throttle the offending flow.

2 Roadmap

§3, DDoS: An Economic Problem

This section gives a feel for the economic nature of the beast we are tackling.

⁵ N_B would, of course, complement this with a fixed charge to cover the balance of capacity and operational costs.

§4, Re-feedback for Economists Here we describe the main features of the proposed ‘re-feedback’ solution sufficiently to highlight its economics, without assuming prior knowledge of Internet technology. We explain how the scheme should prevent DDoS, on the assumption that it will be widely deployed, which is the main question addressed by the rest of the paper.

§5, Will Re-feedback Solve DDoS?

The system has to be built before it will be useful, and someone has to start. Here we consider the most likely initial deployment scenarios, which agents will have to act (OS vendors, network operators, etc), who will have to move first, and why it will be in each of their interests to do so.

But, to solve DDoS, defences have to be near-universally deployed. We explain how, once started, deployment incentives will remorselessly increase in order to press stragglers into joining the club. And how the system is designed to protect those who deploy it from those who don’t, while it is in a (possibly permanent) state of partial deployment.

Finally, we explain the economic process that amplifies the effect of low value incentives between infrastructure providers in order to prevent crimes worth very much more—to the perpetrators. These same processes distinguish DDoS attacks from flash crowds of genuine demand.

The paper ends by surveying related work, and drawing conclusions.

3 DDoS: An Economic Problem

Professional attackers expend no more than the effort needed to achieve their ends. For instance, they could avoid their bots being traced by ‘source address spoofing’—making their floods of data packets appear to come from other computers, preferably computers on other operators’ networks. But they don’t bother. They openly reveal their own addresses, because that merely traces the attack back one step; to thousands of compromised PCs in homes and campuses around the world. In 2004, of 1127 attacks seen on one large ISP’s network, only four spoofed source addresses [Han05].

In response, the Internet security community also expends just enough effort to appear effective. It generally defends against what attackers do, not what they *could* do. State of the art defences against DDoS would be rendered useless if bots made it appear as if they were hopping from address to address around the Internet. Nearer the victim, a carefully crafted attack would just look like new connections from lots of different potential customers. There would be no way to know which ones were genuine. Current filtering defences only block lazy

attackers who send streams of packets from a constant address.

As more of these defences are deployed, the arms race will escalate to its next phase. The code for botnets to spoof source addresses is available if they choose to use it. For instance, in May 2006, 70% of the flood of DDoS traffic that rained down on the anti-spam start-up, Blue Security, was alleged to come from random addresses [Wir06].

The only strategy to defend against address spoofing hangs on the hope that all operators will act for the common cause at their own cost. One operator alone cannot stop spoofing being used in attacks on its own customers. Preventing address spoofing depends on the strength of the weakest link in a wall that must be built around the whole Internet—internal walls aren’t technically feasible.

Certainly, many edge networks are voluntarily building their part of this wall for the common good by checking that their own customers’ packets claim to have been sent from one of the valid range of addresses they would expect their customers to use. But, the ‘Spoofer’ experiment running since Mar 2005 [BB05] found that about 25% of operators around the Internet still allow an attacker to appear as if they are sending from another operator’s network.

These sorry tales show just how pathetic defences against DDoS are, as predicted by economic analysis. Security vendors are following the myopic approach of patching holes as attackers find them. In the face of ever more vulnerabilities, ‘penetrate-and-patch’ might maintain short term sales, but it puts the defending side at an insurmountable disadvantage [And01, §4].

The hope that a complete perimeter wall against address spoofing will be erected by voluntary contributions is contrary to the economics of rational self-interest. Varian’s weakest link model [Var02] shows that voluntary contributions to such a perimeter wall from every operator will be no higher than the contribution from the operator who benefits least relative to their own contribution. So it is unlikely a sufficient perimeter wall will ever be built voluntarily. Nonetheless, the Spoofer project shows that 75% of network operators have acted for the common good. Unfortunately, the reduction in deployment cost necessary to achieve this impressively wide deployment has also reduced the quality. 40% of the 75% that have made the effort to deploy, only limit spoofing to a range of 16 million addresses used within their network. Further, deployment has sadly stuck at the the 75% figure for a year [Bev06]. Unfortunately three walls don’t make a castle, and they are three weak walls too.

But, even if all four walls are built to stop spoofing, we don’t believe address-based throttling of attacks is the correct direction. It will certainly raise the bar and is worth attempting. But we believe it will fall from favour

due to excessive collateral damage from hitting innocent demand.

Our disagreement is about where throttles need to be placed: close to all the potential victims or close to all the attackers? It would seem more cost-effective to throttle close to the potential victims, because there are fewer of them. But our concern is not about cost, it's about effectiveness. Defenders are always at a disadvantage, so there is no point deploying defences that we know the attackers will eventually be able to beat, even if they are cheaper. That is false economy. But it is how the market will always push us—tactical solutions make us look like we're keeping busy. But if we don't resist that temptation, we will always be reacting to the arms race. We need to think strategically. That implies working on the *last* problem in the arms race, not just the next one.

The hardest problem we can think of is discriminating an attack from a flash crowd. We believe we can achieve that, but only if we treat individual users as economic entities (customers), not just addresses, so that we can limit peaks in their demand to what they have paid. But that is private information between the customer and their access network operator. Therefore we have a model in mind where throttles need to be distributed—in front of each customer.⁶

If we deploy policers right in front of attackers, we don't need their address. The policer is right there on their line. It will check their traffic whatever address they say it is coming from. We only need reliable attackers' addresses if we plan to throttle their traffic far away from them. We believe throttles remote from attackers will always be limited to treating individual attackers as mere addresses with little meaning attached to each.

We have a further argument. The above depressing story of spoof prevention concerns just one class of DDoS flooding attack. Granted, it is the one that is commonest and hardest to defend against, but defences still haven't been widely deployed against other classes of vulnerability either (e.g. the numerous TCP receiver attacks on *outgoing* server bandwidth [SCWA99]).

In summary, we believe the current focus on source addresses might succeed eventually, but in the process we will embed more cost and complexity in the Internet, and more and more innocent uses of the Internet will become unreliable. Then in the end it will be beaten anyway. Instead, we should tackle the root causes of the Internet's DDoS vulnerability:

- susceptibility of commodity operating systems to viruses;

⁶And fortunately, this approach rides on the back of the strong desire that networks already have to control the costs their customers can cause with excess but not malicious traffic—more general than DDoS.

- the Internet's inability to control conflicting demands over shared capacity resources.

Both problems are fundamentally economic. This paper concerns a solution to the latter problem—an economic solution to a fundamentally economic problem.

4 Re-feedback for Economists

4.1 Congestion in Networks

DDoS attacks cause an extreme form of congestion. We don't just use the word congestion as a vague indication of overload. It is a precisely defined metric that we will be using for settlements between network operators, so we had better understand what it is.

Congestion is measured as the probability of data being discarded. So 1% congestion means 1% of all the data sent into a path through the network doesn't fit, which leads to 1% of packets having to be dropped (often requiring re-transmission by the sender). In essence, instantaneous congestion, $p = (Y - X)^+ / Y$, where Y is the instantaneous total offered load and X is the available capacity. The terminology $(Y - X)^+ = (Y - X)$ if $Y > X$ or 0 otherwise. During a DDoS attack, ten times more data might be thrown at a link than its capacity, leading to 90% congestion ($p = (10 - 1)/10$).

Although one network might loosely be described as more congested than another, congestion will be different at each link in the network, and each will be in a constant state of flux. Fig 2 shows an example scenario that will help the reader conceptualise congestion. The clouds are networks operated by different economic agents. A selection of the mesh of links that make up the network are shown as fine lines connecting circular routing nodes. Three computers are shown all sending to a fourth, the receiver, R_4 . The flows of data are depicted as heavier, curvy arrows. The flows all converge on the same link in network N_D . It should be understood that the traffic on most links in any of the networks that make up the Internet is typically a mixture of flows from many other networks around the world.

Also it should be understood that low levels of congestion are the norm on the Internet. Whenever well-behaved computers send data, they continually try to seek out the maximum possible rate until they sense congestion, at which point they cut their rate, then seek out more capacity again. The millions of well-behaved computers around the Internet are all adjusting their sending rate every time they send a packet—perhaps a hundred times a second—continually making way for new data flows from other computers in other parts of the world, or taking up the spare capacity when others finish.

But, this behaviour is entirely voluntary. Some uses of the Internet, such as voice (VoIP) or real-time video need a minimum flow rate to be usable. In response to congestion, unlike the flows they are competing with, these flows just don't go any slower. All the voluntarily polite flows that are sharing capacity with them naïvely back down. Playing chicken pays off.

A famous tenet of the Internet's design is that it gives inventors the freedom to use it in unexpected new ways. Not responding to congestion is an innovative use (abuse?) of this freedom. A DDoS attack is just another innovative use of the Internet. In stopping DDoS, should we also stop VoIP? Fortunately, there is a huge difference in degree between the two. But, streaming video—hundreds of times more bandwidth than VoIP—and still not responding to congestion can be seriously anti-social and selfish⁷, though probably inadvertently so. Where do we draw the line? Should we block streaming video? Should we block holographic cinema?

Our answer is that *we*, the designers, should not draw the line. Instead the line should be drawn by the invisible hand of the market [CSWB02]. If we had a properly functioning market, the network (supply-side) would attract enough capacity investment to adequately support video streaming if there was sufficient demand. But if any link became excessively congested, the internal price seen by the network would go stupidly high, automatically drawing the line at the right level, as we shall see. But the basic Internet design lacks the information flows to support this market—a problem we address next.

4.2 Congestion Information Symmetry

We said well-behaved sources sense the congestion on their path. If intervening networks could access this congestion information, wouldn't this support a market in congestion? It would, but the information is inaccessible to all networks except the one with the congested link. Congestion is a measure of the absence of data—data that has had to be discarded. If a sorting office observed the flow of envelopes passing through it, how would it know how many letters were lost or held up in a backlog at another sorting office? This is a fundamental information asymmetry in datagram networks.

So, what information does the source use to sense congestion? Stretching the postal analogy, to send a long letter the sender numbers the pages and puts a few pages in each envelope. The receiver sends envelopes back with messages inside (feedback) saying which pages were received. In this way, the sender works out which pages

⁷It is well-known that electronically mediated communication causes people to treat others less politely than if face to face. Packetisation is just an extreme form of depersonalisation, which if properly anthropomorphised would make video streaming analogous to water skiing in a public swimming pool.

were lost or held up en route. To model the Internet using this postal analogy, recall that whenever a sender senses that a letter has been discarded, as well as re-sending, it voluntarily slows down.

But the important point is that sorting offices cannot infer the absence of letters, because the page number information is inside the envelopes, and anyway only the sender knows which pages it sent.

The solution to this information asymmetry is built on the most recent improvement to the Internet protocol called explicit congestion notification (ECN) that was standardised in 2001 [RFB01]⁸. Staying with the postal analogy, we will describe ECN in terms of coloured stamps.

Sorting offices with ECN capabilities keep a little of their capacity in reserve. The sender puts a grey stamp on each envelope so that it will be entitled to use the reserve. Whenever a sorting office has to bring its reserve capacity into use for an envelope, it covers the grey stamp with a red one. Then the fraction of red stamped envelopes should be a measure of congestion, very much like the fraction of discarded envelopes was before.⁹ The idea is that envelopes don't have to be discarded in order to signal congestion, saving all the duplicate work of re-sending envelopes. The receiver should tell the sender whenever a red stamp has been received and the sender is expected to slow down as if a packet had been discarded (voluntarily, as before).

This is the state of the art before we introduce re-feedback. ECN isn't quite right, but it has a nice side-effect for our purposes. The red stamps are on the outside of the envelope, so congestion is revealed and becomes measurable by other sorting offices. But unfortunately the wrong ones.

We want sorting offices to pay for the congestion they cause (or allow to be caused); we want money to travel in the same direction as the envelopes, from cause to effect. But the red stamps are only seen by sorting offices later in the delivery sequence.

The solution is fairly simple. As before, sorting offices cover grey stamps with red more often, the more congested they are. But in addition we define the re-feedback protocol as follows:

- Two new colour stamps are created: green and black. The sender may put either on an envelope, but black stamps must be placed beside grey ones.
 - Grey stamps are 0 (think free).
 - Red denotes -1 (think debit);
 - Black denotes +1 (think credit);

⁸Implemented in most routers, but not turned on—see §5.1.

⁹Strictly, this explicit measure of congestion relates to congestion of the slightly smaller capacity system—as if it had no reserve.

- Green also denotes +1, but it is special as we shall see (think initial credit or ‘deposit’);
- The sender must label each envelope, as well as addressing it. Conveniently it will use its own address (the source address), but it can choose any label it wants.
- Sorting offices will act in their own self-interest; if a ‘flow’ of envelopes carries insufficient credits to balance the debits, it will discard enough packets to keep it out of debt.
- If a sorting office becomes overloaded, it should discard arriving envelopes by colour (‘colour-preferential drop’), only dropping those stamped with credits (green, then black) as a last resort.
- Sorting offices define a flow of envelopes as all those with the same combination of destination address and sender label.
- A sender can only get a new flow recognised by a sorting office if it starts with a green stamp.¹⁰
- A sorting office accounts for any envelopes that don’t belong to a recognised flow as if all these envelopes (the dregs) were a single bulk flow.
- A sorting office will not hold the account for an inactive flow longer than a set, commonly agreed, period¹¹. So, if a sender wants to continue a flow of envelopes after a longer idle period, it has to use another green stamp effectively to start a new flow.

Later (§4.4) we will consider strategies of malicious senders, but for now we will focus on the intent of the system for rational, non-malicious users and networks. The aim is to force senders to declare the congestion they expect in the chain of sorting offices their data is traversing—the number of black (or green) stamps sent should balance the number of red stamps received.

If the sender wants a flow of envelopes delivered *and the destination wants to receive them*, their best strategy is for the sender to start with a green stamp as an opening credit, then continue sending using grey stamps, but every time the receiver receives a packet stamped red, it feeds a message back to the sender, who adds a black stamp to balance it on the next packet it sends. This

¹⁰The idea of explicitly declaring a flow start is very powerful addition to the Internet protocol in its own right. Then nothing on the Internet that handles flows (TCP servers, firewalls etc) needs to set up any flow state unless a green deposit is ‘paid’, protecting it from malicious flow state exhaustion. It is a similar idea to the state set-up flag proposed by Handley & Greenhalgh [HG04], but also given an important economic twist. This protection is used within the re-feedback system itself; to protect routers that detect negative flows.

¹¹For Internet packet delivery, we have proposed this period should be one second. Of course for letter writing it would have to be longer!

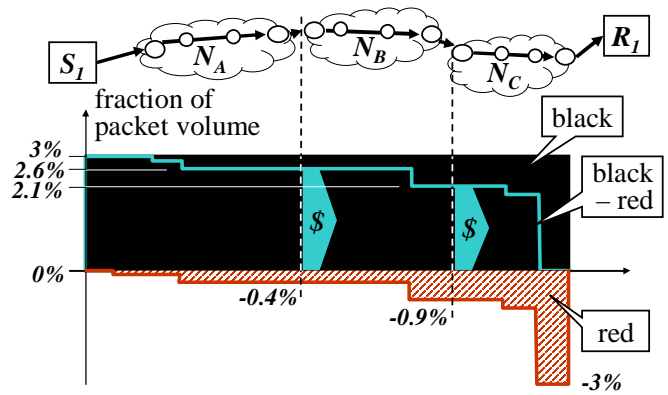


Figure 3: Recursive pairwise congestion charging between networks

is the reason for the name ‘re-feedback’—genuine senders are forced to re-insert feedback from the receiver onto the outside of the envelopes carrying the forward data flow.

Turning from the postal system back to the Internet, increasing numbers of networks are including a usage element in their interconnection settlements [GDL⁺04], largely because of high volumes of peer-to-peer file-sharing traffic. Competitive pressure should drive networks to agree contracts between themselves that include a usage element that tends towards the marginal cost of that usage. Given the cost of network capacity is sunk, usage causes no additional cost unless it causes congestion.¹² So usage charges will tend to the congestion cost. The coloured stamps on packets provide a mechanism for this competitive process to be realised.

The packet colouring scheme has been designed so that it is simple to account for packets between neighbouring networks in such a way that money is automatically distributed from the network allowing congestion to be caused, to the networks suffering as a result, in proportion to how much congestion one causes in the other.

At a network border, for traffic in one direction, the sending network will be expected to pay the receiving network for the bulk volume of all the black & green packets minus the bulk volume of all the red. As is common today with volume charging, they need to agree a fixed price per volume between them and the data volume crossing the border needs to be metered. The only change needed to measure the true marginal cost is for each packet’s charge to be weighted by +1, 0 or -1 depending on its colour. It then becomes a congestion volume charge—effectively a charge for the volume of traffic offered in excess of capacity.

¹²In fact, at the competitive equilibrium, the ratio between usage and fixed capacity charges should be $1/(e-1)$, where e is the elasticity of scale of the cost of capacity (marginal over average cost) [MMV95].

This is illustrated in Fig 3 for just one flow passing through three networks in a row, shown in the top half of the figure. The filled plots in the lower half of the figure show the fraction of black (credit) and red (debit) packets with respect to positions along the path of links through the network. For clarity, we have included green (deposit) packets in with black. The descending staircase superimposed on the black area shows black minus red. And the big arrow heads show settlements between networks determined by this difference. The fraction of black packets remains constant at 3% throughout the path (because black stamps are marked by senders and remain untouched). Tracing the red plot below the x axis, it can be seen that various routers along the way are slightly congested, so each one marks a few passing packets to red, until right near the destination where there is a larger step increase in congestion denoting a more congested router.

It can be seen that the overall effect of paying for the difference between black and red is that N_A pays N_B for all the congestion downstream, whether in N_B or N_C , and N_B in turn pays off the externality N_C suffered. The balance between what N_B gets and receives pays off its own congestion externality. Thus all externalities are correctly internalised.

It might seem that we have given networks a perverse incentive to fake congestion. N_B 's profit depends on how much red congestion marking it introduces. However, inter-domain routing acts as competitive pressure against this incentive. For instance, using Fig 5, if N_B did introduce fake congestion, N_A would find a cheaper route to N_C through another network, perhaps N_E . The perverse incentive to fake congestion of the least congested route would then be limited to the second least congested alternative—the best outcome competition can achieve.

By effectively reversing the information flow, we have cured the information asymmetry that the Internet suffered from previously [CC01]. As Akerlof showed [Ake70], poor market information about quality leads to poor quality services driving out good. With re-feedback, a network must continuously reveal the quality of each route it offers to its neighbours. That is, it not only reveals its own internal congestion, but it combines this information with the quality of its choices of onward routes (its subcontractors) for delivering data to places it doesn't serve directly itself.¹³

¹³Of course, where competition is weak (e.g. at the network edge), and if market regulation is also weak, there may be no other choice of route, or there may be a high cost to switch to another provider, even though this newly symmetric information says it would be beneficial (and therefore a perverse incentive to fake the information could still result). However, these market failures don't affect the way re-feedback turns a DDoS attack into hugely increased costs to the sender's network.

They do result in lower quality, because the immediate intent of congestion notification is to get the sender to slow down. They also result in marginally larger profits for the monopoly network. But these are unsurprising outcomes if natural monopolies are weakly




Figure 4: Aggregation of all downstream congestion externalities

The re-feedback scheme exhibits the elegant property that accounting can be done in bulk at borders, without regard to flows. Fig 4 illustrates this. It shows four representative flows crossing the border between networks N_A and N_B , each one showing along its vertical axis the characteristic downward staircase of downstream congestion from Fig 3. It is only necessary to count the volume of all the black & green packets over say a month and subtract all the red ones, without any regard to which flows they are all in. This balance accounts for the product of the bit rate of each flow and the downstream congestion it contributes to. So if two flows (for example the top two flows) pass through the same congested routers, but one is four times the bit rate of the other (represented by its depth along the axis into the page), within the bulk aggregate four times as much marked volume will correctly be accounted for the faster flow. That is, the bulk aggregate is the sum of all the cross-sectional areas of the flows crossing the border, as shown.

The above re-feedback protocol has been designed to keep the complexity cost of all routers low. It specifies two dropping functions on routers:

- Colour-preferential drop is a very low cost operation that strongly protects genuine users of each router during overload (see §4.4), so an operator deploying re-feedback protections would be well-advised to configure it on every router.
- The function that drops negative flows is more complex (though still very simple), but it need not be

regulated; re-feedback creates no additional problems, but it cannot solve underlying natural monopoly problems in the market supplying the internetwork layer, even though it does fix market failures in its own market layer.

deployed on every router.

We would only expect sampled flow checks for the most negative flows at border routers leaving only the final edge router on any path to do an accurate per flow check, where it is feasible to monitor each flow separately.

It seems wasteful to deliver a flow only to drop it before it gets to its destination. So it might be tempting to send messages upstream to drop it earlier. We advise against such thinking, as it involves over-punishment rather than *proportionate* punishment. Then no-one can exploit the amplifying effect of the punishments to harm someone else. If the flow paid its ‘fare’ to get as far as it did, no-one has been harmed. Anyway, traffic that doesn’t cause congestion costs nothing to transmit. And as soon as it encounters serious congestion, colour-preferential drop preserves service to genuine users.

4.3 Per-user Rate Policing

With re-feedback, the sender is forced to reveal its knowledge of congestion on its path to its own network provider by marking enough packets to black (or green), otherwise they won’t get through the congestion. This cures the information asymmetry we outlined earlier. If the sender under-declares congestion, whether intentionally or because someone else in the feedback loop lied, packets will only traverse the path through as much congestion as has been declared upfront, by which time they will have gone negative and risk being dropped.

These black marks represent the congestion cost the sender is causing to all the networks in its path, so it should be made to pay that cost. But we don’t believe this will happen directly. Odlyzko amassed considerable evidence across all spheres of life to support the common sense view that people are averse to paying unpredictable charges for services [Odl97, §5]. However, as Odlyzko succinctly puts it, the irresistible force runs into the immovable object; people want to act with unpredictable externalities on others, but they don’t want unpredictable charges themselves [Bri02, §3.2].

With re-feedback, we can solve this classic dilemma, because we have balanced the information asymmetry between customer and provider. The ingress network operator can deploy a box to police each user’s traffic¹⁴ to cap the rate at which they cause congestion. But the box can allow a certain degree of give and take. Then the

¹⁴Despite our efforts to the contrary, our work on re-feedback tends to get associated solely with per-flow policing. But we have repeatedly made it very clear that policing is the part of the framework where we expect a market to develop in different policing approaches. We have proposed detailed mechanisms for three points on a spectrum: no user policing, per-user policing and per-flow policing [BJSK06]. Here, *without intending loss of generality*, we focus on per-user policing.

operator can charge a fixed subscription fee in the sure knowledge that its policing box can stop the customer causing more congestion costs than she has paid for. The give and take will lead to gains from some customers and losses to others. This ‘per-user rate policer’ box is also the key to solving the DDoS problem.

The deal the network offers a customer can be simple: In return for a flat subscription of $\$C$ per month, the network operator allows the customer to cause V volume of congestion anywhere on the Internet, spread evenly over the month, but allows an overdraft of V' . This can be conceptualised as a bucket of depth V' continuously filled with tokens at rate V/month . Sending a black packet mark consumes a token from the bucket. If tokens are consumed faster than the bucket is being filled, the bucket level will drop. Once it becomes empty (the ‘overdraft’ is exhausted), pending traffic has to back up, only being released each time one of the regular tokens is added to the bucket. If traffic slows or stops, the bucket gradually fills again with the regular supply of tokens, replenishing the overdraft. If the sender is idle, the bucket eventually starts overflowing, with further tokens irretrievably lost.¹⁵

This type of service offer is simple to understand and weighted proportionally fair [KMT98]. We expect a market in similar service offers to develop (2-stage token buckets etc.). Happily, these types of token buckets and their variants are extremely simple to implement (see [BJSK06, Appx.G]).

4.4 How Re-feedback prevents DDoS

Re-feedback deals with two complementary cases, with and without an intent to communicate:

- Where the sender is trying to communicate with the receiver, the incentive framework traps them between two opposing pressures. They generally want to communicate fast, but the rate policer forces them to go slower, the more congestion (black marks) they declare to the network. So they want to minimise the number of black packets they send. But if they send too few black packets, there won’t be enough to ‘pay the fare’ to get through the path congestion to the destination.
- Where the sender is not trying to communicate, but merely sending dummy attack traffic, it is trapped between the same pressures as above, but it may not care about all traffic reaching a destination—it may be happy to attack the inside of the network. Also, it will be willing to exploit any leeway that the network gives to allow flows to be temporarily negative or to

¹⁵This might seem unfair, but it parallels the fact that communications capacity is a perishable resource—unused capacity cannot be saved for later.




Figure 5: Re-feedback incentive framework

progress a little further than they should once they have gone negative.¹⁶

Fig 5 shows the opposing pressures that re-feedback creates as downward and upward arrows acting to ensure the level of congestion that the sender declares to the network is just sufficient to hit zero at the end of its downward staircase across the network. This staircase represents remaining downstream congestion on the path through the network illustrated below it, which shows the placement of the various functions that contribute to the upward or downward pressures.

Because a router under attack should only drop credit packets as a last resort, if a link is severely congested, the router will give absolute priority to black packets, while dropping most of the incoming packets as well as adding red marks to nearly 100% of those that do get through. The near-100% red marking will cause genuine senders to send near-100% black packets.

If an attacker completely denies the existence of any downstream congestion by sending a flood of grey packets, its ingress policer will allow them through with no throttling. However, as they arrive at the congested link, they will simply all be dropped, because black gets preference.

If an attacker tries to send nearly 100% black packets to get through the congested link, it will quickly exhaust the overdraft in its token bucket policer and be limited to the trickle of black tokens filling the bucket. The problem for the attackers is that the very nature of their DDoS attack causes perhaps 100 to 1000 times more than typical

congestion on the link into which the attack converges. Therefore, the attackers will be forced to use policer tokens very rapidly if they want their attack to reach a target beyond the congestion. This will require the botnet to be about 100 to 1000 times larger than ones seen today to achieve an equivalent force.

In the mean time, genuine users of a server beyond the congested router will sense 100% congestion and also send only black packets. Although these will consume their overdraft, they are unlikely to be sending continuously at line rate like the bots. Therefore, they will enjoy a considerable advantage over the botnet, unless the botnet really can marshal an army 100 to 1000 times bigger than the bigger ones we see today.

The botnet could co-ordinate the dynamics of its attack, with each bot quickly using up its overdraft then holding off to build it up again. This wouldn't achieve any stronger attack if everyone was out of phase, but it could be done in synchrony to generate a pulsed attack. The per-flow variant of the policer mentioned earlier would thwart even this attack.

Yet another strategy the attacker might adopt would be to flood the victim with packets all carrying different source addresses, as if they were all the start of new flows in a flash crowd. Such a 'SYN' attack exhausts the (transport layer) connection memory of a server as well as exhausting (network layer) bandwidth. Servers can return 'SYN cookie' challenges to solve the transport layer problem, but they cannot currently prevent network layer bandwidth exhaustion. If the attacker coloured these packets green, as it should for new flows, it would rapidly consume credits, quickly exhausting its overdraft in the policer, rate limiting further green packets. Any other colour and they would not get through the negative flow dropper, which would treat them all with the dregs.

¹⁶Other, less obvious attacks are open to an attacker as well. As we and others think of these, we have so far been able to harden the system against them without adding complexity. These non-obvious attacks and our defences are described in [BJSK06]

5 Will Re-feedback Solve DDoS?

A pure engineer would answer this question by explaining how it worked. With our amateur economist hat on, we take this question to mean “Will policers be deployed at all, and will policers (or complementary defences) be deployed widely *enough*?”—incentives questions.

5.1 Deployment Bootstrap Incentives

DDoS is just a small part (in terms of value) of the much wider problem space re-feedback addresses. The much higher value reasons for deploying it will be the major factor determining whether it happens. This accords with Ozment & Schechter’s model for reasoning about deployment incentives [OS06], where they pointed out that bundling a security solution with another desirable product can improve uptake. Currently, the costs that one customer can cause to others on the Internet are out of control. The congestion externality of a DDoS attack is just one example of these costs. Although DDoS attacks are of great concern, probably most ISPs would agree that other out-of-control traffic demands are currently of greater day-to-day concern in financial terms, particularly file-sharing.

Two issues are mixed together here. Firstly, ISPs simply want more control over the very wide range of costs that different customers are causing to others. Without this control they can’t extract the low layer value that they know is locked up in the large numbers of more demanding customers. But secondly, ISPs want to sell higher layer value-added services to these more demanding customers, but they cannot compete with other service providers when true costs aren’t being paid for what competitors are currently using. This second issue applies more widely than just file-sharing; it concerns VoIP, video streaming and so on—markets with multi-billion potential that are slipping through the ISP’s fingers.

Most of these higher value products are also demanding to the network. So being able to control demands on the network would prevent customers bypassing the products that ISPs want to offer. At first sight, this seems like a case of the ISPs using a security product to lock-in their customers [And03]. However, re-feedback has been carefully designed to only allow ISPs to prevent theft of their service and no more (‘proportionate punishment’). With re-feedback, ISPs can force true (congestion) costs to be paid when customers use competitors’ services over the ISP’s network. But there is nothing in re-feedback that would stop customers freely choosing another ISP to access the same third party competitor, so there is no lock-in.

It is quite reasonable for ISPs to expect to be able to control demands on their networks. As long as this control

is transparent to the applications being used (‘net neutral’). Re-feedback provides that control, and it is perfectly agnostic to which application is in use. Even if the application is DDoS, re-feedback only controls the congestion caused; it makes no judgement about malicious intent. It is only because the malicious intent *is* purely to cause congestion that re-feedback deals so effectively with DDoS.

Therefore, in the terms of Ozment and Schechter’s deployment incentives model, re-feedback as a solution to DDoS will benefit dramatically from being bundled with re-feedback as a solution to the multi-billion dollar question of ISP cost control. This is a rather special form of intrinsic bundling—‘bundling of a solution with itself’, due to re-feedback being a solution to multiple problems.

But despite such a wonderful advantage in deployment terms, re-feedback suffers from a huge disadvantage too. It requires a change to the Internet protocol, and indeed changes the Internet’s feedback architecture. The Internet community has very little experience of what is required to make a change of that magnitude happen successfully. Differentiated services (Diffserv) is probably the most successful (and perhaps the only) example of a successful architectural change. It actually got taken up quite quickly, but only for enterprise networks—it is nowhere near ubiquitous.

Therefore, to be deployed, re-feedback has been designed to minimise the technical changes required and to make the most of the available strategies for adoption. Focusing on the latter, and borrowing again from Ozment & Schechter, it can also use the ‘co-ordination’ strategy (using alliances), and to a lesser extent the ‘sub-network adoption’ strategy (incremental deployment a sub-network at a time).

Re-feedback gives control over how open or closed an operator wants to be—in the words of David Clark *et al* [CSWB02] it is designed for ‘tussle’. Therefore it is likely to be of interest to those at the closed end of the market first. It hits the big fear and greed buttons, both cost control and revenue defence, so it could become of interest to an alliance of network operators. The most likely would be an alliance of cellular operators who are largely still vertically integrated and threatened most by the openness of Internet technology.

An exhaustive discussion of re-feedback’s incremental deployment incentives is given in [BJSK06, §7.2]. But in very broad overview, initial re-feedback deployment requires two broad technical changes, and a third optional change:

- trivial but essential modification to the sending computer’s Internet software;
- deployment of policing functions at the Internet’s edges and metering at borders;

- optional router modifications to exploit re-feedback’s discrimination against DoS attacks.

Given mobile terminal manufacture and network operation are vertically integrated in the cellular industry, the cellular operators have both the motive and the means to mandate inclusion of re-feedback in mobile terminals and their own network equipment.

If re-feedback does take off in the cellular industry, there will be a pleasant side-effect: most new cellular terminals are being built to be able to roam to other wireless tail technology on fixed networks, such as WiFi and Bluetooth. Therefore roaming could spread re-feedback technology to fixed networks before it appears on the more traditional desktop PC. ‘Cross-infection’ of networks by terminal roaming is an interesting deployment technique that Ozment and Schechter missed from their model.

Another deployment strategy that re-feedback can adopt is also a form of bundling, but rather an interesting one. Re-feedback requires a change to the Internet software on the sender¹⁷. Re-feedback capable senders can use a network alongside non-capable senders. But we recommend that legacy traffic (traffic not marked with one of the colours used in the re-feedback scheme) is rate limited, so it cannot be used to bypass the network’s control of re-feedback traffic. As deployment proceeds, we imagine that network operators will gradually tighten up this rate limit on legacy traffic.

We suspect the ability to degrade the performance of legacy hosts (or equivalently improve the performance of replacements) will be of interest to operating system vendors whose main business model is to earn revenue from encouraging a continual upgrade process. Rather than benefiting from the additional value of a complementary product, it is benefiting from the degradation of a substitute. Given security products tend to restrict, rather than enhance, degrading a substitute is likely to be a generally useful deployment model for security products.

Probably the most important deployment factor in re-feedback’s favour is that it is was designed fundamentally as a technology to align incentives. It is therefore hardly surprising that we have been able to find such strong (and interesting) ways to motivate deployment. And conversely, this most likely explains why security proposals that were not grounded in economics (source address validation, DNSSEC, etc.) are often difficult to get deployed.

5.2 Deployment Closure Incentives

Let us imagine that re-feedback has indeed been implemented in mobile devices. And that major cellular operators have deployed policing boxes around their perimeters, but no non-cellular network has yet followed their

lead. Note, that we use the term ‘cellular network’ to capture the vertically integrated business model aspect of the operator, not the radio aspect, so this term includes the backhaul networks and back-end services run by the cellular operator.

At the border between each cellular and non-cellular operator, the cellular operator charges interconnect usage fees for incoming congestion as described earlier. In the other direction, we assume the non-cellular operators might charge for incoming interconnect data volume, but they would not have the machinery to reliably charge for congestion volume. Further, imagine that numerous mobile devices that comply with the re-feedback protocol are in use on other non-cellular networks because they were designed to be seamless between cellular and non-cellular wireless access.

Now, in common with many security technologies (virus protection etc.), DDoS protection depends on system-wide deployment: Varian’s weakest link model [Var02] is a reasonable approximation. So, the critical question is not just “Will deployment start?” but “Will it finish?” And note that, to prevent DDoS, these other networks don’t have to deploy re-feedback policers - any effective DDoS defence will do (for now).

If one assumes the non-deployers’ inaction was rational, it will have been because they stood to gain less from re-feedback than the cellular operators—less than the cost of deployment. However, once the first movers have acted, deployment costs will reduce considerably. In particular, all the risks of the unknown have been removed and initial research and development costs have largely been recovered.

But the operational finances also seem to start a chain reaction. The non-cellular networks have no re-feedback functions like policers, so they cannot control their outgoing congestion costs. Each cellular network will be profiting from the lack of control the other networks have over their customers. And non-cellular networks will be losing.

For instance, let us focus first on the money movements due to DDoS attacks. When the bots move in they find their attack force is reduced a hundred-fold for attacks that both start and end on a cellular network, because of the rate policers. However, any attack that causes congestion outside the cellular network can succeed. The cellular operators won’t be able to stop incoming DDoS attacks, but they will be able to limit them by prioritising incoming data marked with one of the re-feedback colours. They will also at least recover their costs from their neighbours; attacks from external bots that target high profile sites within the cellular networks (WAP gateways, streaming servers, search engines, caches, location servers) will push very large amounts of money into the cellular networks, through congestion charging at hundreds of times typical levels.

¹⁷Optionally it works best if the receiver is also upgraded.

Without universal deployment, attacks are only prevented if they would have been wholly within the deployment region. Attacks still cross the border between deployment and non-deployment in both directions, but re-feedback is like a valve that only causes money for anomalous peaks in demand to move into the cellular operators, not out. In the outward direction, a DDoS attack will increase volume charges, but the charges won't be hugely amplified by the congestion carried within the specific volume that causes attacks. Because re-feedback forces internalisation of the congestion externality, it has reversed the usual situation: networks that harbour attackers pay victim networks, at least for the infrastructure-related cost of their attacks.

The 'money valve effect' isn't peculiar to the DDoS 'application', it applies for all applications. As we explained earlier, re-feedback deployment will probably be driven more by its ability to control the costs of more prevalent applications (bundling with itself).

We are probably painting a biased picture here, but there is certainly a growing pressure on neighbouring networks to deploy measures to shield the cellular networks from DDoS attacks—to reduce the huge charges cellular networks levy when they are attacked. Importantly, the pressure comes from the company's chief finance officer (CFO). Typically security technology deployment (DNSSEC etc) decisions are pushed by technologists. The chief technical officer (CTO) has an uphill struggle explaining the business case to the CFO. He's never even heard of DDoS. Our goal in proposing re-feedback, is for CFOs of Internet infrastructure companies to be telling the CTO to fix the DDoS problem, not the other way round.

Let us imagine that next generation networks (NGNs) become the next alliance of networks to succumb to the pressure to deploy re-feedback. As ex-telcos, they have similar vertical integration ambitions to those of cellular networks, so the increasing costs of not deploying conspire with their increasing perception of the value of deciding to deploy.

As the deployed bubble grows, the non-deployed bubble shrinks, forcing the bots and other anti-social applications to squeeze into the smaller remaining space (e.g. the 'build it and they will come' ISPs). Therefore, the non-deployers experience ever rising costs. It seems feasible that re-feedback could cause a chain reaction, where the pressure to deploy some sort of defence to DDoS grows inexorably, the more deployment there is. The late adopters may never deploy re-feedback, but they seem to need to deploy at least some effective shield to protect others against bots in their networks, as they become more and more costly to ignore.

5.3 Incentives Not To Be Too Greedy

Finally, we come to the question posed at the start, "Can we discriminate genuine flash crowd traffic from a simultaneous DDoS attack?" The problem here is the profit motive of the infrastructure operator¹⁸. A rate policer is effectively a revenue limiter. Why would an operator limit its own revenue?

The answer to this question gets to the heart of the question in our title. It is in the operators interest to distinguish between the customer's true demand and maliciously faked demand (from a bot). If an operator gets this balance wrong, and a competitor gets it right, the customer will tend to move to the competitor. In other words, the operator risks losing all the revenue from its white market customers just to chase a fleeting temptation that is actually demand from the black market funneled through an innocent customer. Fortunately, our society is still civilised enough that total demand for infrastructure services from the white market dwarfs that from the black, so the white market is not worth sacrificing.

In other words, there is an incentive not to be too greedy *in the short term* (making gains from the black market), due to longer term strategic greed (the risk of losing much greater white market demand).

This line of reasoning implies that there will be competitive pressure to deploy measures against DDoS that are subtle enough to distinguish a flash crowd from a simultaneous DDoS attack—a re-feedback policer¹⁹. We certainly don't believe this pressure exists today. We have shown that re-feedback will encourage deployment of any defences against DDoS, but not necessarily such subtle ones. However, once the bar has been raised against other forms of DDoS, attackers will only be left with the option of attacking during a flash crowd. Only then will competitive pressure push for deployment of the most subtle measure—the re-feedback policer. This is what we mean by working on the *last* problem in the arms race, not just the next one.

This leads to an interesting generalisation. It seems relatively low marginal cost pricing at the infrastructure layer can push back strongly enough to prevent an attack targeted to cause collateral damage that could net thousands of dollars for the attacker at a higher layer. The phone network can't do this, but it seems we might be able to with the Internet. For instance, if Mallory is tendering against Bob for a \$M contract, he can win by tying up

¹⁸Assuming the case of private sector infrastructure, but even public sector operators have to meet budgets.

¹⁹We assume a bot cannot fake a payment to the network operator to make the policer less strict than the one the customer originally bought. We rely on viruses not being able to make financial transactions using the compromised customer's money. Certainly such viruses may be written. But we can expect people to be far more motivated to clear such infections from their machines.

all Bob's phone lines with incoming calls just before the deadline. Mallory would still launch the attack whether or not he has to pay for the calls.

The distinction is the use of marginal cost (congestion) pricing on the wholesale market, which amplifies the distinguishing features of anomalous behaviour, allowing ISPs to discriminate between the black and white markets. They can then isolate the two markets from each other; the black market on which bots are sold for their power to extort money from victims and the white market in Internet infrastructure.

We conclude, admittedly rather speculatively, that self interest will prevent malice; that re-feedback then rate policers will be very widely deployed, all by fixing an underlying market failure, ensuring congestion externalities are correctly compensated.

Of course, this is not a rigorous analysis, for which further work will be necessary. But this is the outline of our intuition from which we intend to build a more solid argument.

6 Related Work

Re-feedback traces its ancestry back to MacKie-Mason and Varian's seminal work on Pricing Congestible Network Resources [MMV95], which was applied as an economic optimisation of the whole Internet in Kelly's celebrated paper on shadow pricing and proportional fairness [KMT98]. Re-feedback uses Kelly's work directly, merely reversing the effective direction of information flow. In this sense, re-feedback has a lot of similarities to MacKie-Mason and Varian's own ideas on how they would have implemented congestion pricing in networks [MMV94], but perhaps re-feedback is more realistic. From another tradition, re-feedback is similar in spirit to Clark's ideas on combining sender and receiver payments in the Internet [Cla96]. Re-feedback also has some broad similarities to Crocker's proposal [Cro04] to solve DDoS.

Many taxonomies have been prepared of DDoS attacks and defences, but if just one were to be chosen, the state of the art is best summarised in Mircovic & Reiher's [MR04]. The research community have proposed a range of novel approaches to various DDoS attacks. A useful summary is provided in the literature review in Yang's paper on receiver capabilities [YWA05].

Bauer *et al* [BFB06] criticise Internet researchers, and specifically our first paper on re-feedback, for omitting to consider malice which would stretch the power of incentive-based approaches beyond their limits. We hope the present paper explains our reasoning better than we managed before.

7 Conclusions

There is no point fixing what DDoS attacks do unless we fix what they could do. DDoS attacks on infrastructure are at their most cost-effective when there is already a flash crowd of genuine demand. Our ambitious aim is to distinguish the two and kill the DDoS attack. We have argued that this requires economic understanding of the two types of traffic sources as 'customers', not just as 'addresses'. This argues for distributed policing by the providers who know their customers better than more remote networks. Conversely it argues that centralised policing will cause unacceptable levels of false negatives as the arms race develops.

We have presented a change to the Internet protocol called re-feedback, that allows congestion externalities to be internalised by removing information asymmetries. We have shown that it encourages network operators to police the congestion response of their own users, and that they are given the correct incentives to deploy the distributed policers we believe will be necessary, both bootstrapping deployment and completing it. Here, completion means winning the last stages of the arms race against network DDoS, at which point near-full deployment will be necessary.

In the closing stages of the arms race, to discriminate between a genuine flash crowd and a DDoS attack on it, we show operators will gain competitive advantage if they use re-feedback policers to forgo the immediate gains from virus-generated demand in favour of genuine demand, because they would otherwise risk losing all the genuine demand to competitors. In short: we have shown how self-interest could indeed be sufficient to prevent malice.

We plan to develop a more rigorous model of the competitive deployment processes described in this paper and would welcome collaboration in this task. We also lay down a general challenge for researchers to try to break the re-feedback protocol, given any proposal aspiring to harden the Internet protocol must receive thorough peer review.

We believe we have struck the right balance that others could learn from. So we also offer our experience as a template for others to consider when securing information infrastructure operated by a large federation of suppliers against unsolicited demand, perhaps including e-mail spam or spam over Internet telephony (SPIT).

Acknowledgements

Thanks are due to Carla Di Cairano-Gilfedder, Ben Strulo, Scott Shenker and Arnaud Jacquet. Also, the anonymous reviewers were particularly helpful in identifying passages where my explanation was poor.

References

- [Ake70] G.A. Akerlof. The market for ‘lemons’: Quality, uncertainty and market mechanisms. *Quarterly Journal of Economics*, 84:488–500, August 1970.
- [And01] Ross Anderson. Why information security is hard — An economic perspective. In *17th Annual Computer Security Applications Conference (ACSAC’01)*, URL: <http://www.acsac.org/2001/abstracts/thu-1530-b-anderson.html>, December 2001.
- [And03] Ross Anderson. Cryptology and competition policy-issues with ‘trusted computing’. In *Proc. 2nd Annual Workshop on Economics and Information Security (WEIS’03)*, URL: http://www.cpppe.umd.edu/rhsmith3/papers/Final_session1_anderson.pdf, May 2003.
- [BB05] Rob Beverly and Steve Bauer. The spoofer project: Inferring the extent of source address filtering on the Internet. In *Proc. Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI 2005)*, pages 53–59, URL: <http://www.mit.edu/~rbeverly/papers/spoofer-sruti05.html>, July 2005. USENIX.
- [Bel03] Steve M. Bellovin. The security flag in the IPv4 header. Request for comments rfc3514.txt, Internet Engineering Task Force, URL: [rfc3514.txt](http://www.rfc3514.txt), April 2003.
- [Bev06] Rob Beverly. State of IP spoofing. URL: <http://spoofer.csail.mit.edu/summary.php> (updated hourly), August 2006.
- [BFB06] Steve Bauer, Peyman Faratin, and Robert Beverly. Assessing the assumptions underlying mechanism design for the Internet. In *Proc. Workshop on the Economics of Networked Systems (NetEcon06)*, URL: <http://www.cs.duke.edu/nicl/netecon06/papers/ne06-assessing.pdf>, June 2006.
- [BJCG⁺05] Bob Briscoe, Arnaud Jacquet, Carla Di Cairano-Gilfedder, Alessandro Salvatori, Andrea Soppera, and Martin Koyabe. Policing congestion response in an internetwork using re-feedback. *Proc. ACM SIGCOMM’05, Computer Communication Review*, 35(4):277–288, August 2005.
- [BJSK06] Bob Briscoe, Arnaud Jacquet, Alessandro Salvatori, and Martin Koyabe. Re-ECN: Adding accountability for causing congestion to TCP/IP. Internet Draft draft-briscoe-tsvwg-re-ecn-tcp-02.txt, Internet Engineering Task Force, URL: <http://www.cs.ucl.ac.uk/staff/B.Briscoe/pubs.html#retcp>, June 2006. (Work in progress).
- [BR05] Bob Briscoe and Steve Rudkin. Commercial models for IP quality of service interconnect. *BTTJ*, 23(2):171–195, April 2005.
- [Bri02] Bob Briscoe. M3I Architecture PtI: Principles. Deliverable 2 PtI, M3I Eu Vth Framework Project IST-1999-11429, URL: <http://www.m3i.org/>, February 2002.
- [CC01] Ioanna D. Constantiou and Costas A. Courcoubetis. Information asymmetry models in the Internet connectivity market. In *Proc. 4th Internet Economics Workshop*, URL: <http://www.m3i.org/papers/ie.pdf>, May 2001.
- [Cla96] David D. Clark. Combining sender and receiver payments in the Internet. In G. Rosston and D. Waterman, editors, *Interconnection and the Internet*. Lawrence Erlbaum Associates, Mahwah, NJ, URL: <http://diffserv.lcs.mit.edu/>, October 1996.
- [Cro04] Steve D. Crocker. Protecting the internet from distributed denial-of-service attacks: A proposal. *Proc. of the IEEE*, 92(9):1375–1381, September 2004.
- [CSI05] CSI/FBI computer crime and security survey. Technical report, Computer Security Institute/FBI Coordination Centre, URL: <http://www.gocsi.com>, 2005. (639 survey respondents).
- [CSWB02] David Clark, Karen Sollins, John Wroclawski, and Robert Braden. Tussle in cyberspace: Defining tomorrow’s Internet. *Proc. ACM SIGCOMM’02, Computer Communication Review*, 32(4):347–356, October 2002.
- [GDL⁺04] Emanuele Giovannetti, Alessio D’Ignazio, Joerge Lepler, Cristiano Ristuccia, and Stefanie Brilon. Initial dataset on transit prices and quality. Deliverable D5-WP1, CoCombine IST project IST-2004-2012, URL: http://www.cocombine.org/pdf/D5_final.pdf, November 2004.
- [GK99] Richard J. Gibbens and Frank P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35(12):1969–1985, December 1999.
- [Han05] Mark Handley. CRN internet architecture WG: DoS-resistant internet subgroup report. URL: <http://www.communicationsresearch.net/object/download/1543/doc/mjh-dos-summary.pdf>, January 2005. (Anonymised digest of presentations at working group launch event).
- [HG04] Mark Handley and Adam Greenhalgh. Steps towards a DoS-resistant internet architecture. In *FDNA ’04: Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture*, pages 49–56, New York, NY, USA, 2004. ACM Press.
- [HR06] Mark Handley and Eris Rescorla. Internet denial of service considerations. Internet Draft draft-iab-dos-05, Internet Architecture Board, URL: <http://www.ietf.org/internet-drafts/draft-iab-dos-05.txt>, July 2006. (work in progress).

- [JKR02] Jaeyeon Jung, Balachander Krishnamurthy, and Michael Rabinovich. Flash crowds and denial of service attacks: Characterization and implications for CDNs and web sites. In *Proc. 11th Int'l Conf on World Wide Web*, pages 293–304, URL: <http://doi.acm.org/10.1145/511446.511485>, 2002. ACM.
- [KMT98] Frank P. Kelly, Aman K. Maulloo, and David K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [MMV94] Jeffrey K. MacKie-Mason and Hal R. Varian. Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. Prentice-Hall, Englewood Cliffs, New Jersey or URL: <http://www.sims.berkeley.edu/~hal/people/hal/papers.html>, 1994.
- [MMV95] Jeffrey K. MacKie-Mason and Hal Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, “*Advances in the Fundamentals of Networking*”, 13(7):1141–1149, 1995.
- [MR04] Jelena Mirkovic and Peter Reiher. A taxonomy of DDoS attack and DDoS defense mechanisms. *Computer Communication Review*, 34(2):39–53, 2004.
- [Od197] Andrew Odlyzko. A modest proposal for preventing Internet congestion. Technical report TR 97.35.1, AT&T Research, Florham Park, New Jersey, URL: <http://www.dtc.umn.edu/~odlyzko/doc/modest.proposal.pdf>, September 1997.
- [OS06] Andy Ozment and Stuart E. Schechter. Bootstrapping the adoption of Internet security protocols. In *Proc. Workshop on the Economics of Information Security (WEIS'06)*, URL: <http://weis2006.econinfosec.org/docs/46.pdf>, June 2006.
- [RFB01] K. K. Ramakrishnan, Sally Floyd, and David Black. The addition of explicit congestion notification (ECN) to IP. Request for comments 3168, Internet Engineering Task Force, URL: [rfc3168.txt](http://www.rfc3168.txt), September 2001.
- [SCWA99] Stefan Savage, Neal Cardwell, David Wetherall, and Tom Anderson. TCP congestion control with a misbehaving receiver. *Computer Communication Review*, 29(5):71–78, October 1999.
- [Sym06] Internet security threat report. Report IX, Symantec, URL: <http://www.symantec.com/enterprise/threatreport/> or http://www.symantec.com/about/news/resources/press_kits/detail.jsp?pkid=threatreport for selected graphs etc, March 2006. (for the period 1 Jul - 31 Dec 2005).
- [Var02] Hal Varian. System reliability and free riding. In *Proc Workshop on Economics and Information Security (WEIS'02)*, URL: <http://www2.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/49.pdf>, May 2002.
- [Wir06] Under attack, spam fighter folds. *Wired* URL: <http://www.wired.com/news/technology/0,70913-0.html?tw=rss.technology>, May 2006.
- [YWA05] Xiaowei Yang, David Wetherall, and Tom Anderson. A DoS-limiting network architecture. *Proc. ACM SIGCOMM'05*, *Computer Communication Review*, 35(4):241–252, August 2005.