GEDDES Fresh thinking

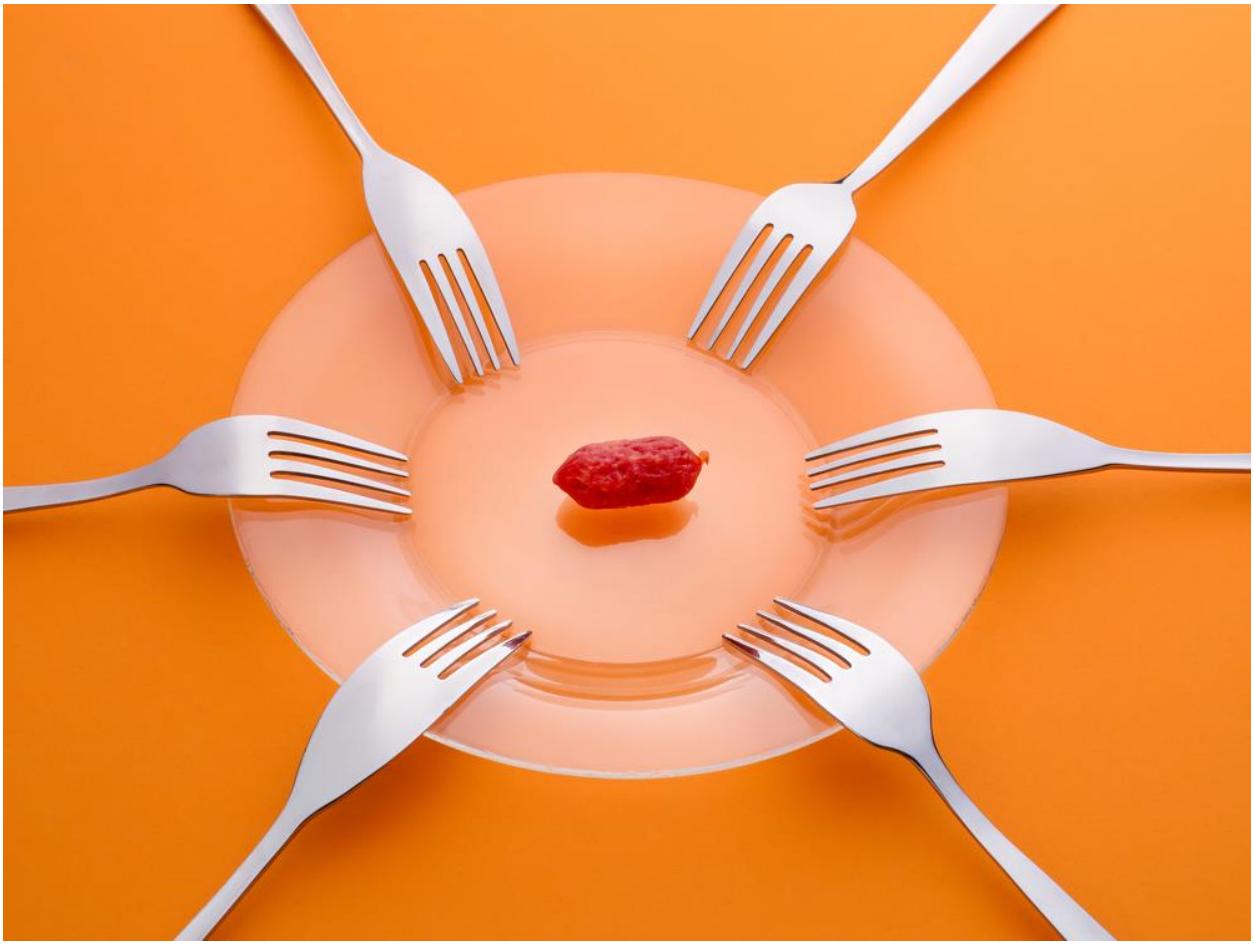

Broadband service quality:
Rationing or markets?

Martin Geddes

 July 2017

About the author

Martin Geddes is a network performance scientist, and a scholar of emerging telecommunications technologies and business models. He holds an MA in Mathematics and Computation from the University of Oxford, and is based in London.

As founder of Martin Geddes Consulting Ltd he advises clients globally on broadband technology, policy, economics and service innovation.

As a co-founder of Just Right Networks Ltd, he is presently working on the launch of the world's first commercial quality-assured broadband service.

martingedd.es

Acknowledgements

I would like to thank Dr Roslyn Layton for her considerable encouragement and advice when preparing this paper. The errors and omissions are my own.

Many of the scientific concepts come from my work in collaboration with Dr Neil Davies and Peter Thompson of Predictable Network Solutions Ltd. The general framework for considering broadband service semantics was jointly developed for a webinar to summarize their scientific advisory work on traffic management detection (TMD). The original report on TMD was commissioned by Ofcom, the United Kingdom telecoms regulator.

Summary

Broadband is increasingly vital to everyday life, as well as to commerce and enterprise. Its importance merits asking the question: **what is the appropriate role of government in managing the market for broadband services?** For historical reasons, the primary agency tasked with answering that question in the US is the FCC.

The FCC has a long history of policing traditional telecom networks, and constraining the actions of players engaged in unfair practices. Over many decades, the FCC has found occasions where dominant carriers to have leveraged their market power to favor or disfavor other players. Whilst this inflates carrier profits, it demonstrably harms consumers.

The subject of "net neutrality" (called "Open Internet" in FCC rulemaking) is about the potential for ISPs to become "king makers" for online services. Its proponents position it as being about the potential for harmful "discrimination" by powerful ISPs. The topic generates diverse viewpoints that are promoted with intense feeling. **This paper proposes that such passion is mostly misguided, as the current policy debate is misframed.**

Consumer advocates have taken a language of fairness for living people, and then misapplied it to packets, whose ethereal nature barely qualifies as inanimate. As a consequence, they have inappropriately focused their "discrimination" concerns on the internal traffic management choices of packet networks. **The resulting regulatory rules unwittingly enforce a grossly unfair rationing of service quality.**

We can successfully reframe the policy problem to achieve a far fairer social and economic outcome. To do so, we must strip purely technical concerns of unwise and irrelevant emotive language. Terms like "discriminate", "throttle" and "violate" are used to fan public outrage, but come overladen with unhelpful semantic baggage. **We need a new policy lexicon**, one that clearly separates mechanistic network processes from orthogonal economic and legal issues.

Richer language and rational inquiry are our newfound friends; they help us to think with crisp precision about fairness concerns in the unfamiliar technology domain of distributed computing. Issues of equitable resource allocation need to be addressed at the human level, where they always belong.

The treatment of packets should only be relevant to the extent that a harm to people exists. Our goal then is to *appropriately* relate fairness to people to that of packet delivery. **This can be achieved by focusing *only* on the external end-to-end service quality.**

This paper proposed an alternative approach that (potentially) meets the needs of *both* consumer advocates *and* free market proponents. It constrains unfair ISP power, whilst also replacing rationing with a fairer market for quality. For this to come to pass, Congress has to express its will that it prefers fair markets over unfair rationing of broadband quality.

Introduction

Broadband is (by definition) the sharing of physical transmission resources using *packet-based statistical multiplexing* (PBSM). It is only affordable as an always-on service because of this statistical sharing of the expensive underlying facilities and capital.

The nature of PBSM means there will often be points in the network where short-term demand exceeds supply. Consequently, resource allocation decisions *must* be made inside the network without recourse to the outside. This means that *traffic management* (TM) is *always* being performed. The operational control over these decisions by the carrier is a power that is intrinsic to broadband service delivery.

As such, all packets networks have the ability to perform *differential* traffic management (DTM). This is TM in which the resource allocation decisions depend upon some aspect of the traffic (source, destination, markings, payload, etc.). This power of *differential* treatment is where "net neutrality" advocates perceive a potential for unfair ISP behavior.

The FCC is confronted with carriers who indeed might conceivably exploit their market power over DTM to harmful ends (not that there is much evidence of such behavior to date). It is unsurprising then that the agency automatically looks to traditional telecom "discrimination" concepts to provide a framework for action.

The core of the "net neutrality" debate thus far been framed in terms of DTM, i.e. the local behavior of network mechanisms that schedule packets. In particular, advocates for regulatory intervention aim to minimize any perceived "discrimination" from DTM. This is achieved by constraining ISP power to engage in it to nebulous "reasonable" use.

The rest of this paper takes on three tasks:

- Firstly, to explain why this aim is a mathematically impossible and technically undesirable objective, even within the false framing adopted by "net neutrality" advocates.
- Secondly, to illustrate why that framing itself is scientifically mistaken and philosophically unsound, since it masks the real underlying issue of the service definition.
- Lastly, and perhaps most significantly, to offer an alternative framing that is scientifically robust and practically implementable, with an execution path via an objective "quality floor".

Why network “neutrality” is not desirable

Is “neutrality” defined, measurable and helpful?

Let us first tackle “neutrality” on its own terms, where it ought to have its best chance of surviving severe intellectual scrutiny.

Any successful regulatory approach to broadband has to support enforcement, and hence survive expert witness challenge. As someone who might someday be summoned in that role, this author would evaluate a “neutral” traffic management policy against (at least) three logical tests:

1. **(Objective)** “Non-discriminatory” (and hence “discriminatory”) behavior must have a clear technical definition.
2. **(Observable)** Any “discrimination” must be operationally measurable as a deviation from “non-discriminatory” behavior.
3. **(Relevant)** The “discrimination” must be relatable to an actual consumer harm, and its mitigation must eliminate the harm.

The approach taken with “neutral” traffic management fails all three tests. This trifecta of fatal flaws may surprise you, so let us digress for a short moment to understand the causal “science gap”. We can then return to examine these issues in more detail.

Physics for spectrum, stochastics for packets

When forming spectrum management policies, we have a deep and well-established body of science to draw upon. The behaviors of electromagnetic waves have been studied for centuries. Physicists have quantitatively modelled the constraints of the real world. For instance, Maxwell’s equations describe radio wave propagation behavior[1].

We can also readily relate the constraints of physics to our policy choices. For example, if we suffer from widespread poor indoor mobile coverage, the answer is not a policy to subsidize building thick stone walls with metal coatings. That would contradict the physics.

In comparison to radio, broadband is a new technology, having been invented during the lifetime of many readers. The science is still being uncovered. Some of the “equations” that define the performance constraints of PBSM have only been discovered in the last

[1] “Maxwell's equations: 150 years of light”, The Guardian, https://www.theguardian.com/science/life-and-physics/2015/nov/22/maxwells-equations-150-years-of-light

decade or so. The algebra of performance is not yet in the textbooks, and its calculus not presently widely known.

Nonetheless, we have many people clamoring for broadband policies that call for magical traffic management outcomes. They are the packet equivalent of faster-than-light communications. In other words, they contradict the constraints of mathematics.

Specifically, the current policy literature on broadband has failed to capture PBSM's quintessential defining feature. In the same way spectrum policy is constrained by the physics of *electromagnetism*, PBSM is constrained by the mathematics of *stochastics*. This is the scientific study of the interaction of random processes, such as statistical multiplexing.

Imagine if the literature on spectrum management omitted any mention of electromagnetism and its properties. You might become concerned that the policy process had disconnected from the science, and thus no longer reflected the reality it represents. Yet that is precisely the situation we face with broadband policy.

Regulating "neutrality" is mathematically impossible

Let us now return to our three problems: why is packet "neutrality" not an objective, observable, or relevant basis on which to prevent harmful "discrimination"?

Firstly, "neutrality" is not an objective phenomenon because the expected behavior of a "non-discriminatory" packet network remains undefined. The user experience is an *emergent* property of a stochastic system. It results from the collective interaction of all those local traffic management processes, together with application and network protocols. In other words, the service quality is the result of "statistical accidents".

Given the service definition is "you got whatever you got" from those accidents, it's rather hard to complain "I didn't get what I expected". When a discontented user of a heavily-loaded network protests "neutrality violation!" to the FCC, the agency is stuck: the user never was entitled to any specific performance for any particular application anyhow. Hence there is no objective "non-discriminatory" baseline against which to compare the supposed "discriminatory" outcome.

It is mathematically impossible to recover "discriminatory" desire from "different" operation (in the general case). That is because you cannot separate out "discrimination-alike" flukes from "deliberate discrimination" faults. It's the general case that the agency is faced with, so will confront overwhelming "neutrality violation" false positives and negatives.

At this point the regulatory vessel is listing alarmingly in heavy stochastics, and its crew ought to be thinking of the location of the lifeboats.

"Neutrality" isn't what matters, anyway

Observing packet "discrimination" is also a technically hopeless task. There is a ubiquitous assumption in the pro-regulation policy literature that there are tools available to detect (undeclared) DTM, and that these tools are fit-for-purpose for regulatory use. Knowing that they were likely to be asked to find such a measurement tool, the UK regulator Ofcom commissioned a scientific study from experts in stochastics[2].

The resulting report[3] concluded that there was no existing technology that had all the required properties: to localize blame in the supply chain, be reliable, and scalable. As a real-world matter, regulating packet "neutrality" cannot be done. That alone is enough to permanently hole this intellectual misadventure below the waterline.

Then there is the humdinger fallacy that finally sinks the whole "neutrality" concept. "Nondiscriminatory" local traffic management is simply not relevant to what matters to users! Measuring it is not merely hopeless, but pointless, too.

Users *only* experience the end-to-end service quality; apart from a few network nerds, they don't care what happens at the traffic management mechanism level. Crucially, the absence of deliberate "discrimination" at the packet level does *not* imply the presence of fairness to applications, and hence users.

Indeed, when "all traffic is equal", you just end up rewarding the greediest users who send the most traffic, and punishing those most sensitive to quality and cost. That's not fair at all. We have observed an ISP service where such "neutral" packet handling has triggered service collapse due to emergent stochastic effects.

Maybe those who proposed these unfit-for-use regulations should sign up for a few undergraduate computer science and statistics courses before they cause more harm to the public.

"Neutrality": seductive framing, scientific folly

As a result of these facts, regulating the "discriminatory" nature of traffic management is a regulatory fool's errand. You cannot recover a user "fairness" doctrine from "neutral" traffic management: it is a practical and philosophical nonstarter. Even if network "neutrality" was objectively measurable, it wouldn't service the needs of protecting users from harm by service providers abusing their market power.

---

[2] "Ofcom publishes scientific report on net neutrality", Martin Geddes, http://www.martingeddes.com/ofcom-publishes-pnsol-scientific-report-on-net-neutrality/

[3] "Traffic Management Detection Methods & Tools", Ofcom, http://stakeholders.ofcom.org.uk/market-data-research/other/technology-research/2015-reports/traffic-management

Traditional telecom discrimination concepts (derived from the carriage metaphor of physical goods) simply don't match the stochastic nature of broadband. It's a new technical category of system, and it requires new concepts for policymakers. The bottom line is that the policy process has gone against nature, and nature isn't changing to accommodate the policy process.

If we care about fairness and non-discrimination, we must abandon packet-level "neutrality" retrofitted into historical policy frameworks. To get out of our undesirable present situation, we have to understand how we got into it in the first place. This demands that we take another little detour, this time into linguistics rather than stochastics

The underlying fallacy of network "neutrality"

"Neutrality" is propaganda for packet equality

The fallacious nature of "neutrality" for packets is located in its implicit framing, as the term "neutral" implies some form of even-handed behavior. This framing in turn is rooted in the loose language used by its advocates. The terms chosen (probably unconsciously) conflate treatment of packets and people, and as a result impute malign intention where none exists.

The starting point is how the word "discriminatory" is used to describe *differential* treatment of packet flows. This takes an objective phenomenon of DTM, and then makes a subjective value judgement about it. In doing so it commits several philosophical and scientific sins.

Firstly, it presumes an intentionality to the specific performance outcome (for that user/application combination), as well as to the service as a whole (for all users and applications). As noted earlier, the emergent nature of the performance of today's broadband offerings means that the service lacks such intentionality; the user experience is all a big, and mostly happy, fluke of fortune.

Sometimes there is an unacceptable emergent user experience[4]. In this case we *may* see a deliberate "bias" by ISPs to the emergence process. This could be by changing TM algorithms or configuration, or using DTM. This drives a *different* emergent outcome, which is still not truly *intentionally* engineered.

Next, the idea of "discriminatory" TM presumes that the intention of any such bias is always a bad one. This is a rhetorical leap of Olympic-winning proportions! If a broadband service provider responds to user needs to direct resources towards a better experience, that is hardly "discriminatory", no matter how extreme the differential treatment.

Lastly, the term conflates technical and economic issues. Any "discrimination" has to be with respect to some context; it is sometimes acceptable to choose movie actors based on skin color to fit a story narrative, but rarely (if ever) those of theatre ticket staff. In our case, the context is the service that the user and edge provider have paid for (if they have paid at all). If you have bought a different service quality that is transparently offered, and then get it, then that's hardly "discriminatory", no matter what TM is applied.

---

[4] The Internet is Broken, and How to Fix It, Jim Gettys, https://gettys.wordpress.com/2012/06/26/the-internet-is-screwed-up-and-how-to-fix-it/

A packet paranoia: who stole my performance?

The term "throttling" is used as the injurious twin of the rather less exciting *rate limiting* of traffic. It conjures up notions of the deliberate depriving of users of their presumed entitlement to reach anywhere and everywhere on the Internet at peak access line speed.

Those who actually engage in the management of real networks, however, will recognize that rate limiting is often used to prevent over-saturation of downstream resources. This *improve* performance and the user experience. So much for the intended value of transparency of TM rules.

This is then often followed up with the idea of a resulting "violation" of the "principle" of neutrality. When we violate people we are engaged in serious criminal acts of violence against individuals. When we "violate" packet neutrality, however, we are transgressing a totemic 1980s academic paper on the "end-to-end" nature of the Internet[5]. These are hardly morally comparable.

This highly-cited paper makes no quantifiable predictions or refutable claims[6]. It is at best an aesthetic (for that is all it offers), one whose prescriptions are mostly observed in the breach in real broadband networks. It is unheard of to use the ideas within it to guide leading-edge network research today.

"Net neutrality" advocates are extrapolating what is little more than a technology folk tale into a doctrine of network design, and giving it the full backing of law and the state's monopoly on violence. If anyone is being "violated", it is innovators and future generations who need a more dependable and affordable digital infrastructure!

Lastly, the result of all these felonies against fragments of data flows is seen as a "distortion" of the marketplace. This again presumes some deviant desire to subvert some natural order, which presumptively advocates of "net neutrality" have special access to.

The messy (and also emergent) nature of free markets and business models is deemed unacceptable. After all, it might result in a system that competes with (and even supplants) their hallowed technical and economic model of the Internet. Just like how the Internet disrupted telecoms before it. That cannot be allowed, now, can it?

---

[5] End-to-end arguments in system design, J.H. Saltzer, D.P. Reed and D.D. Clark, ACM Transactions on Computer Systems, Volume 2 Issue 4, November 1984

[6] A critical review of the end-to-end argument, Martin Geddes, http://www.martingeddes.com/think-tank/the-future-of-the-internet-the-end-to-end-argument/

How to transcend the “neutrality” framing?

The misapplication of intentionality

What unifies all these issues is the additional of subjective values to what are essentially objective technocratic issues. The repeating pattern is one of implying (usually bad) intent to “discriminate”, “throttle”, “violate” and “distort”. This regulatory malice of forethought is remarkably achieved without ever creating a framework for defining it, or its relationship to actual network operation.

The source of this false intentionality to broadband services has an interesting origin. It appears to comes from the misapplication of ideas of (common) carriage to what is really a distributed computing service. ISPs are really a corner case of a bigger space of potential information services.

With physical goods, the fair treatment of the unit of information (like delivering a book) is a strong proxy for fair treatment of its recipient. This held true for the corner case of traditional telecoms circuits. However, this is fundamentally *untrue* of broadband; is categorically different.

For instance, we only care about populations and not individual packets; there is an unfamiliar “high-frequency trading” of resources with two degrees of freedom; its statistical nature negates assumptions of intentionality; elastic protocols have no equivalent in the physical world; packets may be broken apart in flight and later reassembled part-way.

Whilst the legal principles of common carriage have centuries of history, their application needs to be reconsidered for the new distributed computing world. “Net neutrality” has been offered as a construct to bridge these common carriage ideals into broadband and the cloud.

The problem is that its conceptualization of “discrimination” is a naïve translation of “package” to “packet”. This is an anthropomorphic fallacy that (by definition) wrongly treats packets as if they were physical things, rather than as arbitrary divisions of data flows from dueling protocols and applications. We are seduced into believing there is a shared intentionality between the packet and people domains, where there is none.

Nonetheless, the issue of fairness (to people) is undoubtedly a real and pressing one. So if not “neutrality” as a basis for non-discrimination, then what else? What does it mean to be ‘fair’ to broadband users? What is a good or bad intent? How can we reasonably reformulate ‘discrimination’ for distributed computing services?

The real issue: the service is undefined

To answer these, we need to address the fundamental issue: the broadband service is not fully defined. **The core problem is absence of what computer scientists call "intentional semantics".** It is as if the advocates of "net neutrality" have an unconscious belief in a benevolent *deus* in their packet *machina*. This spontaneously produces "good" emergent outcomes, in the absence of our wicked "discriminating" ISPs being allowed to express intent.

If we want to properly define "fairness", the absolute prerequisite is to specify some form of intent to the service that does not involve appeals to magic. Until that is done, it is a moot issue whether the intent is good or bad, or whether a particular user or application was "discriminated" against.

Today we characterize broadband services by their bearer (e.g. cable, DSL, 4G) and service "speed". This does not sufficiently describe the service. **Specifically, measures of average or peak "speed" do not define the *quality* on offer, and it is the quality that determines fitness for purpose in use.**

This lack of clarity over the service definition should concern both side of the policy debate:

- "Net neutrality" advocates are rightly unhappy that the carrier does have an arbitrary and unconstrained power to effectively redefine the service at any point post-purchase. It might be fit-for-purpose one day, and not the next, with no user redress.
- Meanwhile, supporters of free market enterprise would like transparency of fitness-for-purpose, so (given that sufficient competition exists) users can determine which services might be substitutes for one another, and a market may form.

We can unpack this service definition problem from different perspectives:

- As a *political* problem, we need to constrain the power of dominant carriers who are insufficiently help in check by the free market. This means we have a choice over what to constrain. The real issue is fairness and justice *only* to people, and *never* to packets.
- As an *economic* problem, we need to acknowledge that the broadband market is still young. Its delivery models are evolving with new technologies like Internet of Things, 5G, SDN/NFV[7], and RINA[8] to name but a few. Since the "best effort" service

[7] "SDN, NFV, and All That", Yaakov Stein, IETF Journal July 2015, https://www.internetsociety.org/publications/ietf-journal-july-2015/sdn-nfv

is undefined, quality is not being priced. As the price mechanism is not in operation, we are faced with rationing (e.g. data caps). **"Net neutrality" is an appeal to the network gods to perform the rationing in a divine way (that can't be done).**

- As a *technology* problem, broadband is based on the statistical sharing of a resource, so its regulation needs to be anchored in science of randomness (i.e. stochastics). This is currently notable by its absence, both in the policy literature and the resulting rules[9].

[8] "A new kind of network: RINA progress update", Martin Geddes, http://www.martingeddes.com/a-new-kind-of-network-rina-progress-update/

[9] "Beyond 'neutrality' - how to reconnect regulation to reality?", Martin Geddes, http://www.slideshare.net/mgeddes/geddes-beyond-neutrality

Choose markets over rationing

A new "post-neutrality" policy framework

A policy framework that aligns to the reality of PBSM has to address each of these perspectives.

From a *political* standpoint, we need a clear understanding of what "fairness and justice" policy objective(s) are being pursued. There can legitimately be differences of opinion – this is a purely political question. These would then be relatable to a quantitative definition of what the service is supposed to do in pursuit of those objectives, thus turning politics into published policy.

What are the kind of fairness to we seek to promote at the social, individual and application levels? What kinds of network "success" and "failure" match our values? In a finite world, who should shoulder the disappointment of not having access to unbounded resources, and how should this reflect ability and willingness to pay?

Taking the *economic* point of view, we would unquestionably like our digital infrastructure to have sustainable economics. The "all traffic is equal" approach aims to minimize the differential allocation of resources, whereas mathematics tells us that regulators should be seeking to maximize it. The quality needs of application demand are highly diverse, and network supply must reflect this to be viable.

If we are to replace rationing with a true market, where prices reflect actual resource costs, then we need a way of describing the service qualities on offer. How should the ISP service be technically described for experts and B2B buyers? What, if any, consumer "labelling" requirement should there be for quality? How to express technical capability in terms they can relate to? How many qualities should be on offer, whether for a single broadband service, or across the market?

The *technology* point of view requires us to have services that "say what they do, and then do what they say". Broadband networks are man-made worlds, and within the constraints of physics we *can* choose to engineer them to have predictable performance, rather than rely solely on emergent properties that could vary or vanish at any time.

Our ideal world would have an objective service definition with an objective means of measuring the operational behavior to determine if it is compliant. We would be able to assign cause when there is non-compliance to players in the supply chain, and do this with known and bounded false positive and false negative rates. This would turn published policy into a practical system of enforcement.

The "quality floor" approach

There is a clear path forward that addresses these issues. The first step along it is to recognize a simple truth: **the only thing that matters is the customer experience on offer, and not the internal workings that delivered it**. This provokes a new debate around broadband policy, one which can offer far more productive outcomes that the fruitless "neutrality" one.

We call this approach a "quality floor". It is an objectively measurable minimum quality level that is a strong proxy for application fitness-for-purpose.

As a political device, we believe that **"non-discrimination" means simply getting what you paid for, based on a clear service promise**. The carrier may choose to over-deliver to particular users and applications, and good luck to them, but the harm of the arbitrary power to under-deliver is taken away.

With a quality floor approach, regulators should *not* be in the business of defining the broadband services and the shape and structure of those floor(s). Rather, the only consideration is whether abuse of market power exists that leads to consumer harm. In such cases, regulators may choose to impose a *higher* quality floor on some players. The circumstances that might justify this is a question for economists, not us network mathematicians.

As the economic level, **a quality floor allows market pricing to emerge for different quality levels, both at retail and wholesale (interconnect)**. This is important for the rational allocation of resources, and to create incentives to invest in more resources. The cost should fall onto those whose demands for quantity and quality are the greatest. This avoids the harm of rationing, whereby users who might have been willing to pay for better quality (or take a lower price for worse) find their needs are not met.

As a technical approach, a **quality floor introduces the new science of digital experience quality**. We believe there is only one mathematics that describes quality adequately, just as there is only one physics that describes electromagnetism adequately. We call this new quality framework 'ΔQ', and it models what matters most to user experience: the instantaneous probability of packet loss and delay.

By adopting robust science for a quality floor, we avoid the harms associated with weak metrics that do not convey fitness for purpose, or which fail to create a basis for regulating digital supply chains. Any allocation of blame has to stand up in court and survive expert scrutiny. That implies certain technical properties are needed for the metrics and methods of proof of blame for breaking a quality floor.

Congress should investigate the feasibility of "quality floors"

The current regulatory approach of an industrial policy for packets is doomed to failure. The false fairness doctrine of "all traffic is equal" will have serious undesirable customer experience and network cost consequences; it is mathematically unsustainable. The resulting harm to consumers will damage the credibility of the FCC and create demand for Congressional action.

What the FCC, Congress and industry all need to accept is that **broadband cannot ever be retrofitted into legacy regulation designed for circuits**. ISPs offer distributed computing services, and these cannot be tortured into a "carriage" metaphor derived from physical goods.

A new framework is required that engages with the virtualized and statistical reality of the technology. Thankfully, there is an opportunity to reframe of the broadband policy debate around the new science of digital experience quality[10].

A "quality floor" offers the potential to transcend the poisonous feuds around "neutrality". It does this by placing *all* the emphasis on the end user, and *none* on the delivery mechanisms. Our sole concern should be fitness-for-purpose of the service, embodied in a fair and just quality floor for the price paid.

An electricity user gets the voltage and amperage they paid for, regardless of whether the power was generated via solar, coal or nuclear technology. A gasoline buyer gets the octane rating they paid for, no matter what the source of the crude oil. Likewise, a broadband user should get the "quantity of quality" of information exchange that they paid for. This should occur without regard to the bearer technology or its internal configuration.

Replace unfair rationing with fairer markets for quality

By adopting price mechanisms for allocating quality, rather than rationing through randomness, we can achieve far fairer and more efficient outcomes for consumers. The new debate is about the shape and setting of the retail quality floors, and how the supply chain supports that outcome. Appropriate "quality contracts" will have to be in place at wholesale interconnect points, and these should be monitored.

There is a great deal of work to be done to define and explore this emerging regulatory opportunity space of quality floors. For progress to be made, Congress needs to give the FCC the right instructions.

---

[10] The new discipline of Digital Experience Quality", Martin Geddes, http://www.martingeddes.com/the-new-discipline-of-digital-experience-quality/

The three Congressional actions required are:

1. To instruct the FCC to **establish a body of peer-reviewed science** and a reasoning framework in which to discuss broadband quality (including a means of costing and pricing finite resources that are statistically shared).
2. To indicate a legislative desire to abandon the rationing of quality, and instead **institute a transparent market pricing model for quality**. This should explicitly be tied to a quality floor approach, which should undergo a feasibility study.
3. To **set the high-level societal goals for user fairness and justice** (such as protecting the weakest in society), and indicate where intervention is deemed to be justified (e.g. for public safety, abuse of market power, or persistent market failure).

This offers a "win" for everyone. The FCC gets to reassert an unambiguous authority, and takes up a vital impartial role for the technical aspects of regulation. Consumer advocates get a more effective tool with which to hold ISP feet to the service delivery fire. Lovers of free enterprise no longer have to concern themselves with lawyers determined to get into network design. And politicians can sell it as replacing the current "weak neutrality" with something they can market to the public as being a "strong neutrality" alternative.

GEDDES Fresh thinking

To learn more

I recommend reading the presentations Essential science for broadband regulation and Beyond 'neutrality' - how to reconnect regulation to reality?.

I have written a review of the FCC's consultation on Open Internet Transparency, which falls well short of the level of scientific rigor required.

Read more about the science of quality at qualityattenuation.science.

There is a reading list of articles, papers and presentations on my website. This includes a detailed reading list on this topic of broadband quality regulation.

Sign up to my free Future of Communications newsletter and follow me on Twitter.

Get in touch

Thank you for reading this paper. Feel free to reach out to contact me:

mail@martingeddes.com

I offer educational workshops in this subject area, and frequently work as a 'thinking partner' to senior executives.

To learn more about the services on offer visit www.martingeddes.com.

Martin Geddes Consulting Ltd
Incorporated in Scotland, number SC275827
VAT Number: 859 5634 72
17-19 East London Street, Edinburgh, EH7 4BN