

March 2023
Geoff Huston

An Economic Perspective on Internet Centrality

The IETF met in November 2022 in London. Among the many sessions that were held in that meeting was a session of the Decentralised Internet Infrastructure Research Group, (DINRG). The research group's ambitions are lofty: **DINRG** will investigate open research issues in decentralizing infrastructure services such as trust management, identity management, name resolution, resource/asset ownership management, and resource discovery.

So no, this is not the "blockchain research group" per se, but these decentralised systems are relevant to the conversation. The evolution of distributed ledger technologies and the platforms that leverage this model has given rise to the development of decentralized communication and infrastructure systems, and experiments with the same. Some examples include name resolution (Namecoin, Ethereum Name Service), identity management (OneName), distributed storage (IPFS, MaidSafe), distributed applications (Blockstack), and IP address allocation and delegation.

However, we are simultaneously seeing the evolution of use cases (such as certain IoT deployments) that apparently cannot work (or work poorly) in centralized deployment scenarios. So, if centralized systems are supposed to have all these intrinsic problems when they attempt to grow to the size of the global Internet, why is the Internet completely dominated by a small collection of global behemoths? Why are these giants able to defy this reasoning that scale and decentralisation are meant to go hand-in-hand?

Is it because we have been working on the wrong protocol base for the network? Is it some supposed paucity of robust decentralised technologies that inhibits this decentralisation and the associated emergence of diverse service operators? Do protocols such as HTTP or routing protocols lend themselves naturally to centralised operation by a small number of large-scale operators? Should we take Web3 seriously? Can the dominant position of Google, Microsoft and Amazon be attributed to the HTTP protocol? Or the DNS? Could this situation be "fixed" by turning to blockchain-based names and the Web3 framework? Can we find the root cause of centrality in the protocols we use?

A current work in progress, a **draft** on the relationship between Internet standards and the centralization of network functions, exhorts protocol designers to avoid designing protocols that cast a single entity in a role of controller or resource arbiter. However, it's a message that has probably been a feature of Internet protocols since its inception. With perhaps the singular role of the root zone of the DNS we have largely avoided such pitfalls of centralized outcomes based on the communications protocols we use. We've tended towards highly distributed protocol models that attempt to avoid assigning any party into the role of a controller or arbiter. Even when coordinated action might have been of material assistance, such as in technology transitions, as with IPv6, we've avoided such scenarios, relying instead on uncoordinated piecemeal adoption.

I don't think that the level of centrality we see in today's digital world is not directly the result of the communications protocols we use, and trying to adopt a different family of communications protocols would have little bearing on this, unlikely as such a protocol transition would be to achieve in any case. If it's not our protocol base, then are the market pressures that sustain market consolidation largely economic in nature and would such consolidation occur with any technology platform?

The nature of the delivery of goods and services over the Internet are a salient consideration here. The previous limitations to absolute size in the physical world included the ability to amass a large pool of human workers in a single location or gather control of sufficient capital to fund a large-scale operation. In a digital world no such physical constraints apply and the potential for expansion of a digital service is almost unlimited. This probably explains the wild boom phenomenon in the investment sector where the traditional metrics on assessing the future value of an enterprise were based on the assessment of the enterprises' ability to amass and orchestrate physical resources whereas the digital market was essentially one of unconstrained growth potential. In the digital world the marginal cost of additional production is extremely low, so the assessment of the potential market size of a service is not inherently limited by the size of the service provider, whether it's the size of the workforce or the size of available capital.

This situation does not necessarily lead to centrality as an outcome, but the behaviour of the investment sector, but it helped. The venture capital market was looking for opportunities where a digital operator was proposing to disrupt an established activity and provide them with the capital wherewithal to not only disrupt the existing activity base but establish a unique single digital operator in its place. The investors' cherished hope is to create not just the circumstances of centralisation with a digital operator, but, if possible, install a monopoly as the new incumbent. This "winner take all" model of venture capital investor expectation is a pervasive feature of the digital disruptive space. Displacing the existing set of providers with a competitive set of digital players is a second-best outcome. Replacing a competitive structure with a monopoly is very much a preferred outcome.

In classical public sector economics, one of the roles of government is to detect, and presumably rectify, situations where the conventional operation of a competitive market providing services to the public has failed. Of course, a related concern is not just the failure of a market, but the situation where the market collapses and simply ceases to exist. Perhaps markets are more than enablers of simple transactions between a buyer and a seller. Karl Marx was one of the first to think about the market economy as a global entity, and its role as an arbiter of resource allocation in society. When we take this view, and start looking for potential failure points, one of the signs of incipient failure is that of "choke points" where competitive investment levels fall and there is a single critical service point that is the enabler for the entire market space. Any study of an economy involves understanding the nature of these choke points. Telecommunication services are not an isolated case but can be seen as just another instance of a choke point in the larger economy. Failure to keep the market functioning efficiently and effectively can have implications across many other areas of economic activity.

From this perspective, there is this perception that today's centralised digital behemoths have seized far too much social and economic power on today's world. But exactly what is it that gives these entities such powers? How these powers manager to position themselves outside of effective competition and outside of effective regulatory curbs? If it did not take the accumulation of a large pool of labour, capital and resources to create the current incumbent then why are there any inherent barriers to competitive entry into the same market? Why isn't the position of today's digital behemoths subject to constant competitive erosion? What sustains their size and control in today's digital space?

Drivers of Centralisation and Monopolies

It might seem to be the case that a market that is constantly changing does not lend itself to consolidation and domination by a few actors. The Adam Smith model of commodity markets has goods being traded at the marginal cost of production. Each producer must adjust the market price of their goods and services to match their competitors, yet also use a sustainable price point. Excessive margins open up opportunities for competitors, while negative margins are unsustainable. New market entrants are motivated to use marginal improvements in the means of production to establish their products at a new price point but are disinclined to take radical steps that would destroy the market completely. Yet this does not appear to be the case today and we are living in a post-Adam Smith economic world. What are the drivers that tend to work against the balance of competitive interests in digital markets? Let's look at a few here.

Intellectual Property Rights

It seems that the twentieth century has broadened many of our concepts of what is *property* and how this property is *owned*. Here I mean *own* in the sense of exclusive exclusionary rights to exploit the property. Thomas Jefferson *wrote* in the early nineteenth century that, once articulated, ideas inscribe themselves into the thoughts of others. No one can rid themselves of an idea once they have it, and everyone can access the idea once it comes into existence. As he put it: "Inventions then cannot, in nature, be a subject of property."

However, during the twentieth century in the US primary production and manufacturing industry sectors have been gradually supplanted by information processing as the principal source of jobs and wealth creation. The impact of this fundamental change was an increase in the perceived need for intellectual property rights that took the precepts of property and ownership and applied them to non-tangible forms including techniques and processes. This created a new market, not of ideas per se, but of the property of these ideas, and reinforced exclusive ability to exploit these ideas in the market for goods and services.

This process dramatically accelerated in the latter part of the twentieth century and remains with us today. For example, Microsoft filed a handful of intellectual property patents in its first years of operation, while now it files some 6,000 to 10,000 patents annually and a significant proportion of its corporate asset valuation is based on this patent portfolio. Google's 2011 \$12.5 billion purchase of Motorola was considered to be motivated by the latter's 17,000 patents, and its subsequent sale to Lenovo for just \$2.9 billion could be partially explained by Google retaining Motorola's patent portfolio.

However, it's unclear as to the extent to which these patent filings act as disincentives to competitive entry and support the further entrenching of monopoly positions. They certainly support centralisation but are probably not the sole determinant factor.

Large Data

Markets with a high reliance on data have positive feedback loops. The more data an enterprise can gain through operating its product, the better the product. This leads to strong data-driven network effects. A search engine like Google is able to 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 we observed that: "We don't have better algorithms than anyone else. We just have more data." The end result is that consumers may be locked into using a dominant service by reason of better accuracy. Large data has a much higher utility value.

This applies today in the formative models of ChatGPT and similar AI services. There is no essential *intelligence* in these models, artificial or otherwise, 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 role of both labour and capital in the digital economy. Those entities who are able to amass extremely large data sets are in a 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 off smaller data sets.

Network Effects

A *network effect* occurs where the benefit to a user of a product or service increases as the number of users increases. This was evident in the 1980's with the introduction of the fax machine into the corporate world. While the benefit of owning the first fax machines was minimal, the benefit increased as the number of users increased. The multiplier is often quoted as the square of the number of users, as this is a rough approximation of the number of pair-wise transactions. More recently we have seen the same effect in social network platforms where size, as in the number of active users, begets greater size.

The concept of a *network effect* dates back to Theodore Vail's promotion of AT&T at the start of the twentieth century, noting that maximal value of a telephone service for a given number of subscribers was achieved with a single telephone service operator. It was restated as "Metcalfe's Law", conceived in 1983 by a co-inventor of Ethernet, Robert Metcalfe, and given a more quantified formulation in that the value of a network rises in proportion to the square of the number of its users. It has been variously debated in the ensuing forty years, but the general concept that value rises as the number of users rise is not disputed.

Digital platforms have a high propensity to undergo network effects because they function on interoperability and communication with other similar products and users. A novel service or application needs to adequately meet some need at the outset. Once a critical mass of users has been established, this user base becomes the major driver for its further growth. For example, in the case of Microsoft's Office product, once it became the dominant word processor, user familiarity and software compatibility meant that more and more users were encouraged to adopt the product, creating a defacto "standard" for office productivity tools. While a network effect does not inevitably result in a "winner takes all" outcome, it appears that for digital platforms, the network effect is a major determinate of market dominance.

Capital Investment Strategies

Even taking into account the regulatory risk factor, services that operate on a monopoly or highly restricted competitive basis generally offer significantly better returns on invested capital. If this position also has high barriers to entry, then the incumbents can translate this position into an enduring position that further supports the position of the incumbent or incumbents. This lack of competitive pressure allows the service price to include a monopoly premium, allowing the price to be set far above the marginal cost of service production.

This is often expressed as a "winner take all" market position, but this description can misattribute cause and effect. It's not that the market leader has the ability to remove competition and thereby establish itself as a monopoly, but the capital investment strategy is to select incipient monopolies and support their efforts to create a monopoly position. For the investor or venture capital provider the rewards from their investment are potentially far higher.

As the tech investor Peter Thiel has **observed**, "competition is for losers", and "if you want to create and capture lasting value, look to build a monopoly."

But the caveat is that one should never confuse size with value.

As Peter Thiel also writes on this topic: "This means that even very big businesses can be bad businesses. For example, U.S. airline companies serve millions of passengers and create hundreds of billions of dollars of value each year. But in 2012, when the average airfare each way was \$178, the airlines made only 37 cents per passenger trip. Compare them to Google, which creates less value but captures far more. Google brought in \$50 billion in 2012 (versus \$160 billion for the airlines), but it kept 21% of those revenues as profits—more than 100 times the airline industry's profit margin that year. Google makes so much money that it is now worth three times more than every U.S. airline combined. The airlines compete with each other, but Google stands alone. Economists use two simplified models to explain the difference: perfect competition and monopoly."

The result is what we see today where the capital investment profile is weighted heavily in favor of enterprises and service platforms that are establishing a unique position or market that has no existing digital competitor. By assembling a technology platform that is buttressed by patents, assembling a large-scale data collection to support the service and then assembling a pool of early adopters to start the process of building a networked effect, then the result is sustained high value returns from the initial capital investment.

In short, the preferred path for capital is to a highly centralised outcome. Little wonder that this is exactly what has ensued for this industry.

Responses to Centrality - A Brief Historical Perspective

Given all that has happened in the past thirty years is sometimes hard to concede that the nineteenth century saw a similar wave of dramatic social change. The industrial age transformed the means of production from small scale cottage industries to large scale production factories, with its associated massive displacement of people, occupations, and wealth. We saw the transformation of transportation systems with the railway boom, and the transformation of communications by the telegraph and then the telephone. These technologies facilitated corporate entities to project power over longer distances and a new set of industrial enterprises supplanted local activities with regional and national enterprises. By the end of the century, we saw the emergence of a number of industrial giants, many of whom are still with us today. In the US there was US Steel, J.P. Morgan, Standard Oil, and General Electric, to name a few.

At the start of the twentieth century a member of the US Supreme Court, Louis Brandeis, argued that big business was too big to be managed effectively in all cases. He argued that the growth of these very large enterprises that were at the extreme end of the excesses of monopolies, and their behaviours harmed competition, harmed customers, and harmed further innovation. He observed that the quality of their products tended to decline, and the prices of their products tended to rise. When large companies can shape their regulatory environment, take advantage of lax regulatory oversight to take on more risk than they can manage, and transfer downside losses onto the taxpayer, we should be very concerned.

It is hard to disagree with Brandeis if this outcome is an inevitable consequence of simply being big and given the experiences of the 2008/2009 financial meltdown, we could even conclude that Brandeis' observations apply to the financial sector today. But do these systemic abuses of public trust in the financial sector translate to concerns in the ICT sector?

Brandeis' views did not enjoy universal acclaim. Others at the time, including President Theodore Roosevelt, felt that there were areas where there were legitimate economies of scale, and that large enterprises could achieve higher efficiencies and lower prices to consumers in the production of goods and services without undermining the salaries and working conditions of their labour force, by virtue of the economies achieved through volume of production. The evolution of the auto manufacturing industry in the early twentieth century, and the electricity industry both took exotic and highly expensive products and applied massive scale to the production process. The results were products that affordable by many of not all, and the impact on society was truly transformational.

This perspective that there are reasons why we might decide to tolerate very large enterprises was further reinforced by the events of the so-called panic of 1910-1911, where a 25% fall in the US stock market index was attributed to the breakup of the Standard Oil Company and the American Tobacco company under the application of the Sherman Anti-Trust Act. The US administration of the day moved to implement regulatory oversight over these corporate behemoths, but not necessarily act to dismantle their monopoly position, and the Sherman Act was not applied in the US for some decades thereafter.

The prevailing view in the twentieth century was that market forces would act as a sufficient balance against corporate abuse of monopoly market powers. Moreover, the trends in the nineteenth century that supported the projection of power at a national level continued through the twentieth century into the projection of power at an international or even at a global level. The benefits to the country that hosted the corporate domicile of these multi-nationals was such that corporate behaviours that would be excessively anti-competitive within a national market were seen as enhancing national commercial interests in an international setting. The global dominance of the US automobile industry in the 1930's is probably a good case in point here. The rewards for expanding the projection of corporate power from a national scope to a global enterprise were such that many nations tolerated domestic market domination by a small set of incumbents in order to capture market share in international market spaces, and thereby transfer value back into the domestic economy. It is readily apparent that the United States domestic economy has been the major beneficiary of the emergence of the global digital economy.

Centrality in the Digital World

Business transformation is often challenging. What we have today with the rise of content and cloud providers into dominant positions in this industry is a more complex environment that is largely opaque to external observers. What matters for consumers is their service experience, and that depends increasingly on what happens inside these content distribution clouds. As these content data network (CDN) operators terminate their private distribution networks closer to the customer edge, the role of the traditional service providers, which used to provide the connection between services and customers, is shrinking. But as their role shrinks then we also need to bear in mind that these carriage networks were the historical focal point of monitoring, measurement, and regulation. As their role shrinks so does our visibility into this digital service environment.

It is a significant challenge to understand this content economy. What services are being used, what connections are being facilitated and what profile of content traffic are they generating, and just how valuable is it?

This brings to the forefront the same question from Brandeis over a century ago: Is *big* necessarily *bad*? There is little doubt that the digital environment is dominated by a small number of very big enterprises. The list of the world's largest public companies, as determined by market capitalisation, includes the US enterprises Apple, Microsoft, Alphabet, and Amazon, and the Chinese enterprise Tencent. Admittedly, there are other metrics of size that includes metrics of revenues, profits, customers and social impact, but the considerable market capitalization of these five companies places them in the global top ten, which makes them extremely *big*. But are they *bad*? Or to use less pejorative terms, when does their size generate risks to the balanced and stable functioning of the larger economy in which they operate? When is an enterprise so big that failure is untenable in terms of regional or global economic stability? At what point does the level of overarching size of these centralised entities present an unacceptable economic risk?

The global financial crisis of 2008 explored the concept of "too big to fail" in the financial world. Do we have a similar situation with some or all of these digital service enterprises?

Regulation?

If the only oversight mechanism is national (and in some cases regional) regulation, then have we allowed the major multinational actors in the digital service sector to become too big to regulate? Any company that can set its own rules of engagement and then behave in a seemingly reckless fashion is potentially damaging to the larger economy and the potentially threatening to the stability of social democracy. One need only mention Facebook and US elections in the same sentence to illustrate this risk of apparently reckless behaviour. To quote Brandeis again: "We believe that no methods of regulation ever have been or can be devised to remove the menace inherent in private monopoly and overwhelming commercial power."

But if we choose to reject Brandeis' view and believe that regulation can provide the necessary protection of public interest, then it is reasonable to advance the proposition that we need to understand the activity we are attempting to regulate. Such an understanding might be elusive. In the digital networking world, we are seeing more and more data traffic go 'dark'. Content service operators are using their own transmission systems. They are using encrypted transport over the public Internet to ensure that the transactions between the service portal and the user are not exposed to casual inspection by third parties. This withdrawal of traffic from the shared public communications platform is now not only commonplace, but the limited visibility we have into this activity suggests that even today the private network traffic vastly overwhelms the volume of traffic on the public Internet, and the growth trends in the private data realm also is far greater than growth rates in the public Internet.

How can we understand what might constitute various forms of market abuse by these service providers, such as dumping, exclusions, deliberate efforts to distort a market, or discriminatory service provision, when we have no real visibility into these private networks? Yet these private networks are important. They are driving infrastructure investment, driving innovation, and indirectly driving the evolution of the residual public communications realm. Are we willing and able to make an adequate case to expose, through various mandatory public filings, reports and measurements, the forms of use of these privately owned and operated facilities and services that provide services to the public? Should these large data sets that these enterprises have collected be opened up and exposed to all forms of public inspection? Do we have regulatory power to do so considering the size and powers of the entities we are dealing with?

We've seen in the past situations where national regimes have attempted to avoid the test of abuse of market power by handing the problem to another jurisdiction. The anti-trust action against Microsoft was eventually undertaken in Europe and even then, the result was largely unsatisfactory. Even if we might believe that greater public exposure of the traffic carried by these private networks, and public exposure of the data that their applications are collecting might be in the public interest, we might simply not have the capability to compel these service operators to undertake such public reporting in any case.

Consolidation?

The internet has been constructed using a number of discrete activity areas, and in each area appeared to operate within a framework of competitive discipline. Not only could no single actor claim to have dominate or overwhelming presence across the entire online environment, but even in each activity sector there was no clear monopoly position by any single actor.

Carriage providers did not provide platforms, and platform providers did not provide applications or content. The process of connecting a user to a service involved a number of discrete activities and different providers. The domain name being used can from a name registrar, the DNS lookup was an interaction between DNS resolver application and a DNS server host, the IP address of the service was provided by an address registry, the credentials used for the secured connection came from a domain name certification authority, the connection path provided by a number of carriage providers, and the content was hosted on a content delivery network, used by the content provider. All of this was constructed using standard technologies, mostly, but not exclusively defined through standards developed in the Internet Engineering Task Force (IETF).

This diversity of the elements of a service is by no means unique, and the telephone service also showed a similar level of diversity. The essential difference was that in telephony the orchestration of all these elements was performed by the telephone service operator. In the Internet it appears that there is no overarching orchestration of the delivered composite service. It would be tempting to claim that the user is now in control, but this is perhaps overreaching. Orchestration happens through the operations of markets, and it appears that the market is undertaking the role of resource allocation. However, the user does have a distinguished role, in that it is the users' collective preference for services that drives the entire supply side of this activity.

But this is changing, and not necessarily in a good way. Services offered without cost to the user (I hesitate to use the term “free” as this is a classic two-sided market instance where the user is in fact the goods being traded to advertisers) have a major effect on user preferences. However, there is also the issue of consolidation of infrastructure services. As an example, Alphabet not only operates an online advertising platform, but also a search engine, a mail platform, a document store, a cloud service, a public DNS resolver service, a mobile device platform, a browser, mapping services to name just a few. It is also the owner of a considerable portfolio of transmission systems and is perhaps the single largest operator of undersea cable systems. It appears that in this case it is one enterprise with engagement in many discrete activities. The issue with consolidation is whether these activities remain discrete activities or whether they are being consolidated into a single service.

There is a more insidious form of consolidation in the Internet than we’ve seen to date with various corporate mergers and acquisitions. It’s not the individual actors that are consolidating and exercising larger market power, but the components within the environment that are consolidating. Much of this is well out of normal regulatory oversight, but the result of accumulation of broad market power are not dissimilar to the outcomes of corporate consolidation.

Future Tense

The Internet was purported to be a poster child for so many potential positives in the communications world. It was not positioned as a public sector monopoly, in contrast to its telephone company forebears in many parts of the world. It was not implicitly rigid and unchanging in its technology and its service base, but could sustain continued innovation and evolution. Indeed, this innovation was not necessarily orchestrated across all parts of the network at once, but was intended to take place in a way that was termed *permissionless innovation*. What was meant to sustain this environment was a deregulated space that enabled vibrant competition between providers at every level. The focus of all of this effort was the end user, and delivery of goods and services that best met their needs.

Obviously, that situation not what we have today. The acute centralisation at the upper levels of the network protocol stack has sucked most of the oxygen from everywhere else and concentrated the vast wealth of this sector into the hands of a few global behemoths. But they do not sustain their globally dominant position by employing millions of workers. They have not commandeered such vast amounts of capital that no other prospective competitor can enter into the market space. They do not have the traditional control over physical and capital resources that were the previous essential elements of a monopoly. The traditional barriers to competitive entry do not exist here.

But, as we've explored here, there are other less tangible aspects of this digital environment that constitute equally effective barriers to competitive entry. The use of IPR as a trade weapon, the networked effect that acts as positive feedback to market share are two such factors. But there are two more factors that are probably the more critical ones: The first is big data, where large data sets provide an enterprise with a clarity of knowledge of the user and the user's preferences that cannot be accurately induced from smaller data collections. The second is the agenda of capital investment. There is little appetite to financially support competitive entry into a given market when the alternative prospect of establishing a new digital monopoly is on offer. Existing activities have been captured by this digitisation and in its wake large scale centralised entities have been established, whether it’s the taxi industry and Uber, the accommodation industry and AirBNB, or television and the digital streamers. Many of these areas are now captured by monopolies, while others have at their core a small cartel of centralised players. This combination of the power of extremely large data sets and the lucrative financial rewards that await the establishment of an effective monopoly that drive investments markets are perhaps the unhappy coincidence of circumstance that sustain today's digital behemoths.

It’s challenging to create a regulatory framework that can alter the preferences of the investment market, and by and large the regulatory sector has not tried to go down this path. The chosen path is perhaps the equally forbidding task of attempting to change the inherent value of big data collections. Here the various actions of the EU in the form of GDPR, and prospective measures to impose mandatory data

sharing and the more fundamental question of who *owns* data that directly is about me as a user are perhaps the most effective actions that we can do to try and rebalance this space. Or at least that's the thought bubble coming from the regulatory sector. If the data is all about *me*, then surely I should be aware of where this data is, I should be able to amend it, and I should control how it is traded and exploited. It's all about *me*, after all. I suspect that these debates over the rights relating to the collection and use of large data, particularly relating to data about users and the way in which such data is exploited will remain as the most pressing of regulatory matters in the coming years.

My inner cynic is not overly optimistic about the outcome here, but there are few other available options for the regulator in attempting to restore a more equitable competitive balance back into this digital space.

It's not a case of the protocols we use, and no, the ledger-less world of blockchain-based technologies will have very little practical bearing on the issue of centralisation in the digital world and the unique position of the small clique of behemoths that occupy its core. It's perhaps a replay of the circumstances that occurred in the industrial revolution, where a small set of parties were able to amass large fortunes by committing early on to realise the competitive advantages that the industrial age technology was offering early adopters. In this case it's the marriage of computing and communications that have had a transformative impact on our society, and once more a small set of parties have been able to amass large fortunes by committing early in realizing the competitive advantages basing the delivery of goods and services on digital technology.

Perhaps this large-scale centralisation was always going to happen and in this case the fault, whatever it may be, is not within our failings but within the very nature of fundamental transformational change. To misquote Shakespeare's Julius Caesar, *the fault, dear Brutus, is indeed in our stars, that we are underlings*.

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