


Diffusion of Networking Technologies




Sharon Goldberg
Boston University

Electronic Commerce (EC'12)
Valencia, Spain
June 7, 2012



Diffusion in social networks: Linear Threshold Model

[Kempe Kleinberg Tardos'03, Morris'01, Granovetter'78]



A node's utility depends only on its neighbors!



I'll adopt the innovation if θ of my friends do!

- $\theta = 1$
- $\theta = 2$
- $\theta = 3$
- $\theta = 4$
- $\theta = 6$

Optimization problem [KKT'03]: Given the graph and thresholds, what is the smallest seedset that can cause the entire network to adopt?

Seedset: A set of nodes that can kick off the process.



Marketers, policy makers, and spammers can target them as early adopters!

Today's questions:

- Where did this theory come from?
- Can it be applied to networking technologies?
- What are these networking technologies?



Tutorial Plan

- Classical foundations of diffusion modeling
 - Diffusion of Innovations (Social Sciences) [Everett '62, '03]
 - “Bass Model” and extensions (Marketing) [Bass '69]
 - Network externalities or effects (Economics)[Farrell, Saloner '86], [Choi '94]
- Quick interlude - Internet economics
- Networking technologies
 - IPv6 and the challenge of adopting an incompatible technology
 - IPsec – a success story
 - BGPsec and the challenge of coordinating independent agents
 - DNSSEC – quick overview of a rollout this happening right now.



Why should you care?

1. EC community has expertise in diffusion problems on graphs.
 - Most of these problem involve network `externalities' with graph structure.
2. This is a real problem the practitioners care about **right now**.
 1. DNSSEC rollout is ongoing since 2005.
 2. BGPsec is currently being standardized and will be rolled out in ~ 5 years.
 3. World IPv6 launch day happened yesterday!



As academics, we can help answer policy question of how to rollout these technologies. There is surprisingly little work in this area.



Diffusions of Innovations Theory (social sciences)



“Diffusion is the process by which an **innovation** is communicated through certain **channels** over **time** by members of a **social system**.” [Rogers 2003]

Social system:

- Social norms
- “Opinion leadership” (power & position in social network)

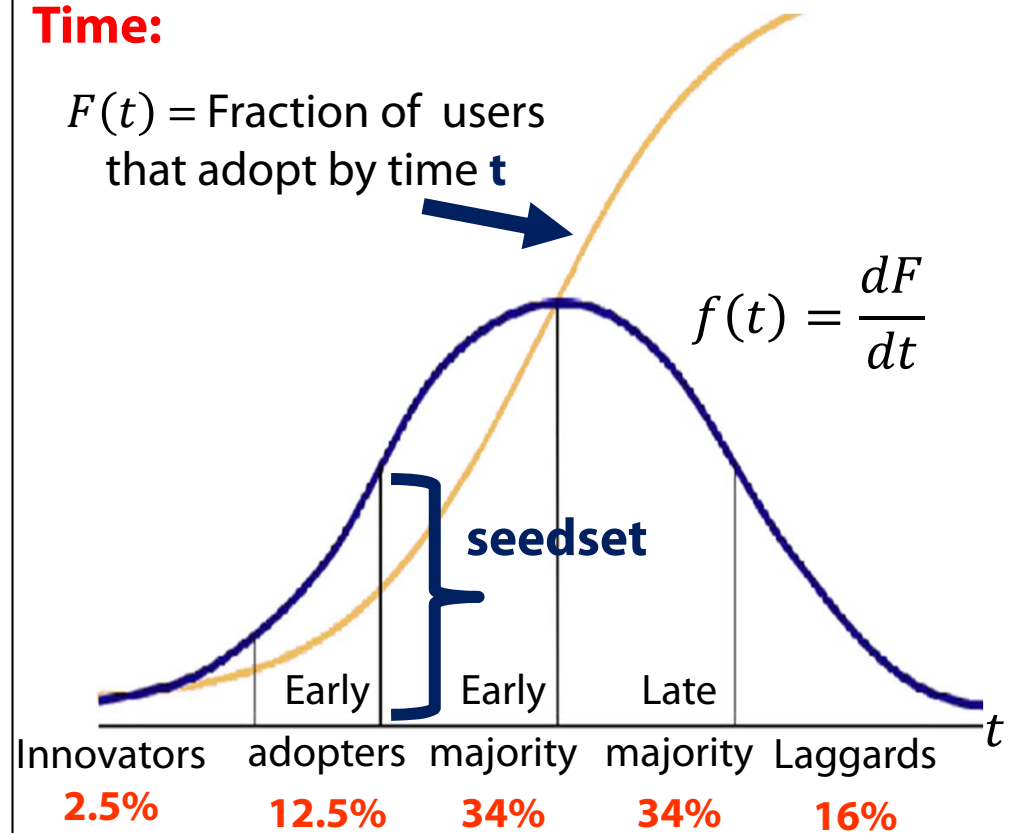
Communication channels:

- Mass media
- Interpersonal communication

Innovation characteristics

- Relative advantage
- Compatibility
- Complexity
- Trialability
- Observability

Time:



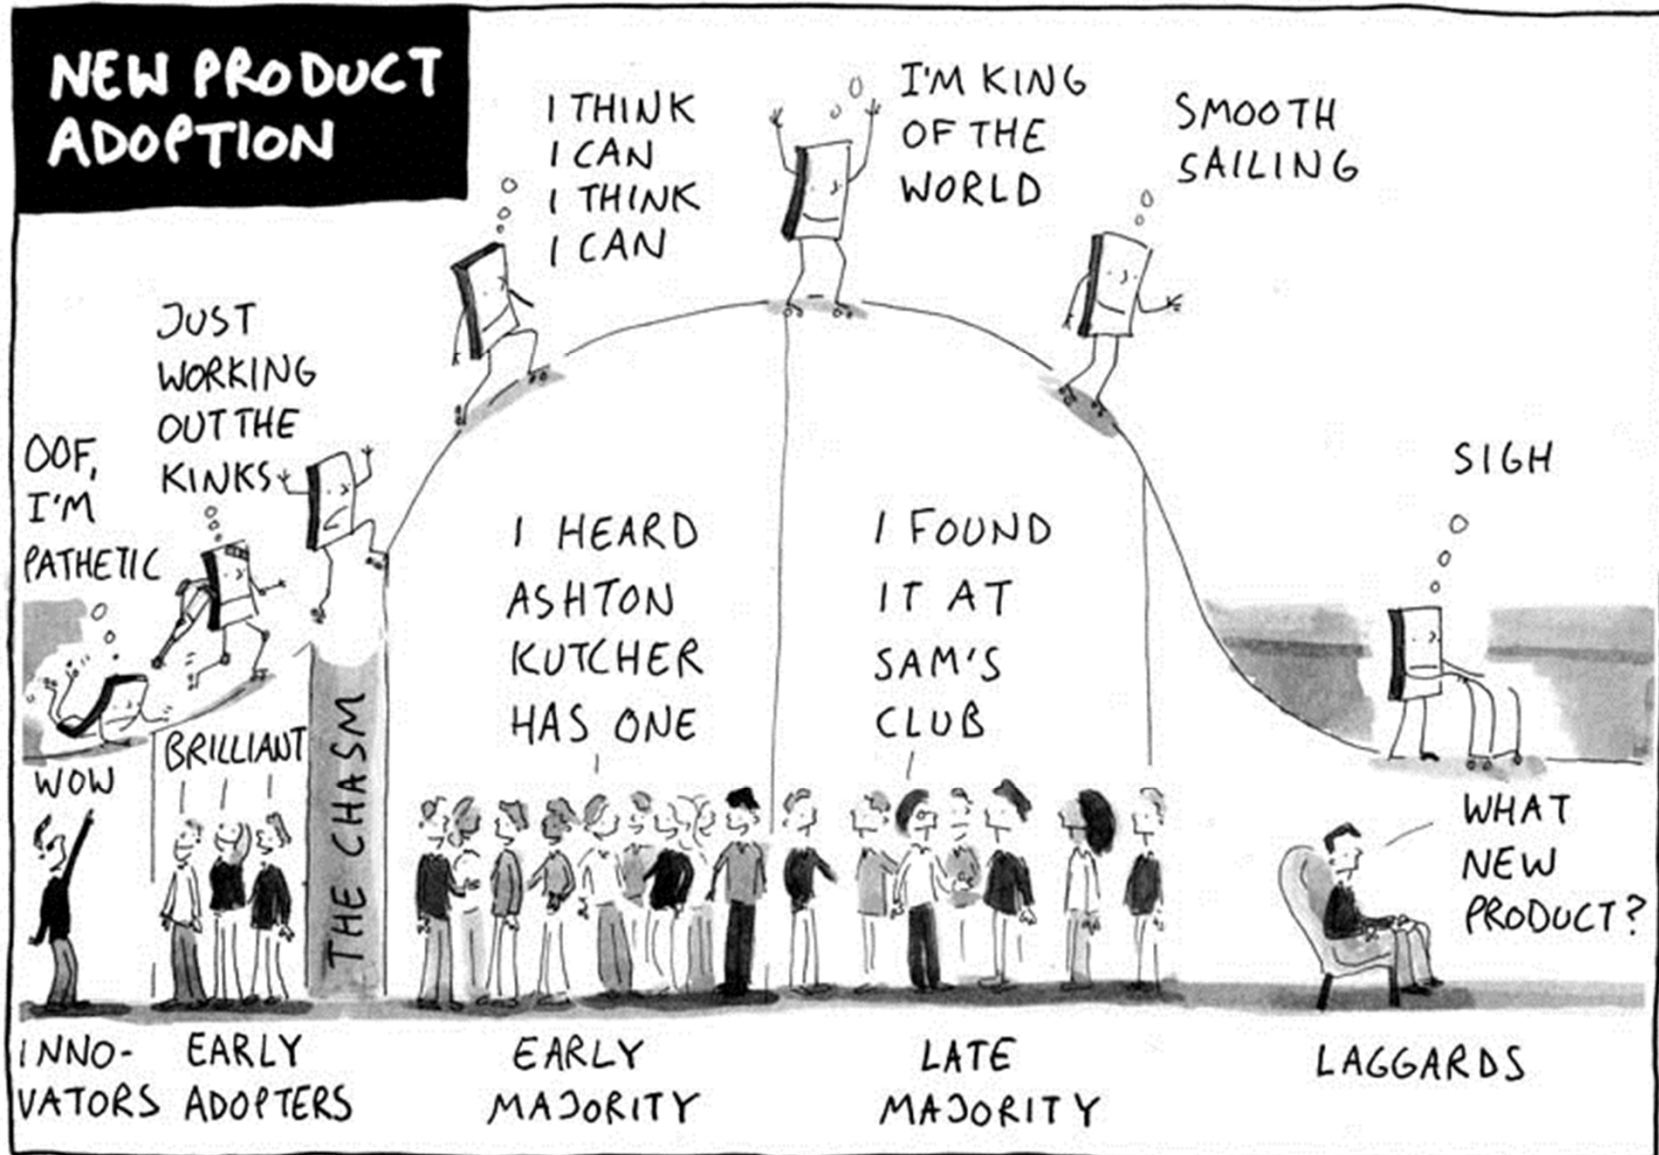
- **Usually:** an individual's decisions are influenced by past adopters decisions

Diffusions of Innovations Theory (social sciences)



BRAND CAMP

by Tom Fishburne



© 2007 Thanks to G. Moore

SKYDECKCARTOONS.COM



Diffusions of Innovations Theory (social sciences)



- Pioneered by rural sociologists **[Ryan&Gross 1949]**
- Uses to understand why some innovations take off, but others don't
- ... to understand impact of communication channel use
- ... to categorize "innovativeness" of organizations / individuals, etc
- 100's of studies of innovations e.g. seed corn, new drugs, HIV prevention, sanitization, photovoltaics, fax, computers, Internet, video games, ...

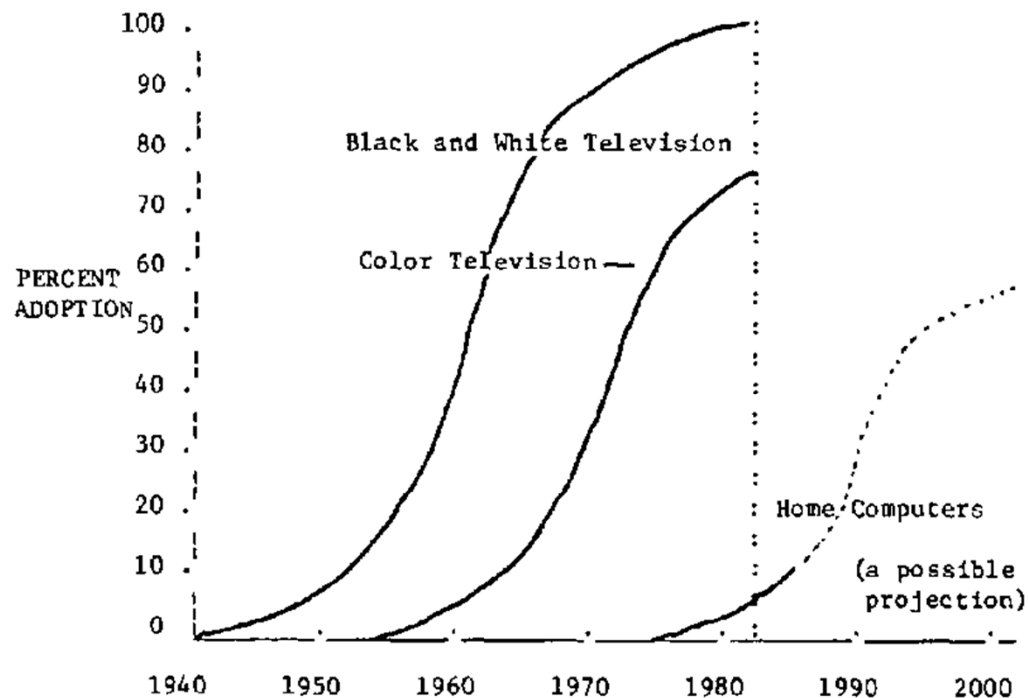


Figure 1. Diffusion Curves for the Adoption of Three Household Communication Technologies.

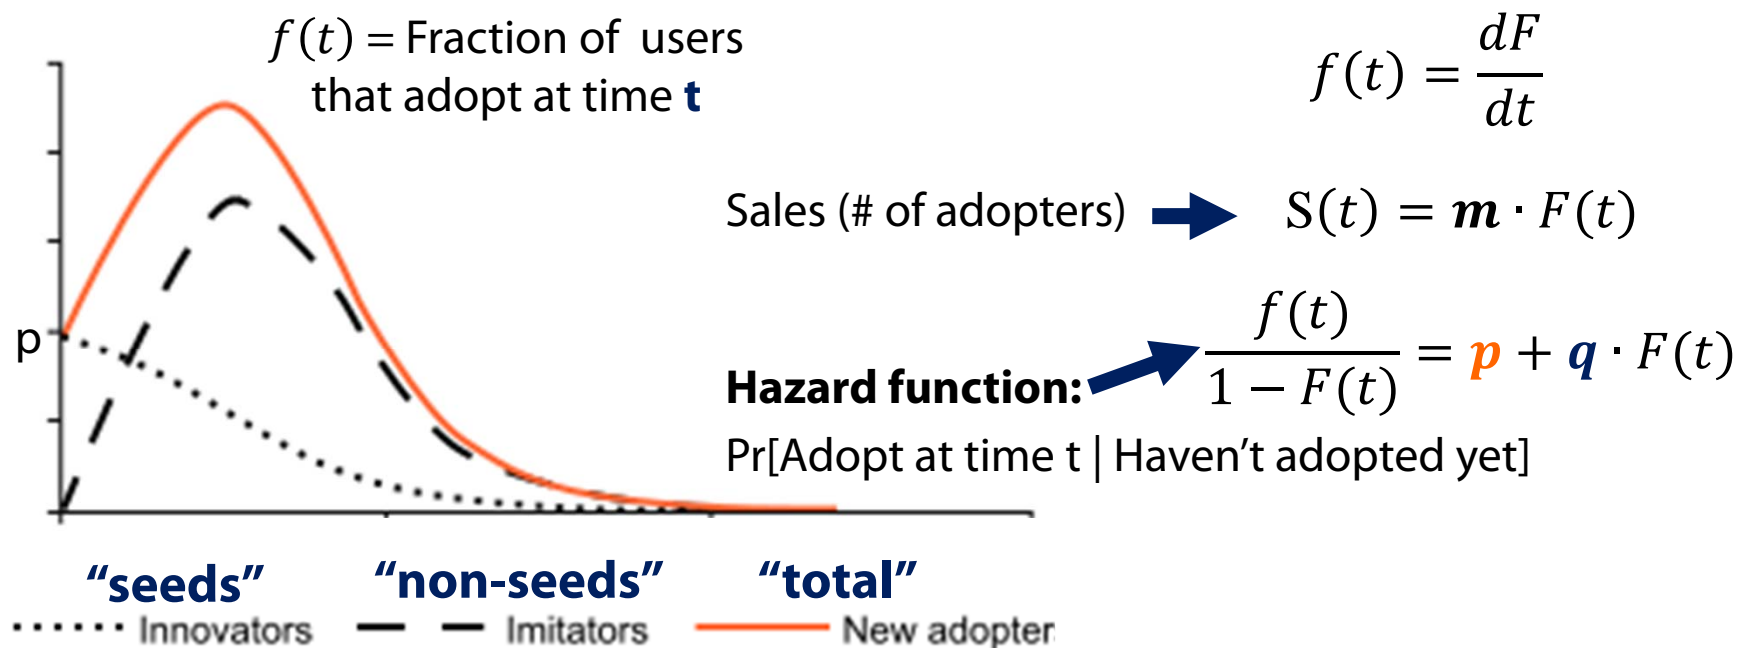
Source: Rogers. "The Diffusion of Home Computers Among Households in Silicon Valley" (1985)



The Bass Model (marketing)

- **[Bass 1969] Three parameter diffusion model:**

- **p** (coefficient of innovation / external influence / external marketing)
- **q** (coefficient of imitation / internal influence / “word of mouth”)
- **m** (market potential / max number of possible adopters)



- “The most popular model in the field of marketing” **[Dekimpe]**
- Used to forecast extent of diffusion, and how pricing, marketing mix effects it
- ... and for normative and descriptive purposes (e.g. pricing, timing strategies)



Norton-Bass Model (marketing successive generations)

- Both diffusion and substitution of new products (e.g. IC feature sizes, OS version)

1070

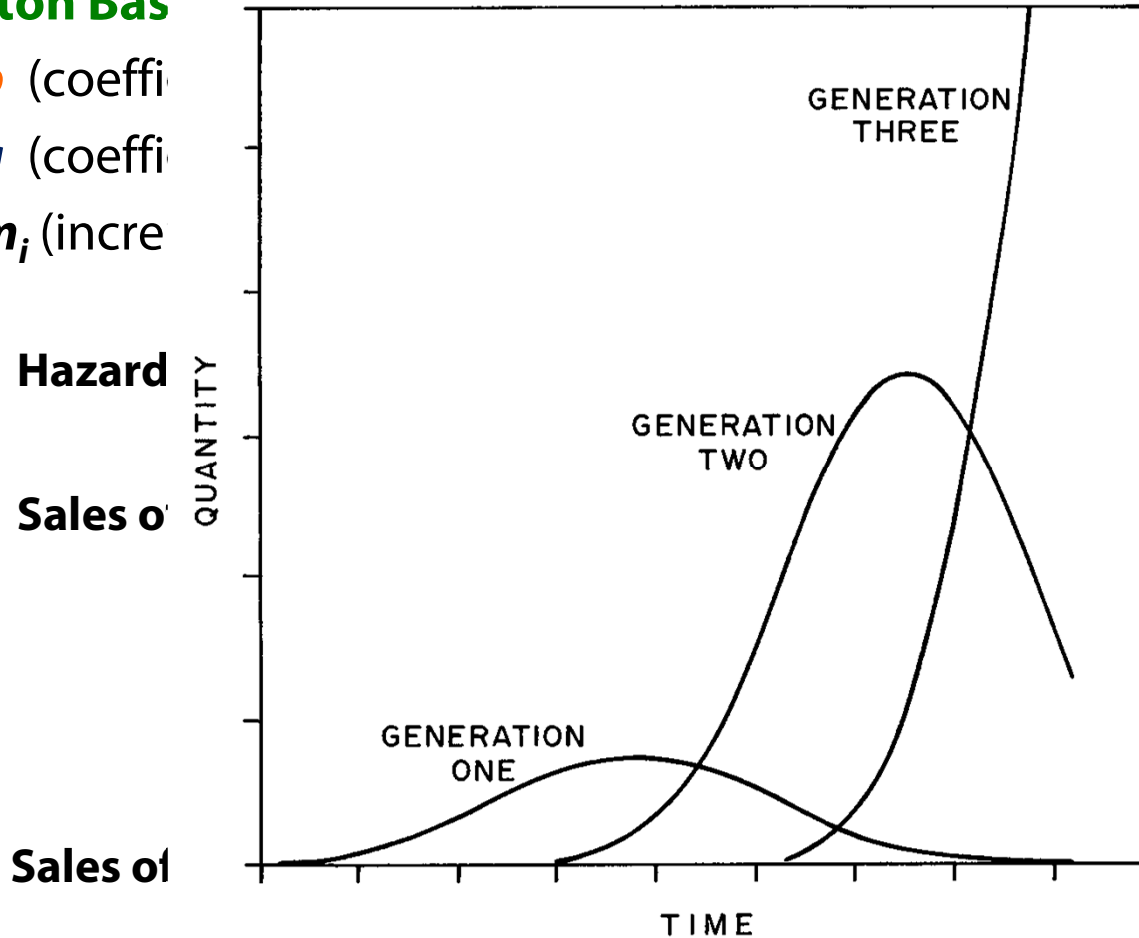
JOHN A. NORTON AND FRANK M. BASS

[Norton Bass

- p (coeffi
- q (coeffi
- m_i (incre

el:

nal marketing)
 | of mouth")
 ration)



Hazard
 Sales of
 Sales of

FIGURE 1. A Series of Technological Generations.

$$[1 - F(t - \tau_3)]$$

etc.

Cumulative fraction of adopters

Market potential of 2nd gen

Fraction lost to 3rd gen.



Network Externalities/Effects (economics)

“The utility that a given user derives from the good depends upon the **number** of other users who are in the same “network” as he or she.” **[Katz & Shapiro 1985]**

- **Direct network effects:**

- Increased direct usage leads to direct increases in value
- Classic examples: phone, fax, videoconferencing

**Impact on
compatibility
& standards ?**

- **Indirect network effects:**

- Increased direct uses increases the value of complementary goods

- **Two-sided network effects:**

- “two sets of agents interact through an intermediary or platform, and the decisions of each set of agents affects the outcomes of the other set of agents, typically through an externality.” **[Rhysman 2009]**
- Videogames, search engines, credit cards, dating services, etc.

- **Local network effects:**


- Users influenced by decision of their neighbors (e.g. **[Kempe et al '2003]**)



Rogers calls some of these “interactive innovations” : an individual’s decisions are influenced by the decisions of **future adopters**




Installed base & compatibility [Farrell&Saloner 1986] (2)



- Utility of choosing **U** at $t=0$ given equilibrium is **adoption**
$$u^-(0) = \int_0^{\infty} (\alpha + \beta I) e^{-rt} dt = (\alpha + \beta I)/r$$
- Utility of choosing **U** at $t=0$ given equilibrium is **nonadoption**:
$$u^+(0) = \int_0^{\infty} (\alpha + \beta I + \beta t) e^{-rt} dt = (\alpha + \beta I)/r + \beta/r^2$$
- Utility of choosing **V** at $t=0$ given equilibrium is **nonadoption**:
$$v^-(0) = \int_0^{\infty} \alpha' e^{-rt} dt = \alpha'/r$$
- Utility of choosing **V** at $t=0$ given equilibrium is **adoption**:
$$v^+(0) = \int_0^{\infty} (\alpha' + \beta' t) e^{-rt} dt = \alpha'/r + \beta'/r^2$$



Installed base & compatibility [Farrell&Saloner 1986] (3)



Utility of choosing **U** given equilibrium is **adoption**: $u^-(0)$

Utility of choosing **U** given equilibrium is **nonadoption**: $u^+(0)$

Utility of choosing **V** given equilibrium is **nonadoption**: $v^-(0)$

Utility of choosing **V** given equilibrium is **adoption**: $v^+(0)$

What should user arriving at time 0 do?


Choose **adoption** if $u^+(0) < v^-(0)$

Choose **nonadoption** if $v^+(0) < u^-(0)$

Else, both **nonadoption** and **adoption** are equilibria



Installed base & compatibility [Farrell&Saloner 1986] (4)



Utility of choosing **U** given equilibrium is **adoption**: $u^-(0)$

Utility of choosing **U** given equilibrium is **nonadoption**: $u^+(0)$

Utility of choosing **V** given equilibrium is **nonadoption**: $v^-(0)$

Utility of choosing **V** given equilibrium is **adoption**: $v^+(0)$

Improved welfare of **adoption** equilibrium vs **nonadoption** equilibrium **at time 0**:

$$\underbrace{\int_0^{\infty} (v^+(t) - u^+(t))e^{-rt} dt}_{\text{Change in utility for users arriving after time 0}} - \underbrace{(u^+(0) - u^-(0)) \cdot I}_{\text{Loss in utility for installed base.}} = \underbrace{\frac{(\alpha' - \alpha)}{r^2} + \frac{(\beta' - \beta)}{r^3} - \frac{2\beta I}{r^2}}_{\text{+ve or -ve!}}$$

Change in utility for users arriving after time 0

Loss in utility for installed base.

+ve or -ve!

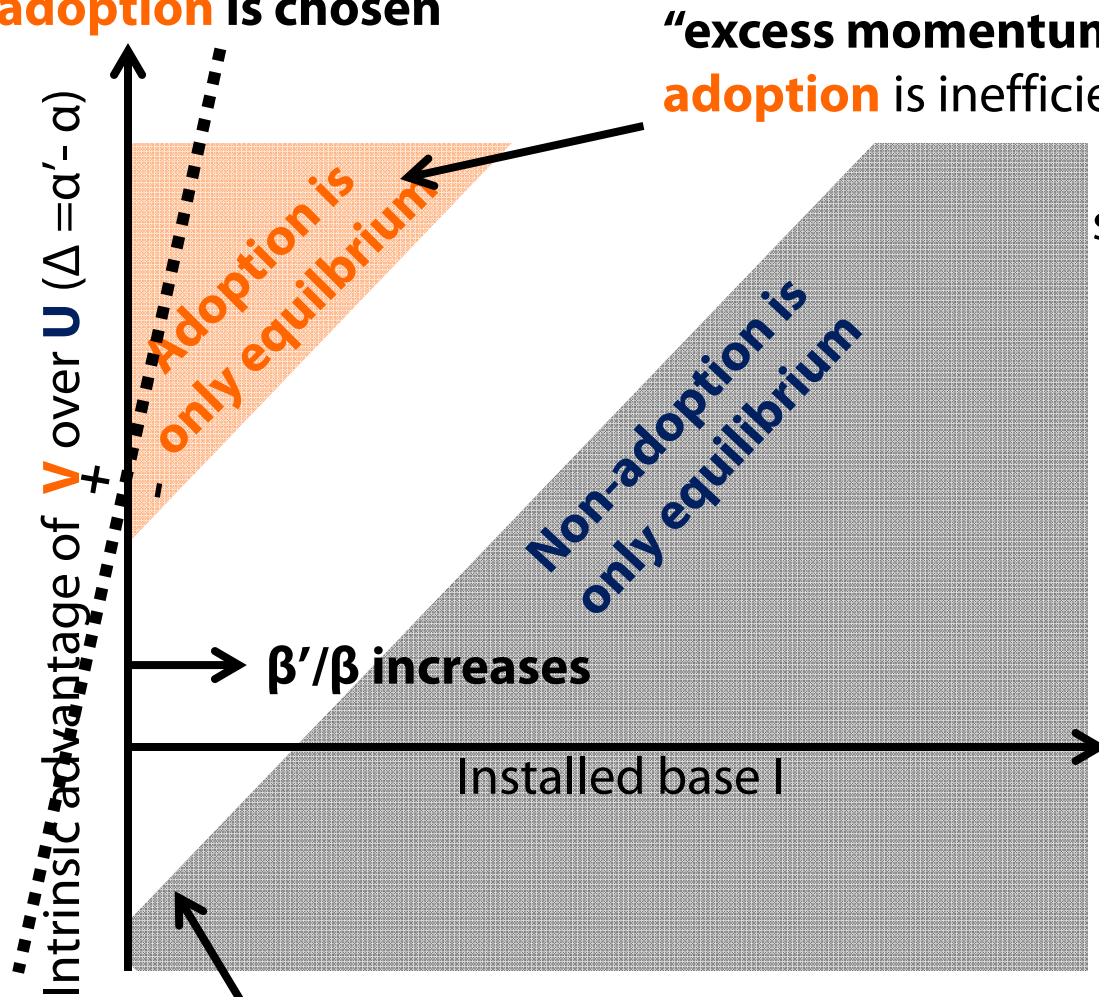


"Relative advantage" [Rogers'62]



Installed base & compatibility [Farrell&Saloner 1986] (5)

Change in welfare
if **adoption** is chosen



"excess momentum"

adoption is inefficient but is an equilibrium

Why? Installed base is stranded and loses utility.

β'/β increases

Installed base I


"excess inertia"

adoption is efficient but **nonadoption** is an equilibrium

Why? Externality kicks in late, no incentive to adopt at $t=0$



Do converters facilitate the transition? [Choi 1994] (1)



$$u(t) = \alpha + \beta \cdot (\# \text{ U users at time } t) + q \cdot \beta' \cdot (\# \text{ V users at time } t)$$

$$v(t) = \alpha' + \beta' \cdot (\# \text{ V users at time } t) + q \cdot \beta \cdot (\# \text{ U users at time } t)$$

$q < 1$ is the cost of conversion

Recalculate $u^-(0)$, $u^+(0)$, $v^-(0)$ and $v^+(0)$ using equations above....

As before:

Adoption is only equilibrium if $u^+(0) < v^-(0)$

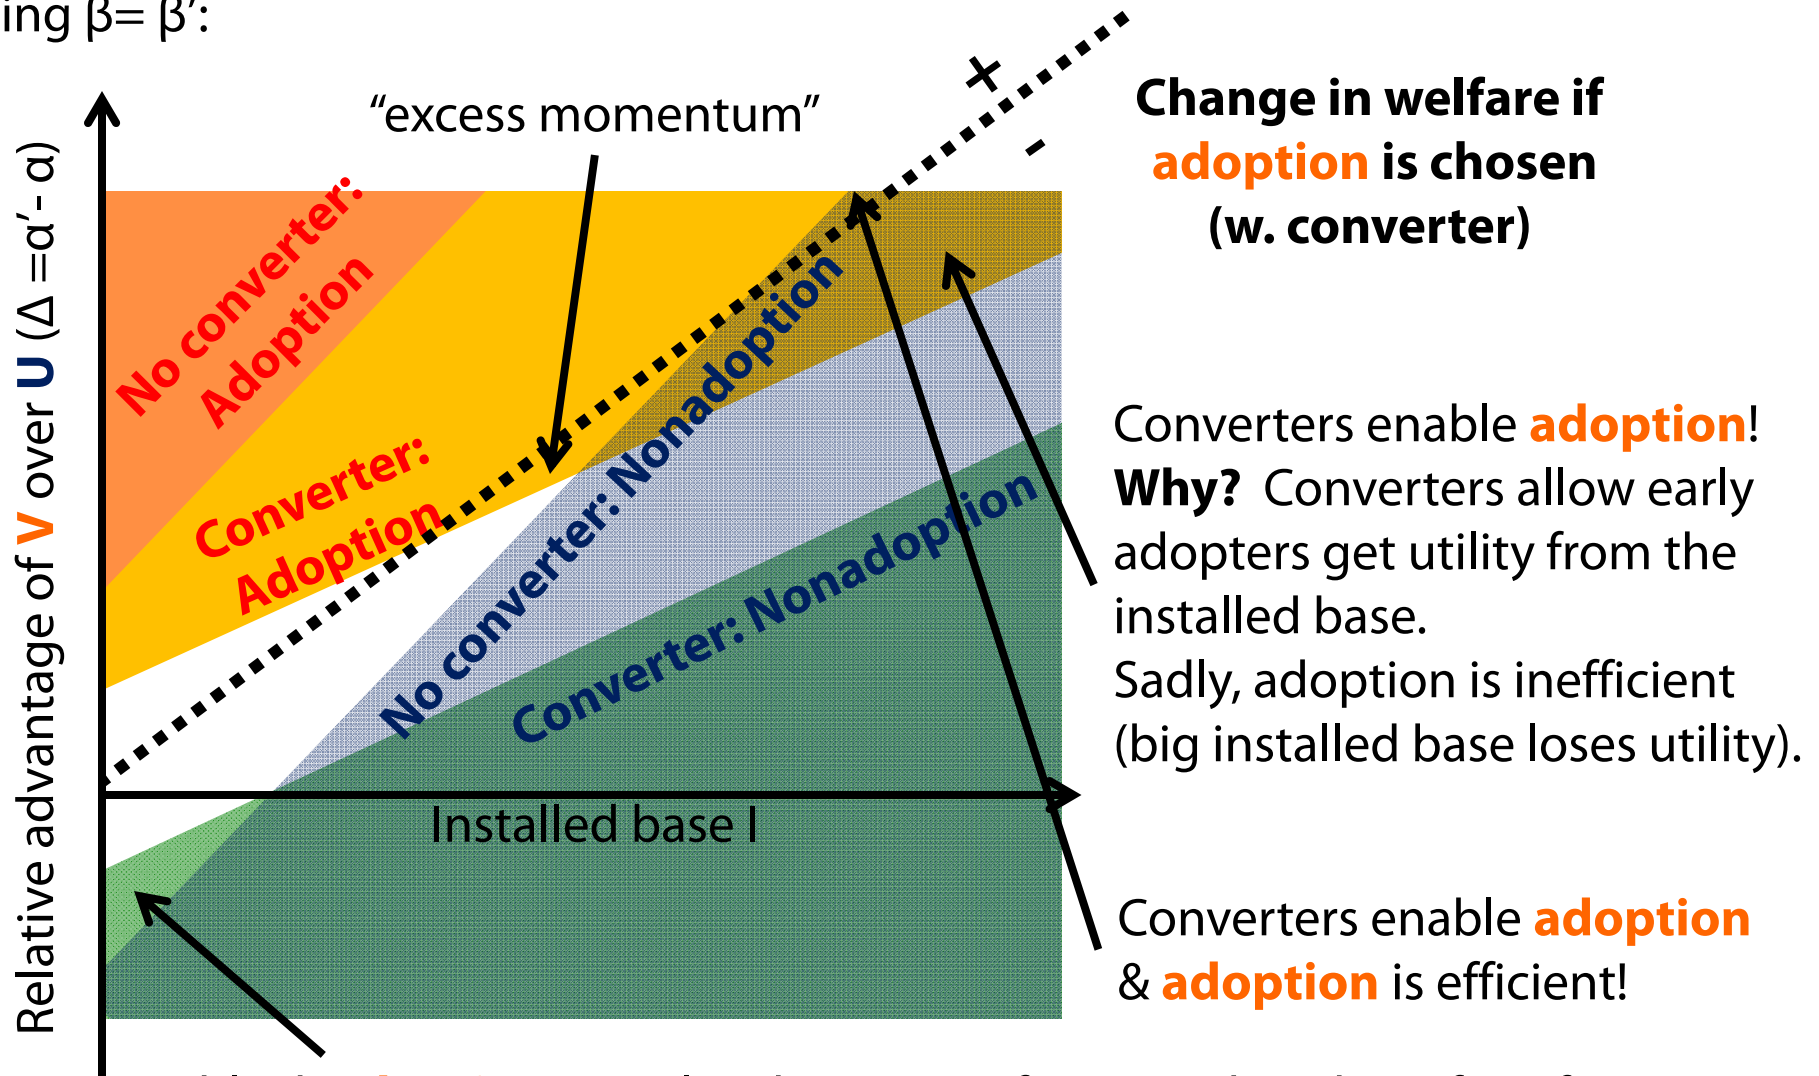
Nonadoption is only equilibrium if $v^+(0) < u^-(0)$

Else, both **nonadoption** and **adoption** are equilibria



Do converters facilitate the transition? [Choi 1994] (2)

Letting $\beta = \beta'$:

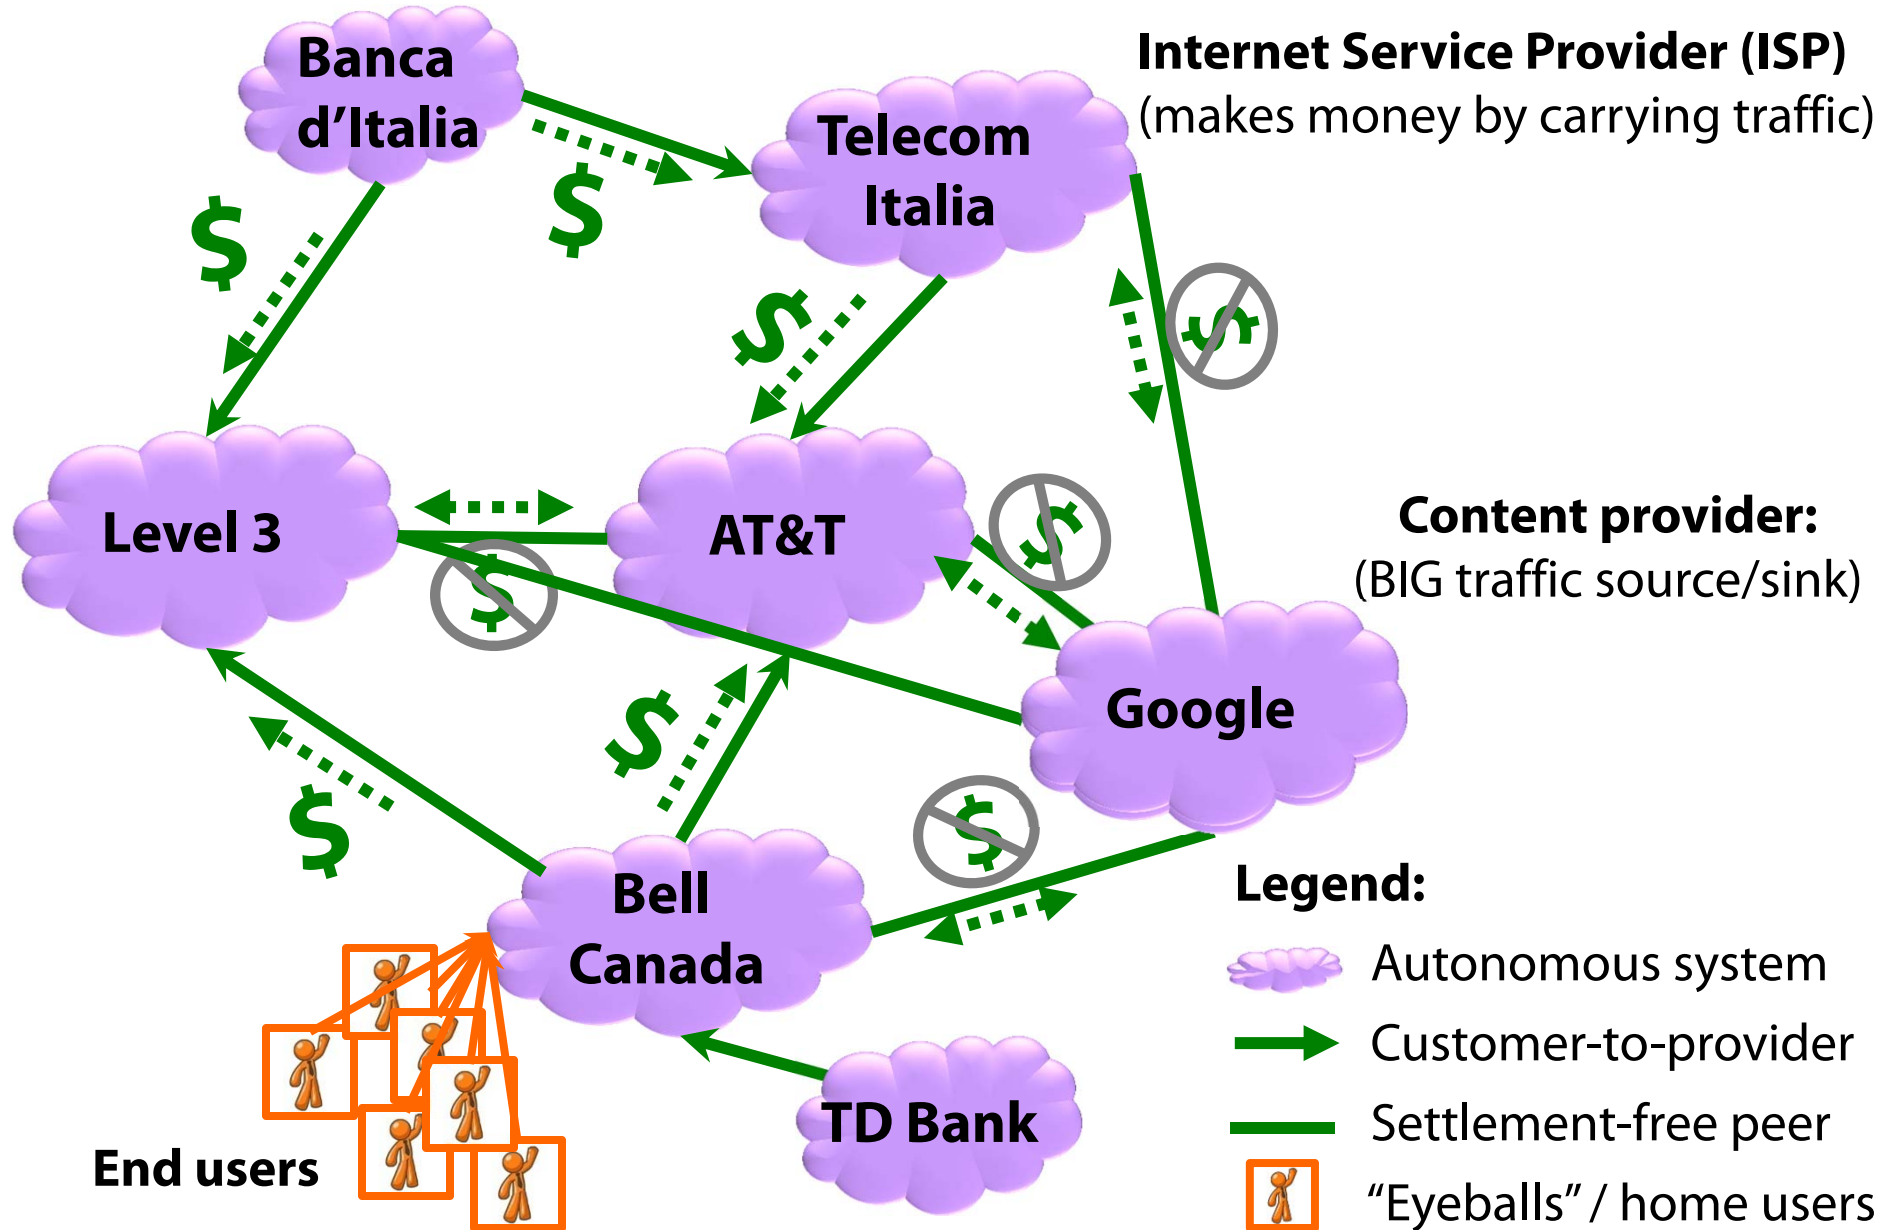


Converters block **adoption**! (Early adopters prefer immediate benefits of joining installed base, instead of long-term gain they preferred w/o converters.)
Fortunately, adoption is inefficient.




Internet Economics 101

Stub AS (traffic source/sink!) – **85% of ASes!** (Fictional topology, for illustrative purposes only!)



The (looming?) upgrade to IPv6



A technology that is:

- 1) not compatible with the installed base, and
- 2) imposes a network externality.



Internet Protocol 101: IPv4



Tx: 65.215.1.12

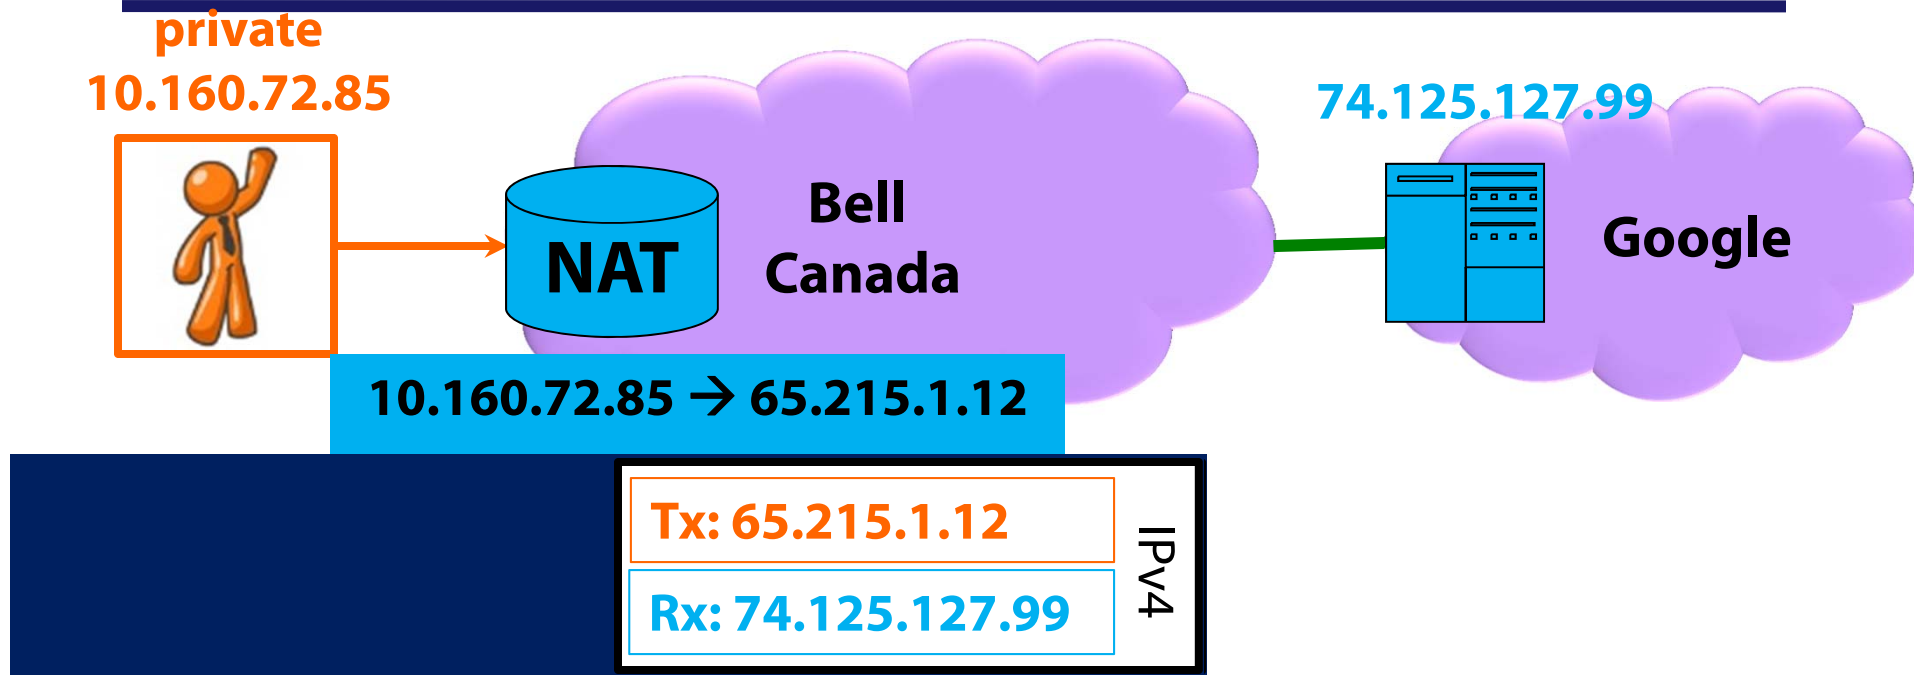
Rx: 74.125.127.99

IPv4

- **The length of the IP address is 32 bits.**
 - This is hard coded into all browsers, laptops, routers, switches, everything!
- **To put this in perspective: $2^{32} = 4$ Billion. # of people in world = 7 Billion**
 - So how do we survive?
- **NATs (Network address translation):**
 - A box that translates a *private* IP address into a *public* one.



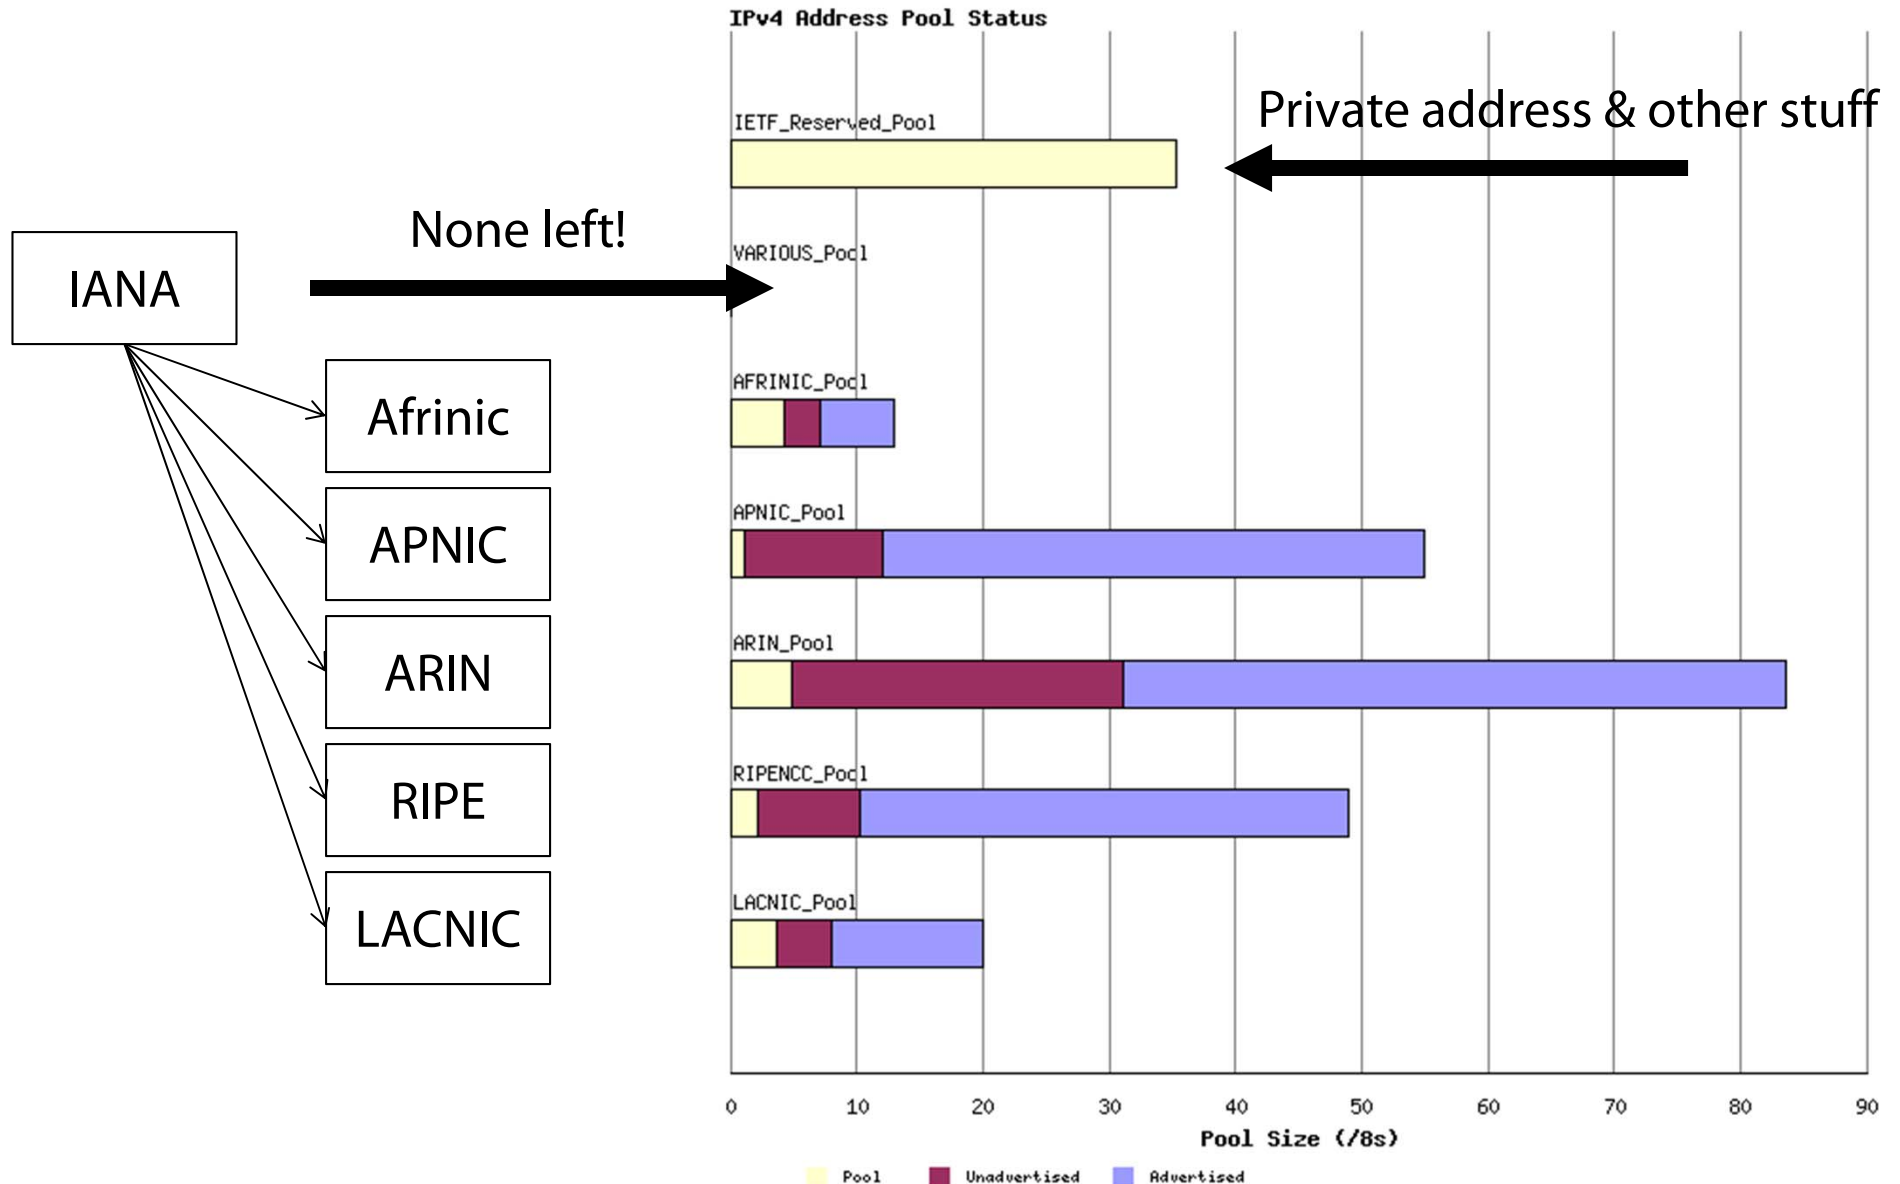
Internet Protocol 101: IPv4 & NAT



- **The length of the IP address is 32 bits.**
 - This is hard coded into all browsers, laptops, routers, switches, everything!
- **To put this in perspective: $2^{32} = 4$ Billion. # of people in world = 7 Billion**
 - So how do we survive?
- **NATs (Network address translation):**
 - A box that translates a *private* IP address into a *public* one.



We have run out of “unallocated” IPv4 addresses





Internet Protocol 101: IPv4 & IPv6



Tx: 10.160.72.85

IPv4

Rx: 74.125.127.99

Tx: 2001:4898:4030:3010:eda7:9a75:49f4:3010

IPv6

Rx: 2001:4860:a002::68


- **IPv6 was standardized in 1998. It increases address length to 128 bits.**
 - Now we have $2^{128} = 3 \times 10^{38}$ addresses. And that's about all it does...
 - except maybe get rid of NATs once **everyone** uses it.
- **IPv6 is not compatible with IPv4, because the headers are different!**



Why is this transition so difficult? (1)

Problem 1: Everyone on Internet must be able to talk to everyone else.

- How to achieve this with incompatible technologies?
- **Solution 1:** Use conversion. The catch: Performance degrades.



- **Solution 2:** “Dual stack” devices run both IPv6 and IPv4.
 - The catch: This doesn’t save addresses.
 - The catch: IPv4-only device perform better & can still talk to everyone!



Why is this transition so difficult? (2)

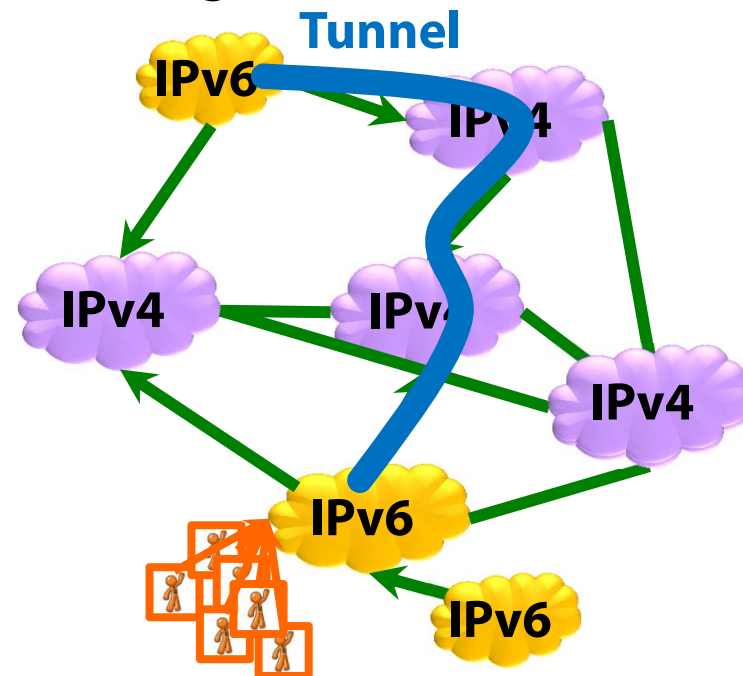
Problem 1: Everyone on Internet must be able to talk to everyone else.

- How to achieve this with incompatible technologies?
- **Solution 1:** Use conversion. The catch: Performance degrades.
- **Solution 2:** “Dual stack” devices run both IPv6 and IPv4.
 - The catch: This doesn’t save addresses.
 - The catch: IPv4-only device perform better & can still talk to everyone!

Problem 2: There may not be a IPv6 path through the network

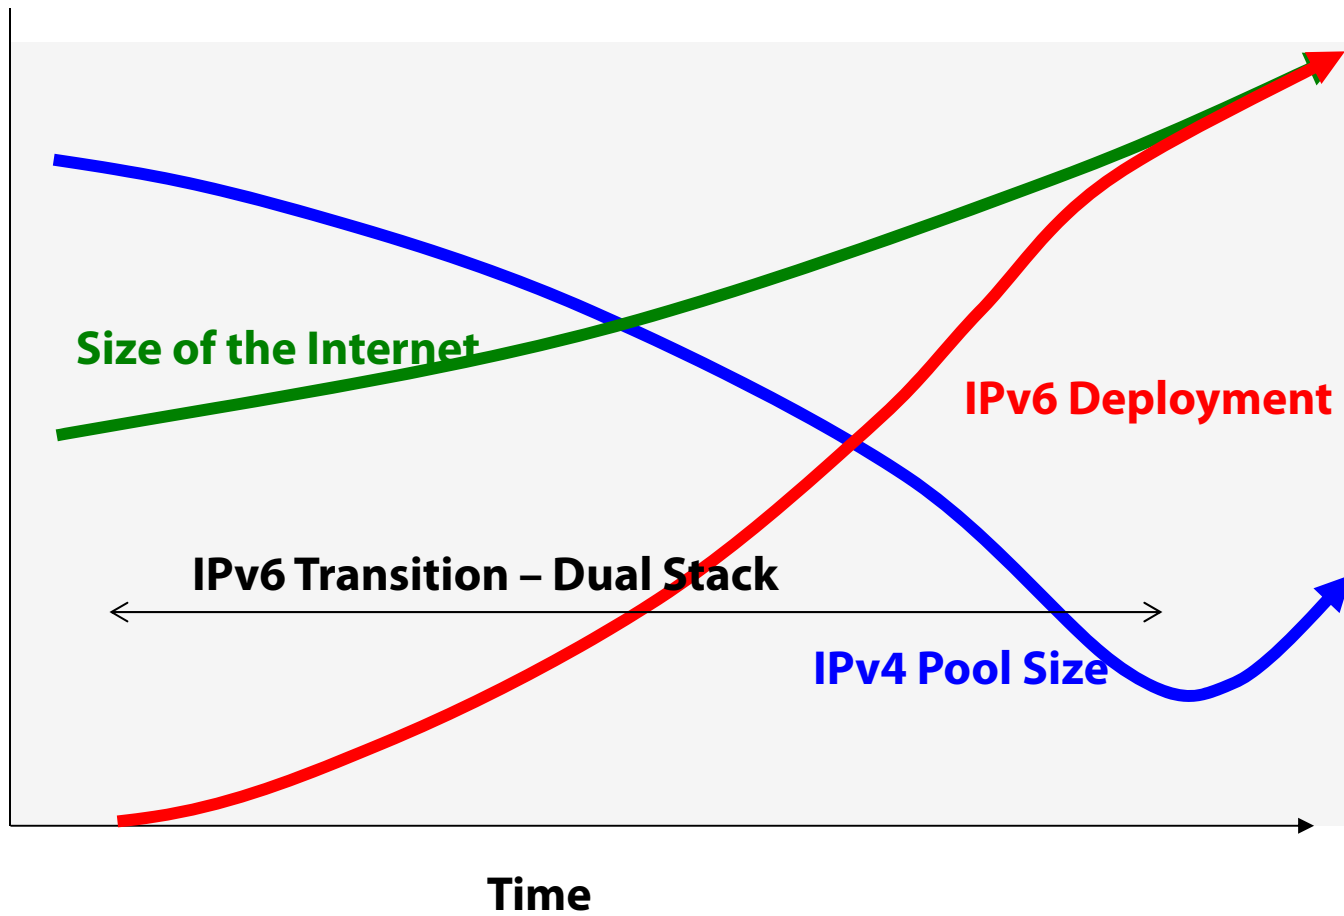
- Solution: Use a tunnel.
 - The catch: Performance degrades.

Are conversion technologies speeding/slowing the transition?





The IPv6 Transition Plan





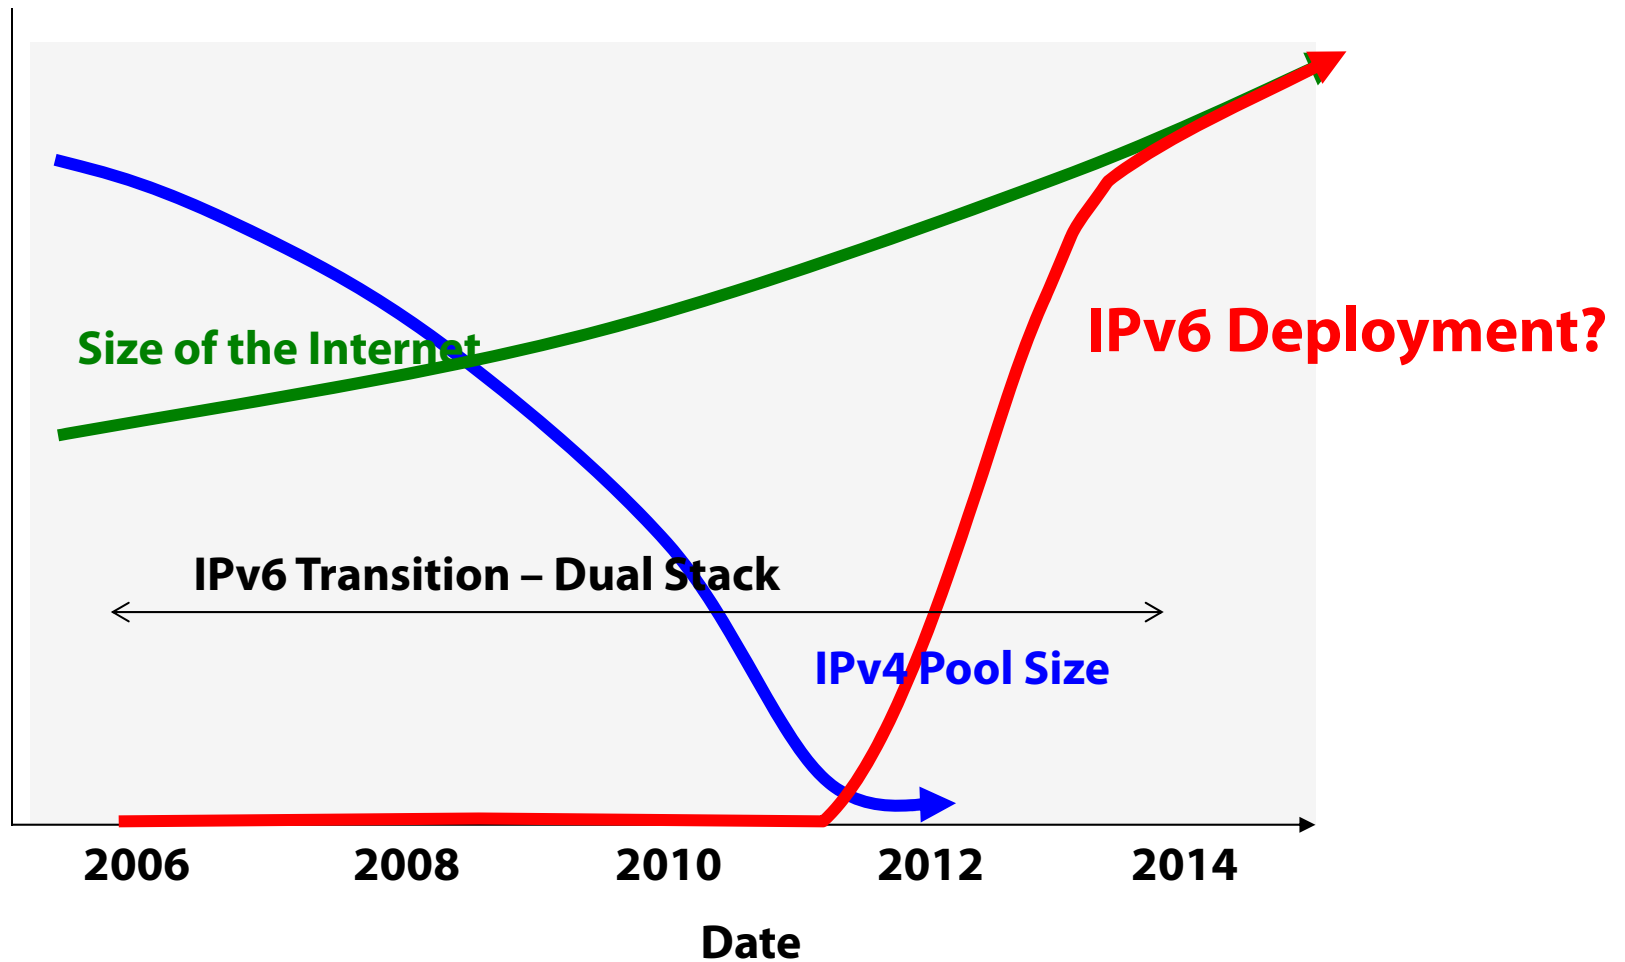
IPv6 Adoption as seen by Google (June 3, 2012)



<http://www.google.com/intl/en/ipv6/statistics/>



The IPv6 Transition Plan v 2.0



Or maybe we'll just keep using IPv4 indefinitely?



Can this transition be managed via market mechanisms?

“The minister for communications and information technology does not believe that regulatory intervention is appropriate. Adoption of IPv6 needs to be lead by the private sector. The private sector must recognise that adopting IPv6 is in their own best interests to protect their investment in online capabilities into the future. Issues of advantages and disadvantages, costs, risks, timing, methodology etc, have to be for each enterprise to assess for itself.”

New Zealand Minister for Communications, Aug 24, 2009

- **IPv4 vs IPv6. What's the difference?**
 - Cost? Not right now.
 - Functionality? No.
 - Performance? No. Actually IPv6 usually performs worse.
 - Consumer visibility difference? No.
 - Consumer demand? No.
 - Competitive differentiators? Only future risk.

[Slide: Geoff Huston]



Overview of results on the transition from IPv4 to IPv6

How do converters, quality & price affect the transition:

- [Choi'94] & [Sen, Jin, Guerin, Hosanger'10]

Transit ASes should act as 2-sided market [Guerin, Hosanger']

- Give customers & content providers incentives to use IPv6.

Markets aren't going to work here – IPv6 is public good.

- Geoff Huston subsidies and regulation could help.
- Or maybe just peer pressure & publicity?



13D 02:30:13
THIS TIME IT IS FOR REAL
6 JUNE 2012

AKAMAI	AT&T	CISCO
COMCAST	D-LINK	FACEBOOK
FREE TELECOM	GOOGLE	INTERNODE
KDDI	LIMELIGHT	MICROSOFT BING
TIME WARNER CABLE	XS4ALL	YAHOO!

**DO YOUR PART
JOIN THE LAUNCH!**

Federal Government Imposes New IPv6 Deadlines

'Native' IPv6 by FY 2012 for public-facing servers and services

Sep 30, 2010 | 04:03 PM |

By Kelly Jackson Higgins

The Obama administration's CIO has informed federal agencies that they must run native IPv6 on their Web, email, ISP, and DNS servers and services by the end of fiscal year 2012, and their internal client applications by fiscal year 2014.

In a Sept. 28 memo (PDF) sent to federal CIOs, Vivek Kundra, the nation's federal CIO, outlined the

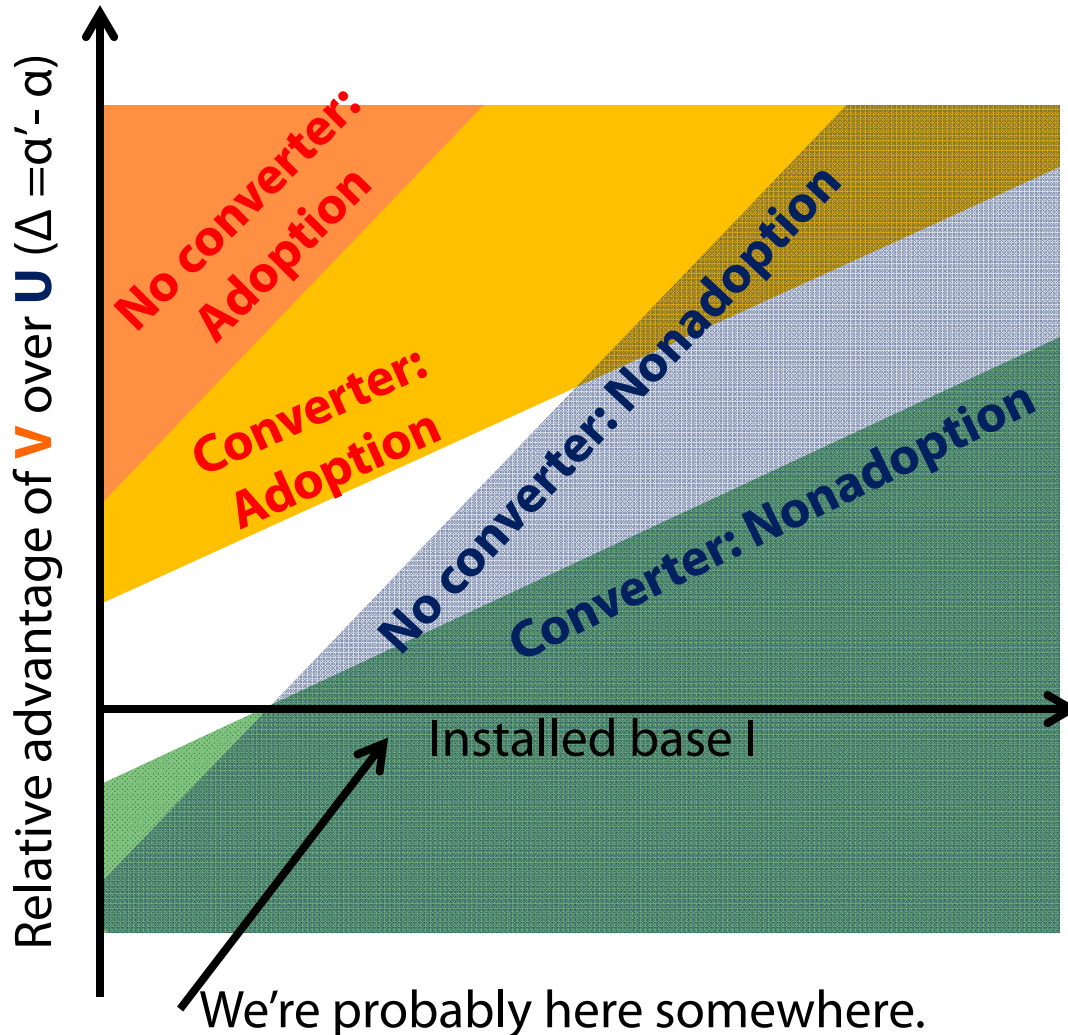
Forget it, it won't happen. Treat IPv4 address as scarce commodity & move on!

- [Edleman Swartz '11] Let's just auction off IPv4 address space.



Were do we fall in the [Choi 1994] model?

It's hard to say...



One way to fit IPv6 in this model is to have a very nonlinear network externality function, so that the positive benefit only kicks in when **most** people adopt.

Also, installed base should also be able to switch to IPv6.



Another model of the transition [Sen, Jin, Guerin, Hosanger'10]

Technology 1: $\mathbf{U}_1(\theta, \mathbf{x}_1, \mathbf{x}_2) = \theta \mathbf{q}_1 + (\mathbf{x}_1 + \alpha_1 \beta \mathbf{x}_2) - \mathbf{p}_1$

Technology 2: $\mathbf{U}_2(\theta, \mathbf{x}_1, \mathbf{x}_2) = \theta \mathbf{q}_2 + (\mathbf{x}_2 + \alpha_2 \beta \mathbf{x}_1) - \mathbf{p}_2$

Cost (recurrent) of each tech (\mathbf{p}_i) and intrinsic technology quality (\mathbf{q}_i)

Linear Network Externalities ($\mathbf{0} < \mathbf{x}_1 + \mathbf{x}_2 < \mathbf{1}$)

- $\alpha_i, \mathbf{0} \leq \alpha_i \leq \mathbf{1}, i = 1, 2$, captures converters performance

User sensitivity to technology quality (θ)

- Private information for each user, but known distribution

If $\mathbf{H}_i(\mathbf{x}_1(\mathbf{t}), \mathbf{x}_2(\mathbf{t}))$ is fraction of users that prefer tech \mathbf{i} at time \mathbf{t}

$$\mathbf{x}'_1(\mathbf{t}) = \boldsymbol{\gamma} [\mathbf{H}_1(\mathbf{x}_1(\mathbf{t}), \mathbf{x}_2(\mathbf{t})) - \mathbf{x}_1(\mathbf{t})]$$

$\boldsymbol{\gamma}$ is hazard rate, ie. $\mathbf{Pr}[\text{User adopts at time } \mathbf{t} \mid \text{It didn't adopt at time } < \mathbf{t}]$

At equilibrium $\mathbf{H}_1(\mathbf{x}_1^*(\mathbf{t}), \mathbf{x}_2(\mathbf{t})) = \mathbf{x}_1^*(\mathbf{t})$

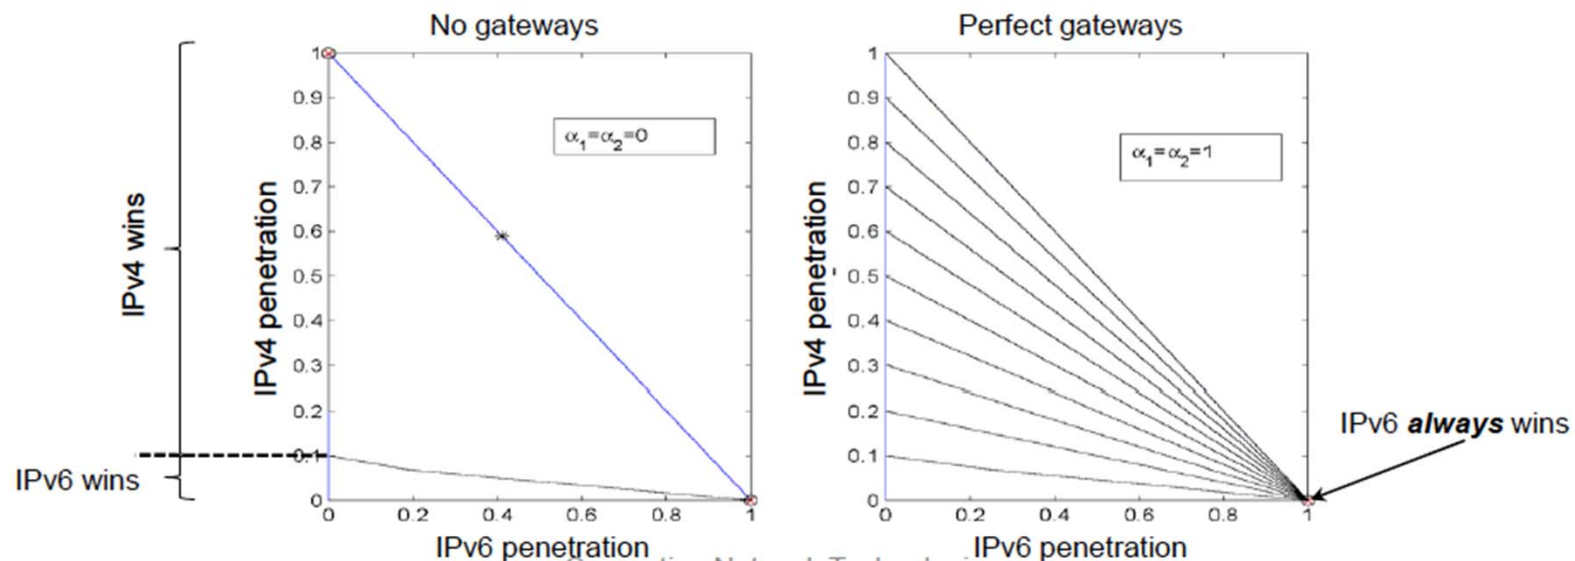


Sample Results [Sen, Jin, Guerin, Hosanger'10]

Question: Given a starting point $(x_1(0), 0)$, what is (x_1^*, x_2^*) at equilibrium?

IPv4 Slightly “Better” than IPv6

- In the absence of gateways, IPv6 never takes off unless IPv4 initial penetration is very low...
- After introducing “perfect” gateways ($\alpha=100\%$), IPv6 eventually takes over, irrespective of IPv4 initial penetration
 - There is a “threshold” value (**80%**) for gateway efficiency below which this does not happen!





What if ISPs act as 2-sided market? [Guerin, Hosanger'11]



Number of users is
 $x_4 + x_6(t)$
exogenous, fixed

Chooses qualities: q_4, q_6, q_{46}
Choose α : fraction of IPv4 users
that also get a IPv6 address .
Subject to b : budget for fraction of
converted traffic.

Chooses whether to make
server available via IPv6
based on the quality IPv6
traffic receives.

Findings:

1. If $q_6 < q_{46}$ then Google has no incentive to become available via IPv6.
2. If q_{46} is "high quality" then keeping converted traffic $< b$ requires $q_4 < q_6$.
3. If $q_6 > q_4$ then choose bigger α , Google adopts IPv6 & converted traffic $< b$.

Issue: Not clear ISP has incentive to be a platform!



Is the transition to IPv6 a market failure?


Huston and others suggest that the transition is a market failure.

- “situation in which ... market equilibria cannot be relied on to give Pareto optimal outcomes” [Mas-Collel, Whinston, Green 1995]
- Usually requires subsidies or government intervention.

Market failures are often caused by public goods:

- “Consumption of a unit of the good by one agent does not preclude its consumption by another” [Mas-Collel, Whinston, Green 1995]
- **Huston** suggests IPv6 is a public good.
- **Non-excludable?** Not possible to prevent people who have not paid for it from having access to it.
- **Non-rivalrous?** For any level of production, the cost of providing the good to a marginal (additional) individual is zero.

Maybe its just a case of very nonlinear network externalities? $u(x)$





Or, forget the transition & make IPv4 allocation efficient!

Lots of addresses are allocated but not used – so redistribute using auctions!

- This is already happening, without a formal auction mechanism (i.e. ARIN)
- **[Edleman & Swartz 2011]** an auction for IPv4 address space.

NETWORKWORLD

News | **Blogs & Columns** | Subscriptions | Videos | Events | INS

◀ Previous Post

Next Post ▶

Microsoft pays Nortel \$7.5 million for IPv4 addresses

Bankrupt Nortel finds a buyer for 666K of its legacy IPv4 addresses, raising questions if the IPv4 black/grey market has arrived.

By [Microsoft Subnet](#) on Thu, 03/24/11 - 3:35pm.

Subject: [apnic-talk] need help
From: Rajeev Garg <rajeev@aninetwork.in>
To: apnic-talk@lists.apnic.net
Wed, Jun 6, 2012 at 5:13 PM

Dear sir..

Could you please help us in the matter of locating a user of APNIC who like to transfer its unused IPv4 IPs. As we are short of IPv4 and under APNIC transfer policy we need more IPv4 series...

So please reply as soon as possible..




Or, forget the transition & make IPv4 allocation efficient!

Lots of addresses are allocated but not used – so redistribute using **auctions!**

- This is already happening, without a formal auction mechanism (i.e. ARIN)
- [Edleman & Swartz 2011] an auction for IPv4 address space.

Aggregation in routing impose constraints on the auction:



Routing table

Dest. IP Prefix	Out link
74.125.*.*	B
74.125.127.*	C
74.125.1.*	C
74.200.*.*	B

Tx: 10.160.72.85

Rx: 74.125.127.99

IPv4

...routed the same way

Fewer prefixes → shorter routing table → cheaper & faster routing

Today, with 2^{32} addresses, only about 400K prefixes used

So the auction should break up address space into as few prefixes as possible.




[Edleman & Swartz 2011] IPv4 address auction

Sellers can sell to multiple buyers; buyers can buy from only 1 seller.

“Spartan rule”: After each trade, one agent becomes **extinguished**.

An extinguished agent cannot engage in further trades.



Goal: Find minimal allocation (fewest outgoing edges from each block).

Thm [E&S'11]: A spartan allocation with N buyers has at most N cuts.

Thm [E&S'11]: There is a spartan allocation for every minimal allocation.

Note: No algorithm given to find spartan allocation.



What will happen next?



13D 02:30:13

THIS TIME IT IS FOR REAL
6 JUNE 2012

**AKAMAI
COMCAST
FREE TELECOM
KDDI
TIME WARNER CABLE**

**AT&T
D-LINK
GOOGLE
LIMELIGHT
XS4ALL**

**CISCO
FACEBOOK
INTERNODE
MICROSOFT BING
YAHOO!**

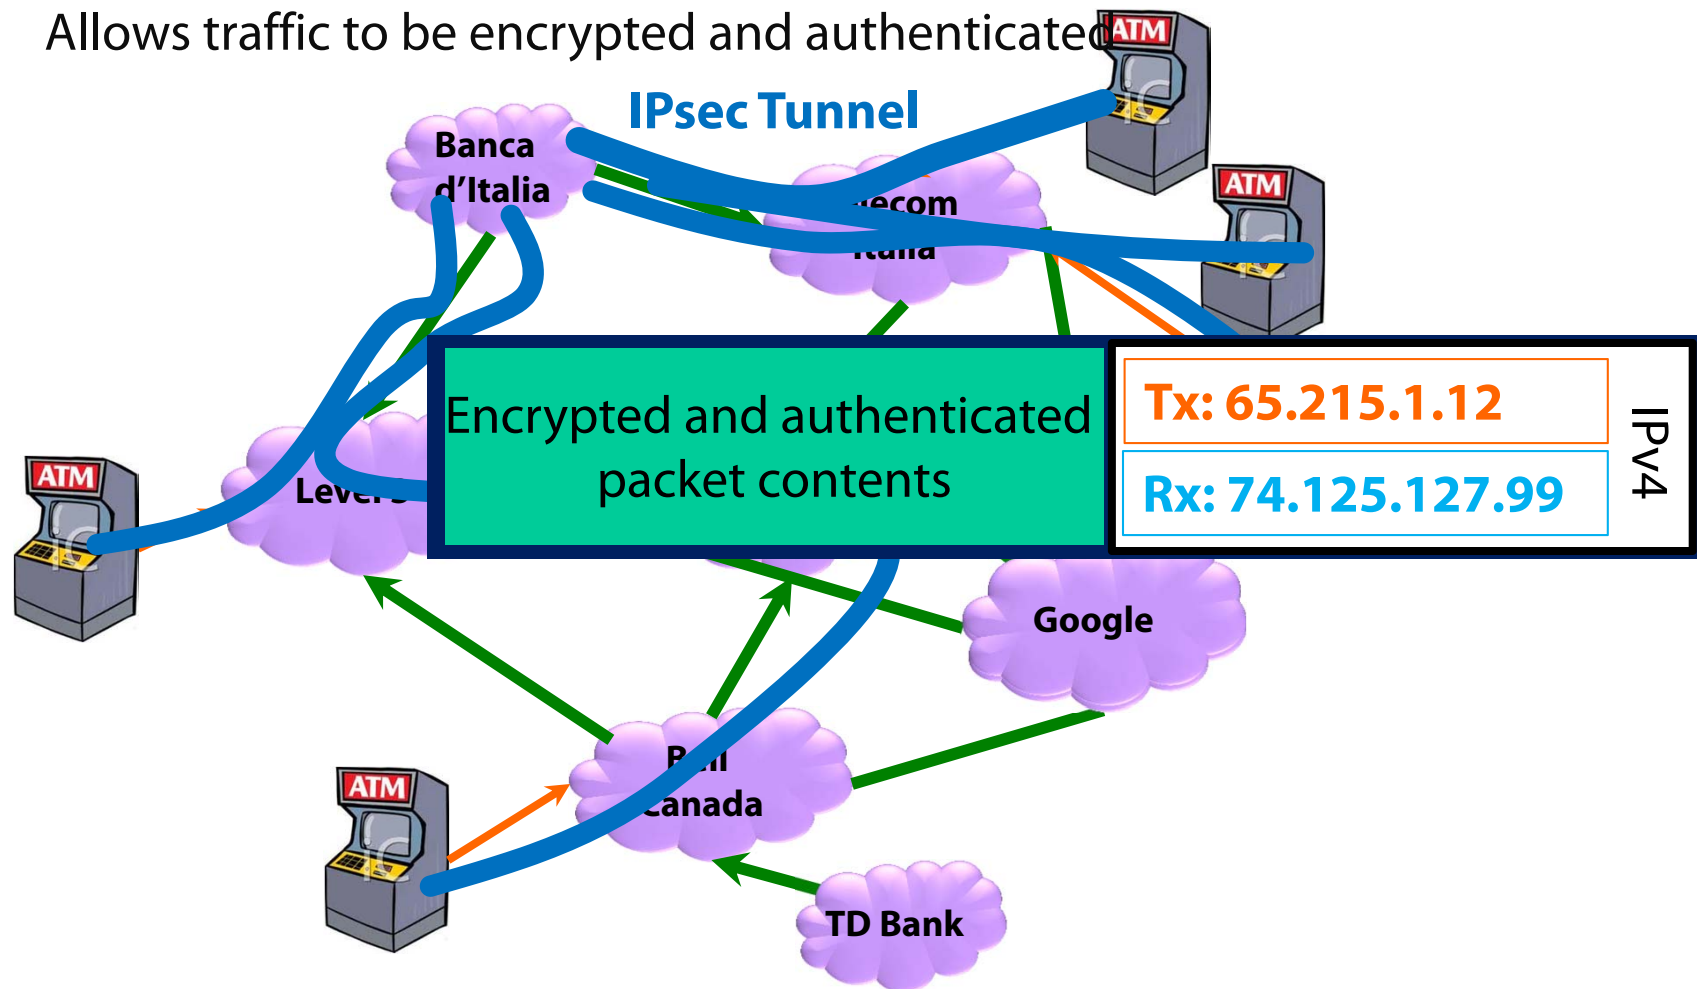
DO YOUR PART
JOIN THE LAUNCH!



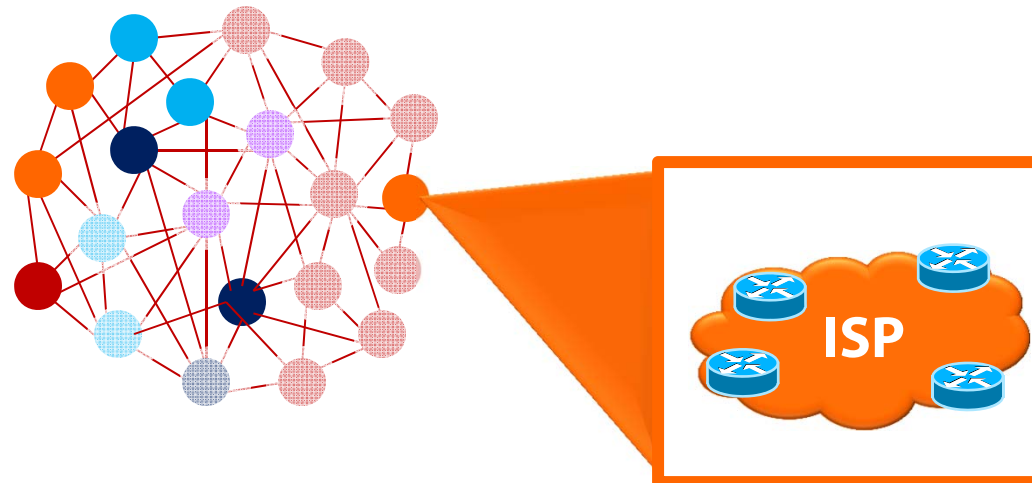
Interlude: IPsec - a success story

IPsec was standardized at same time as IPv6 but **has** been adopted. Why?

- It's compatible with IPv4 and IPv6 – no changes to the packet headers.
- Has incentives for adoption: security & VPN (virtual private network).
- Allows traffic to be encrypted and authenticated



Adoption of routing (BGP) security



A technology that may be rolled out in ~5 years:

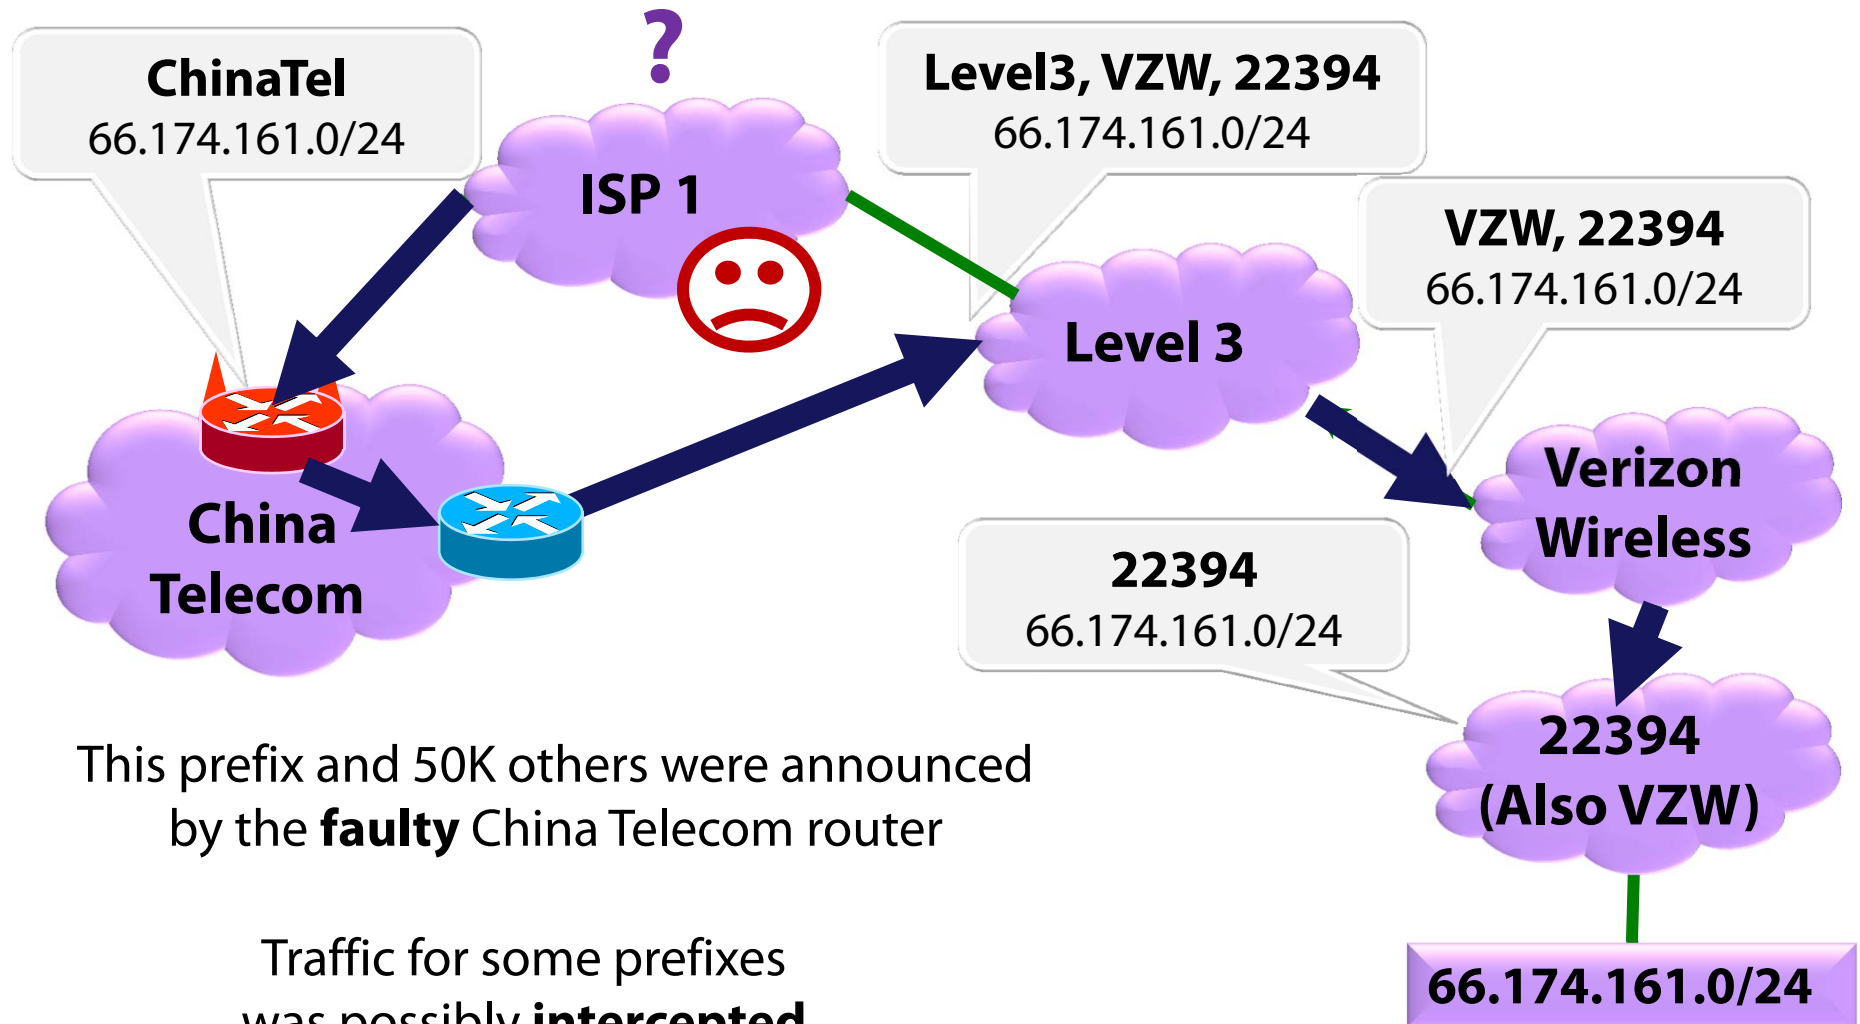
- 1) It's compatible with the installed base, but
- 2) It imposes a network externality on a graph, and
- 3) Interacts with the payment structure in the Internet



Traffic Attraction & Interception Attacks on BGP

An interesting incident from April 8, 2010

ChinaTel path is shorter



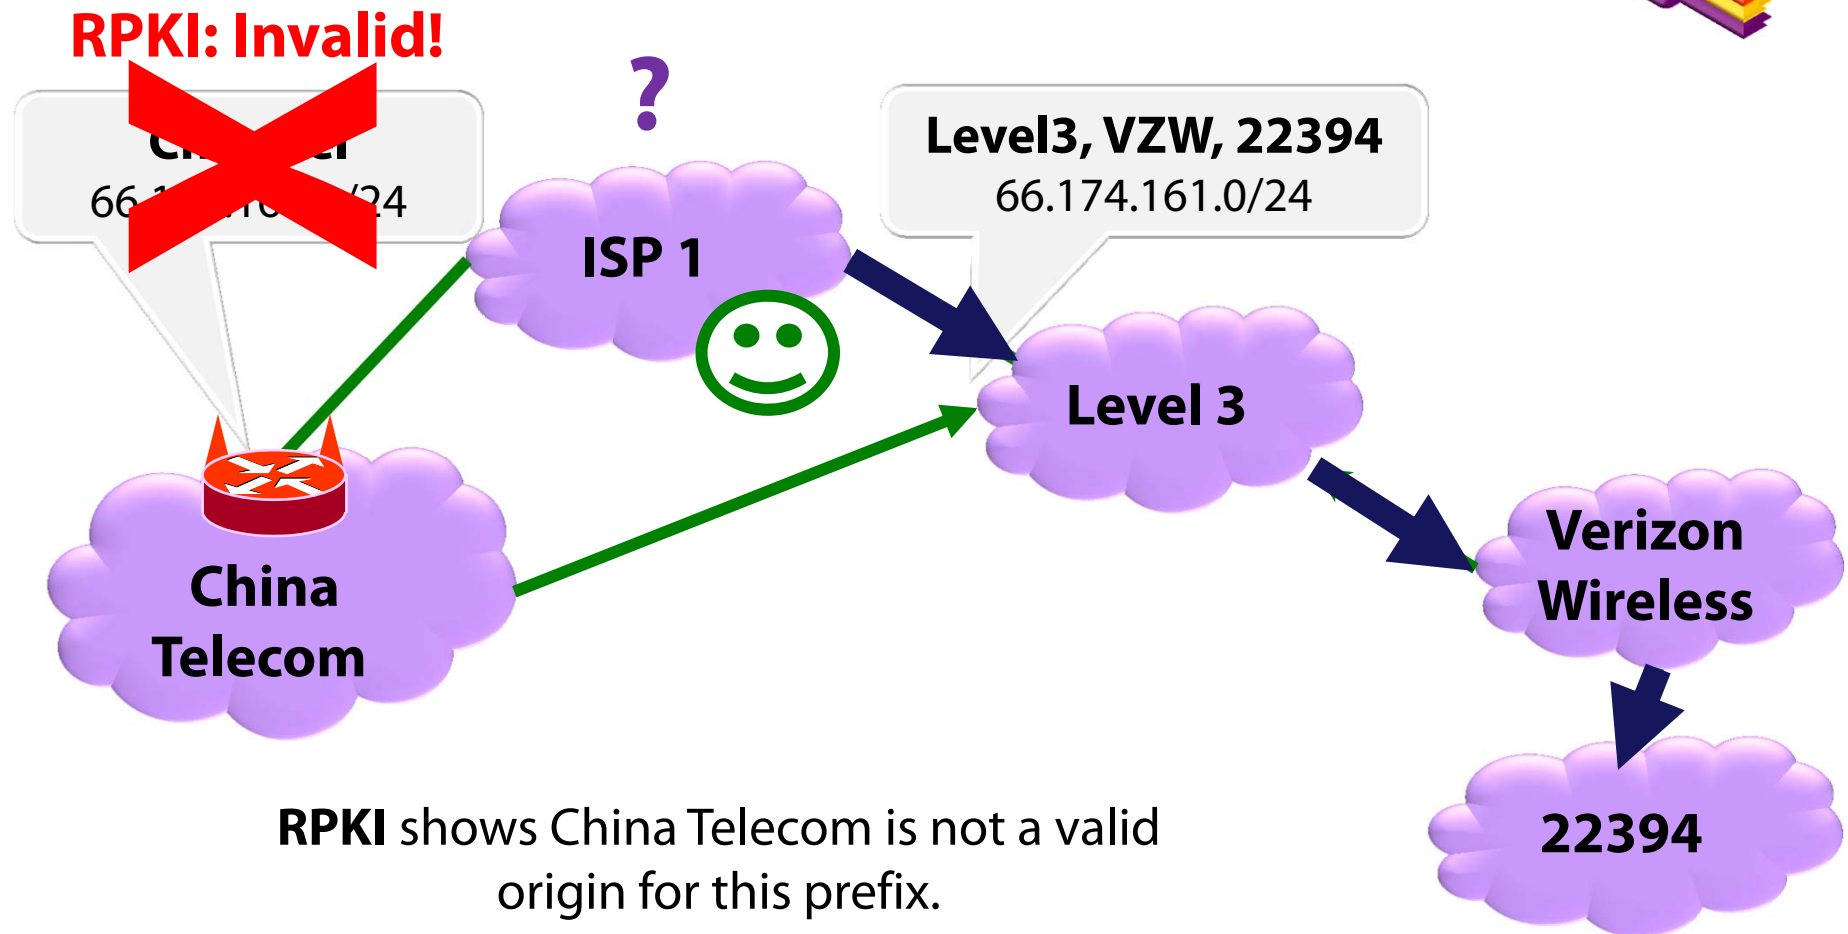

This prefix and 50K others were announced by the **faulty** China Telecom router

Traffic for some prefixes was possibly **intercepted**



Currently under deployment : The **RPKI**

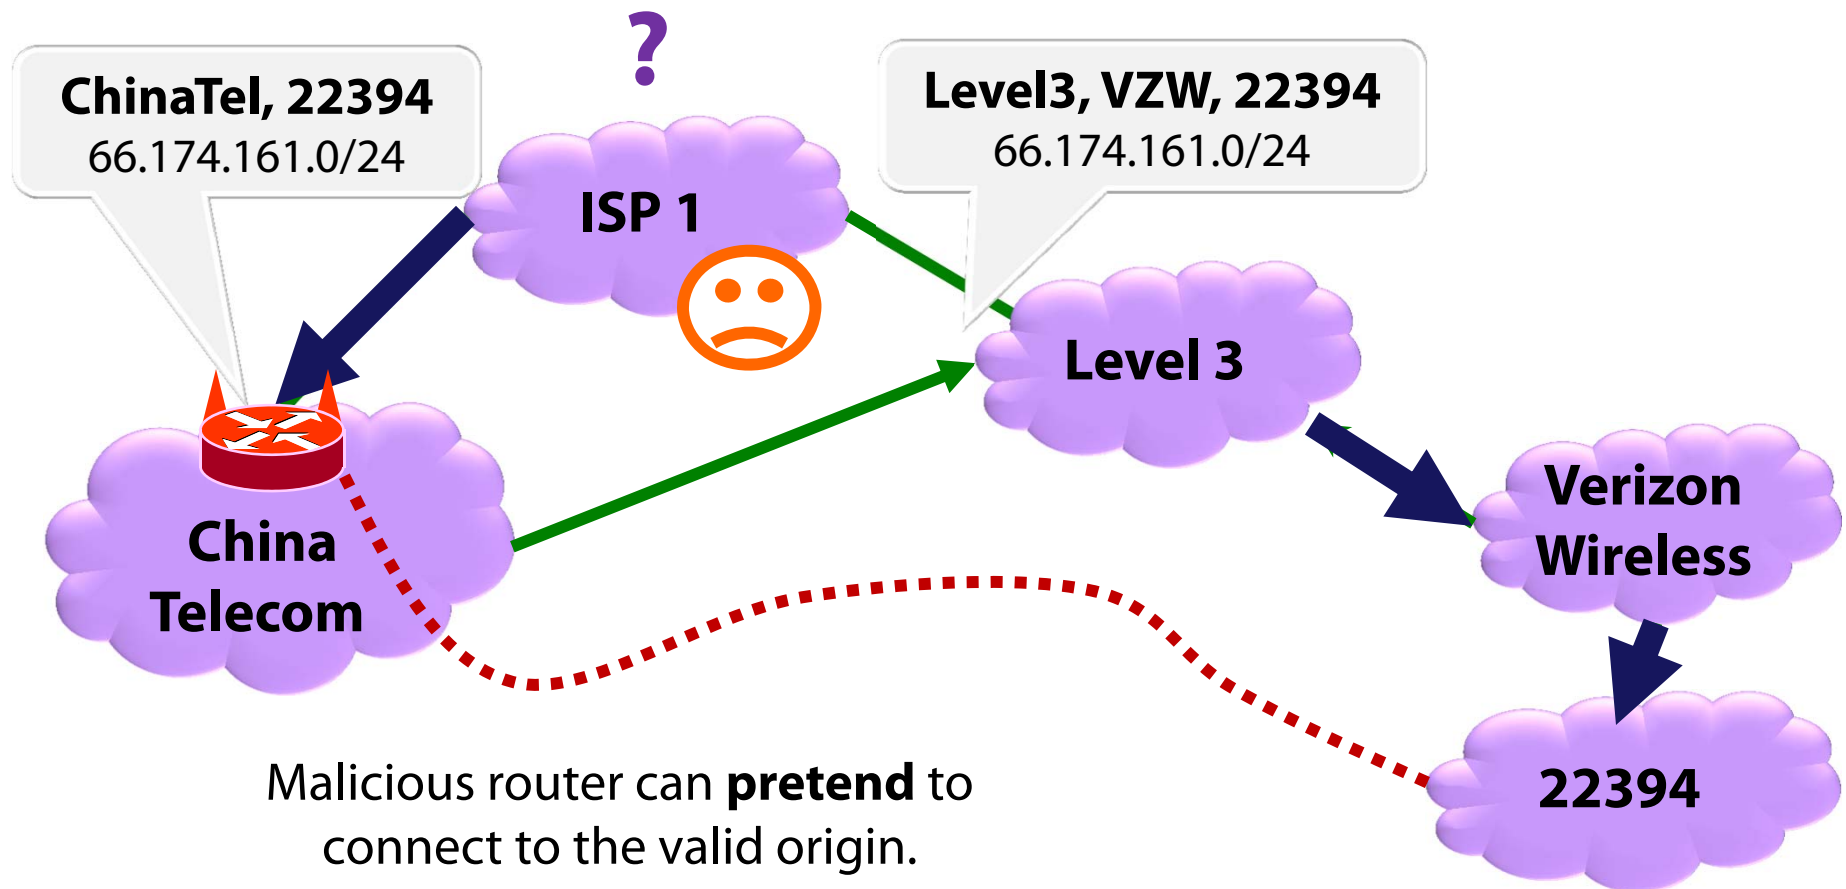

Resource Public Key Infrastructure (RPKI): Certified mapping from Autonomous Systems to public keys and IP prefixes.





But **RPKI** alone is not enough!

Resource Public Key Infrastructure (RPKI): Certified mapping from Autonomous Systems to public keys and IP prefixes.

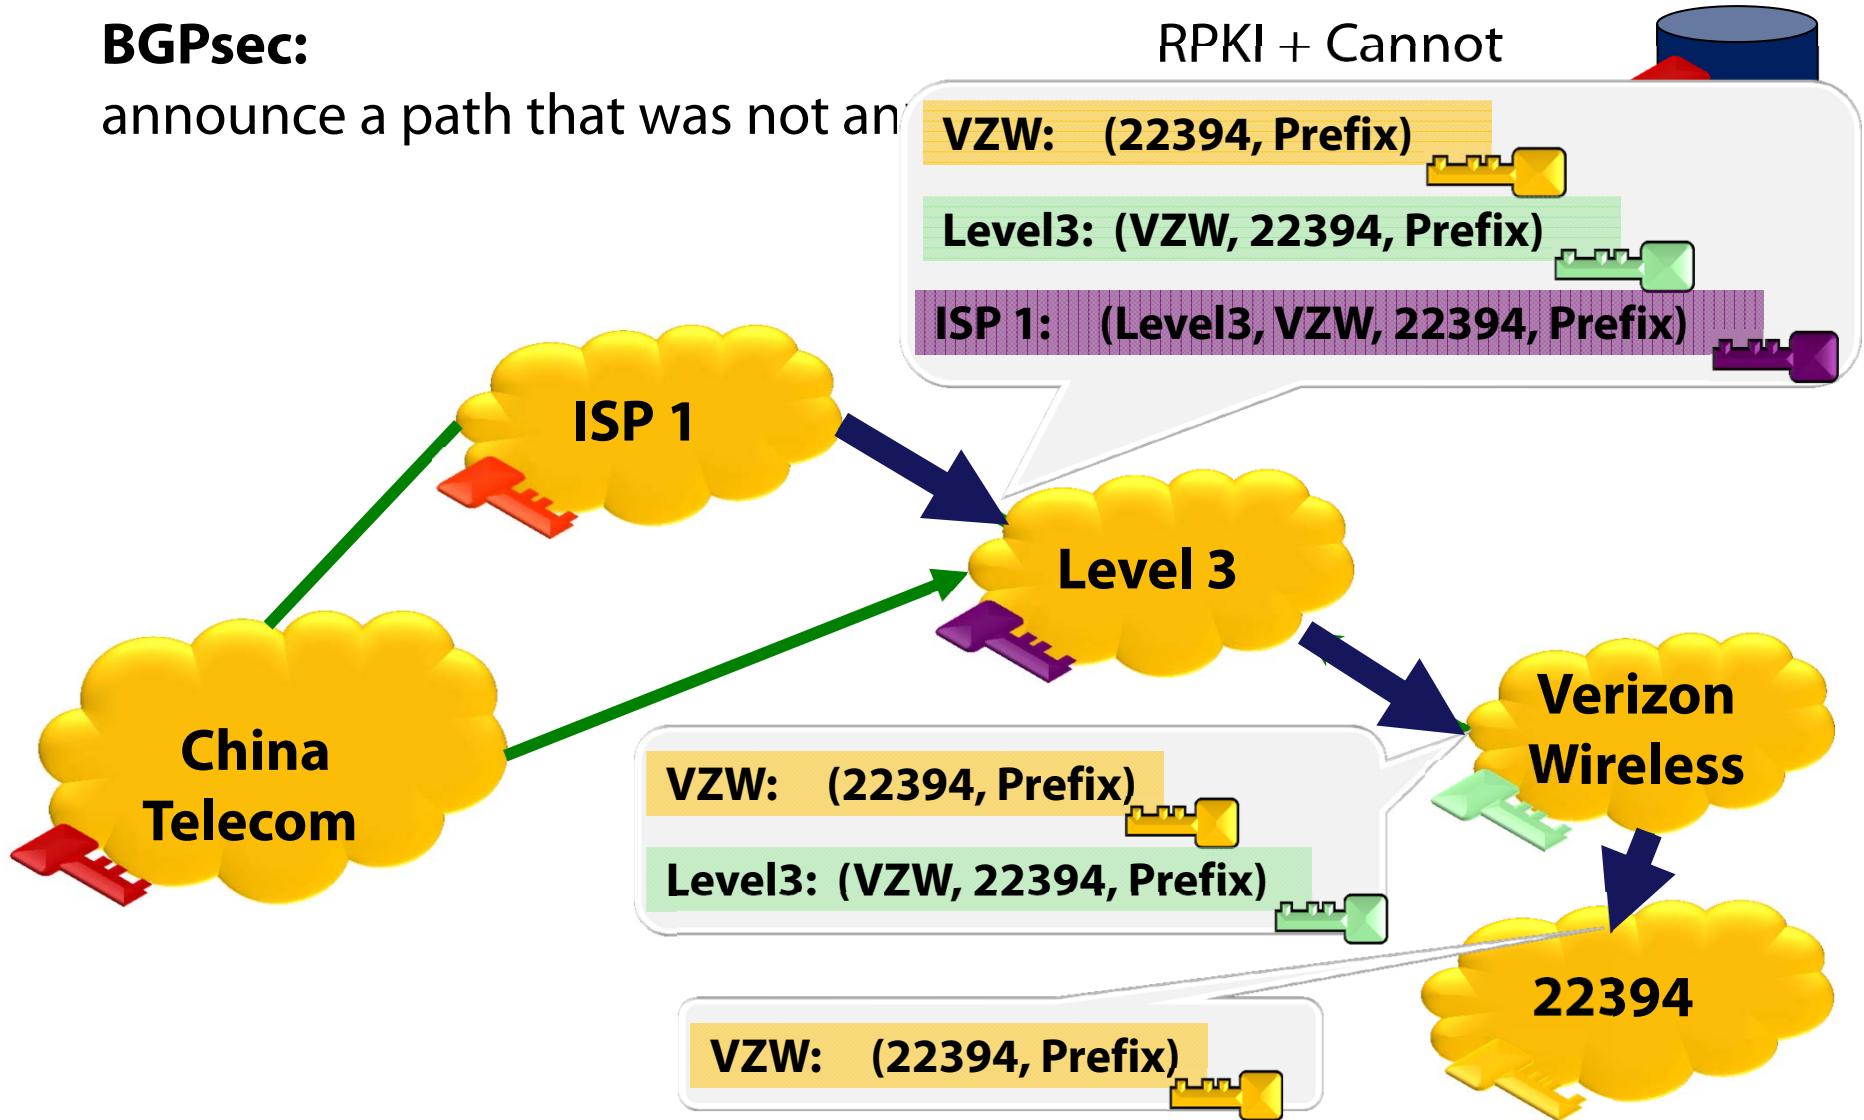




To stop this attack, we need **BGPsec** (1)

BGPsec:

announce a path that was not an



Public Key Signature: Anyone with 22394's public key can validate that the message was sent by 22394.

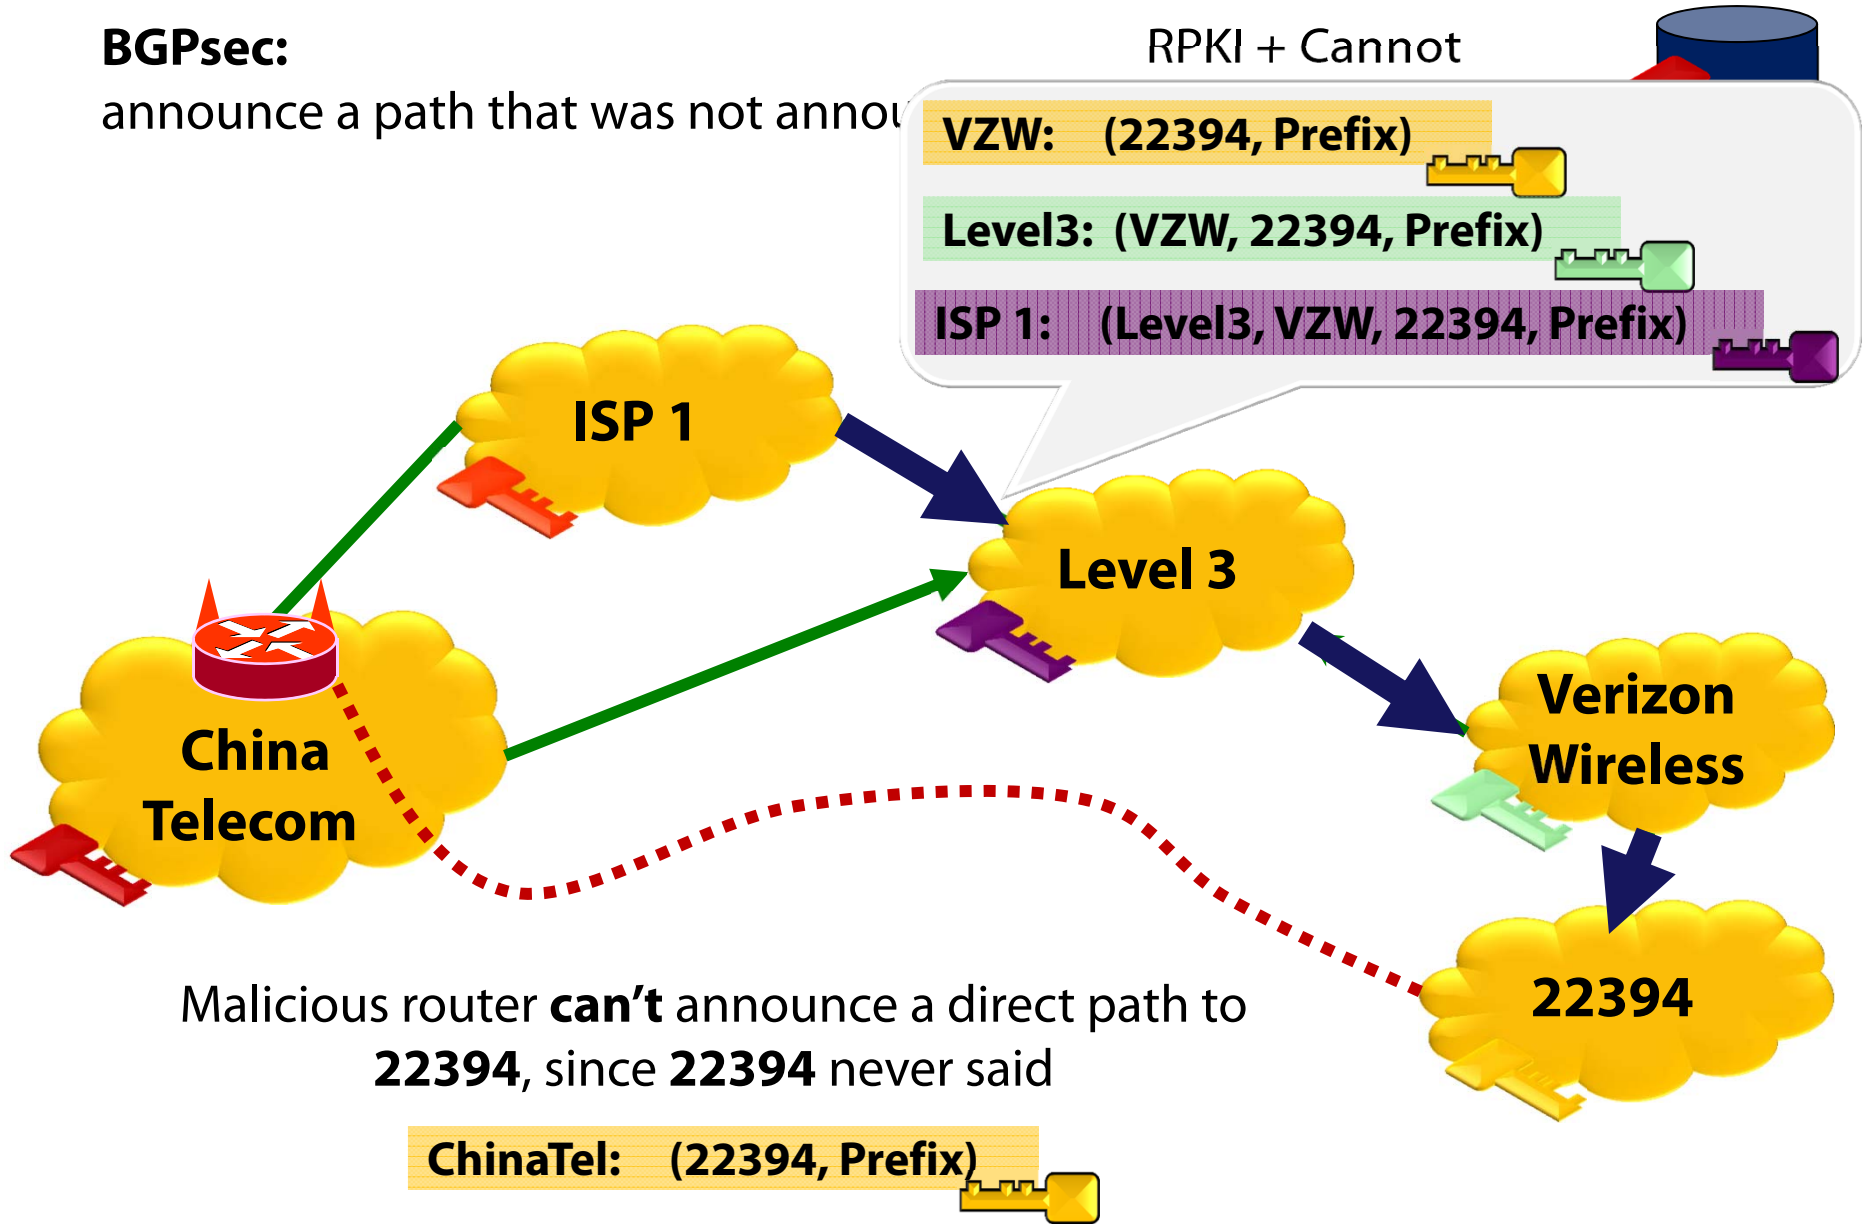




To stop this attack, we need **BGPsec** (2)

BGPsec:

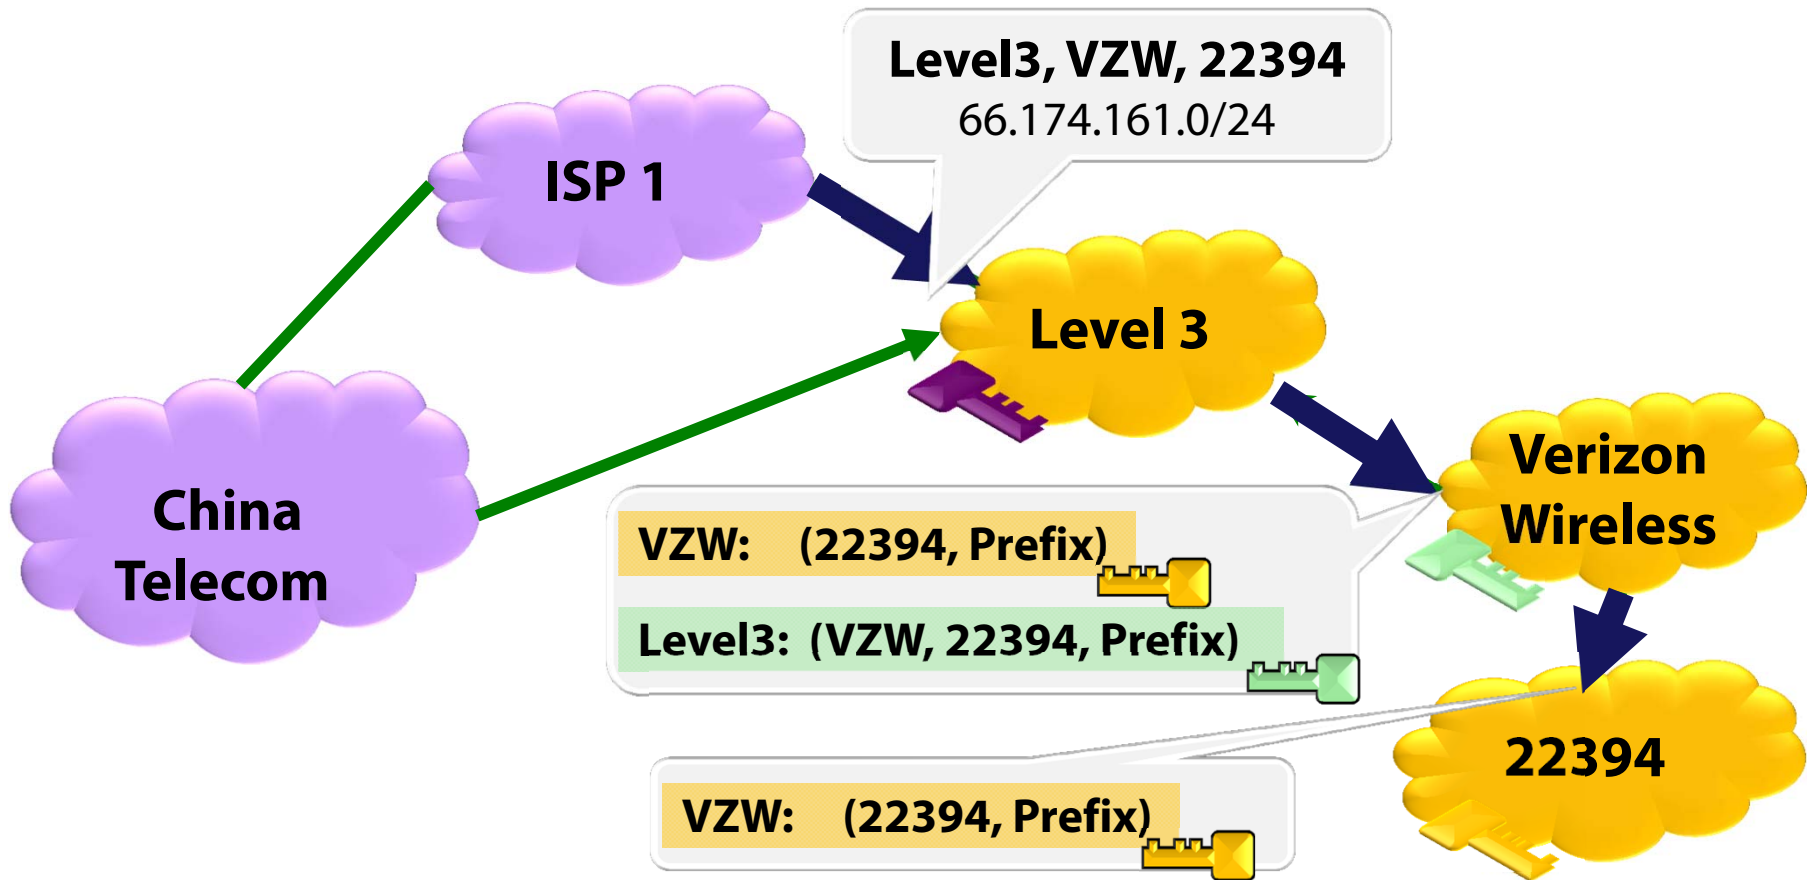
announce a path that was not announced





We learned a lesson on backwards compatibility

Any device that doesn't have BGPsec will see a plain old BGP announcement!





Why is it taking so long to deploy BGPsec?

RPKI is a necessity: But now it's finally happening! (Slowly.)

Local incentives for deployment of BGPsec?

- ASes are economically-motivated agents.
- Security benefits only kick-in when **all ASes on a path** deploy
 - As with IPv6 where quality degrades if not all ASes on path handle IPv6.

We've seen similar problems before:

- Technology diffusion in social networks [**Morris'00**], [**Kempe et.al. '03**]
 - But, utility only depends on immediate neighbors
 - Here it depends on full paths

Why upgrade if
(security) benefits
don't kick in unless
everyone else does?

43284



Overview of work on diffusion of BGPsec

Goal: Develop guidelines for BGPsec deployment.

- Which early adopters lead to cascading BGPsec deployment?
- How should BGPsec interact with routing decisions?

How to evaluate these guidelines?

1. Develop model:

- Model ISP utility. Model routing (Shortest-path? Economics?)
- Game: ISPs myopically upgrade if utility $>$ threshold

2. Analyze model: (Tractability? Convergence?)

3. Simulations:

- Use empirical graphs **$G(V,E)$** of Internet [**UCLA Cyclops**]
- Simulate the deployment process using guidelines + model


Why would ISPs adopt BGPsec?

[**G. & Liu 2012**] Because they can use it to communicate with other ISPs

[**Gill, Schapira, G. 2011**] Because they want to make money!



Technology diffusion in internetworks [G. Liu 2012]



Captures why its so hard to deploy new technologies like (**IPv6**, and **secure BGP**,




I'll adopt the innovation if I can use it to communicate with at least θ other Internet Service Providers (ISPs)!

- $\theta = 2$
- $\theta = 3$
- $\theta = 12$
- $\theta = 15$
- $\theta = 16$

Our model of node utility: Node u 's utility depends on the size of the connected component of active nodes that u is part of.

eg. $utility(u) = 5$



Seedset: A set of nodes that can kick off the process.  Policy makers, regulatory groups can target them as early adopters!

Optimization problem: Given the graph and thresholds, what is the smallest seedset that can cause the entire network to adopt?



Social networks (Local) vs Internetworks (Non-Local)

Minimization formulation: Given the graph and thresholds θ , find the smallest seedset that activates every node in the graph.



Local influence: Deadly hard!

Thm [Chen'08]: Finding an $O(2^{\log^{1-\epsilon}|V|})$ -approximation is NP hard.



Non-Local influence: Much less hard.

[G. Liu 2012]: An $O(r \cdot k \cdot \log |V|)$ approx algorithm

Maximization formulation: Given the graph, assume θ 's are drawn uniformly at random. Find seedset of size k maximizing number of active nodes.



Local influence: Easy!

Thm [KKT'03]: An $O(1-1/e)$ -approximation algorithm.

How? 1) Prove submodularity. 2) Apply greedy algorithm.



Non-Local influence: ?

[G. Liu 2012]: The usual submodularity tricks fail.



[G. Liu 2012] Results

Minimization formulation: Given the graph and thresholds θ , find the smallest seedset that activates every node in the graph.



Main result: An $\mathbf{O}(r \cdot k \cdot \log |V|)$ approx algorithm

r is graph diameter (length of longest shortest path)
 k is threshold granularity (number of thresholds)



Lower Bound: Can't do better than an $\mathbf{\Omega}(\log |V|)$ approx.
(Even for constant r and k .)



Lower Bound: Can't do better than an $\mathbf{\Omega}(r)$ approx. with our approach.



What if ISPs value revenue above security?

Pessimistic view:

- No local economic incentives; only security incentives.
- Similar to IPv6 (except we have backwards compatibility)

“ISPs would be the ones forced to upgrade all of their equipment to support this initiative, but how would it benefit them?

As commercial companies, if there is little to no benefit (potential to increase profit), why would they implement a potentially costly solution?

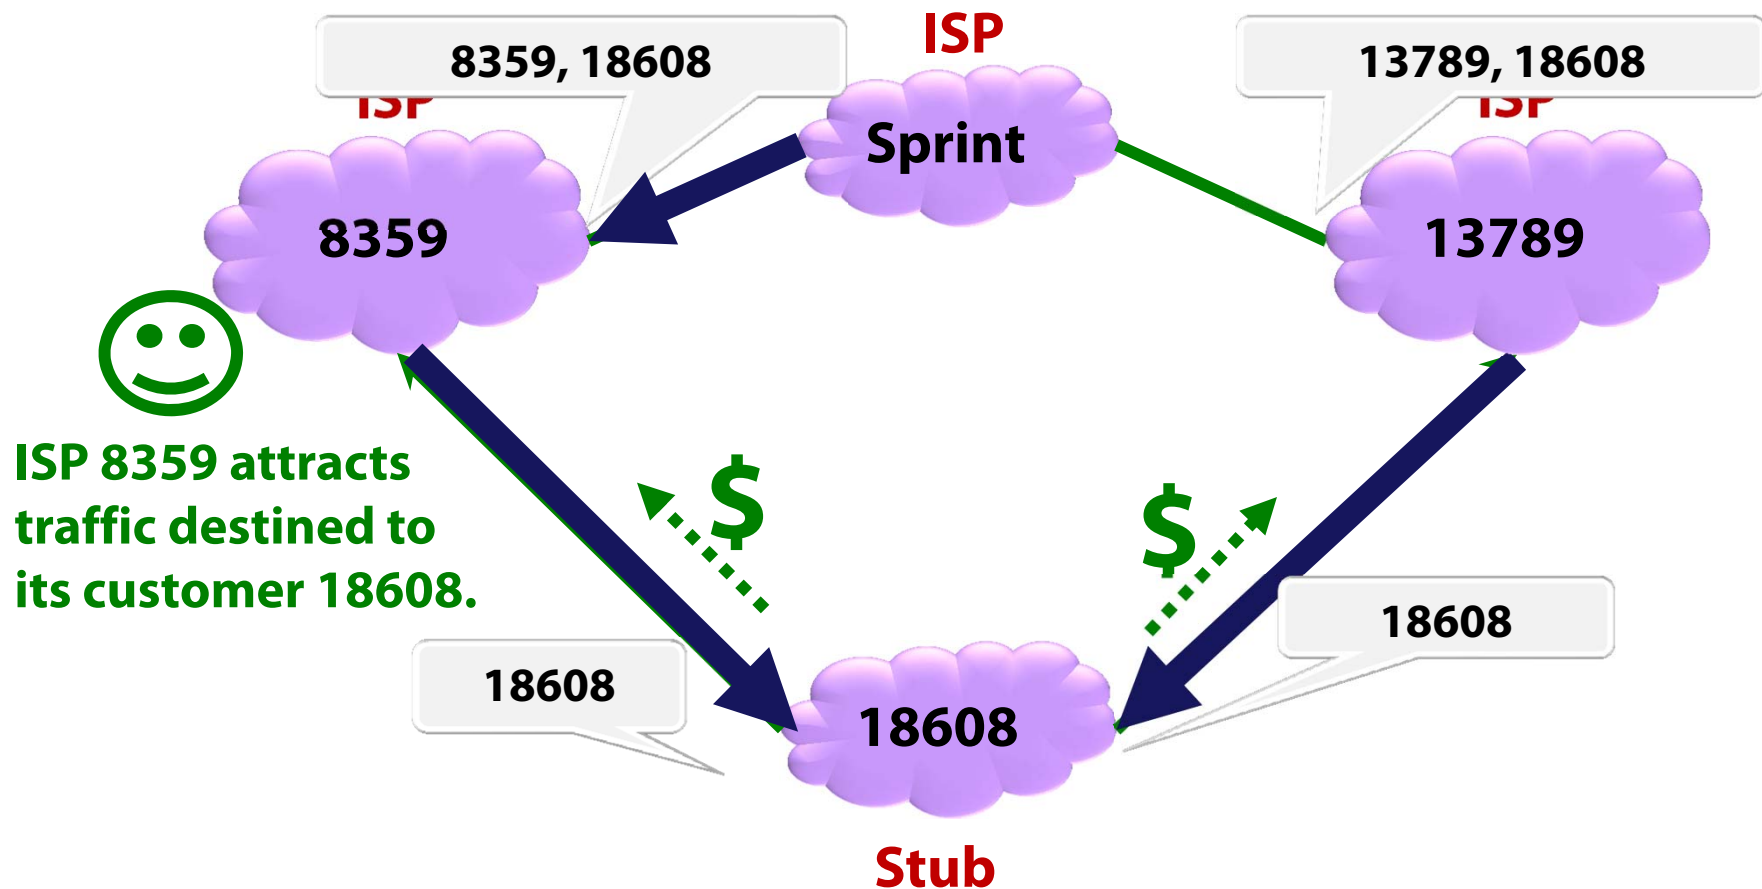
The answer is they won't.”

[http://www.omninerd.com/articles/Did_China_Hijack_15_of_the_Internet_Routers_BGP_and_Ignorance]



What if ISP utility depends on attracting traffic? (1)

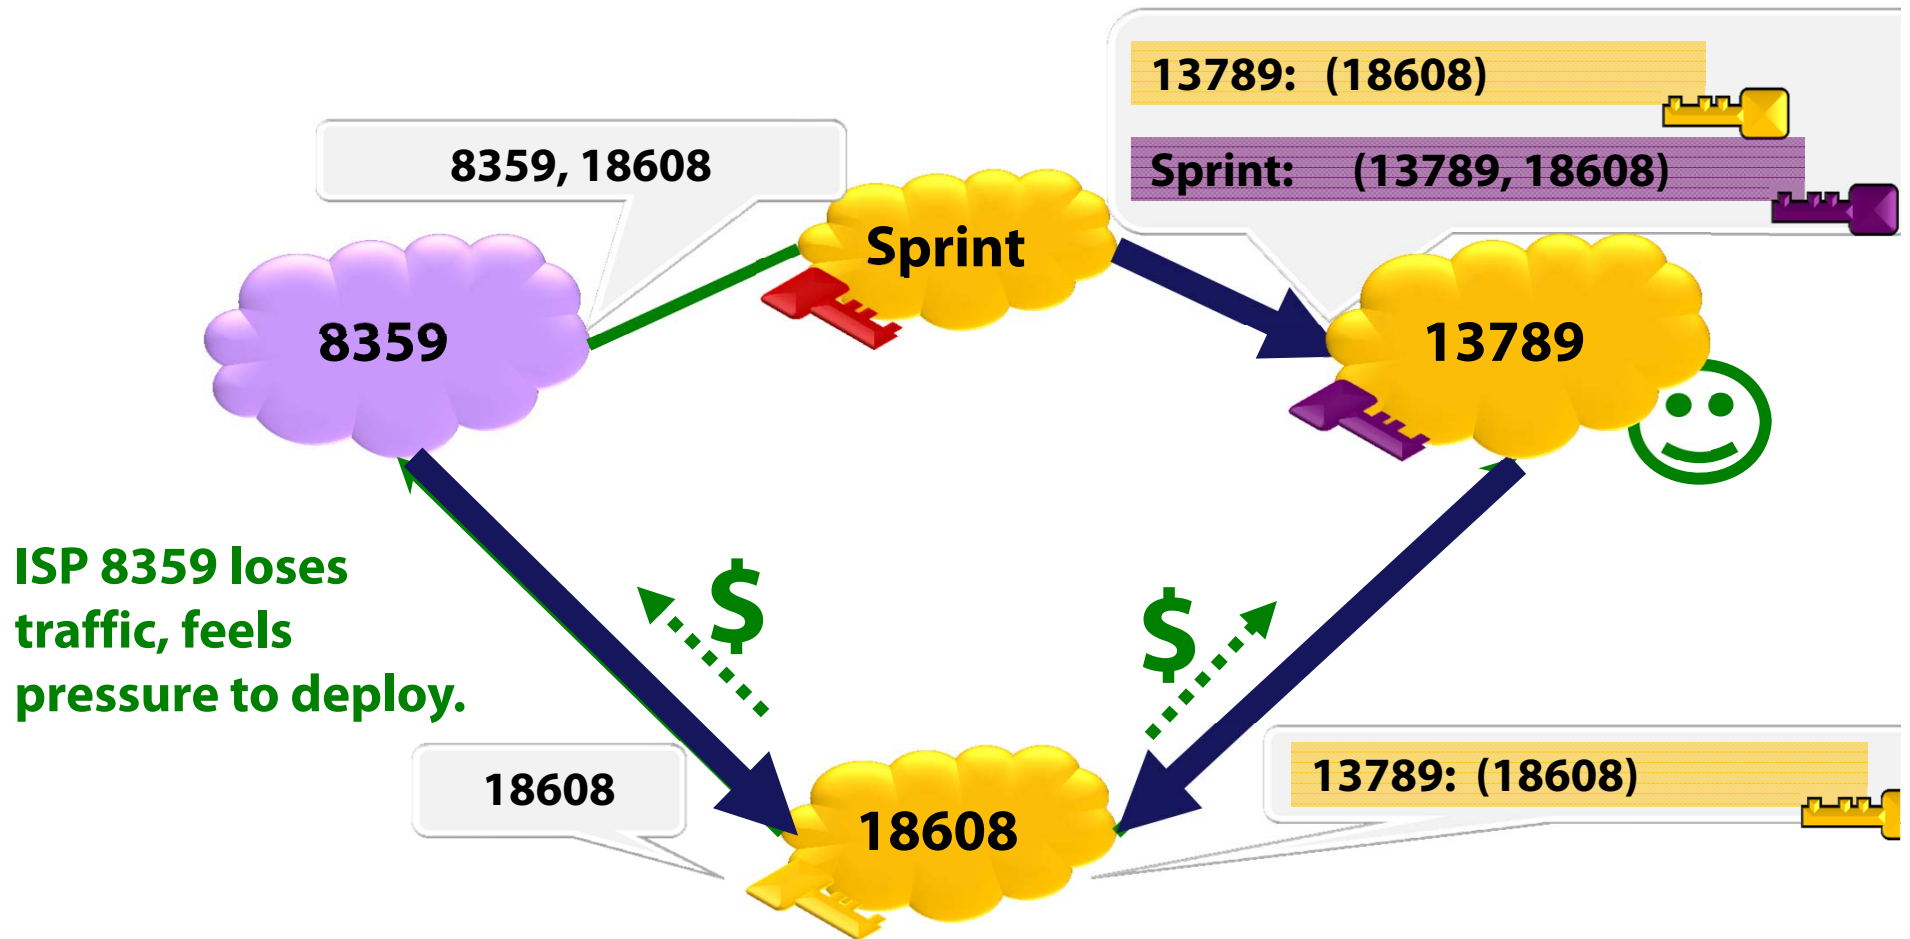
A motivating example. Suppose all are insecure...





What if ISP utility depends on attracting traffic? (2)

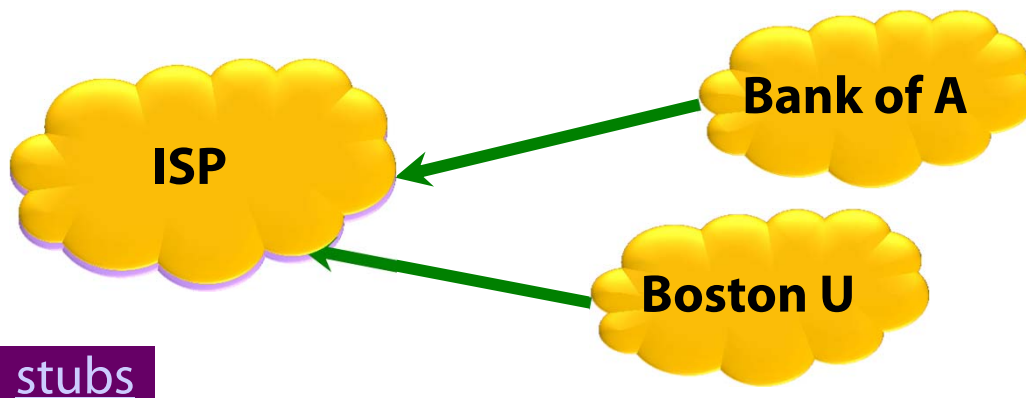
Assume that secure ASes *break ties* on secure paths!





[Gill, Schapira, G. 2011] Guidelines for deploying BGPsec

1. **Secure ISPs *at least* break ties in favor of secure paths**
2. **A few early adopters initially deploy BGPsec (A least 5 biggest ISPs)**
3. **ISPs deploy simplex BGPsec in their stub customers**



stubs

(Gov'ts should subsidize ISPs that do this.)

Stub with Simplex BGPsec:

- Need only sign; trusts provider to validate.
- Minor security impact
- No hardware upgrade!

**Crucial, since
85% of ASes are stubs!**




[Gill, Schapira, G. 2011] model

- **To start the process:**
 - Early adopter ASes become secure
 - Their **stub customers** become secure (e.g. simplex BGPsec)
- **Each round:**
 - Compute **utility** for every insecure ISP (Number of source ASes routing through ISP n to all customer destinations).

BGP Routing Policy Model:

1. Customer > peer > provider paths
2. Prefer shorter paths
3. **If secure, prefer secure paths**
4. Arbitrary tiebreak



- If **ISP n** can increase utility by more than $\theta\%$ by becoming secure
 - ... then **ISP n** secures itself & all its stub customers
- **Stop when no new ISPs decide to become secure.**



[Gill, Schapira, G. 2011]: Results

Thm: An ISP has no incentive to undeploy BGPsec.

Cor: The game terminates.

Thm: Choosing the optimal set of early adopters in NP Hard. (Reduction from set cov)

Use heuristics & simulations instead...

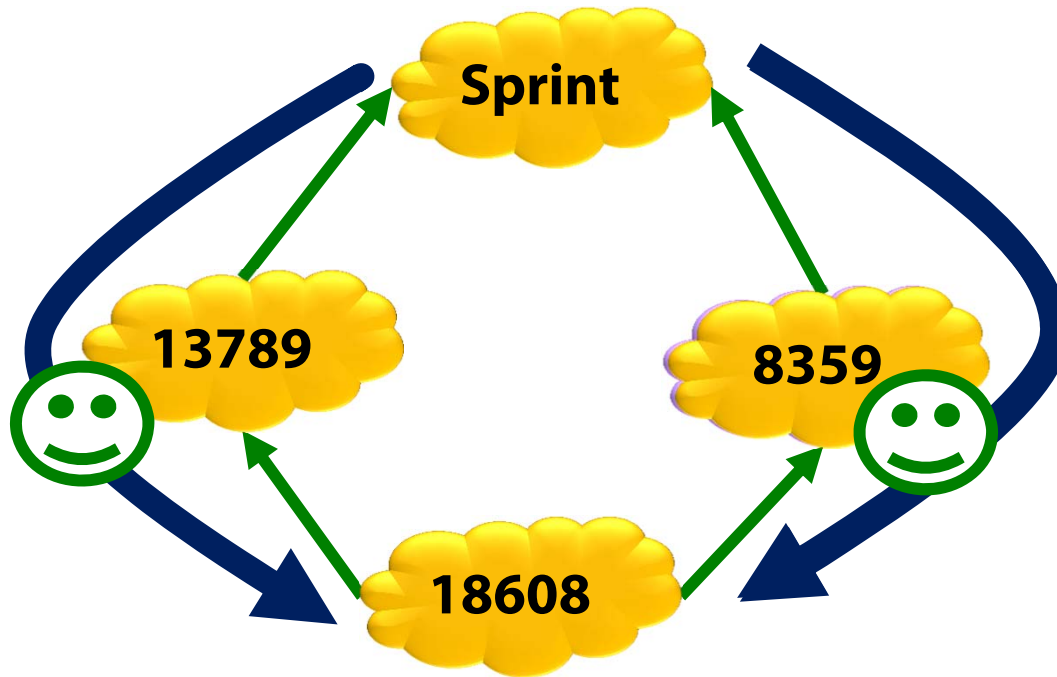
Case study: Ten early adopters, heuristically chosen:

- Five high-degree ISPs: (Sprint, Verizon, AT&T, Level 3, Cogent)
- Five big Content Providers (Google, Microsoft, Facebook, Akamai, Limelight)
- The five content providers source 10% of Internet traffic
- All nodes have the same threshold $\theta = 5\%$.
- **Leads to 85% of nodes to deploy BGPsec!**



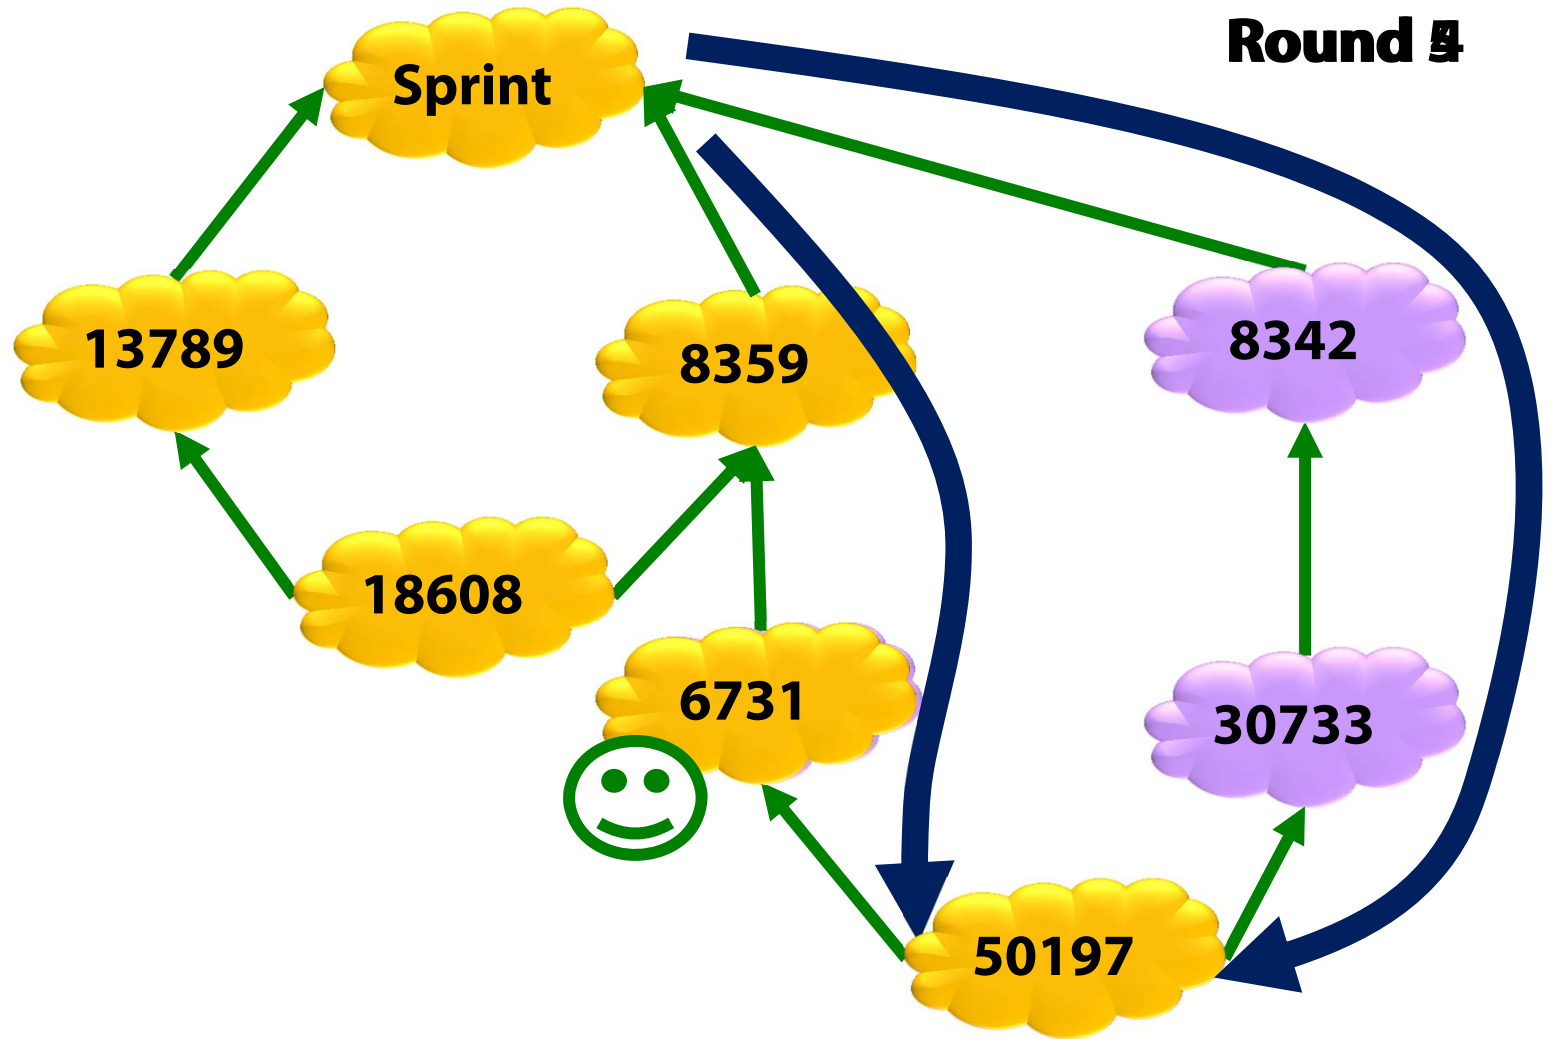
[Gill, Schapira, G. 2011]: Simulations (1)

Round 0



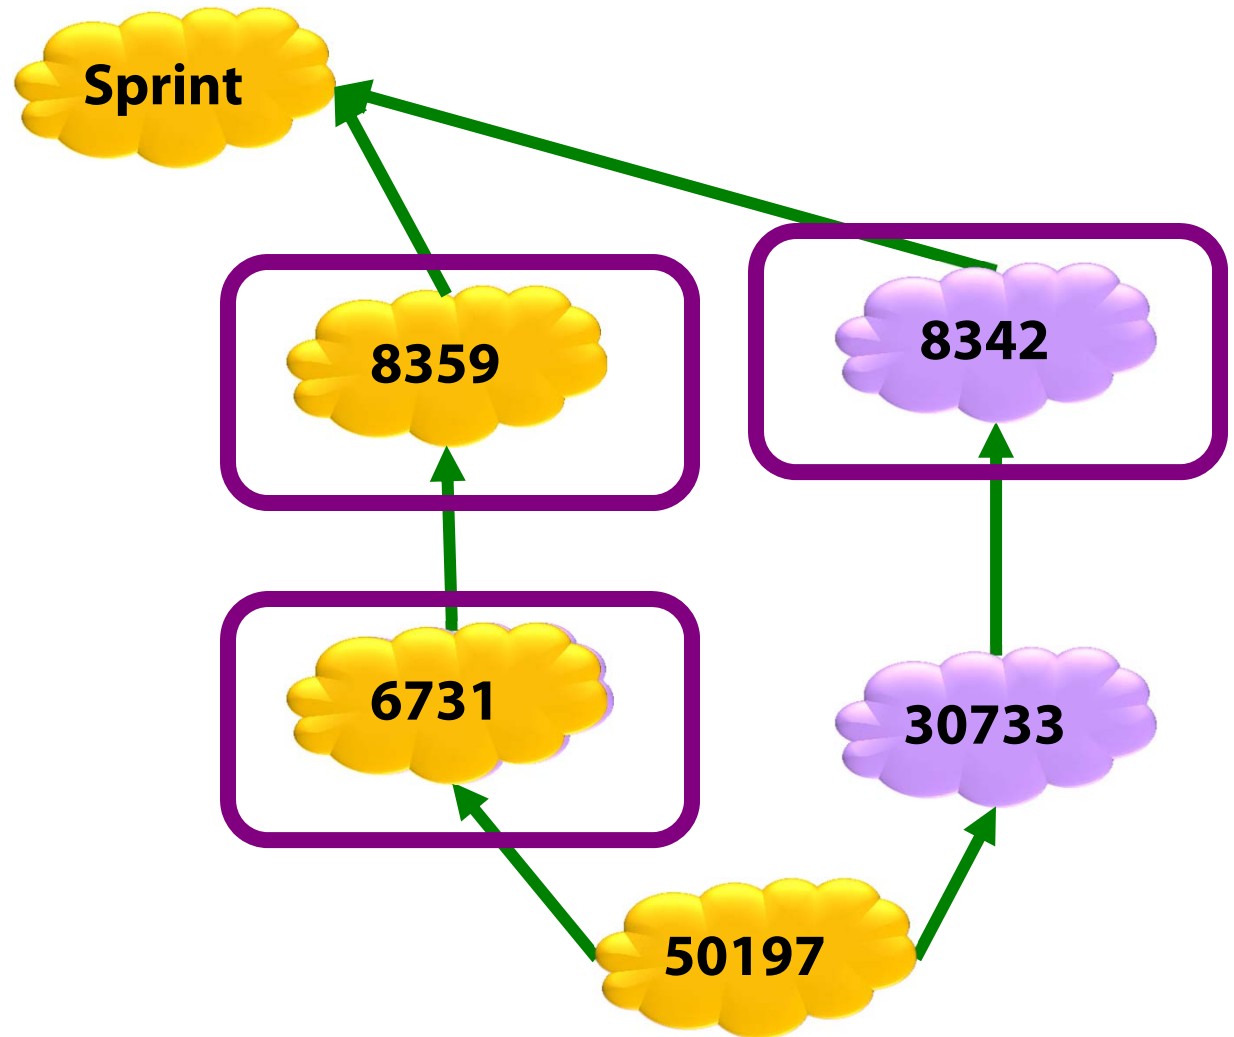


[Gill, Schapira, G. 2011]: Simulations (2)





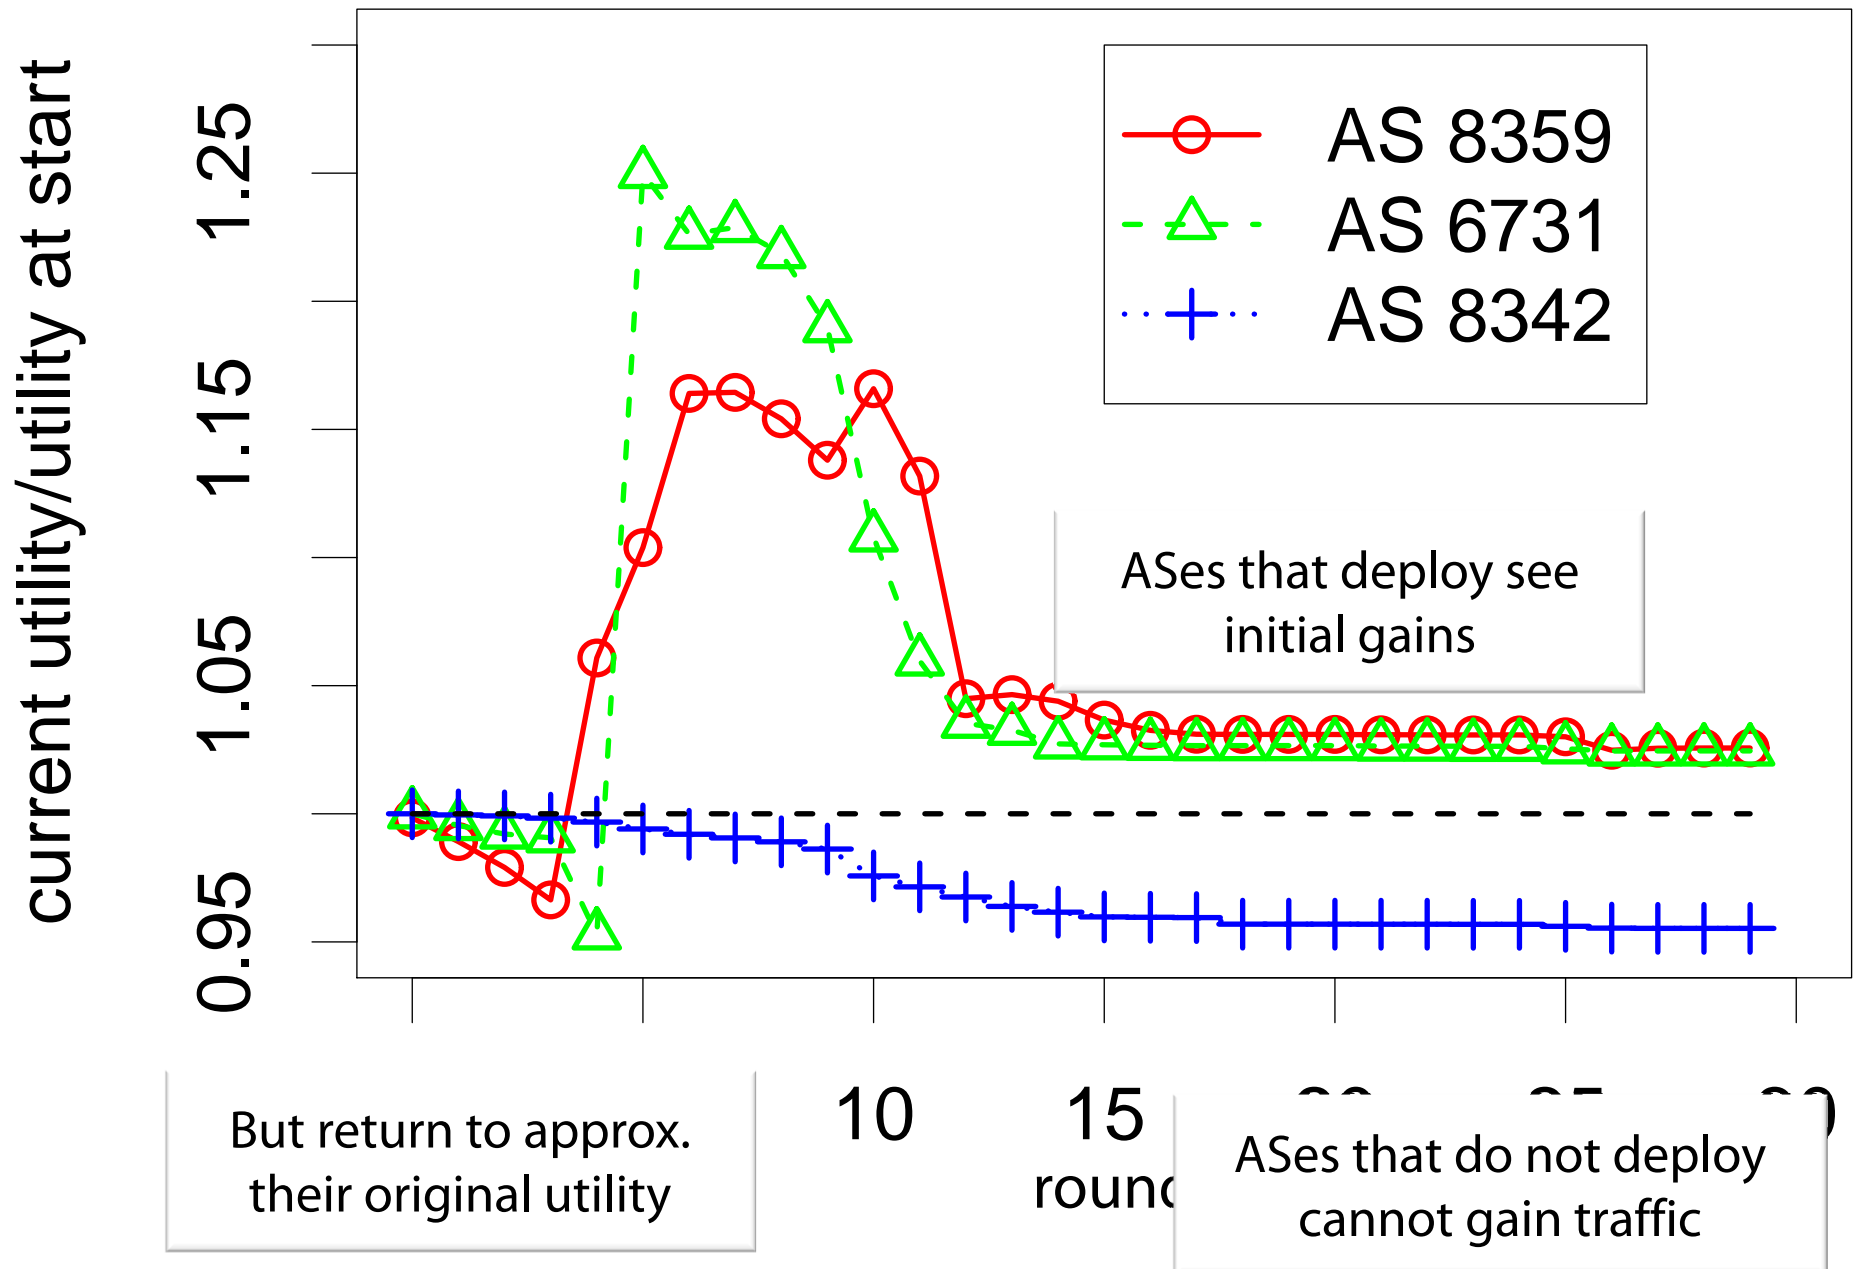
Changes in Utility as Deployment Progresses (1)




Zoom in on utility of each of these three ISPs...



Changes in Utility as Deployment Progresses (2)



Postlude: DNSSEC



A technology is being rolled out today!

- 1) Both hierarchical and distributed network externality



Postlude: DNSSEC – a protocol we're rolling out now

DNSSEC was standardized at around the same time as IPv6

- Basic idea: Take DNS and add cryptographic digital signatures.
- They started thinking about this in mid-1990. Standards 1999, 2005, 2008.

What's DNS?

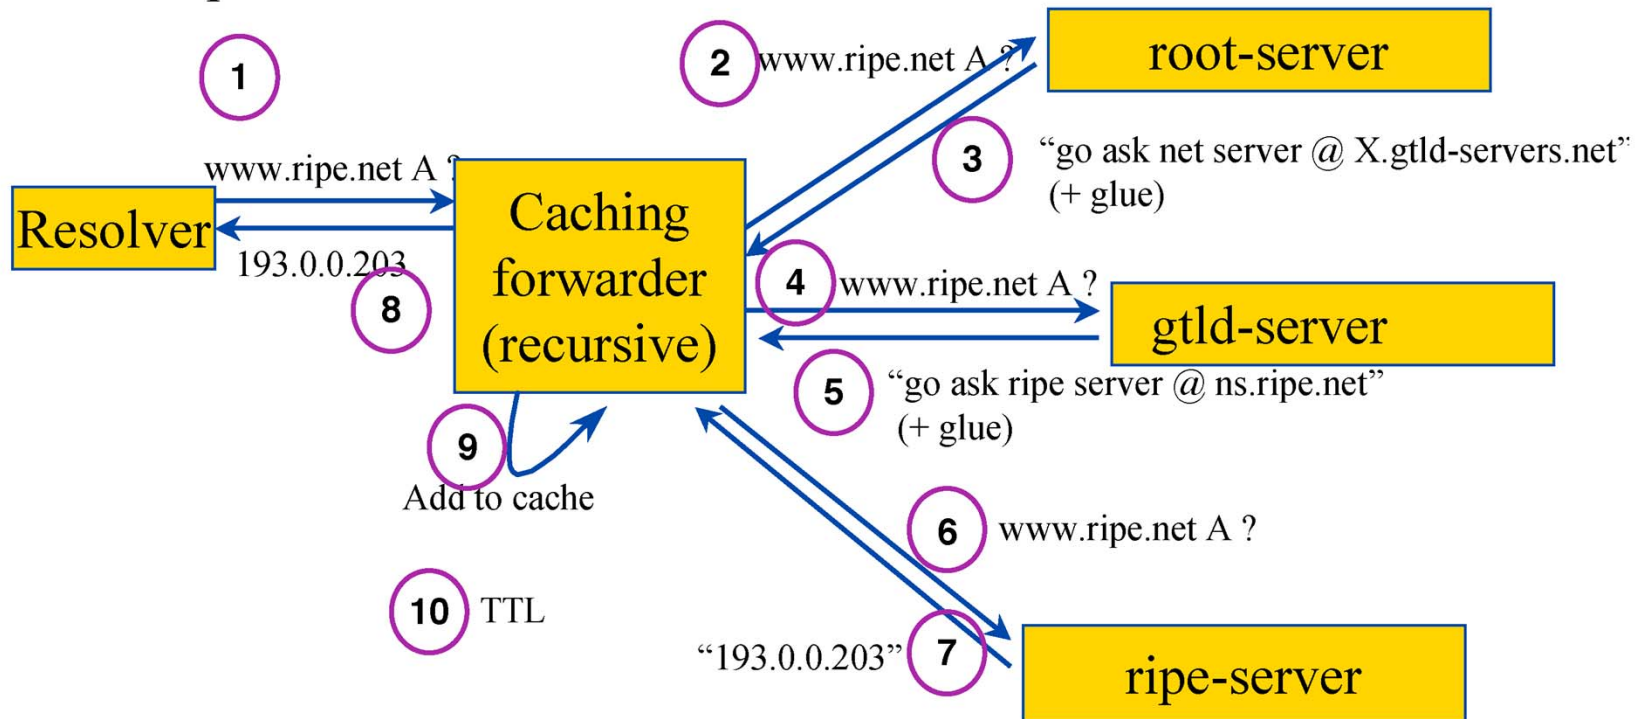
- A distributed and hierarchical database mapping URLs etc to IP addresses.
- How to map www.bu.edu to an IP address?
- Recursively! Ask "root" how to find **.edu**
- Then ask **.edu** how to find **bu.edu**
- Then ask **bu.edu** how to map **www.bu.edu** to an IP address!



Reminder: DNS Resolving

Question:

www.ripe.net A





Traditional DNS





Add a public key to zone





Sign zone with private key





Give hash(pubkey) to parent





Rinse and Repeat



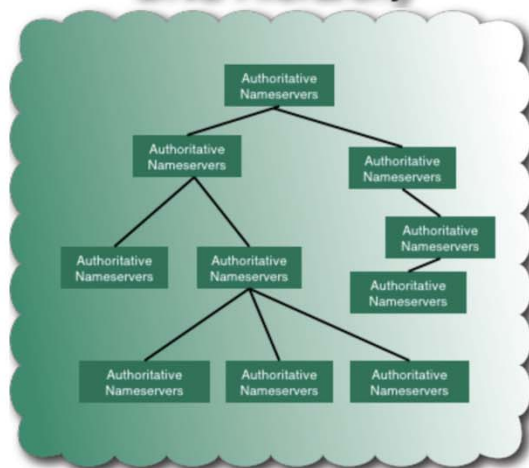
DNSSEC's Chicken & Egg

*Why invest in signing
while the signatures are
not going to be validated*

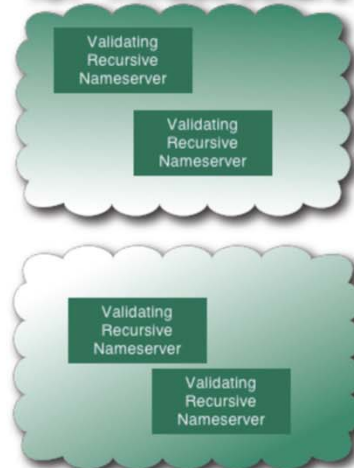
*Why invest in validation
while there is nothing to
validate*

*Why invest in
development while there
is no infrastructure?*

DNS Hierarchy



ISP infrastructure



OS and Application
Support





Postlude: DNSSEC – a protocol we’re rolling out now

DNSSEC was standardized at around the same time as IPv6

- Basic idea: Take DNS and add cryptographic digital signatures.
- They started thinking about this in mid-1990. Standards 1999, 2005, 2008.

Challenges of DNSSEC deployment:

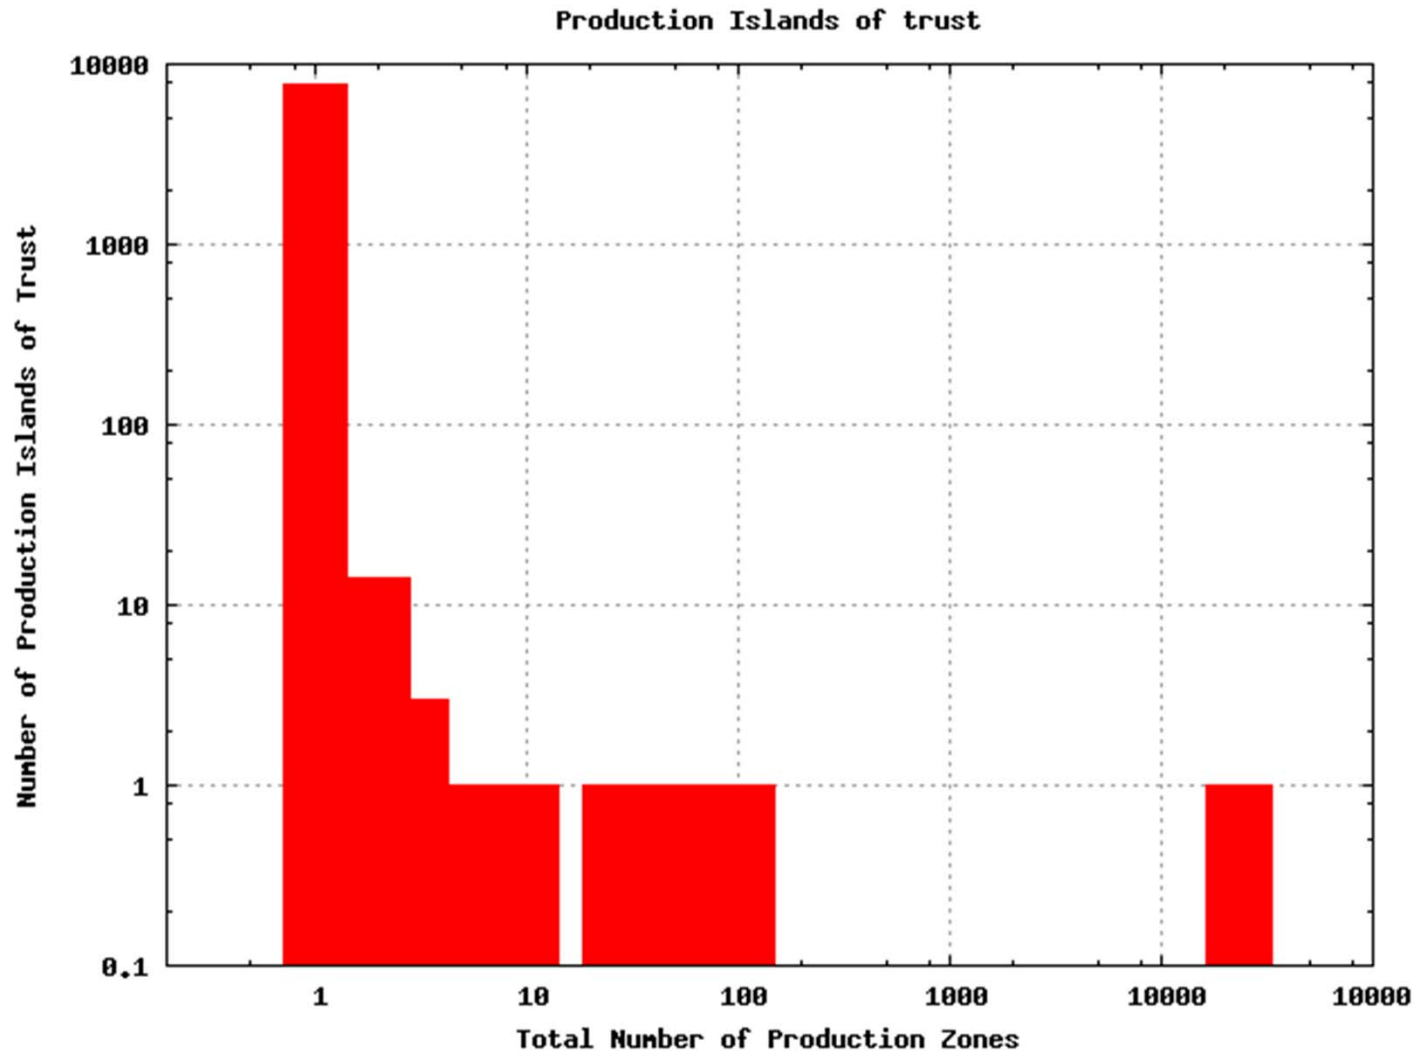
- Hierarchical: A nameserver can’t deploy until its parent nameserver does.
- Distributed: Economically-motivated agents need incentives for deployment.
 - Agents are **nameservers** and **resolvers**
- Political: Who owns the keys to the root zone?

Where we are today?

- First “zone” to deploy in 2005 is Sweden, **.se**
- Root zone signed in 2010 (with key shared by 14 Internet “personalities”!)



Deployment as of June 5, 2012 per UCLA's SecSpider




Note: there are millions of DNS servers in the world. So we aren't there yet.



A few open questions

- Models to understand IPv6 deployment & IPv4 address allocation
 - Descriptive models to understand the market failure
 - Conversion/incompatibility interaction with network structure (graphs)?
 - Two sided markets that create incentives for IPv6 adoption?
 - Auctions for IPv4 address space?
 - Effects of IPv4 scarcity on adoption?
- Non-local network externalities?
 - More realistic models of the network externality that map accurately to BGPsec and IPv4 deployment
 - Characterizing equilibria with certain graphs or threshold models
- Descriptive models of DNSSEC & RPKI deployment
 - Way to leverage or avoid dependence on hierarchy? Interactions with resolvers?

Thanks!



<http://arxiv.org/abs/1202.2928>