

Pricing and Efficiency in the Market for IP Addresses

Benjamin Edelman* Michael Schwarz
Harvard Business School Yahoo Research

September 26, 2011

Abstract

We consider market rules for the transfer of IP addresses, numeric identifiers required by all computers connected to the Internet. Excessive fragmentation of IP address blocks causes growth in the Internet's routing table, which is socially costly, so an IP address market should discourage subdividing IP address blocks more than necessary. Yet IP address transfer rules also need to facilitate purchase by the networks that need the addresses most, from the networks who value them least. We propose a market rule that avoids excessive fragmentation while almost achieving social efficiency, and we argue that implementation of this rule is feasible despite the limited powers of central authorities. We also offer a framework for the price trajectory of IPv4 addresses. In a world without uncertainty, the unit price of IPv4 is constant before the first time when all blocks of IPv4 addresses are in use and decreasing after that time. With uncertainty, the price before that time is a martingale, and the price trajectory afterwards is a supermartingale.

1 Introduction

Every device connected to the Internet – from PCs to tablets, printers to cash registers – needs an IP address. The current addressing standard, IPv4, uses addresses with 32 binary digits, allowing approximately 4 billion IP addresses. The world's centralized supply of unused IP addresses reached exhaustion in February 2011, and networks in most countries will soon find they cannot easily obtain additional IPv4 addresses. While addresses may now be bought and sold, the institutions and rules of these transfers are not yet well-developed. Nor have economic models examined the unusual characteristics of this market. In this paper, we seek to speak to the latter gap.

*Edelman advises ARIN's counsel on matters pertaining to IPv4 exhaustion, v6 transition, and associated ARIN policy. This paper expresses only his personal view.

Low-cost electronic communications have sparked the transformation of many markets – both in creating the need to accommodate large communities of market participants, and in providing tools to make centralized markets mechanisms logistically viable. For example, keyword auctions let tens of thousands of buyers bid for millions of distinct objects – a context ill-suited to the negotiations traditionally used to sell advertising. (Edelman et al. (2007)) At the same time, the Internet facilitates market mechanisms that would have been too cumbersome without modern information technology. For example, the spectrum auctions considered in McMillan et al. (1998) require quick transmission of other bidders’ activities. In this paper, we consider the market for IPv4 addresses – another example of the Internet expanding feasible market structure, and also a market that is important for the continued growth of the Internet. Meanwhile, in the spirit of papers which let the unusual requirements of a particular context guide the design of the corresponding allocation rules and institutions (e.g. school choice, Abdulkadiroglu et al. (2005), and kidney allocations, Ashlagi et al. (2011)), we consider novel rules in service of special concerns in IP markets.

We proceed as follows: In Section 2, we present the technical and institutional details that inform understanding of this market. In Section 3, we consider a one-period model of the market for IPv4 addresses and a mechanism that yields a (near) efficient outcome (notwithstanding the negative externalities we identify in Section 2.4). In Section 4, we offer an equilibrium framework for modeling the time trajectory of IP addresses prices.

2 The technologies and institutions of IP addressing

2.1 The institutions of IP addressing

Evaluating feasible rules for IPv4 markets requires understanding responsible institutions. We summarize relevant details here; Edelman (2009) offers more details.

Like most of the Internet’s infrastructure, the IP addressing system is largely private. IP communication systems were initially developed pursuant to US government contract, but private parties have always administered address assignments. Under current allocation procedures, the Internet Assigned Numbers Authority (IANA), a California non-profit, assigns addresses to five regional internet registries (RIRs), which in turn assign addresses to ISPs and large networks. (In some regions, RIRs first assign addresses to country-specific internet registries, which in turn assign addresses to networks.) Large networks tend to get their addresses directly from RIRs, while smaller networks and residential end-users get addresses from ISP’s.

Because RIRs are private, they have little ability to penalize networks that flout their rules. Indeed, networks seem to value the limited powers of their RIRs, treating this as a philosophical matter (analogous to “limited govern-

ment”). That said, RIRs do have some important powers. Notably, most ISPs check an RIRs records before accepting a new customer purporting to use a given block of IP addresses. Through “WHOIS” listings, RIRs report the network authorized to use a given block of IP addresses. Historically, WHOIS has been sufficiently desirable – and RIR rules sufficiently unobjectionable and well-regarded – that networks have been willing to comply with RIR rules in order to obtain official address assignments, WHOIS listings, and associated benefits.

To date, RIRs have examined each network’s request to confirm the need for the requested addresses; addresses are issued only to networks that can justify their requests, e.g. through business plans, equipment purchases, and customer lists. Meanwhile, RIRs charge low fees; the North American RIR, the American Registry of Internet Numbers (ARIN), charges just \$18,000 per year to its largest ISPs.

With modest fees and rapid Internet growth, IPv4 resources are in short supply. IANA announced in February 2011 that it had exhausted its reserves. APNIC, the RIR for the Asia/Pacific region, finished its supply in April 2011, and the North American and European RIRs are expected to run out by early 2012.

2.2 Technical responses to IPv4 scarcity

Facing limited availability of IPv4 addresses, growing or newly-created networks have several options:

- Networks can move devices to an alternative addressing system called *IPv6*, which offers 128 binary bits of numbering capacity and therefore is in abundant supply. However, a device with only an IPv6 address cannot directly communicate with IPv4-only devices; for example, it cannot directly access a web site that has only an IPv4 address. Translation devices are not yet widespread, and some protocols may be difficult to translate between IPv4 and IPv6, making IPv6 a more appealing solution once others have adopted it also. Furthermore, until IPv6 is in widespread use, IPv6 communications will take longer routes (limited to passing among IPv6-enabled routers), and some have found that IPv6 systems less reliable than v4 (typically for lack of around-the-clock monitoring and troubleshooting).
- Alternatively, networks can seek to *make do with fewer* IPv4 addresses, typically by installing network address translation (NAT) devices which allow multiple devices to share a single public-facing IP address. By rewriting packet headers, NAT maintains the illusion that all local devices have a single address. NAT devices have become common in the home gateways that many residential end users install in order to connect multiple computers to a single Internet DSL modem or cablemodem. While most protocols function as expected through NAT, others do not or require significant customization, and NAT impedes certain kinds of innovation. (Blumenthal and Clark (2001)) Furthermore, large-scale “carrier-grade” NAT would raise weightier questions of reliability and scalability, as well

as reducing the likelihood of transition away from NAT in the coming years.

- Finally, networks can *buy* IPv4 addresses from others. Some networks may have more than they need – whether due to overly generous initial allocations, reduced requirements, or migration towards address-conserving alternatives (such as IPv6 or NAT). Meanwhile, others need more than they have – for example, new and growing networks.

In the long run, it is widely expected that networks will move to IPv6. But in the short run, v6 transition has been slow. For one, network effects encourage use of IPv4: with most users and sites on v4 only, it is natural to want v4 connectivity and unusual to want v6. Limited and untested translation systems – untested in part due to lack of customer demand – have further hindered transition. Meanwhile, large-scale NAT has not been deployed on the required scale and, even in the best case, adds complexity to network structure. For these reasons, it is currently anticipated that many growing networks will, in the short run, turn to IPv4 purchases to meet their v4 needs – prompting questions of the design of markets and institutions to facilitate such transfers.

2.3 The prospect of paid transfers of IPv4 addresses

Facing demand that exceeds supply, it is natural to use prices to reveal valuations and to transfer resources to the networks that most value them. Historically, paid transfers of IP addresses were impermissible; after all, with sufficient addresses available directly from RIRs, there was no need for sales between networks, and any buyer was either a scofflaw (e.g. a spammer needing limitless new addresses to send unsolicited email – a group networks did not seek to assist) or a fool (who could be better served by a strong norm of obtaining addresses directly from RIRs). But as scarcity loomed, RIRs revised their policies to allow sales. For example, the ARIN transfer policy is codified in Number Resource Policy Manual (NRPM), <https://www.arin.net/policy/nrpm.html#eight3>.

2.4 A negative externality from IPv4 sales

While trade in IPv4 addresses promises various benefits, including as sketched in preceding sections, transfers also prompt concerns. Most common is the fear that certain address transfers might threaten the Internet’s routing system – the systems that transfer data from one network to another. The Internet’s routing architecture requires that each “default-free zone” (DFZ) router on the Internet keep a record of each block of addresses used anywhere else on the Internet. These records are stored in each router’s routing table – a high-speed memory bank limited in size and limited further by the need to process and search the routing table exceptionally quickly. Importantly, a small block of addresses requires just as much routing table capacity as a large block. If large networks begin to acquire many small address blocks, rather than a few large blocks, the routing table could grow sharply. (Consider: a network might elect

to buy 16 blocks of 256 addresses, rather than one block of 4,096, if the former is slightly cheaper. In that case, growth would be an order of magnitude more rapid.)

At best, growth in the routing table would require that networks upgrade their routers more frequently. But particularly rapid growth could exceed router manufacturers' ability to improve their offerings – destabilizing the Internet's routing system.

No single entity – neither public nor private – exerts meaningful control over the routing system. As a result, no one can easily impose rules on what routes may be added to the routing table, nor is there anyone to collect a fee for each addition. In North America, ARIN established rules to require each buyer to satisfy the entirety of its short-run need in a single transfer – disallowing multiple small transfers, and thereby avoiding unnecessary growth of the routing table. (See ARIN NRPM rule 8.3, allowing transfers only to buyers which “can demonstrate the need for such resources, as a single aggregate, *in the exact amount* which they can justify under current ARIN policies” (emphasis added). The “exact amount” language is understood to require that a buyer satisfy its entire need, not merely a portion thereof.) However, other RIRs took a different approach. For example, APNIC's transfer rules included no restrictions to prevent many small purchases.

In Section 3.2, we examine the welfare and efficiency properties of a rule that generalizes ARIN's “exact amount” requirement.

3 One period model and market mechanism

In most asset markets, it is typical for large buyers and large sellers to split their purchases and sales into small pieces to be bought or sold separately. An unusual feature of market for IPv4 addresses is that such splits are socially costly for the reasons set out in Section 2.4. We therefore seek to devise market rules to discourage unnecessary splits, and in this section, we develop a formal model to evaluate the performance of such rules.

3.1 Notation and definitions

Consider N networks, each with an endowment of IPv4 addresses. We will use the size of the smallest tradable block of IPv4 addresses as the unit of measurement, so when we say that network k bought x_k addresses, we mean that k bought x_k times the smallest block. $x_k < 0$ means that k was a seller. Denote by $f_k(x)$ the value that k derives from increasing (or the cost from decreasing) its IPv4 address holdings by x , and normalize $f(0) = 0$. We assume that $f(\cdot)$ is increasing. We also assume decreasing differences (that the marginal benefit of an address decreases in the number of addresses the network holds), i.e. $f(x) - f(x - 1) \leq f(x - 1) - f(x - 2)$ for all x .

Definition 1 *An outcome is a vector specifying how much each buyer purchases and how much each seller sells.*

Definition 2 An outcome is feasible if $\sum_k x_k = 0$ and $x_k \in I$ for every k .

Definition 3 An outcome is efficient if $\sum_k f(x_k)$ is maximized subject to the feasibility constraints.

Due to the negative externality presented in Section 2.4, we will consider restrictions on trading procedures to prevent excessive fragmentation of IPv4 blocks without inhibiting efficiency. The following definitions will assist in developing and analyzing such restrictions:

Definition 4 Under the spartan rule, the agents engaging in a bilateral trade must designate one of the two agents as “extinguished.” The spartan rule prohibits an extinguished agent from trades with other extinguished agents; an agent can be extinguished at most once.

Notice the relationship between our concept of an “extinguished” agent and the “exact amount” requirement from Section 2.4: Whereas “exact amount” required that every buyer’s first trade be his last, our spartan rule allows either buyer or seller to be “extinguished” by a transaction, such that either the buyer or the seller may cease to trade further after that transaction. In Section 3.2 we examine the welfare and efficiency implications of this modified rule.

Because we offer a one-period model, an agent is extinguished forever. However, in practice IP allocations typically operate on a rolling basis, i.e. providing a network with the addresses it can justify for a six month period. As a practical matter, it would be logical for extinguished status to last for that same duration.

A sequence of trades satisfies the spartan rule if and only if each agent is extinguished at most one time and if each agent has at most one trade with an agent extinguished by someone other than himself. Consequently, if a sequence of trades satisfies the spartan rule, then the same trades in a different sequence also must satisfy the spartan rule. Thus, we can define spartan allocations without considering the order of trades.

Definition 5 An allocation is an outcome and the set of bilateral trades that lead to that outcome. In particular, an allocation specifies which seller(s) each buyer was matched with, the size of the transaction between each buyer and sellers, the price of each such transaction, and which party was designated as extinguished.

An allocation is spartan if each agent is extinguished at most one time. In that case, regardless of the order of trades, the spartan rule is respected.

Definition 6 If a seller is matched to N buyers, we will say that $N - 1$ cuts are required for that seller,

Remark 7 If a seller sells a positive quantity of resources to N different buyers, then $N - 1$ cuts are required for that seller.

3.2 Welfare and efficiency under spartan allocations

Lemma 8 *A spartan allocation involving N buyers never entails more than N cuts.*

Proof. Suppose there are N buyers and K sellers involved. The maximum total number of transactions is $N + K$. (Either a buyer or a seller must be extinguished in each transaction, and each entity is extinguished at most once, so the number of trades cannot be more than the number of participating entities, $N + K$.) Suppose seller k was matched to n_k buyers. Since the total number of transactions is no more than $N + K$, it must be the case that $\sum_{k=1}^K n_k \leq N + K$. By definition the total number of cuts is $\sum_{k=1}^K (n_k - 1) = \sum_{k=1}^K n_k - K$. Combining, $\sum_{k=1}^K (n_k - 1) = \sum_{k=1}^K n_k - K \leq N + K - K = N$. ■

Proposition 9 *For any feasible outcome, there exists a spartan allocation that induces that outcome.*

Corollary 10 *Let N denote the number of buyers. There exists a spartan allocation with no more than $N - 1$ cuts that induces an efficient outcome.*

Proof. Consider an outcome involving N buyers and K sellers. Denote by V_i^0 the quantity obtained in that outcome by buyer i , and denote by W_j^0 the quantity sold by seller j . Consider a sequential mechanism where there is exactly one transaction in each period. Denote by V_i^t and W_j^t the number of units that buyer and seller respectively will acquire (or provide) in transactions that take place after period $t - 1$. Consider the following sequence of trades satisfying the spartan rule: In period t , find the buyer with the smallest remaining demand. (Buyer k has the smallest remaining demand if $V_k^t \leq V_j^t$ for any j .) Similarly, find the seller with the smallest remaining supply at period t . ($W_m^t \leq W_j^t$ for any j .) In period t , match this buyer k with seller m . If $V_k^t > W_m^t$ (buyer k wants more than seller m can offer), they transact amount W_m^t - exhausting m 's supply, leaving $W_m^{t+1} = 0$, while k still seeks additional resources $V_k^{t+1} = V_k^t - W_m^t$. Alternatively, if $V_k^t \leq W_m^t$ (seller m wants to sell more than buyer k requires) then in period t the transacted amount is V_k^t . In other words, in each period the smallest buyer and the smallest seller transact with each other the largest possible amount, thereby extinguishing either the buyer or the seller in each transaction. The sequence of trades constructed through this process can implement any outcome (including the efficient outcome) with $N - 1$ cuts. ■

Corollary 11 *Any coalition-proof spartan allocation is efficient.*

3.3 Properties of minimal allocations

The preceding section characterizes the maximum number of cuts required for spartan allocations and efficient outcomes. In this section, we examine the *minimal* number of cuts in order to minimize the externality flagged in Section 2.4.

Definition 12 *A minimal allocation is an efficient allocation with weakly fewer cuts than any other efficient allocation.*

Proposition 13 *There exists a spartan allocation that is also minimal.*

Proof. Suppose any efficient spartan allocation results in more cuts than an efficient allocation with the smallest number of cuts (a minimal allocation). First, note that there are many minimal allocations because without the spartan rule, the identity of the party extinguished in each trade can be assigned arbitrarily. An allocation can be viewed as a graph where agents who trade are connected and arrows points away from extinguished agent. Without loss of generality we can consider an allocation that has only one connected component. Indeed, if there are multiple connected components, the problem becomes separable because there is an outcome corresponding to each component and for each component it must be impossible to reduce the number of cuts without changing the outcome. Thus, without loss of generality we will assume that minimal allocation has only one connected component.

Suppose there does not exist a spartan allocation that leads to as few cuts as a minimal allocation. Then for any minimal allocation, there must exist an agent who is extinguished more than one time. (Otherwise there would exist a spartan allocation that coincides with a minimal allocation.)

Consider a minimal allocation with the smallest possible number of non-extinguished agents. Observe that a non-extinguished agent must be directly connected to an agent who is extinguished only one time. (If a partner of a non-extinguished agent were extinguished more than once, then we could reduce the number of none extinguished agents, a contradiction.)

The number of extinguishings equals the number of trades, so the total number of extinguishings is at least $N + K + 1$ where N and K are the numbers of buyers and sellers respectively. One seller participates in every trade, and the number of cuts is the total number of trades minus the number of sellers, thus we showed that the number of cuts in the minimal allocation is at least $N + 1$. But this contradicts Lemma 8 which showed that there are at most N cuts in a spartan allocation. Hence, there must exist a spartan allocation that is minimal. ■

Corollary 14 *The number of cuts in the minimal allocation equals the number of buyers minus the number of (disjoint) components in the graph representing the minimal allocation.*

Proof. It follows from Corollary 10 that the k -th component has at most $N_k - 1$ cuts. Summing over all components, we obtain the upper bound on the number of cuts. Now let us show that a component of a graph induced by minimal allocation cannot have less than $N - 1$ cuts. From Proposition 13, it follows that without loss of generality we can consider a minimal allocation that is spartan and hence all agents are extinguished no more than once. Suppose a component has fewer than $N - 1$ cuts. In that case the total number of links is at most $N + K - 2$ and hence at least two agents remain unextinguished. The trading

partners of the unextinguished agents are extinguished (by transactions with non-extinguished agents). Denote by Ω_{1h} and Ω_{2h} the sets of agents who are h links away from non-extinguished agents 1 and 2. (Ω_{10} is agent 1 himself, and Ω_{11} is the set of agents who are connected to agent 1 and hence extinguished by that trade.) Note that agents in a set Ω_{1k} are extinguished by agents in a set Ω_{1k-1} . Consequently, any agent connected via a chain with agent 1 is extinguished by agents closer to agent 1, and the same is true for those surrounding other unextinguished agents. Thus, the clusters around unextinguished agents can never be connected to each other, which contradicts the assumption that they belong to the same component of the graph. ■

The preceding results establish the favorable characteristics of the spartan rule: it achieves efficiency and achieves the fewest possible number of cuts. At the same time, the spartan rule generalizes the current ARIN rule presented in Section 2.4. Relative to the mechanism ARIN uses, the spartan rule is more flexible in that it allows large buyers to more readily find the addresses they seek, yet it does so without unnecessarily constraining trades.

4 Price trajectory

The preceding section ignored dynamic aspects of the market for IPv4 addresses. But in fact IP addresses are long-lived assets, and networks' needs will change over time. In this section, we consider the dynamics of the v4 market, characterizing the price trajectory of v4 addresses.

Although the value of IPv4 addresses may be high in the near future, it is expected that scarcity will eventually diminish as networks migrate to IPv6. In particular, once much of the Internet supports IPv6, there will be less need for IPv4 addresses. As a result, at some time in the future, we expect the value of an IPv4 address to be zero. But what happens to prices until then?

We propose to treat IPv4 addresses as inputs into the production process – much like ordinary capital assets. On this view, the price of an IPv4 address should reflect the present discounted value of its future rental prices. (We recognize that there may never be a rental market for IPv4 addresses – there are good reasons why such rentals might be infeasible and/or undesirable – but the framework of rental prices helps clarify prices over time.)

We begin by considering the period when all IPv4 addresses are not yet in use. During this period, the Lagrange multiplier on the constraint associated with scarcity of IPv4 addresses is zero, so the rental price of IPv4 must be zero. In other words, as long as there are blocks of IPv4 addresses that are not in use, equilibrium behavior of market participants requires that the rental price of IPv4 must be zero. (Of course, a zero *rental* price does not mean the prevailing market price of IPv4 should be zero; the market price of an IPv4 address incorporates expectations about future rental prices of that address.) Importantly, the efficient market hypothesis implies that an asset with zero per-period rental income must yield a market rate of return. Hence if the interest rate is zero, the market price of an IPv4 address must be martingale as long as




Figure 1: Price Profile of IPv4 Addresses

there are some unused IPv4 blocks. In a world without uncertainty, the price of IPv4 addresses must remain constant as long as there are some unused IPv4 blocks.

We now turn to the period when there are no unused IPv4 blocks available. In this period, true scarcity begins, and the implied rent of an IPv4 address becomes positive. Notably, the market price of an IPv4 address declines each period by an amount equal to the implied rent of that addresses. In a world with uncertainty, the price of IPv4 addresses becomes a super-martingale once all IPv4 blocks are in use.

Using the preceding analyses, Figure 1 presents plausible price profiles for the complete information case (without uncertainty).

We believe our analysis of IPv4 price profiles is likely to be counterintuitive to non-economists. Non-economists might be surprised to hear that prices are at their highest before scarcity actually sets in, and that prices only decrease as networks search for sources of the scarce resources at issue.

5 Concluding remarks

The impending scarcity of IPv4 addresses calls for economic research to facilitate suitable institutions and market rules. In many markets, participants had the benefits of a period of years to design such systems, and often geographically-isolated markets can serve as parallel laboratories to invent and test alternative approaches. In contrast, IPv4 addresses have been issued at de minimis price (and typically zero marginal price) and will continue to be issued that way until RIRs have fully exhausted their supply – delaying the development of market institutions. Then, IPv4 scarcity will arrive in one fell swoop, raising the stakes for any adjusting of rules and offering less opportunity to refine rules over time. Furthermore, because IPv4 resources are inherently portable, experiments in one region would tend to affect behavior elsewhere – limiting the potential for regional experiments. In these circumstances, economic theory can provide particularly valuable insights on market design.

Our proposed spartan rule builds on ARIN’s current restriction – seeking to address the same negative externality, but offering somewhat more flexibility in order to avoid unnecessary deviations from efficient allocations. We believe ARIN could adopt our rule within its existing transfer framework and with minimal administrative burden.

As IPv4 markets develop, there will also be an opportunity for empirical economic research. For example, will prices in fact be linear in block size, or will large blocks carry disproportionately higher (or lower) prices? Combining such observations with relevant economic theory could yield prompt diagnoses of market malfunctions and timely interventions to preserve efficiency while limiting negative externalities.

6 References

- Atila Abdulkadiroglu, Parag A. Pathak, and Alvin E. Roth. The new york city high school match. In *American Economic Review, Papers and Proceedings*, pages 95–364, 2005.
- Itai Ashlagi, Duncan S. Gilchrist, Alvin E. Roth, and Michael A. Rees. Nonsimultaneous chains and dominos in kidney paired donation – revisited. *American Journal of Transplantation*, 11:984–994, 2011.
- Marjory Blumenthal and David Clark. Rethinking the design of the internet: the end-to-end arguments vs. the brave new world. *ACM Transactions on Internet Technology*, 1:70–109, August 2001. ISSN 1533-5399.
- Benjamin Edelman. Running out of numbers: Scarcity of ip addresses and what to do about it. In Sanmay Das, Michael Ostrovsky, David Pennock, and Boleslaw K. Szymanski, editors, *AMMA*, volume 14 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 95–106. Springer, 2009. ISBN 978-3-642-03820-4.
- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- John McMillan, Paul Milgrom, Bradley Miller, Bridger Mitchell, Daniel Vincent, and Robert Wilson. Simultaneous ascending auctions with package bidding. 1998.