July 2023
Geoff Huston

On Centrality and Fragmentation

I attended a workshop on the topic of Internet Fragmentation in July. The workshop was attended by a small collection of Australian public policy folk, some industry representatives, folk from various cyber-related bodies and those who have a background in Internet Governance matters. It was a short meeting, so the perils of fragmentation were not discussed at length, as they often can be, but the concerns about the breakup of the essential bonds that keep the Internet together was certainly palpable in that meeting.

On the other hand, there is a concern, that no doubt is voiced in different venues at different times, that not only are the bonds that tie the Internet together too strong already, but these bonds are exclusively controlled by a handful of digital behemoths whose collective agenda appears to be based more on the task of ruthless exploitation of everyone else in the singular pursuit of the accumulation of unprecedented quantities of capital and social power.

It's proved impossible for me not to share my thoughts about this apparent contradictory state of affairs!

There seem to be two dominant themes in the enumeration of potential perils that face the Internet these days, and oddly enough they seem to me to be opposite in nature.

The first is the theme of ***centrality*** that points to the current set of Internet behemoths, the so-called group of 5, the online giant enterprises which dominate the Internet, namely, in order of market capitalisation, Apple ($3.0T), Microsoft ($2.5T), Alphabet ($2.0T), Amazon ($1.9T), and Meta ($1.0T). If the Internet was ever built upon the foundation of vibrant open competition, then those days are truly over, and the current Internet landscape contains an elite cartel of enterprises which are at a global scale that appears to defy most national efforts to impose regulatory constraint upon their actions. The Internet is, in the eyes of a cynical commentator, just a set of five web sites that do little more than just reference each other in an endless cycle.[1]

The second is the theme of ***fragmentation***, which appears to be based on the premise that the Internet is in danger of splitting itself asunder, creating a set of mutually disconnected segments[2]. The premise of Internet fragmentation points to the struggles at a technical level in maintaining a coherent and single naming system in the face of ongoing fragmentary efforts such as alternate root domains, blockchain based name systems and various forms of obfuscation, coupled with the issues with address fragmentation and routing incoherency. The second theme in this space is termed commercial fragmentation which points to various commercial practices including peering and interconnection arrangements, traffic discrimination, content dissemination constraints. The term fragmentation also applies to the diversity of the various public policies in the area of Internet governance with often

[1] https://www.abc.net.au/radionational/programs/futuretense/cory-doctorow-enshittification-platform-capitalism/102492918
[2] https://icannwiki.org/Internet_Fragmentation

conflicting national efforts in the areas of content curation and censorship, privacy and data protections, intellectual property rights, national security and regulatory purview. They all are specific barriers to the larger concept of the Internet as a global, open, universally accessible space for collaboration and communications.

Centrality

We have witnessed the emergence of a small set of massive commercial enterprises that appear to have a commercial imperative to impose a uniform regime on the Internet environment. There are a few throwaway statistics to illustrate the level of centrality in today's internet. Google's Chrome browser is used by more than 62% of the world's users, and if you include Edge, which is built upon the Chrome engine) the number rises to 68%. The next largest is Apple's Safari with a 21% market share.[3] If the Internet is just a collection of web pages, then the browser application rules, and from this perspective it's a Google-coloured Internet.

These days the mobile device is the dominant way of interacting with the Internet, so maybe we should look at the market share of the operating platforms for mobile devices. Here the picture of centrality is even more pronounced, with Google's Android platform claiming just under 71% market share and Apple's iOS second with just under 29% market share. All other platforms squeeze into the 0.5% that's left.[4]

What about the other side of this picture? In the web hosting content world GoDaddy appears to the dominant provider, with 44M distinct web sites, or some 33% market share. The next largest in terms of number of distinct websites is Cloudflare with 12M sites, or 9% market share, and Google with 8%.[5] In the domain name system, we can look at the ranking of domain names by query count and then look at who hosts these highly used domain names. A similar picture of market concentration emerges, with 35% of queried named hosted by Amazon-O2, then Cloudflare with 9%, Google with 8% and Akamai with 4%.[6] These days streaming video dominates network use profiles, and this market sector is dominated by Netflix, Amazon Prime Video, Disney+, and HBO Max. Just 8 content platforms, namely those operated by Akamai, Google, Microsoft, Disney, Netflix, Meta, Amazon and Apple account for more than 70% of the volume of traffic on the network.[7]

Google have 93% of the search market. The next closest is Bing, with a market share of under 3%. Do users really care about tracking and privacy? DuckDuckGo's market share of 0.51% would tend to suggest that this is not a widely shared concern.[8]

What do these observations say about an open competitive Internet? Not very much, unfortunately.

A conventional view of aggregation in a market is based around the exposure of economies of scale, where the unit costs of the production of goods and service drops as the volume of production increases. Competitors who enter such markets were at a natural disadvantage, as they have to bear significant losses at the outset to maintain a competitive price for their goods as compared with a high-volume incumbent. Regulation of such markets is typically focussed on the market share of the dominant provider, and once the dominant provider passes some threshold value of market share, then intervention of some form is triggered, with the objective of restoring some semblance of competitive access into the market by competitive providers.

---

[3] https://www.oberlo.com/statistics/browser-market-share

[4] https://www.bankmycell.com/blog/android-vs-apple-market-share/

[5] https://firstsiteguide.com/web-hosting-stats/

[6] https://www.potaroo.net/ispcol/2022-11/dns-ctl.html

[7] https://ripe86.ripe.net/presentations/49-The_New_Encyrpyted_Stack_RIPE.pdf

[8] https://www.oberlo.com/statistics/search-engine-market-share

But maybe this conventional view of the operation of markets and the influence of economies of scale in the cost of production misses the key characteristics of digital markets. Markets with a high reliance on data have positive feedback loops. The more data an enterprise can gain through operating its product, the more effective the positioning of the product. This leads to strong data-driven network effects. A search engine like Google can improve its search results by using the data in its search database that it continually collects from its billions of users.

Large data sets tend to be broader, in so far as such large data sets tend to provide detail in a larger spread from the average. Even a poorly designed algorithm can find more valuable information and insights in high volumes of various data than a superior algorithm can when working with a more coherent, but smaller, dataset. Google's chief scientist, Peter Norvig, admitted as much in 2011 when he observed that: "We don't have better algorithms than anyone else. We just have more data."[9] The result is that consumers may be locked into using a dominant service by reason of superior relevance through better accuracy. Large data has a much higher utility value.

This applies today in the formative models of ChatGPT and similar generative language services. There is no essential intelligence in these models, but a formidable data set combined with mechanistic rules to summarize and classify the data. This then allows the system to produce new outputs which are not a clone of any single source, but rather a pastiche of the collection of related texts reassembled in a novel way. The larger the original data set the greater the ability of the system to generate responses which match our intuition. In that way the system appears to exhibit intelligence, but this ability is not based on a deep model of cognitive processing of deduction and inference from a limited data set, but a model of pattern recognition and word prediction across massive data sets.

It could be argued that data is taking the roles of both labour and capital in the digital economy. Those entities who are able to amass extremely large data sets are in as uniquely privileged position to create digital service outcomes that are customisable to individual transactions, an option that is not available to entities who are working with smaller data sets.

The challenge for the public sector is to formulate an effective framework to counter the inherent tendency of such highly centralised data-driven markets to exploit their users. As Cory Doctorow observes: "First, they are good to their users [to attract both more users and business advertisers]; then they abuse their users to make things better for their business customers; finally, they abuse those business customers to claw back all the value for themselves."[10]

Fragmentation

Fragmentation is not a new or unique aspect of the computing and communications realm. Back in the 1980's each computer vendor adopted a unique approach to their equipment that did not necessarily interoperate with that of any other vendor. The differences were often very fundamental, such as in the number of bits in an addressable "word" in the computer's memory, and the ordering of the bytes in a word from left to right or right to left. The result was an inevitable vendor lock-in for customers, as switching vendors was a highly disruptive and expensive process. Unsurprisingly, customers quickly lost their tolerance for this situation, and there was conscious effort by customers to force the industry to build computing devices that were based on open standards such that they could interoperate with products from other vendors. In this way customers had some level of assurance that a multi-vendor computer environment was feasible, and that the various devices would interoperate if not seamlessly then at least in a functionally adequate manner.

There were no regulations that enforced such outcomes, but the existence of various industry standards assisted customers to state their desire to purchase standards-based products and to guide providers as to the required attributes of their products. The adoption of these standards was a market-based

[9] https://www.forbes.com/sites/scottcleland/2011/10/03/googles-infringenovation-secrets/?sh=368cad7830a6

[10] https://www.wired.com/story/tiktok-platforms-cory-doctorow/

movement where producers and consumers were both motivated to abide by these standards as an act of self-interest. But adoption of standards intended to facilitate interoperation and admit the potential for product substitution are not the same as the imposition of regulations over behaviours. The Internet was a direct product of the progressive deregulation of the world of telephony, where the actions of a strongly competitive sector was meant to act as a counter-measure for the excesses of monopolistic exploitative behaviours by both the private and public sector operators of these national telephone monopolies. The Internet entered the picture at a time when consumers had lost patience with what we might characterise as fragmentation of the computing industry, and it emphasised the value proposition of coherence and interoperability for consumers.

But these phases appear to be cyclic, and at some point fragmentary pressures surface in otherwise coherent and consistent environments. The Internet appears to be no exception here. As noted in the introduction, fragmentation of the Internet can occur in a number of ways, at the technical level, at the commercial level, and at the public policy level.

At the level of the Internet's common infrastructure, the topic of fragmentation arises in the administration on the Internet's naming and addressing schemes, and in the operation of the Internet's routing system. In this case the fragmentation discussion is about defining norms and safeguards that apply to every service provider in operating this common infrastructure, to enable uses to experience a single seamless Internet on top of this diverse collection of individual component services. The assignation of a DNS name or an IP address prefix to be used by an entity is intentionally an Internet-wide unique assignation. Fragmentation in this context of common infrastructure directly confronts the assumptions about uniqueness and global scope. In the case of the Internet's namespace, the DNS, there is an implicit assumption that we all resolve domain names in the context of the unique DNS root zone. If some applications or hosts were configured to use different name resolution contexts, then effective communication is impaired, as names would no longer be useful as a unique and consistent reference label. The same name would be associated with potential different services, depending on who is asking, when and where.

The story of Internet addresses has been one of some unexpected twists and turns. While the exhaustion of the supply of IPv4 addresses was well understood, and extensively analysed, the implications of this situation were not. The widely held expectation was that the Internet would migrate to IPv6 and its 128-bit address capability. The most optimistic expectation was that this migration would be largely complete before the IPv4 address pools had been fully exhausted. What has happened is that the network largely completed a transition into a client/server architecture and dispensed with the requirement for unique persistent addressing for client hosts. At the same time servers turned toward name-based service addressing, allowing a large number of distinct services to share a common endpoint IP address. This shift in the use of IP addresses has taken much of the pressure off with address exhaustion and the migration to IPv6 has been a protracted exercise. The IPv4 address space has been fractured with the widespread use of private addresses and NATs, but the use of name-based frameworks for service identification and service integrity has meant that this fracturing of the end-to-end model of address coherency has not been of major consequence. The side into what one could call a name-based network service model has appeared to be relatively seamless, and oddly enough has been facilitated by strong shifts to centrality in the roles of content provisioning.

Fragmentation in the name space started with various efforts to augment the name space through the use of alternate root zones, and while they attracted some attention from time to time achieved very little in the way of mainstream acceptance. More recent initiatives have centred around alternate name resolution structures, looking at blockchain-based distributed ledgers for name registration and shift away from a single root name hierarchy into distributed hash tables. It is debatable whether such efforts have been directed towards fragmenting the structure of the name system or are more focussed towards testing innovations in name system using the at-scale Internet as the test platform! IN any case the result so far has been largely the same, in that such alternate name systems have still achieved very little in the way of mainstream acceptance. We've had to go to some extraordinary lengths to resist name fragmentation in some cases. The efforts to pull Unicode into the name system to support non-Latin scripts in the name

space has been a challenging exercise. The Unicode character set is everything ASCII is not. Its ambiguous, in that there are multiple ways to encode the same visual appearance, it's not predictable, in that the same glyph can be presented in multiple ways, it does not have a clear upper/lower case distinction and the DNS itself is not 8-bit clean. The choices were to adopt a new DNS environment for each script collection, or somehow fold Unicode into the constrained ASCII character set. The first solution was clearly going to present fragmentation issues for the name infrastructure, while the latter solution, mapping Unicode into an ASCII representation, enrols application behaviour into the DNS, something the DNS had tried hard to avoid until then. The "solution" of IDNs is far from ideal. It adds various vulnerabilities into the user presentation of the application environment., by virtue of permitting visually identical name strings to be represented by distinct Unicode glyph sequences which in turn encode to distinct domain names. But the main factor in favour of this approach is that is keeps the DNS as a cohesive and unfragmented space. So far, the effort to resist fragmentation of the name space has been largely successful. So far, we've been able to operate the Internet in a mainstream of the name space that provides the essential attributes of universality and cohesion. But there is a cost of this outcome. The incumbents in the provision of name infrastructure are few in number and large in size, and the longer-term trend is fewer and larger. The price of this outcome is centrality. The price is an ever-increasing level of dominance by incumbents and an increased level of resistance to all forms of evolutionary change. Obviously, there are no clear answers as to which is the preferred outcome here.

The routing platform is also experiencing some fragmentation pressures. The role of the network in teleporting clients to services has been reversed with the shift to content distribution systems that are intended to bring content ever closer to the client. The ever-increasing proportion of traffic that is delivered from these content distribution systems point to the decreasing relative commercial importance of shared transit infrastructure, and as long as local clients can be connected with content delivered from local servers, then the majority of client requirements can be met without reliance on long distance transit services, nor on the routing system that holds this framework together. The routing system shows some minor signs of fragmentation, and the span of the routed address space is not the same in every default-free network. However, such fragmentation is rarely visible, in that it manifests itself in the situation where some clients cannot directly reach other clients, which in a client/server network is a marginal circumstance most of the time in any case.

This consideration impacts on the issues of the inter-provider commercial agreements in the carriage space. Don't forget that it was the asymmetric nature of the inter provider agreements in the telephone world that provided much of the impetus for US position to withhold the Internet from ITU-T oversight in the first place, so it was no surprise to see the Internet adopt a position where inter-provider arrangements were an outcome of market processes. Universal connectivity in the Internet was not an assured outcome, but an aspirational objective arising from a disparate collection of bilateral arrangements between providers. However, as already noted, the rise of content distribution systems and the move by the largest of the digital giants to use their own transmission infrastructure wherever feasible has meant that the level of relative importance of this particular topic of inter-provider arrangements has receded in commercial terms. These days it is far more important for an access provider to be serviced by well-populated nearby content data centres than it is to secure high-capacity low-cost transit arrangements.

The third aspect of fragmentation considered here is the area of governance and public policy. Like many booms, the initial pace of change with the explosive appearance of the Internet took much of the public sector (and everyone else) by surprise. The transformation from a public national telephone monopoly to a largely unregulated space dominated by venture capital and entrepreneurial disruption was chaotic at best. The initial public policy aspiration here was to treat the Internet as a poster child for the beneficial power of market-based disciplines, where competition between providers would be aimed at efficient production of goods and services that were able to meet users' expectations of desirable and affordable service. The public policy mantra of the day was probably along the lines of: "Deregulation works! We've removed structural inefficiencies and enabled a new generation of competitive providers to focus on the current needs of consumers in innovative and effective ways. Just look at the Internet's success!"

And for a while that may have been the case, but the market quickly skewed, and aggregation of providers quickly became a dominant theme. As Peter Thiel, a veteran venture capital funder, has observed "competition is for losers!"[11] Many market sectors were exposed to digital transformation, such as entertainment, retail and even social interaction, and were quickly dominated by either a single actor or a select clique operating as an effective cartel. National regulatory frameworks that were intended to counter such overarching levels of market dominance and the attendant risks of abuse of consumers were simply inadequate to effectively engage with these large-scale digital enterprises and the global nature of their scope. However, this imbalance was a temporary situation, and the public sector has reacted by increasing its skill set and acting with a greater level of confidence in rule making with the intention of curbing such excesses.

For example, the entire issue of consumer privacy was subordinated by the demands of a rapacious model of surveillance capitalism. What privacy and consumer protection frameworks that did exist appeared to be covered by the simple advice to "do whatever you want, but don't lie about what you intend to do!" The often cavalier attitudes to the collection and protection of personal data[12] prompted the European Union to adopt the General Data Protection Regulation (GDPR) to force the industry to act more responsibly, at least with the data relating to EU citizens. There is also the EU Digital Markets Act (DMA) and Digital Services Act (DSA) that attempts to force large online platforms to interact with others on a fair and open basis, at least in the context of their activities as they relate to the EU markets.

Will other nation states adopt precisely the same regulatory measures? If history is anything to go by the answer will be an approximate yes, and each national regulatory regime will craft their own customised set of similar measures. The handful global behemoths are doubtless more than capable of customising their activities to meet each of these local constraints and obligations, but the ideal of the Internet as an open, accessible global marketplace for a diversity of providers, large and small, tends to fall aside in this complex fragmented policy compliance environment.

Pick your Peril

It seems that we are dealing with a set of diametrically opposed potential perils at the same time.

On the one hand we have the issue of unprecedented levels of centrality in the digital space, where a handful of massive enterprises are now stifling many forms of competition, diversity, and innovation. Their objective is to maximise the extraction of value from end users, commercial customers, and anyone else who is unfortunate enough to be in range, including sovereign nations. When social media giants take on nation states in a competition for mind share, then so far, it's the social networks that appear to exercise a greater level of societal influence. In terms of the prospects for open, diverse, and supportive societies that make up our national communities, it's hard to see beyond the deliberate efforts to ferment division and social tension as part of the ruthless exploitation of our attention span to maximise platform value above everything else. Centrality is simply not going to end well if this is where we've got to!

Fragmentation has its own pitfalls and risks. The moment in time when the Internet was a coherent space when anyone could reach into a global community, and not founder on the rocks of geolocation blocks, national censorship, arbitrary content blocks, arbitrary application of variant intellectual property rules, incompatible technical diversity appears to have been a very fleeting one. Without a compelling framework to advocate interoperation, consistency, and coherency in this digital world then it's likely that entropy will simply increase, and thus unique coherent communication platform will disintegrate of its own accord. That may well cause more serious and deeper fractures in our world.

Maybe the underlying problem here is that the Internet has been just too successful. As Douglas Adams observed in the Hitchhikers Guide to the Galaxy: "Meanwhile, the poor Babel fish, by effectively

[11] https://www.wsj.com/articles/peter-thiel-competition-is-for-losers-1410535536

[12] Such as the situation in the US, for example. See https://www.nytimes.com/wirecutter/blog/state-of-privacy-laws-in-us/

removing all barriers to communication between different races and cultures, has caused more and bloodier wars than anything else in the history of creation."

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net