INTERNET EXPANSION, REFINEMENT AND CHURN

Vocal and lyrics:

Andre Broido, Evi Nemeth, kc claffy & elves
(broido, evi, kc) @ caida.org

CAIDA/SDSC/UCSD

NANOG, Miami 12 feb 02

(Eur.Trans.Telecomm., Feb.2002)

www.caida.org/~broido/nanog200202.egr.pdf
www.caida.org/~broido/nanog200202.egr.ps
www.caida.org/presentations/nanog0202/

Acknowedgements

David Meyer

Geoff Huston

Sean McCreary

Bill Woodcock

Cengiz Alaettinoglu

Sean Doran

CAIDA elves: Marina Fomenkov, Brad Huffaker, David Moore, Colleen Shannon

Part I

Motivation

Routing Concerns

- Network stability
- AS number exhaustion
- IP addreses depletion
- IP space fragmentation

Router constraints

Processing power limitations

router in a fridge?

Memory limitations

router in a wristwatch?

Table growth strains both

Affects: equipment costs

Timely route computation

Network stability

Router stability

BGP updates per 15 min

Black dots: route announcements
Red: withdrawals, 22 routers

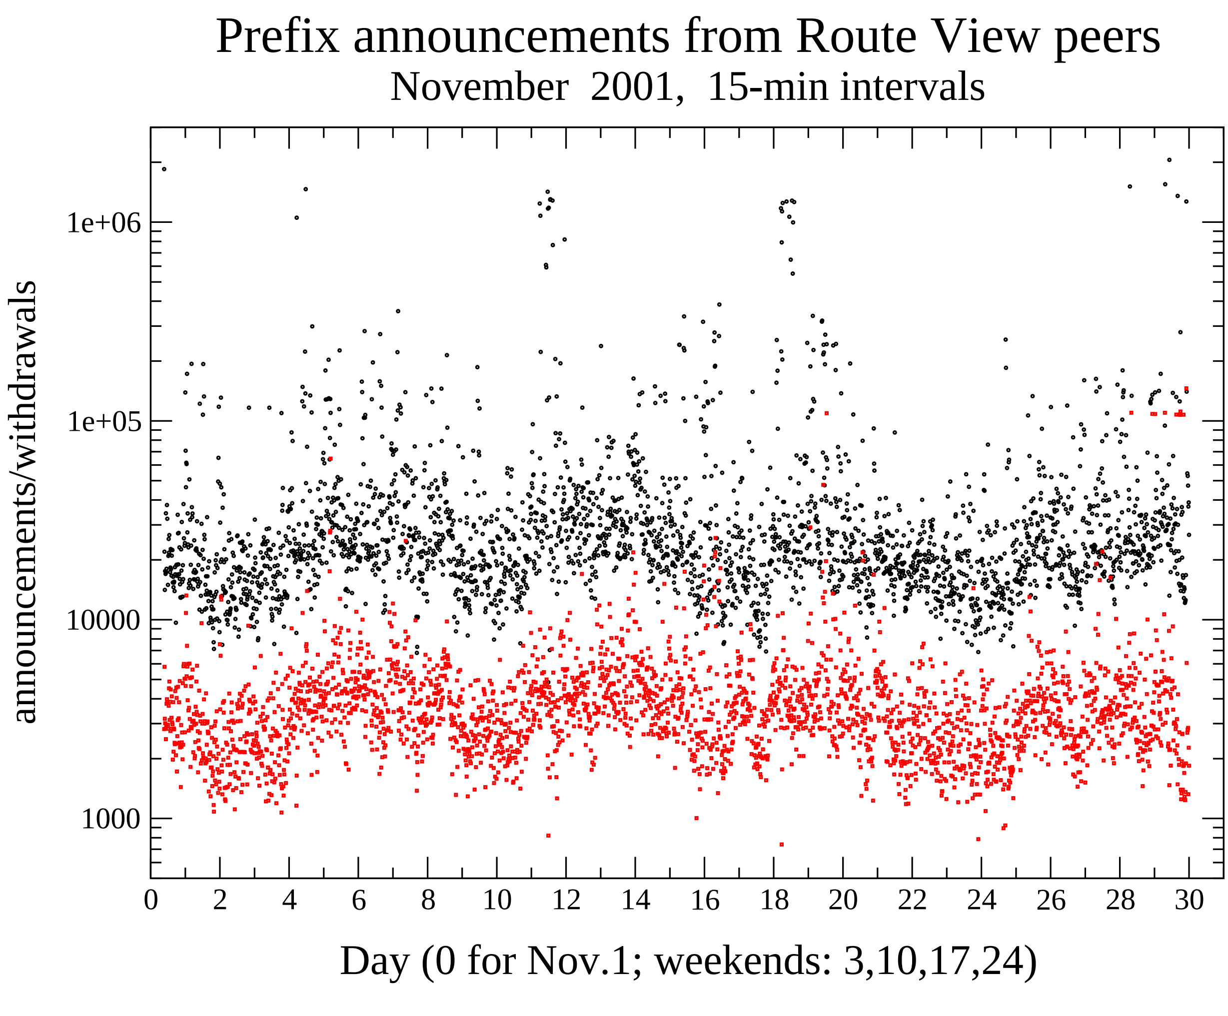

Figure 1:

Two operating modes

Diurnal/weekly background noise

Occasional full table uploads

50K..2M prefixes announced in 15 min

Meltdown scenario

Neighbours' session resets

Many full tables arrive at once

Surge in processing load

Router crash

More session resets

"Domino effect"

The core "burns out"

The likelihood grows with table size

Part II

Facets of change

Spectrum of change

Expansion

Refinement

Sophistication

Churn

$+$ *stability*

Trends & patterns

- Network invariants
  - Several ratios change very slowly
  - Proportionate growth in many respects
- BGP table size
  - Prefix & AS growth slowed in 2002
  - Small ISPs do not contribute growth/churn
  - Route updates per prefix decrease (Cengiz, NANOG 23)
  - Internet demand changes with economy
- IP addreses

  Uneven, mostly slow growth since 1997

  Fast growth of more specific space

  Reasonable IP space fragmentation

Data Analysis

Timeframe: 1997-2002

RouteViews BGP tables

Sampling: every 6 months

Novembers, Mays

Full-size tables only

Prefixes carried by most peers, i.e. semiglobal prefixes

Unstable 60/8 to 69/8 prefixes removed

Table Size

#routes times per route state:

#ASes (path length + prepending)

communities, other attributes

Table can grow even when #routes is stable

NANOG mailing list, Aug 2001:

150-180 bytes/route

Estimate: 16-20 Mbytes per RIB

Part III

Counting routes

Factor1: # routes

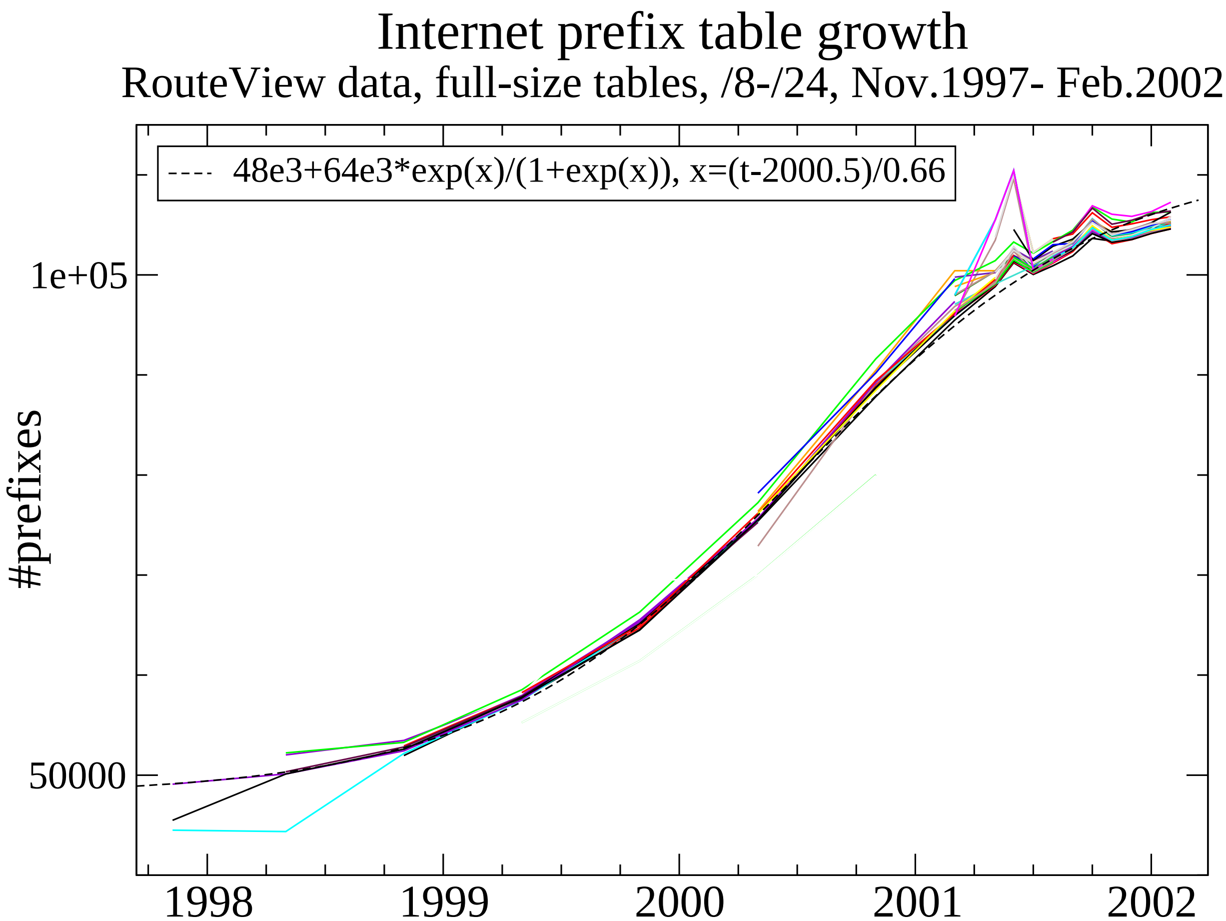

Figure 2: Routes count in XX-XXI c.

Number of routes

Grows much slower in this century

Close to logistic curve:

$$\exp(x)/(1+\exp(x)) + C$$

Generic shape for tech transfer

steam engines to locomotives etc.

growth = product of haves to have-nots

"contagious gadgets"

Most likely cause of slowdown:

high prefix dropout rate

See churn numbers in IETF slides

Summands of prefix count

Standalone prefixes

More specific (subset) prefixes

Root prefixes (aggregates)

Top prefixes = standalones + roots

More specifics and /24s, multihomed ASes and more specifics are independent

Having one property does not increase the probability of another

Summands of prefix count

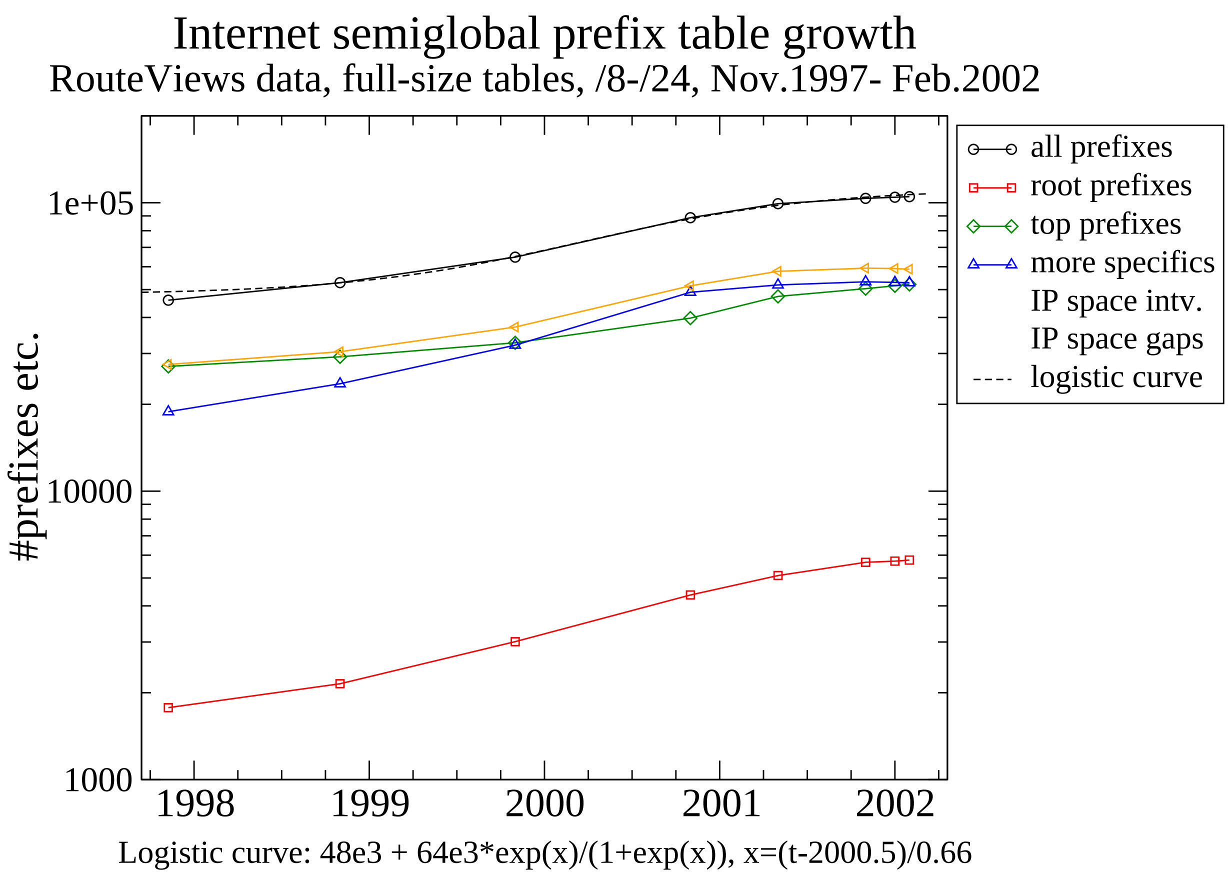

Figure 3:

#more specifics = #top prefixes in 2002

Prefix growth slowdown

Caused by economy?

"Bubble" burst in November 2000

Many 'Net ventures failed in 2001

Compounded by Sept.11 events

Is it a coincidence?

Prefix growth slowdown (cont'd)

Business customers connect in thousands

Average 6,000 new prefixes per month

It is the drop rate that maxed

Oct.2001 count = Feb.2002

More specifics stable since Sept.2001
even decrease slightly

Smallest global blocks (/24s)
stable/down since Aug.2001

Churn: drop rate = 1/2 growth in 1990's

Now they are equal

Part III

Topological state growth

ASes

#AS paths

AS path lengths

Average, variance

Prepending

AS edges

Measuring growth by ratios

AS paths / ASes ratio in the table
related to topological diversity (redundancy)

AS path length average
=#non-prepended AS tokens / routes
related to memory consumption

Variance of AS path length

BGP path selection is largely AS path length-based

Numbers of AS and AS paths

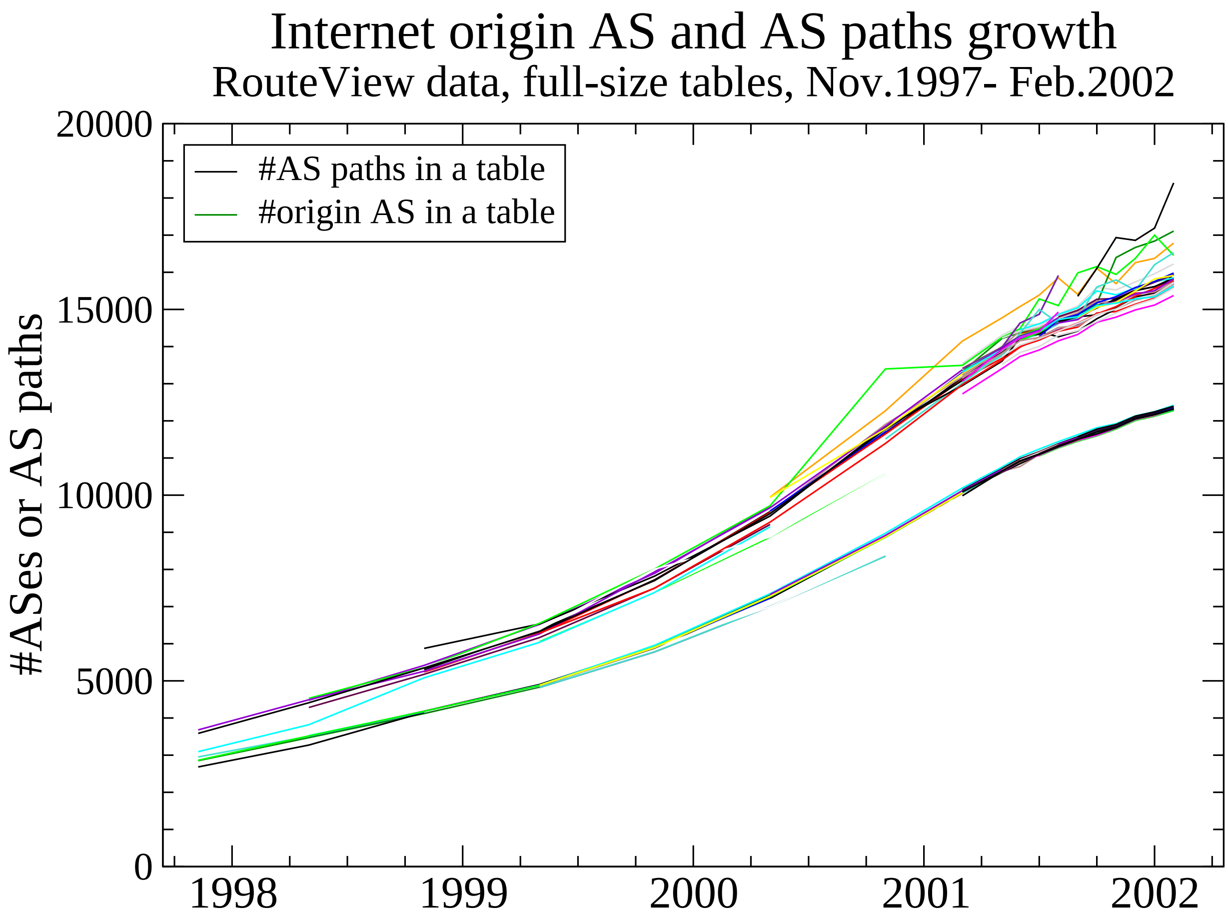

Figure 4: AS paths (top curve) and origin AS counts

Slowdown

ASes still on the rise

Not as fast as in 1990s

Would last for 10 years

If linear growth persists

AS path counts / ASes

- Appear to decrease since 2000
- Close to 1.25-1.3, almost invariant

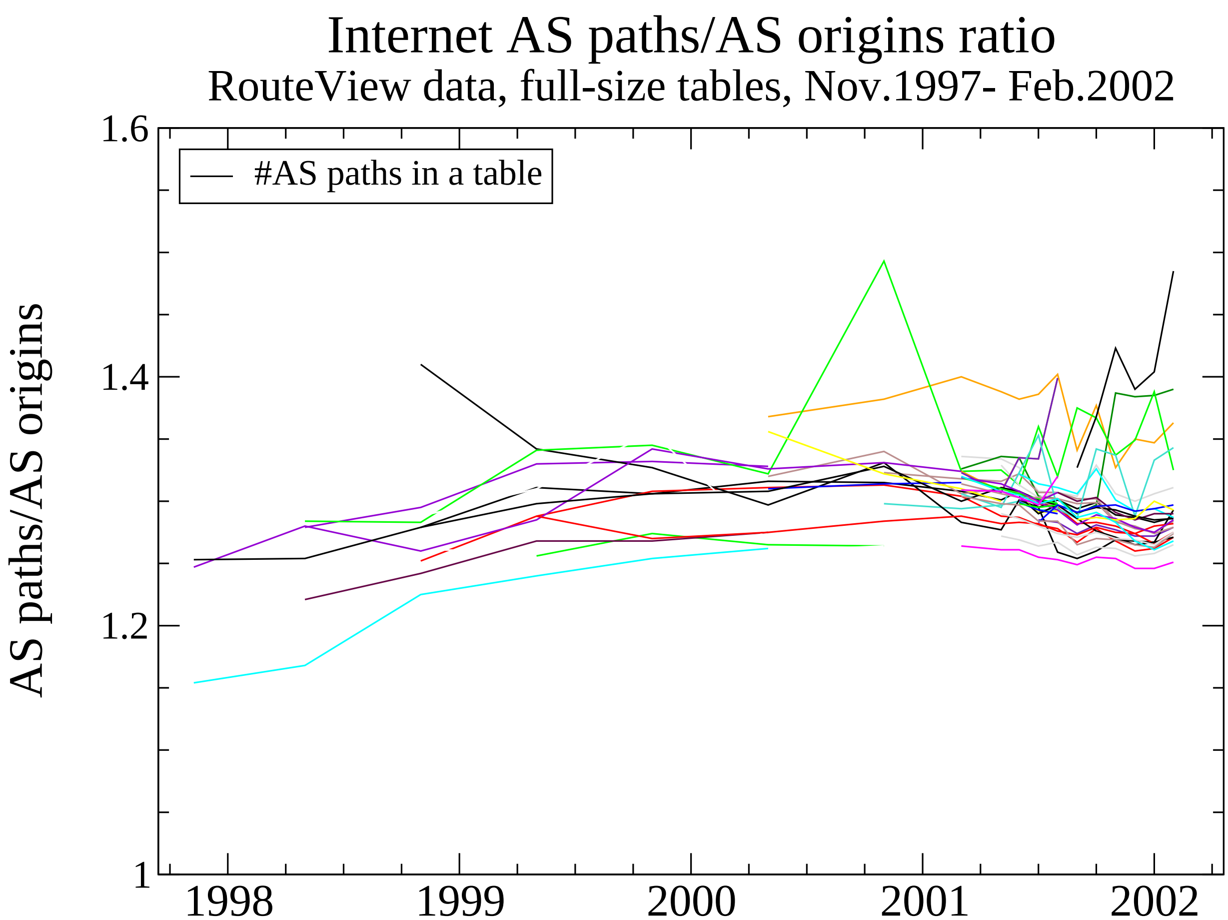

Figure 5: AS Paths/origins per peer

Stability of ratios (1)

Best path selection results in low path redundancy

Connectivity policies contribute moderately

May delay reaction to topology changes

Average number of AS neighbours grows slowly, now about 2

Stability of ratios (2)

AS path length can change when upstream provider is changed; usually doesn't. Many AS have stable path length distributions.

AS path length diversity: standard deviation, path length entropy (binary choice) near constant, appear to decrease slowly

Prepending is reasonably low (5.5% AS tokens and lines)

Routing table size: Factor2

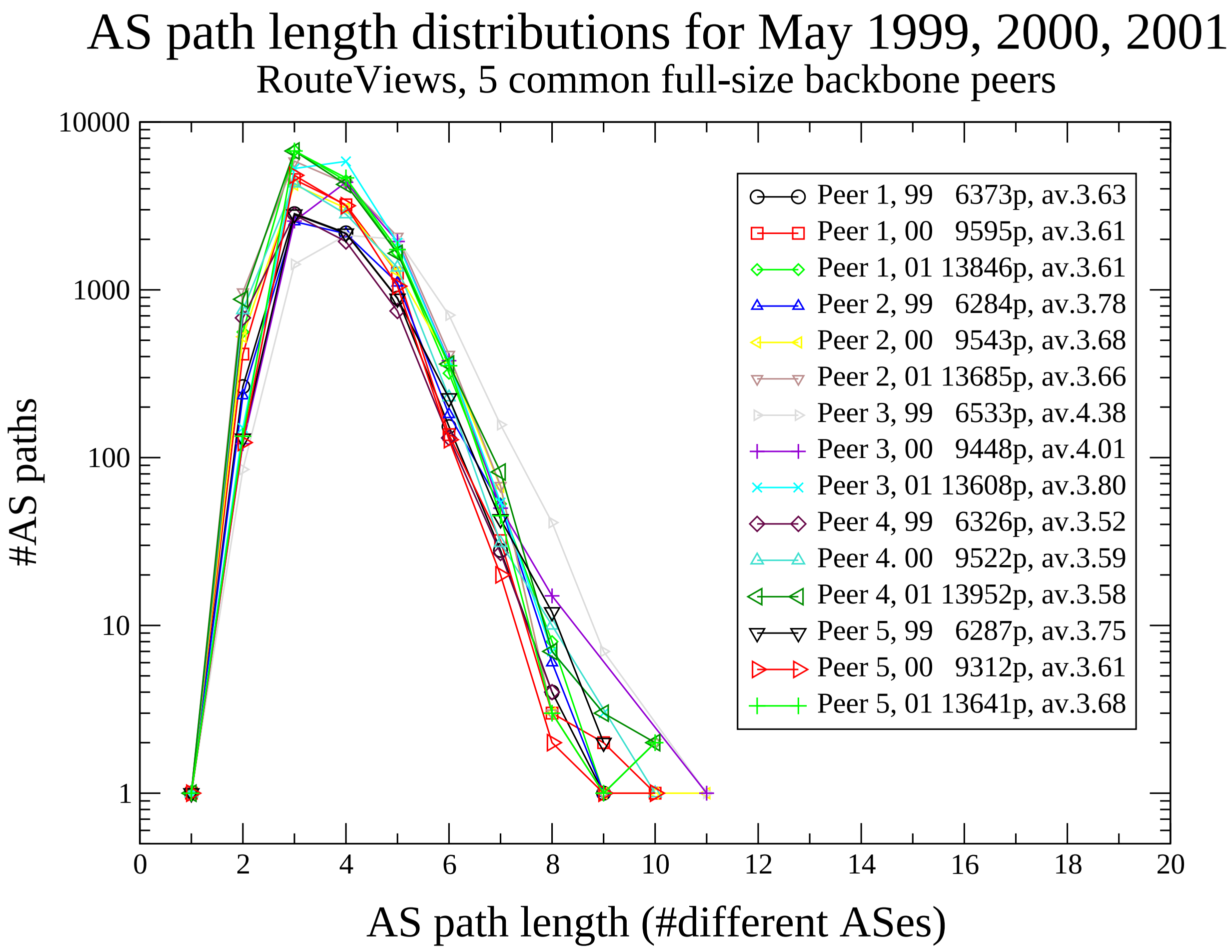

Figure 6: AS path length distributions (RouteViews 1999-2001)

Proportion of AS paths with given length is constant; the number of paths grows

Average AS path length for 40+ peers

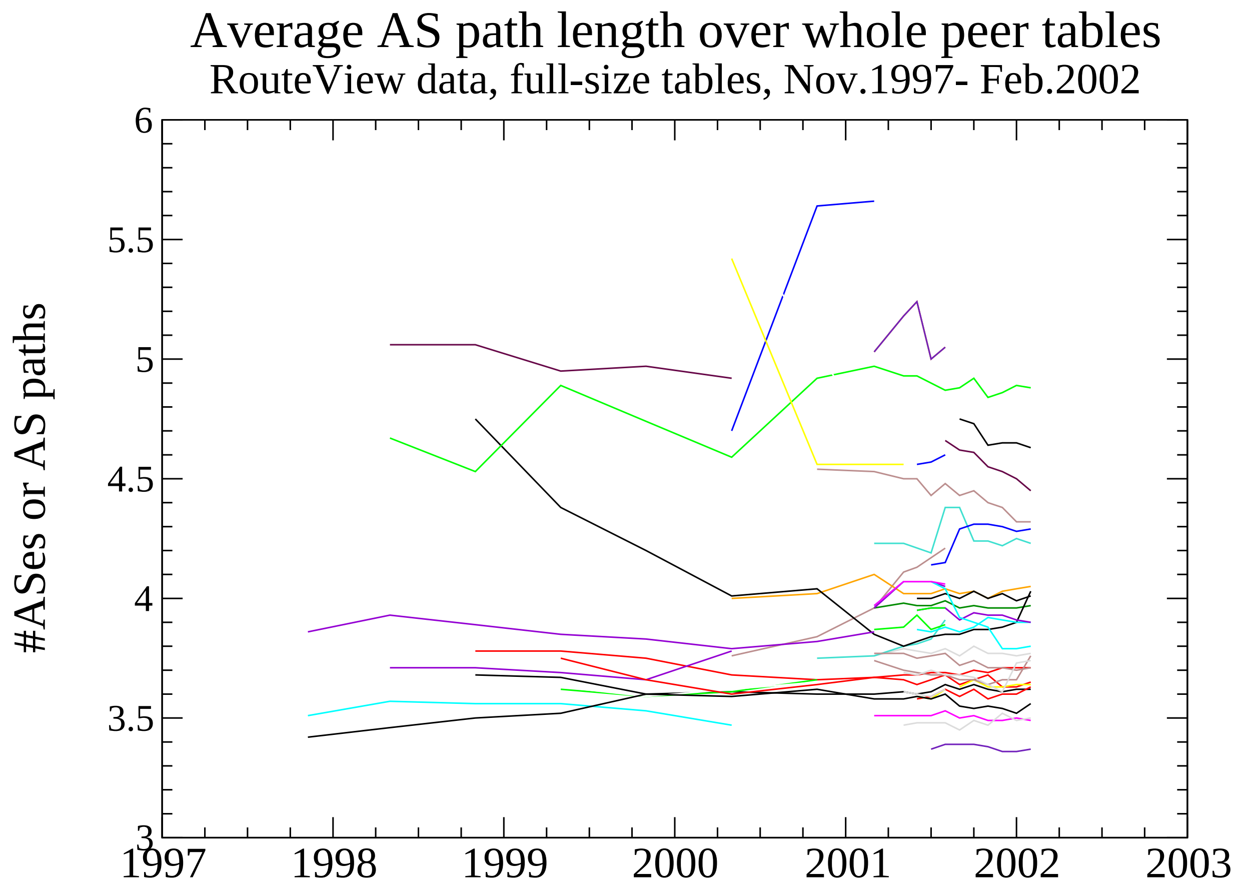

Figure 7: Average AS path length per peer

No visible trend, close to contant

Standard deviation of AS path length

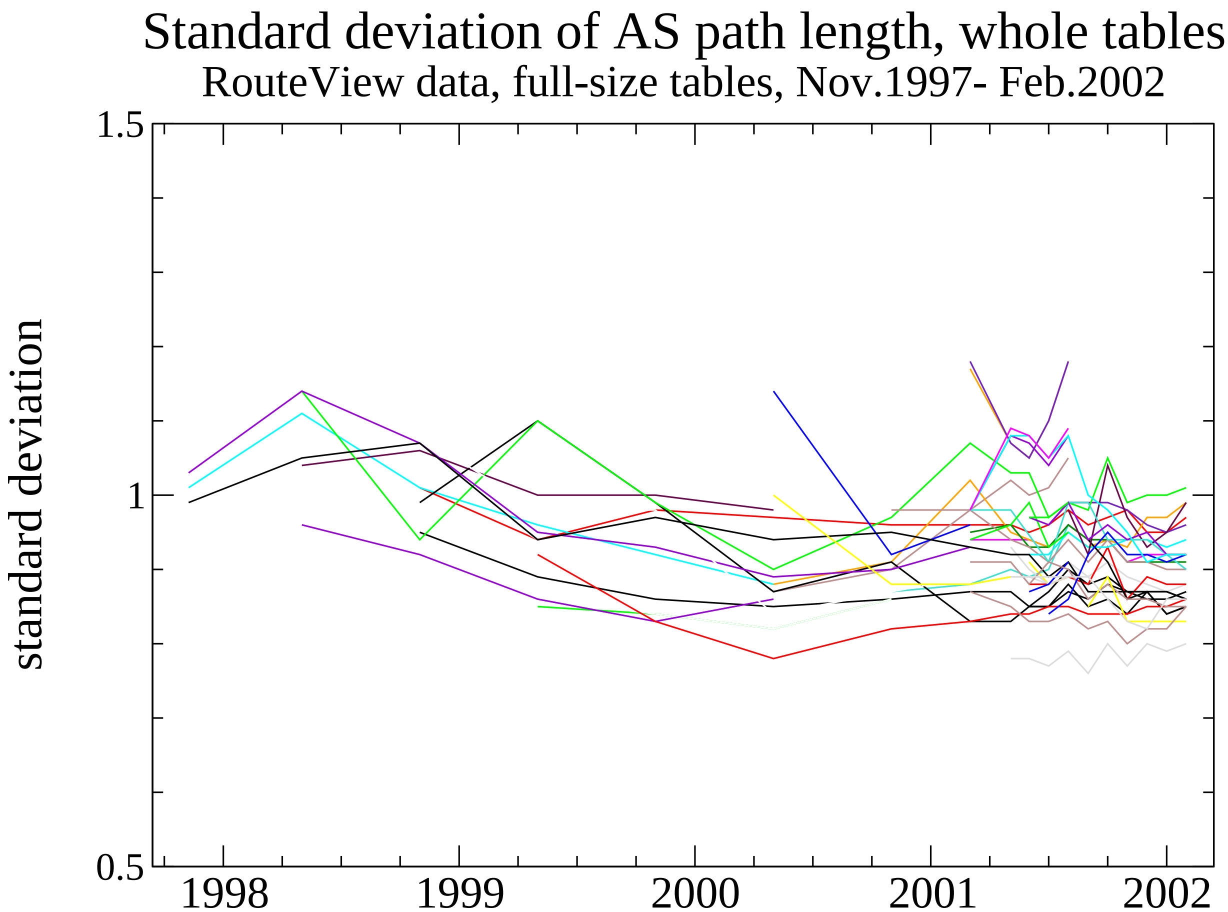

Figure 8: Standard deviation of AS path length

Appears to decrease in 1997-2000
no visible trend in 2000-2002

AS number exhaustion

Largest AS number is close to 25,000

Half of those are routed

5 new ASes per day

64,512 maximum

Will suffice for a few years (4-10)

Eventual implementation of 32-bit AS numbers

Prefix/AS refinement

Prefix growth rate was 2/3 of AS rate

Currently less than 1/2

Prefixes/ASes ratio steadily decreasing

Currently about 8

Will eventually have less than 1 prefix/AS
About 80 ASes are already transit-only

Benefits of granularity reduction from prefixes to origin ASes and similar concepts gradually disappear

Prefix/AS refinement (cont'd)

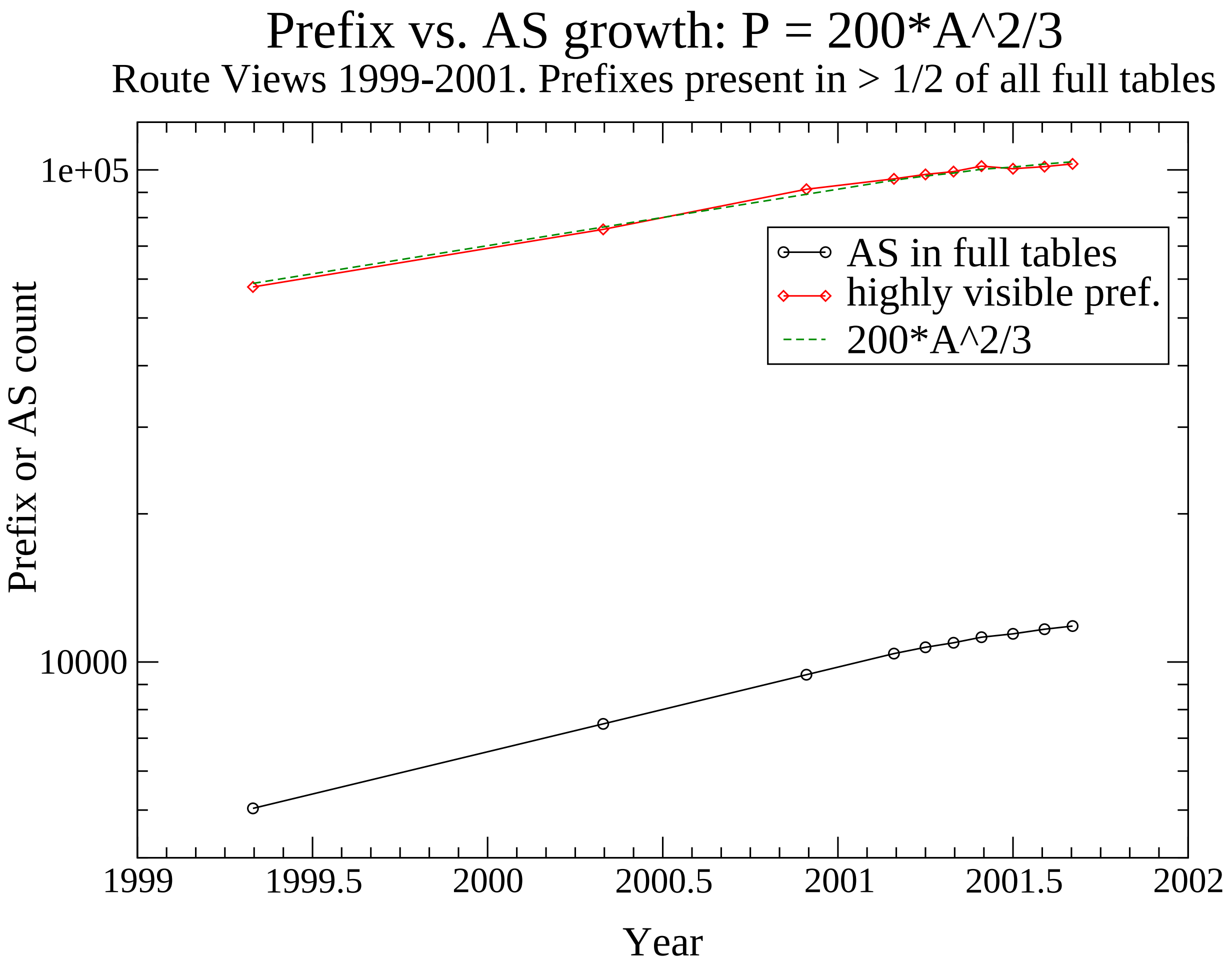

Figure 9: AS refinement)

Route origination

40% ASes originate 1 prefix; contribute 5% prefixes

1% ASes originate 100 or more prefixes; contribute 32% prefixes

Same stats in Nov.2001 and Feb.2002

Small ISPs do not contribute to table growth

This disparity changes slowly

Relatively fewer ASes are origins of many prefixes

Reflects AS refinement

Route origination

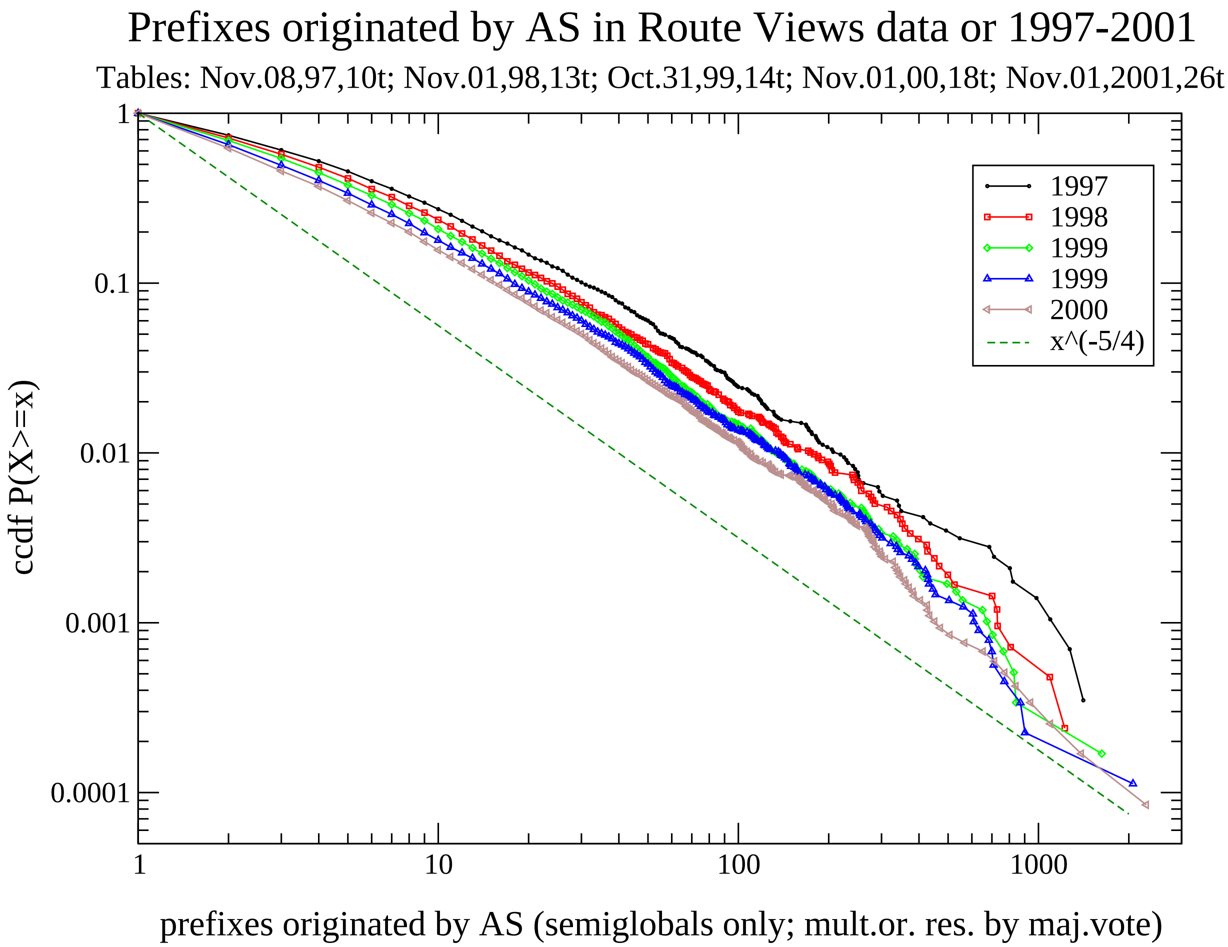

Figure 10: BGP AS origin distributions (RouteViews 1997-2001)

Part IV: Trends in IPv4 address space use

- Consumption
- Refinement
- Sophistication
- Fragmentation

Importance for prefix growth

New top prefixes, Dec.2000-May 2001:

50% cover new address blocks

37% deaggregations of existing prefixes

+aggregation & expansion

IP addresses' consumption

Slowed in mid-90s

Currently 27% of potential addresses routed

Close to 1/2 are allocated/assigned

2/5 of allocated/assigned addresses are not globally routed

Yearly growth in 1998-2001:
10%, 7%, 1.2%, 4.6%

Assume yearly growth of 7%
40% overhead of local addresses
Current semiglobal coverage
IPv4 space can potentially last 10 years

IP space evolution

Addresses in standalone blocks decrease

Addresses in roots blocks grow

Addresses in more specifics grow fastest

Sophistication: nested prefixes replace standalones

IP address space growth

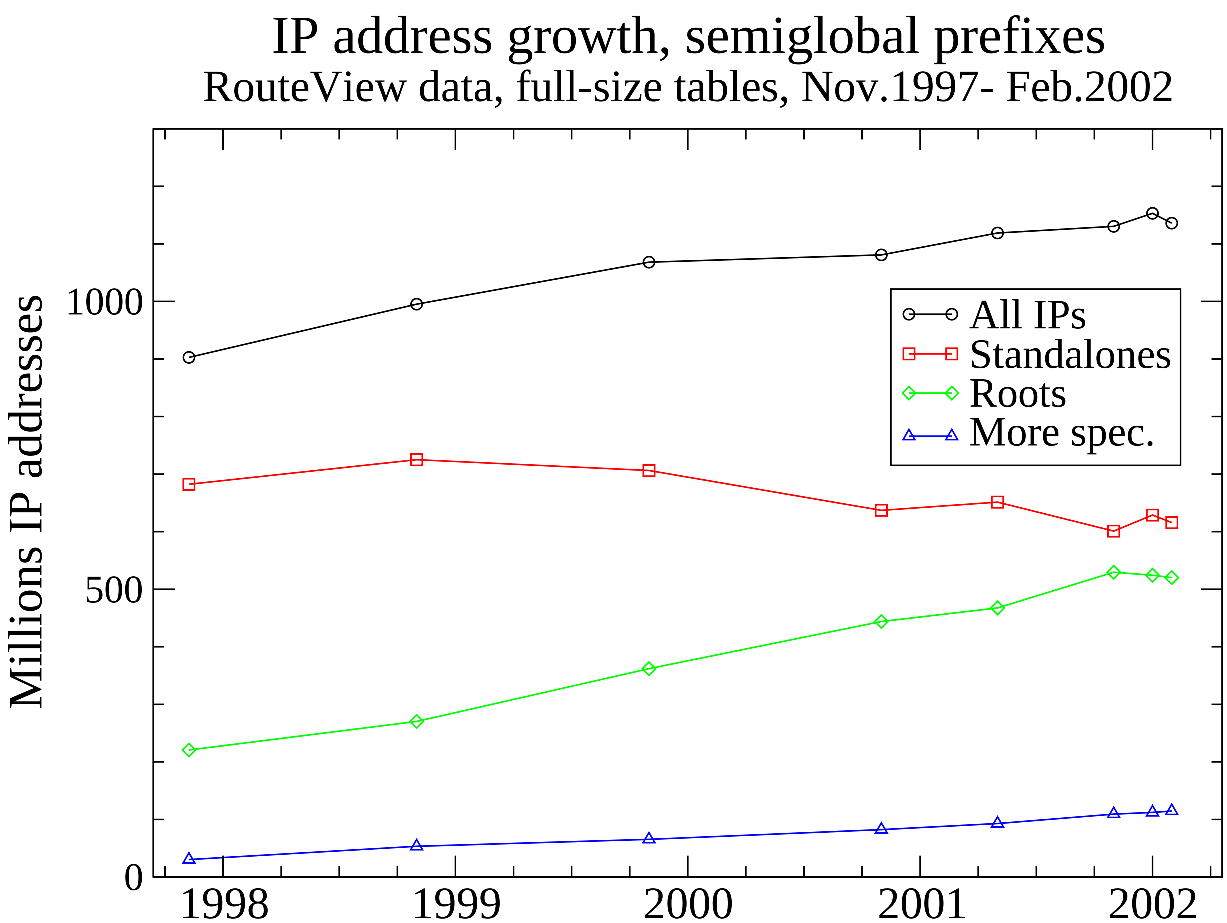

Figure 11:

Churn

Nov.2000-Nov.2001:

120M new IP addresses

70M dropped

Net growth 50M addresses

Drop rate higher than net growth rate

IP address space
refinement, sophistication, fragmentation

Address space refinement

End of 1990s:

Prefixes grew much faster than addresses

Average #addresses per prefix steadily decreased

2001-02:

Two growths comparable, no significant refinement

Sophistication

routing-induced subdivisions
Intervals in IP space

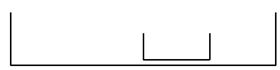

A "hole" is punched in a larger prefix

Customer changes provider, keeps prefix

Provider has a reason to deaggregate

Should subdivide pieces to CIDR blocks

Prefix tower intervals in IP space

Contiguous sets of IP addresses

Sharing common set of covering prefixes

"Prefix tower"

An IP address can be in up to 32 prefixes

in reality, maximum is 6

Intervals = lower bound for #prefixes arising in total deaggregation which makes all more specifics standalones

Generated by shifted and skipped more specifics

Prefix tower intervals (2)

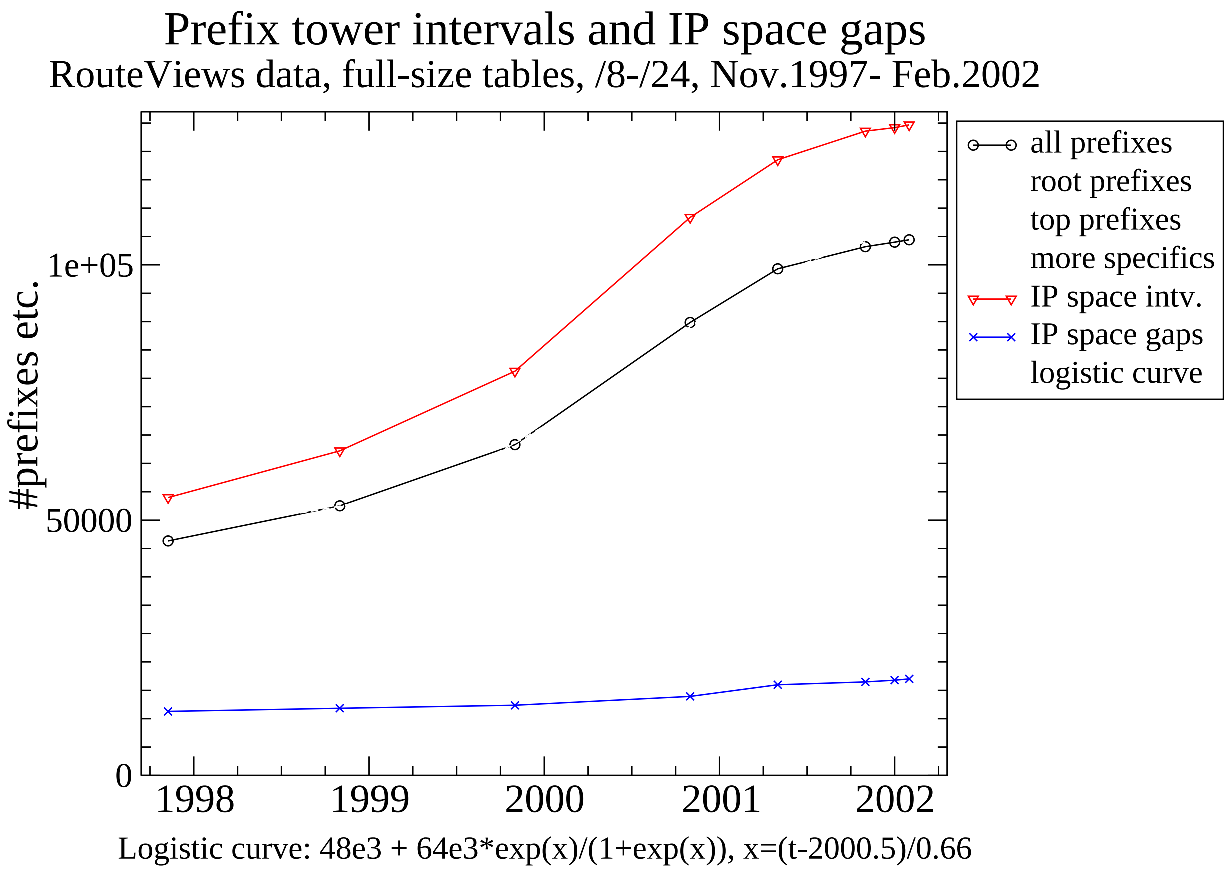

Figure 12:

Growth of intervals and prefixes, linear scale

Prefix tower intervals (3)

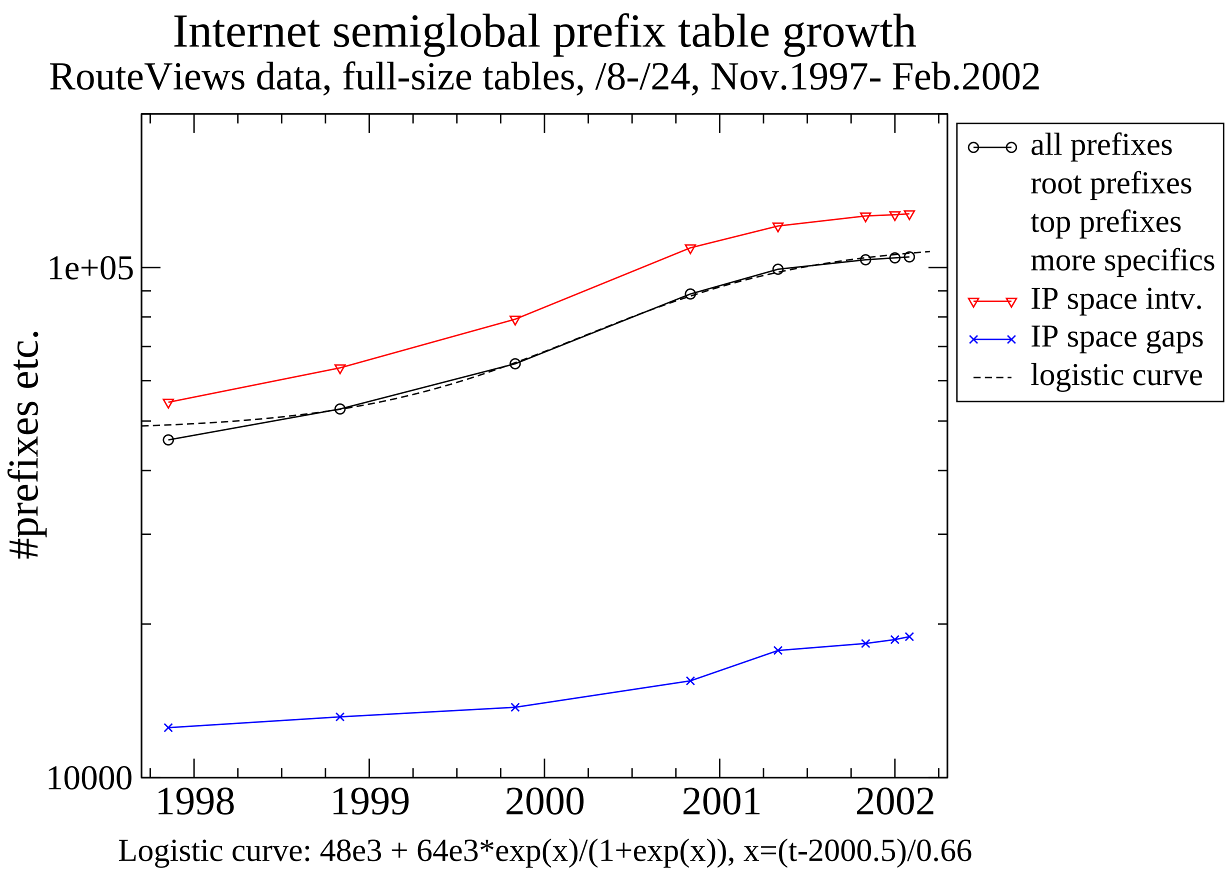

Figure 13:

Growth of intervals and prefixes, log scale

Prefix tower intervals (4)

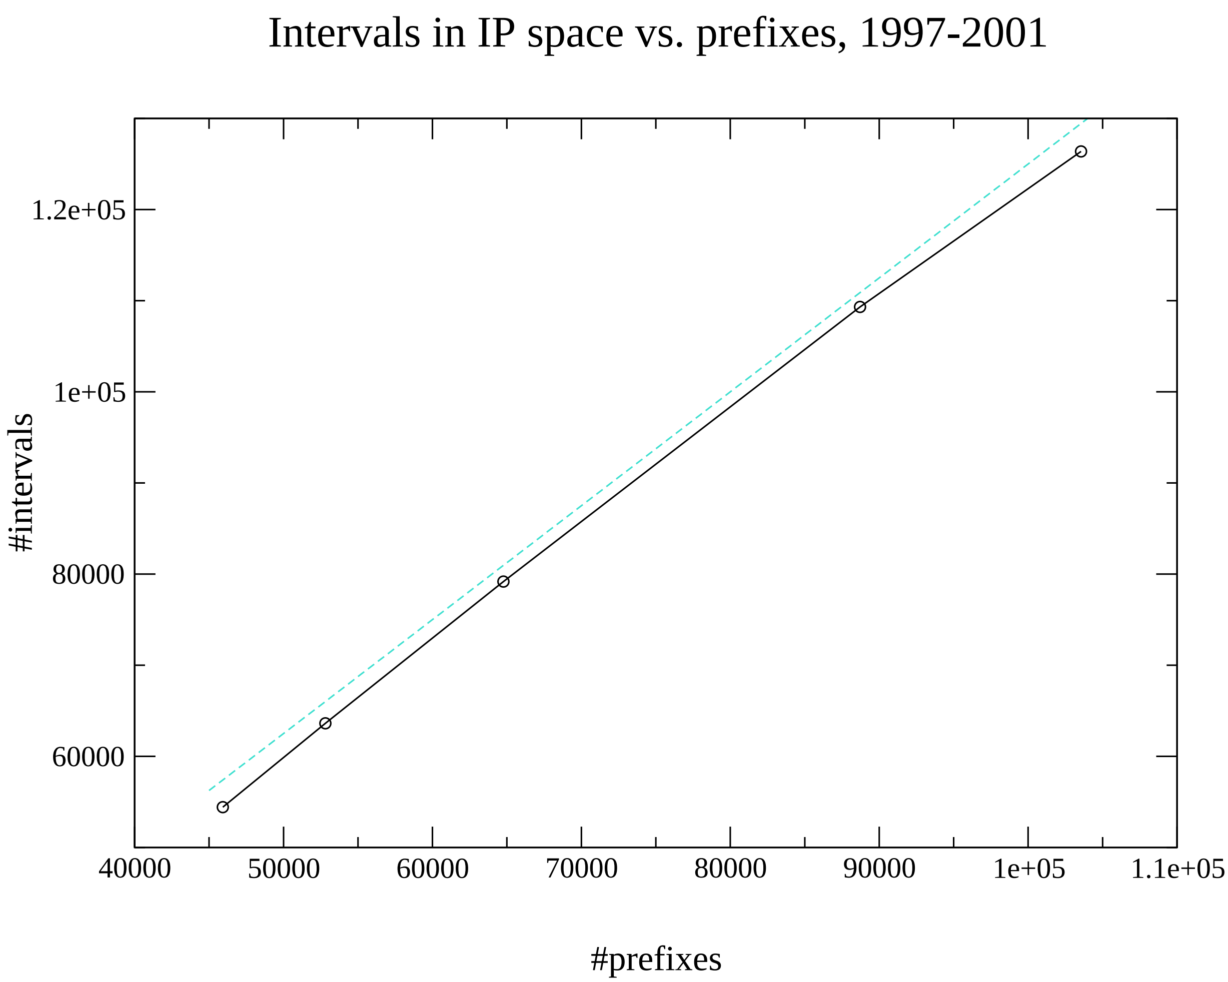

Figure 14: Intervals in IP space vs. prefixes (RouteViews 1999-2001)

IP space usage sophistication

Intervals grow in concert with prefixes

constantly 1/4 more than prefixes

Potential number of CIDR blocks is large though

75K new routes if all holes punched by more specifics' were deagregated in the most economic way

Prevailing deaggregation is linear
a /16 to 256 /24s

More specifics save routing table from explosion

Fragmentation: gaps in IP space

Make address space management harder

effectively make fragments of IP space unusable.

Depend on ways address space is managed

Delegation, assignment, allocation, subdivision,
reuse and return of address blocks

Excessive gaps growth in Nov 2000 - May 2001 6 months same as 2 previous years

due to high death rate of smaller blocks

Local IP space can also pose as gaps

Internet growth: conclusions

Slowed in 2001-2002

Economy slowed down, too

Growth masked by high drop rate (churn)

Concerted growth: constant ratios

Granularity refinement: AS, root prefixes, address blocks

A breathing space – for how long?

Suggestions

Make routing more network-friendly

Avoid bandwidth waste by circuitous routes

Make delay, loss, jitter, bandwidth BGP parameters

In-band micro-payments beyond peering links

#routes as an economic indicator?

INTERNET EXPANSION, REFINEMENT AND CHURN

Vocal and lyrics:

Andre Broido, Evi Nemeth, kc claffy & elves
(broido, evi, kc) @ caida.org

CAIDA/SDSC/UCSD

NANOG, Miami 12 feb 02

(Eur.Trans.Telecomm., Feb.2002)

www.caida.org/~broido/nanog200202.egr.pdf
www.caida.org/~broido/nanog200202.egr.ps
www.caida.org/presentations/nanog0202/