

Two economic perspectives on the IPv6 transition

Acknowledgements

The author would like to acknowledge the helpful and constructive feedback and suggestions from an anonymous reviewer.

Abstract

Purpose – IPv6 is the replacement for the Internet’s incumbent protocol, IPv4. IPv6 adoption is required to allow the Internet to continue to grow; however, there has been almost no uptake since its standardisation in the late 1990s. This paper seeks to explain how this non-adoption may be a consequence of current policies paradoxically intended to promote IPv6.

Design/methodology/approach – Economic theories of exhaustible resources and permit markets are used to provide an explanation for the lack of adoption of IPv6.

Findings – The current policy approach will not yield a significant adoption of IPv6 until after the IPv4 address space is exhausted and may also constrain Internet growth after IPv4 exhaustion occurs.

Practical implications – Current policies intended to promote IPv6 diffusion through the Internet must be reconsidered. The economics of permit markets in particular can inform discussions about IPv4 address transfer markets.

Originality/value – Economic analyses of IPv6 adoption are almost non-existent and very few prior studies are known. This paper helps to rectify this important gap in the literature.

Article Type: Viewpoint

Keywords: IPv4, IPv6, technology diffusion, exhaustible resources, permit markets.

1. Introduction

“The problem of address scarcity is as severe for the Internet economy as the oil shocks and gasoline shortages of the 1970s were to the industrial economy” (Mueller, 2008).

The above quotation paints a dire picture for the Internet, and by extension for economies in which the Internet is currently, or expected to be, a critical piece of infrastructure. The core of the problem is that IP addresses – which are to the Internet what telephone numbers are to the telephone network – are fixed in number and rapidly running out. Exhaustion is currently projected to occur between 2011 and 2012[1], at which point future growth of the Internet will be constrained.

The vast majority of previous literature surrounding this issue has focused on technological aspects. This paper departs from this pattern and investigates the problem from the perspective of economics. The remainder of the paper is structured as follows: first, a brief orientation to the history and background of the problem is provided. Second, current obstacles to a solution are discussed. This is followed by an application of two different economic theories to the problem. Finally, the paper finishes with discussion and conclusions.

2. Background to the problem

All computers connected to the Internet need a unique IP address. The protocol currently used as the basis for Internet communications is Internet Protocol version 4 (IPv4), and the addresses in use today are thus referred to as IPv4 addresses. Although IPv4 addresses are typically expressed as four numbers separated by dots, they are actually a single, 32-bit number. Thus, a total of 2^{32} – approximately 4.3 billion – possible addresses exist.

This is clearly inadequate for future needs. The Internet has grown exponentially to its current point and this trend seems likely to continue due to economic development and a growing global population, and the fact that more and more technologies are being designed with Internet connectivity in mind. This includes those that are still in their infancy, such as sensor-webs, and those that do not yet exist. It is possible that demand for Internet addresses could run into the tens or even hundreds of billions.

There were warnings about the eventual exhaustion of the 32-bit IPv4 address space even before the massive wave of commercial Internet adoption that took place during the mid-1990s. The RFC1287 memorandum, issued by the Internet Engineering Task Force (IETF) in 1991, describes routing and addressing as “the most urgent architectural problem, as it is directly involved in the ability of the Internet to continue to grow successfully”.

Classless Inter-Domain Routing (CIDR) was standardised in 1993 by RFC1518 and RFC1519 as a medium-term measure to slow the rate at which IPv4 address space was depleted, thus extending the useful life of IPv4 to the projected dates mentioned above. CIDR tackled supply by enabling more flexible address allocation policies and thus minimising the proportion of wasted – allocated but unused – addresses. Today, CIDR is universal and there are now few gains to be made from further improvements to address supply.

Meanwhile, Network Address Translation (NAT), standardised in 1994 by RFC1631, tackled demand by allowing a single address to be used to connect an entire network, where previously a single address had been required for each device on that network. The introduction of NAT has introduced technical problems for a wide range of protocols and services, which in turn has led to the development of a range of “workarounds”, the most common of which are NAT-PMP (NAT Port Management Protocol), TURN (Traversal Using Relay NAT), STUN (Standard Traversal of UDP through NAT) and ICE (Internet Control Exchange). Unfortunately, none of these are universal and each introduces its own technical problems.

Nevertheless, NAT is now the de facto means for connecting networks to the Internet and further impact on demand is unlikely. While implementing multiple NAT systems, sometimes called “layered NAT”, could in theory be used to reduce demand for IPv4 addresses, this is not a desirable solution to future address shortage as it compounds the problems NAT introduces, as well as introducing performance problems.

Concurrently with the development of CIDR and NAT, a new version of Internet Protocol was developed as a long-term solution to the problems with IPv4. IP version 5 (IPv5) had already been created as an experimental version to test media streaming features, and therefore the name given to the version set to replace IPv4 was IP version 6 (IPv6). IPv6 was officially standardised in 1995 in RFC1883, followed by an updated version in 1998 (RFC2460).

The principal advantage of IPv6 is that it has a much larger address space – 128 bits, as opposed to IPv4’s 32 bits – and this provides an extraordinarily large number of addresses: 3.4×10^{38} addresses, or 6.7×10^{23} addresses for every square metre of the Earth’s surface. Indeed, it has been said that this is sufficient for a unique address for every grain of sand on Earth (Wiljakka, 2002).

Nevertheless, it is now at least 19 years since the initial warnings about IPv4 address space and more than a decade since the final version of IPv6 was published, and yet IPv6 adoption is still negligible (Domingues *et al.*, 2007; Joseph *et al.*, 2007). Indeed, a projection based on current trends anticipates that 80% diffusion of IPv6 could take between another 8 and 22 years (Elmore *et al.*, 2008). Of course, trends can change and such projections may thus prove to be inaccurate; however, such projections do show

that it is unlikely that significant adoption will occur before the IPv4 address space is exhausted, and that there are still significant barriers to IPv6 adoption. The following section examines these barriers in more detail.

3. Current barriers to IPv6 adoption

Hovav *et al.* (2004) proposed the Internet Standards Adoption (ISA) model, which describes individual adoption decisions in terms of how useful the technology is to the organisation and how conducive the organisation's environment is to adoption. Both usefulness (UF) and environmental factors (EC) can be high or low, leading to four categories: Status Quo (low UF and low EC), Replacement (low UF and high EC), Niche (high UF and low EC) and Full Implementation (high UF and high EC). Given the almost complete lack of adoption of IPv6 to date, it is clear that almost every organisation is firmly within the Status Quo quadrant.

~~Hovav *et al.* (2004) proposed the Internet Standards Adoption (ISA) model, which describes individual adoption decisions in terms of how useful the technology is to the organisation and how conducive the organisation's environment is to adoption. Both usefulness (UF) and environmental factors (EC) can be high or low, leading to the 2x2 matrix shown in Figure 1. Given the almost complete lack of adoption of IPv6 to date, it is clear that almost every organisation is firmly within the Status Quo quadrant.~~

~~[FIGURE 1 AROUND HERE]~~

That IPv6 is not perceived as useful is hardly surprising. Technologically superior features of IPv6, such as improved security, mobility and Quality of Service (QoS), do

not offer the same level of advantage over IPv4 that IPv4 offered over its predecessor (Huston, 2007) – the benefits are primarily long-term rather than immediate (Bohlin and Lindmark, 2002), and the belief that there is no business case for IPv6 is widespread (Roberts, 2009). Indeed, the primary benefit of IPv6 – a vastly expanded address space – is not relevant to organisations that already have sufficient IPv4 address space and have no immediate need for expansion. Further, the impending IPv4 run-out is not likely to occur until approximately 2011-2012 – very soon from a “whole of Internet” point of view but far enough into the future to be less worrying for individual firms concerned only with their own networks. Thus, rightly or wrongly, it is unlikely that most organisations will perceive IPv6 as useful.

The current environment for IPv6 is also extremely un conducive to adoption. IPv4 is ubiquitous, and this would create high drag, inertia and conversion costs should an organisation decide to adopt IPv6 (Bohlin and Lindmark, 2002; Hovav *et al.*, 2004). Further, most organisations have little access to IPv6 skills and experience (c.f. Warfield, 2003; Dell *et al.*, 2007) and there have also been few monetary incentives or opportunities for sponsorship available in most countries.

Even in the albeit rare case where an organisation might fall in the coexistence, replacement or full implementation quadrants, the opportunities to obtain an IPv6 connection from an Internet Service Provider (ISP) are extremely limited. As at January 2009, there were only 44 ISPs or similar organisations worldwide that provide native IPv6 services to their customers, details of which are provided in Table 1.

Country	Companies
Japan	NTT Communications KDDI IIJ Nifty Dream Train Internet Powerdcom Japan Telecom
	JENS Media Exchange Freebit Plala Networks OnDemandTV Softbank Japan Sustainable Community Centre
Germany	Easynet IDKOM Networks Individual Network Berlin rh-tec
	Space.Net SpeedPartner TAL.DE Klaus Internet Service Titan Networks
USA	Epik Networks Citynet ipHouse
	Lava.net Spectrum Networks Transaria/Cutthroat Communication
UK	Claranet Andrews and Arnold Bogons
	Entanet Goscomb
Switzerland	Cyberlink Jaguar Network
	Init Seven Nexellent
France	Jaguar Network Nerim
	Proxad / Free SAS Wanadoo France
Italy	ITGate
	Panservice
Ireland	Airwire
Canada	Epik Networks
Ukraine	NetAssist
Finland	Nebula
Estonia	Linxtelecom
Netherlands	BIT
Australia	Internode

Table 1: ISPs offering native IPv6 to customers

Source: SixXS (2009), IPv6Style (2009)

Examination of the ISPs in Table 1 shows that 14 are Japanese, with few options available in other countries. The majority do not have a significant market share and therefore it is presumed that they are capable of only limited capacity. Indeed, Ladid (2008) reported that only seven of the world's biggest (by revenue) 21 telecommunications companies were IPv6 ready, and only two of those – NTT and KDDI – had retail IPv6 offerings. Thus, it is reasonable to conclude that although

native IPv6 services do exist, they are extremely rare and that IPv6 is not yet significantly diffused.

Even if an organisation is willing to use one of the very few providers from whom IPv6 services are available, they must also consider which of a number of “gateway technologies” (David and Bunn, 1988) they intend to use to provide interconnection between IPv6 and IPv4. Three such technologies have been proposed over the years. First, “dual-stacking” was suggested, in which computers had both IPv4 and IPv6 installed. This is no longer relevant as a mainstream transition mechanism, as the looming exhaustion of IPv4 address space will render it impossible for new connections to run native IPv4 in parallel with native IPv6.

A second gateway technology involves the implementation of protocol translation systems such as NAT-PT (Network Address Translation – Protocol Translation) and NAPT-PT (Network Address Port Translator – Protocol Translation) in order to connect their own IPv6 network to IPv4 service providers, or passing IPv6 traffic through IPv4 “tunnels” to IPv6 tunnel brokers elsewhere on the Internet. NAT-PT and NAPT-PT standards were first put forward in 2000 (RFC2766), however after several years this approach still faced serious technical challenges and was abandoned as a transition mechanism in 2007 (RFC4966).

A third gateway technology that has been proposed is the use of “tunnels” to connect IPv6 networks across an IPv4 backbone. To the minimal extent that IPv6 has been implemented – typically in experimental, hobbyist, and test environments – tunnels have been the approach most often used.

Regardless of which gateway is used, gateway technologies are typically only adopted when the benefits derived from improved system performance and beneficial network effects outweigh the costs imposed on those who adopt them (David and Bunn, 1988). In the case of the IPv6 and its associated gateways, the costs have almost invariably outweighed the benefits, and significant adoption of IPv6 has not occurred. Tunnelling imposes extra costs and degrades performance and thus is likely to be considered unacceptable for important or critical services. Consequently, it remains difficult – if not impossible – to put forward a business case for IPv6 adoption.

4. Economic perspectives of the IPv6 transition

Clearly, initiatives to promote update of IPv6 have largely not succeeded, a fact which may well be due to an absence of relevant theoretical or empirical research to inform them. Despite a reliance on market economics to promote the diffusion of IPv6 in many countries, there have been very few serious discussions to date of the economics of IPv4 and IPv6.

First, Bohlin and Lindmark (2002) considered the individual firm's incentive to invest in IPv6 and view the issue as an example of the Boiteux problem, in which the decision to invest in maintenance of an existing system appears rational when considering local issues, but is irrational when one takes in the global perspective. Thus, individual firms continue to invest in stop-gap measures such as NAT to extend the life of IPv4; however, this is inefficient when one considers the Internet as a whole.

Second, Mueller's (2006) examination of IP addressing policy concluded that there is significant room for variation and called for further research to identify optimal address policies. Any effort to introduce change had been blocked by the existing policy regime; however, subsequent to Mueller's (2006) analysis, three proposals for address transfer markets of various forms have been put forward by various RIRs (ARIN 2008-2, RIPE 2007-08 and APNIC Prop50).

It is possible that a market mechanism might encourage IPv6 adoption by giving the holders of IPv4 address space the incentive to relinquish their IPv4 space and make the transition to IPv6 (Edelman, 2009). However, others do not share this confidence in the ability of a market to promote the IPv6 transition. Mueller (2008) asserts that a successful market would not prevent the eventual transition to IPv6 but could prolong it by extending the life of the IPv4 address space. On the other hand, Lehr *et al.* (2008) argue that the transition may stall if IPv4-holding incumbents were to become sufficiently powerful to be able to preserve the status quo. Finally, Lehr *et al.* suggest that it is not guaranteed that such a market would even work at all, for example if sufficient buyers or sellers cannot be attracted to the market.

In summary, economies around the world are increasingly reliant on a global communications network that is fast reaching a point at which further growth will be problematic. Despite being a major economic issue, economic analyses of the transition from IPv4 to IPv6 are few. Three RIRs have recently put forward proposals for a market for the transfer of IPv4 address space and these currently form a focus of debate, however it may be difficult to evaluate the impact of such proposals on the transition to IPv6 in the absence of related literature.

To help rectify this situation, this paper now turns to such an examination of the IPv6 transition in which two approaches are considered: the economics of exhaustible resources and the economics of permit markets. If no transfer market ends up being implemented, the economics of exhaustible resources can explain the behaviour of IPv4 address space consumption. In the case that an IPv4 address transfer market is adopted, the economics of permit markets would seem to have explanatory power. These two theories are discussed in the following sections.

4.1. Economics of exhaustible resources

Economic analysis of exhaustible resources is not a new discipline; there is extensive literature on the subject from Hotelling's (1931) ground-breaking paper to the present day. The study of exhaustible resources typically has tended to focus on natural resources, particularly minerals and energy, but can be applied to any exhaustible resource. Considering IPv4 address space as an exhaustible resource which is being depleted in much the same way as the world's oil reserves are being depleted gives a valuable insight into any transition from IPv4 to IPv6 under market forces.

The essence of the Hotelling's analysis – now known as the “Hotelling Rule” – is that efficiency and competitive market forces will result in an increasing scarcity rent of an exhaustible resource that is equal to the interest rate. In other words, increasing scarcity of a resource contributes to increasing market price. Eventually, the market price of the exhaustible resource reaches the backstop price, and the backstop technology (Nordhaus, 1992) becomes more economical. It has been argued by some that for this reason,

“exhaustible resources” are never actually exhausted. In the case of IPv4 addresses, depending on the address allocation policy that prevails at the time it may be that the market price does not reach the backstop price, but there may be administrative difficulties in obtaining new addresses. In other words, it may be that the administrative burden for obtaining an IPv4 address will be more costly than that required for obtaining an IPv6 address.

In monopolistic situations, production is equal to demand at the current market price. If the net price rises too slowly, production shifts forward due to increased demand, leading to earlier exhaustion. Similarly, if the net price rises too quickly, the producer is inclined to keep its resources in the ground to maximise return (Solow, 1974).

The former is exactly what is happening in the case of IPv4 addresses. The market price of IP addresses cannot be observed because there is no IP address market. As Mueller (2006) observes, government agencies are not subject to market forces, so actual prices for IP address allocation – levied indirectly through other fees, but effectively the price one must pay to obtain address space – almost certainly do not reflect the price that would be obtained in an IP address market. The artificially low price thus promotes increased consumption of IPv4 addresses.

It is also possible to see evidence of economics of exhaustible resources in the “production” of IPv4 addresses. The monopolistic “producer” of IPv4 addresses is the Internet Assigned Numbers Authority (IANA), which is operated by the Internet Corporation for Assigned Names and Numbers (ICANN) and is responsible for the global coordination of IP address space allocation. The IANA allocates IPv4 addresses

to the various Regional Internet Registries (RIRs) according to their needs. RIRs in turn allocate address space to Local Internet Registries (LIRs), which are effectively customers and are usually Internet Service Providers (ISPs).

Address allocation policies typically prevent stockpiling of addresses and require the LIR be able to demonstrate a need for the address space. In other words, the production of IPv4 addresses is done in response to demand, just as Hotelling's analysis of exhaustible resources predicts.

If one accepts that IP address space consumption obeys the Hotelling Rule, it follows that significant uptake of IPv6 will not occur until IPv4 address space becomes more expensive than IPv6 address space. However, given that the "market price" of IPv4 address space is artificial it may well be the case that it never reaches the backstop price, in which case IPv6 uptake will only begin when IPv4 address space is completely exhausted, as illustrated in Figure 12. In any case, economic theory of exhaustible resources predicts that meaningful IPv6 diffusion does not commence until the point where further IPv4 deployment ceases, resulting in a sudden – as opposed to gradual – transition.

[FIGURE 2-1 AROUND HERE]

This model assumes that IPv6 addresses can be directly substituted for IPv4 addresses, albeit at a higher cost; however, this assumption may be incorrect. Roberts (2009) suggests that in the event of IPv4 address space exhaustion, most organisations may simply increase their use of NAT, while very few would adopt IPv6. If this eventuated,

deployment of new services and Internet diffusion in developing countries would be constrained.

Regardless of which of these two scenarios one regards as more likely, neither is desirable. While the latter case has obvious disadvantages – the inability to serve new markets and offer new services – the former case is also undesirable. To conduct the inevitable IPv6 migration only as a response to the complete exhaustion of IPv4 address space would be a risky strategy as the inevitable time pressures in such a situation would lead to ad hoc and unplanned actions. Indeed, serious problems with interoperability may result in a fragmented Internet if the transition to IPv6 is not properly managed, posing significant reliability issues for critical infrastructure (SIFT, 2007), and likely significantly higher cost of conversion than if the process was able to be planned in advance (Klensin, 2002).

4.2. Permit markets

Permit markets have been increasingly used as a means for environmental regulation, particularly since the 1980s. Issues such as carbon trading, regulating water consumption, sustainable fisheries, regulating the taxi industry, the phasing out of leaded petrol and limiting the emissions of chlorofluorocarbons and sulphur dioxide have all been subject to various schemes around the world, although not all schemes are alike. They are designed to meet specific objectives and will have varying rules concerning permit pricing, allocation and subsequent trading, and are subject to a variety of different monitoring and enforcement regimes.

Since it is impossible to connect to the Internet without an IP address, these can be considered analogous to a permit to connect to the network. There are a fixed number of “IPv4 permits”, and IP address allocation policies typically prohibit on-selling or subletting of addresses; thus, IP addresses are currently analogous to non-tradeable permits. Indeed, non-tradable permits can be conceptualised as a method for allocating a scarce resource (Leffevre, 2005), and this objective is an explicitly stated policy goal of the RIRs and is recommended in RFC2050.

Current policy is that IPv4 addresses are technically free and in practice require only the payment of a nominal annual membership fee to the RIR and to be able to meet the most basic of eligibility criteria – essentially to be able to demonstrate a need for the address space. As long as IPv4 permits continue to be allocated on this basis, firms will continue to have very little motivation to adopt IPv6 – in fact, it is likely that they will remain more motivated to continue using IPv4 due to network effects.

If this policy does not change, firms will continue to use IPv4 address space until it is completely exhausted, after which point they will have no choice but to adopt IPv6 or to restrict the number of connections to the Internet. Of course, this is how permit markets are intended to operate – a maximum number of permits is decided upon that allows an industry to function while restricting the impact of negative externalities to a satisfactory level. The only significant difference with IPv4 “permits” is simply that the maximum number is determined by the technology rather than by the consideration of externalities.

The exhaustion of the IPv4 address space will lead to the creation of scarcity rents for IPv4 addresses. It therefore may be desirable to allow the trading of IPv4 permits, as recommended by Mueller (2006; 2008) and Edelman (2009). The three proposals put forward by ARIN, RIPE and APNIC thus effectively transform IPv4 addresses into tradeable permits. If issues concerning address aggregation can be dealt with adequately, allowing IPv4 address trading via an appropriate market design would seem to be potentially advantageous for a number of reasons.

First, it is extremely unlikely that a central entity that issues non-tradeable permits will have access to enough information to allocate permits in a way that yields the greatest benefit. It is thus extremely unlikely that the IANA, RIRs, and so on can effectively judge which applicants for IPv4 address space will use addresses in ways that maximise public benefit. Applicants themselves are more likely to have the best information about their needs than the RIRs, and applicants could also have an incentive to misrepresent this information to the RIRs if RIRs were to use it to make allocation decisions. Further, such information may be commercially sensitive and applicants may be unwilling to provide it to RIRs (Lehr *et al.*, 2008).

Second, non-tradeable permits tend to be static and unresponsive to changing social values (Howe, 1996). This is evident in the way IPv4 address space has been allocated – many address blocks that were inefficiently allocated early in the development of the Internet remain with the original applicants, resulting in large amounts of allocated but un-used address space.

Third, non-tradeable permits imply that a consumer cannot benefit directly from trade even if there are opportunities to reduce consumption of the regulated resource[2] (Vollebergh *et al.*, 1997). Thus, even if an organisation could reduce their use of IPv4 address space, there would be little incentive to do so (Edelman, 2009).

In contrast, permits move to applications with the highest value in competitive markets in which trading is allowed (Montgomery, 1972). Indeed, this is true regardless of whether permits are auctioned or issued for free. When permits are issued for free and permit trading is allowed, total expenditure on permits is equal to the cost of the permit system objective, i.e. the cost of emission control or restricting fishery catches to sustainable levels (Montgomery, 1972), or managing the transition to IPv6 in the current context.

If IPv4 addresses were allocated free of charge and subsequently allowed to be traded, as is practically the case in the proposals currently put forward by the RIRs, any expenditure on such trades would equal the cost of restricting the (IPv4) Internet to a maximum size. However, the US and Canadian experiences in allocating cellular licenses in the 1980s, in which there were many frivolous and speculative licence applications (McMillan, 1995), suggests that the current first-come-first-served approach to allocation would not necessarily be the most appropriate. Thus, if address trading is to be allowed, auctioning may be a more appropriate method for initial IPv4 address allocations. While it may seem that there will be no remaining unallocated IPv4 address space to be auctioned, this may not be the case. This issue is examined in more detail below.

A counterargument to allowing the trading of IPv4 addresses is that trading permits may not obtain the maximum value from a resource if there are high transaction costs (Stavins, 1995), which may apply to the case of IPv4 address space. The effort involved in renumbering networks – as would be required by trading IPv4 addresses in a market – is immense. Further, trading IPv4 addresses could result in a highly fragmented IPv4 address space, leading to backbone routing performance problems. This paper thus seconds Mueller's (2006) and Lehr *et al.*'s (2008) call for further research into the implications of such trading.

A last consideration is that tradeable permits can be inequitable if the poor are forced to sell their permits to the rich in order to fund their own development. Thus, it can be argued that permit trading consolidates the power of the rich (Leffevre, 2005). Thus, any policy to allow trading of IPv4 address space would need to include measures to minimise this outcome (through initial address allocations, for example).

Lehr *et al.*, (2008) note that all three of the mechanisms currently put forward by the RIRs address this by disallowing inter-regional trading. Meanwhile, Mueller (2008) disputes whether inter-regional trading would harm developing countries at all. Indeed, as developing countries are likely to experience stronger demand than developed countries, it is likely that exhaustion – without a mechanism to re-use previously allocated addresses – would hurt developing countries more than those already developed.

I return here to the question of auctioning of permits. Although the free allocation vs. auction question may seem a moot point in the event of IPv4 exhaustion, this is not

necessarily the case. IPv4 address space is not allocated in perpetuity; rather, it is “licensed” to address holders for a specified period (typically one year), at the conclusion of which it may be renewed. For example, ARIN’s policy is as follows:

“An annual fee for registered space is due by the anniversary date of the ISP’s first allocation from ARIN. ISPs should take care to ensure that their annual renewal payment is made by their anniversary due date in accordance with the Registration Services Agreement. If not paid by the anniversary date, the address space may be revoked” (ARIN, 2010).

Similarly, APNIC (2010) requires address space to be renewed on an annual basis: “APNIC will allocate and assign Internet resources on a ‘license’ basis, with such licenses to be of specific limited duration (normally one year).” RIPE and LACNIC do not make the annual renewal explicit; however, membership of the RIR is required in order to hold address space, and such membership requires the payment of an annual fee, creating a de facto annual renewal requirement.

This has some similarities to Singapore’s Vehicle Quota System (VQS), which has operated in Singapore since 1990 (Olszewski and Turner, 1993). Under the VQS a fixed number of Certificates of Entitlement (COE) are offered by tender each month. A COE is required in order to register a new vehicle, and therefore a COE is effectively a “vehicle ownership permit”. COEs remain valid for a fixed period (10 years), after which they must be renewed at the prevailing price. If a consumer buys a second-hand vehicle they must also pay the prevailing price when that vehicle reaches 10 years of age. As some consumers choose not to renew the COE, new COEs can be auctioned

without increasing their overall number. This system has contributed to reduced car ownership in Singapore (Lam and Toan, 2006).

If it were possible to reserve some unallocated IPv4 address space for auction after the remainder is exhausted, it would be possible to identify a prevailing price. Henceforth, in a similar way to the VQS, one possible IPv4 market mechanism might be to require incumbent address holders to renew IPv4 address permits at the prevailing price. Thus, stronger demand for IPv4 address space would be at least partly counteracted by greater incentive for incumbents to make the transition to IPv6. IPv4 address space not renewed could be available for subsequent auctions. The possibility of this or a similar model is deserving of further investigation.

Finally, the impact of different market-based instruments on technology diffusion is subject to only limited agreement. Some contend that auctioning permits also results in greater diffusion of alternative technologies than other instruments (Milliman and Prince, 1989, 1992), while others argue that taxes provide greater incentive (Denicolò, 1999). Jaffe *et al.* (2002: 53) review the literature and decide that “unambiguous exhaustive ranking of instruments is not possible on the basis of theory alone”; indeed, each case needs to be considered on its own merits.

Figure 23 illustrates two possible scenarios for IPv6 diffusion through the lens of the economics of permit markets. It is likely that significant IPv6 diffusion will not commence until IPv4 exhaustion is reached. At that point a range of trajectories are possible, depending on the policy that is adopted at that point – whether IP address trading is allowed, and whether IPv4 addresses are taxed, auctioned, grand-fathered or

freely allocated on a first-come-first-served basis. It is assumed that the maximum demand for IP addresses would not exceed exponential growth, but that different policy frameworks will yield different levels of diffusion.

[FIGURE 23 AROUND HERE]

This is clearly a similar result to the previous analysis based on the economics of exhaustible resources and is undesirable for the same reasons.

5. Discussion and conclusions

Economic theory of exhaustible resources suggests that a safe and effective transition to IPv6 will not occur if current policies do not change. The prevailing policy approach will only yield significant IPv6 diffusion after IPv4 exhaustion has already occurred, which would likely create reliability and connectivity problems.

Three market models have recently been put forward for discussion by the RIRs. If a market for trading IPv4 addresses were to be implemented it could be considered as a permit market for Internet connectivity, and thus lessons from the economics of permit markets have illustrative power. In particular, such a market could allow permits to move to applications of the highest value and to respond to changing social values.

However, determining the impact of permit markets on technology diffusion requires analysis on a case-by-case basis and cannot be determined by theory alone, and it is perhaps for this reason that there is disagreement in the literature about the impact

address trading would have on IPv6 diffusion. Given that none of the three models have ever been implemented, discussion of their impact remains conjecture.

One permit market which can be observed is Singapore's VQS. This may serve as a useful example for the design of an appropriate mechanism for the management of IPv4 address space.

Whatever mechanism is implemented, it is clear that something must be done in order to achieve the outcomes deemed necessary by the OECD:

“Deploying the newer IP version 6 address blocks is necessary to enable growth in use of the Internet. But making the switch is difficult and it takes time and resources as well as a commitment by all stakeholders, including governments” (OECD, 2008: 3-4).

Clearly there is a need for change to the prevailing policy. There are undoubtedly other options to be discussed, and this paper therefore calls for an urgent and vigorous debate of new policy ideas.

6. Notes

1. See <http://www.potaroo.net/tools/ipv4/index.html> (accessed 3 February 2010).
2. In other tradeable permit markets – emissions trading, for example – selling of permits can be a source of additional capital before relocating to another country where

such permits are not required, negatively affecting the effectiveness of the permit system (Vollebergh *et al.*, 1997). However, in the case of IPv4 addresses such trade would be pointless since the IPv4 address “permit market” is global: it is not possible to simply move to another country where IPv4 addresses are unnecessary.

7. References

APNIC (2010), “Policies for IPv4 address space management in the Asia Pacific region”, Available at: <http://www.apnic.net/policy/add-manage-policy#9.1> (accessed 9 March 2010).

ARIN (2010), “ARIN Number Resource Policy Manual”, Available at: <https://www.arin.net/policy/nrpm.html#four212> (accessed 9 March 2010).

Bohlin, E. and Lindmark, S. (2002), “Incentives to Innovate with Next Generation Networks”, *Communications and Strategies*, Vol. 48, No. 4, pp. 97-117.

David, P.A. and Bunn, J.A. (1988), “The Economics of Gateway Technologies and Network Evolution”, *Information Economics and Policy*, Vol. 3, pp. 165-202.

Dell, P., Kwong, C., Syamsuar, D., François, S. and Choy, K. (2007) “A comparison of attitudes to IPv6 in three countries”, in: Cooper, R. and Madden, G. (Eds), *Conference proceedings – corporate strategy and the path to a new economy*, Curtin University of Technology, Perth, Australia.

- Denicolò, V. (1999), “Pollution-reducing Innovations Under Taxes or Permits”, *Oxford Economic Papers*, Vol. 51, pp. 184–199.
- Domingues, M., Friças, C. and Veiga, P. (2007), “Is global IPv6 deployment on track?”, *Internet Research*, Vol. 17, No. 5, pp. 505-518.
- Edelman, B. (2009), “Running Out of Numbers: Scarcity of IP Addresses and What To Do About It”, in: Das, S., Ostrovsky, M., Pennock, D. and Szymanski, B. (Eds), *Auctions, Market Mechanisms and Their Applications*, Springer Berlin Heidelberg.
- Elmore, H., Camp, L.J. and Stephens, B. (2008), “Diffusion and Adoption of IPv6 in the ARIN Region”, paper presented at the Workshop on the Economics of Information Security 2008, June 25-28, Hanover, New Hampshire.
- Hotelling, H. (1931), “The Economics of Exhaustible Resources”, *Journal of Political Economy*, Vol. 39, No. 2, pp. 137-175.
- Hovav, A. Patnayakuni, R. and Schuff, D. (2004), “A model of Internet standards adoption: the case of IPv6”, *Information Systems Journal*, Vol. 14, pp. 265-294.
- Howe, C.W. (1996), “Sharing water fairly”, Available at: <http://dlcvm.dlib.indiana.edu/archive/00003054/> (accessed 12 February 2010).

Huston, G. (2007), “Addressing a Fundamental Part of the Internet”, paper presented at the NSF/OECD Workshop – Social and Economic Factors Shaping the Future of the Internet, 31 January, Washington, Available at: <http://www.oecd.org/dataoecd/30/15/38221103.pdf> (accessed 27 February 2009).

IPv6Style (2009), “IPv6 Services in Japan”, Available at: <http://www.ipv6style.jp/en/statistics/services/index.shtml> (accessed 15 January 2009).

Jaffe, A.B., Newell, R.G. and Stavins, R.N. (2002), “Environmental Policy and Technological Change”, *Environmental and Resource Economics*, Vol. 22, pp. 41-69.

Joseph, D., Shetty, N., Shuang, J. and Stoica, I. (2007), “Modeling the Adoption of New Network Architectures”, in: Kurose, J. and Schulzrinne, H. (Eds), *Proceedings of the 2007 ACM CoNEXT Conference*, ACM, New York.

Klensin, J.C. (2002), “A Policy Look at IPv6: A Tutorial Paper”, available at: <http://www.itu.int/oth/T0501000013/en> (accessed 28 January 2009).

Ladid, L. (2008), “IPv6 Readiness: Facing the Challenge”, paper presented at the 2008 Australian IPv6 Summit, 17-19 November, Canberra.

- Lam, S.H. and Toan, T.D. (2006), "Land transport policy and public transport in Singapore", *Transportation*, Vol. 33, No. 2, pp. 171-188.
- Leffevre, J. (2005), "Greenhouse Gas Emissions Trading: A Background", In: Bothe, M. and Rehbinder, E. (Eds.), *Climate Change Policy*, Eleven International Publishing, Utrecht, The Netherlands.
- Lehr, W., Vest, T. and Lear, E. (2008), "Running on empty: The challenge of managing internet addresses", in: *TPRC '08: Proceedings of the 36th Research Conference on Communication, Information and Internet Policy*, TPRC, Farnham, VA, available at: <http://www.tprcweb.com/images/stories/2008/lehrtprc08.pdf> (accessed 11 February 2010).
- McMillan, J. (1995), "Why auction the spectrum?", *Telecommunications Policy*, Vol. 19, No. 3, pp. 191-199.
- Milliman, S.R. and Prince, R. (1989), "Firm incentives to promote technological change in pollution control", *Journal of Environmental Economics and Management*, Vol. 17, pp. 247-265.
- Milliman, S.R. and Prince, R. (1992), "Firm Incentives to Promote Technological Change in Pollution Control: Reply", *Journal of Environmental Economics and Management*, Vol. 22, pp. 292-296.

- Montgomery, W.D. (1972), "Markets in Licenses and Efficient Pollution Control Programs", *Journal of Economic Theory*, Vol. 5, No. 3, pp. 395-418.
- Mueller, M.L. (2006), "IP addressing: the next frontier of internet governance debate", *Info*, Vol. 8, No. 5, pp. 3-12.
- Mueller, M.L. (2008), "Scarcity in IP addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries", Available at http://www.internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf (accessed 9 February 2010).
- Nordhaus, W.D. (1992), "The ecology of markets", *Proceedings of the National Academy of Sciences*, Vol. 89, pp. 843-850.
- OECD (2008), "The Future of the Internet Economy", available at <http://www.oecd.org/dataoecd/20/41/40789235.pdf> (accessed 27 February 2009).
- Olszewski, P. & Turner, D.J. (1993), "New methods of controlling vehicle ownership and usage in Singapore", *Transportation*, Vol. 20, No. 4, pp. 355-371.
- Roberts, P. (2009), "Internet Society Organization Member IPv6 Study", available at <http://www.isoc.org/pubs/2009-IPv6-OrgMember-Report.pdf> (accessed 23 March 2009).

SIFT (2007), “Future of the Internet (FOTI) Project: Reliability of the Internet”, available at http://www.dbcde.gov.au/_data/assets/pdf_file/0004/75676/FOTI-Reliability-FinalReport.pdf (accessed 16 December 2008).

SixXS (2009), “Where can I get native IPv6?”, available at <http://www.sixxs.net/faq/connectivity/?faq=native> (accessed 15 January 2009).

Solow, R.M. (1974), “The Economics of Resources or the Resources of Economics”, *American Economic Review*, Vol. 64, No. 2, pp. 1-14.

Stavins, R.N. (1995), “Transaction Costs and Tradeable Permits”, *Journal of Environmental Economics and Management*, Vol. 29, No. 2, pp. 133-148.

Vollebergh, H.R.J., De Vries, J.L. and Koutstaal, P.R. (1997), “Hybrid Carbon Incentive Mechanisms and Political Acceptability”, *Environmental and Resource Economics*, Vol. 9, pp. 43-63.

Warfield, M.H. (2003), “Security implications of IPv6” Available at <http://www.its.ms.gov/docs/IPv6.pdf> (accessed 2 February 2010).

Wiljakka, J. (2002), “Transition to IPv6 in GPRS and WCDMA Mobile Networks”, *IEEE Communications*, Vol. 40, No. 4, pp. 134-140.