Observed Relationships between Size Measures of the Internet

or

Is the Internet really just a star network after all?

Brian Carpenter
Department of Computer Science
The University of Auckland

June 2009

THE UNIVERSITY OF AUCKLAND
NEW ZEALAND
Te Whare Wānanga o Tāmaki Makaurau

Work partially supported by HUAWEI

Of course, the Internet is doomed

- The particular doom of interest in this talk is:
  The wide-area BGP4 routing system grows beyond any reasonable scaling limit.
- The limits that might matter include
  - size of the routing table
  - lookup time in the forwarding tables
  - propagation and computing time of routing updates
  - power density & energy consumption of core routers
- There's a sad lack of hard data about these issues
- But at least we have historical growth data

How many hosts on the Internet? (1969-2008, log scale)

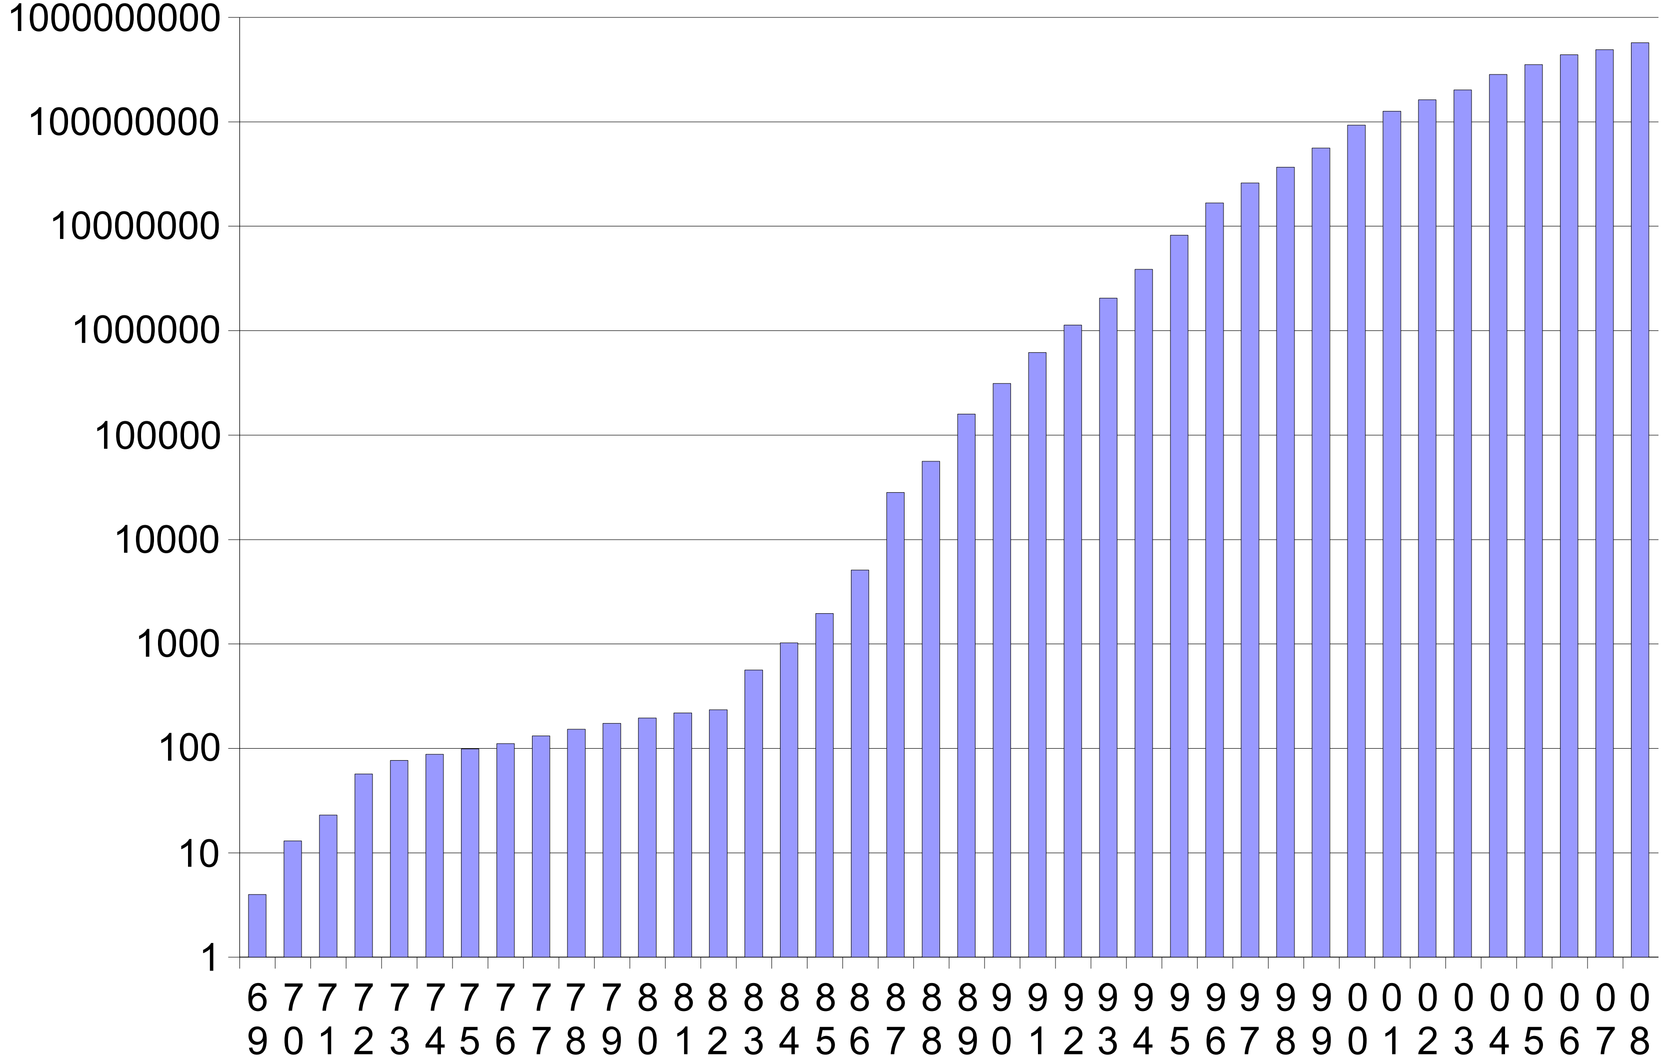

Data from http://www.isc.org/ and other sources

Host growth (1994-2008, linear scale)

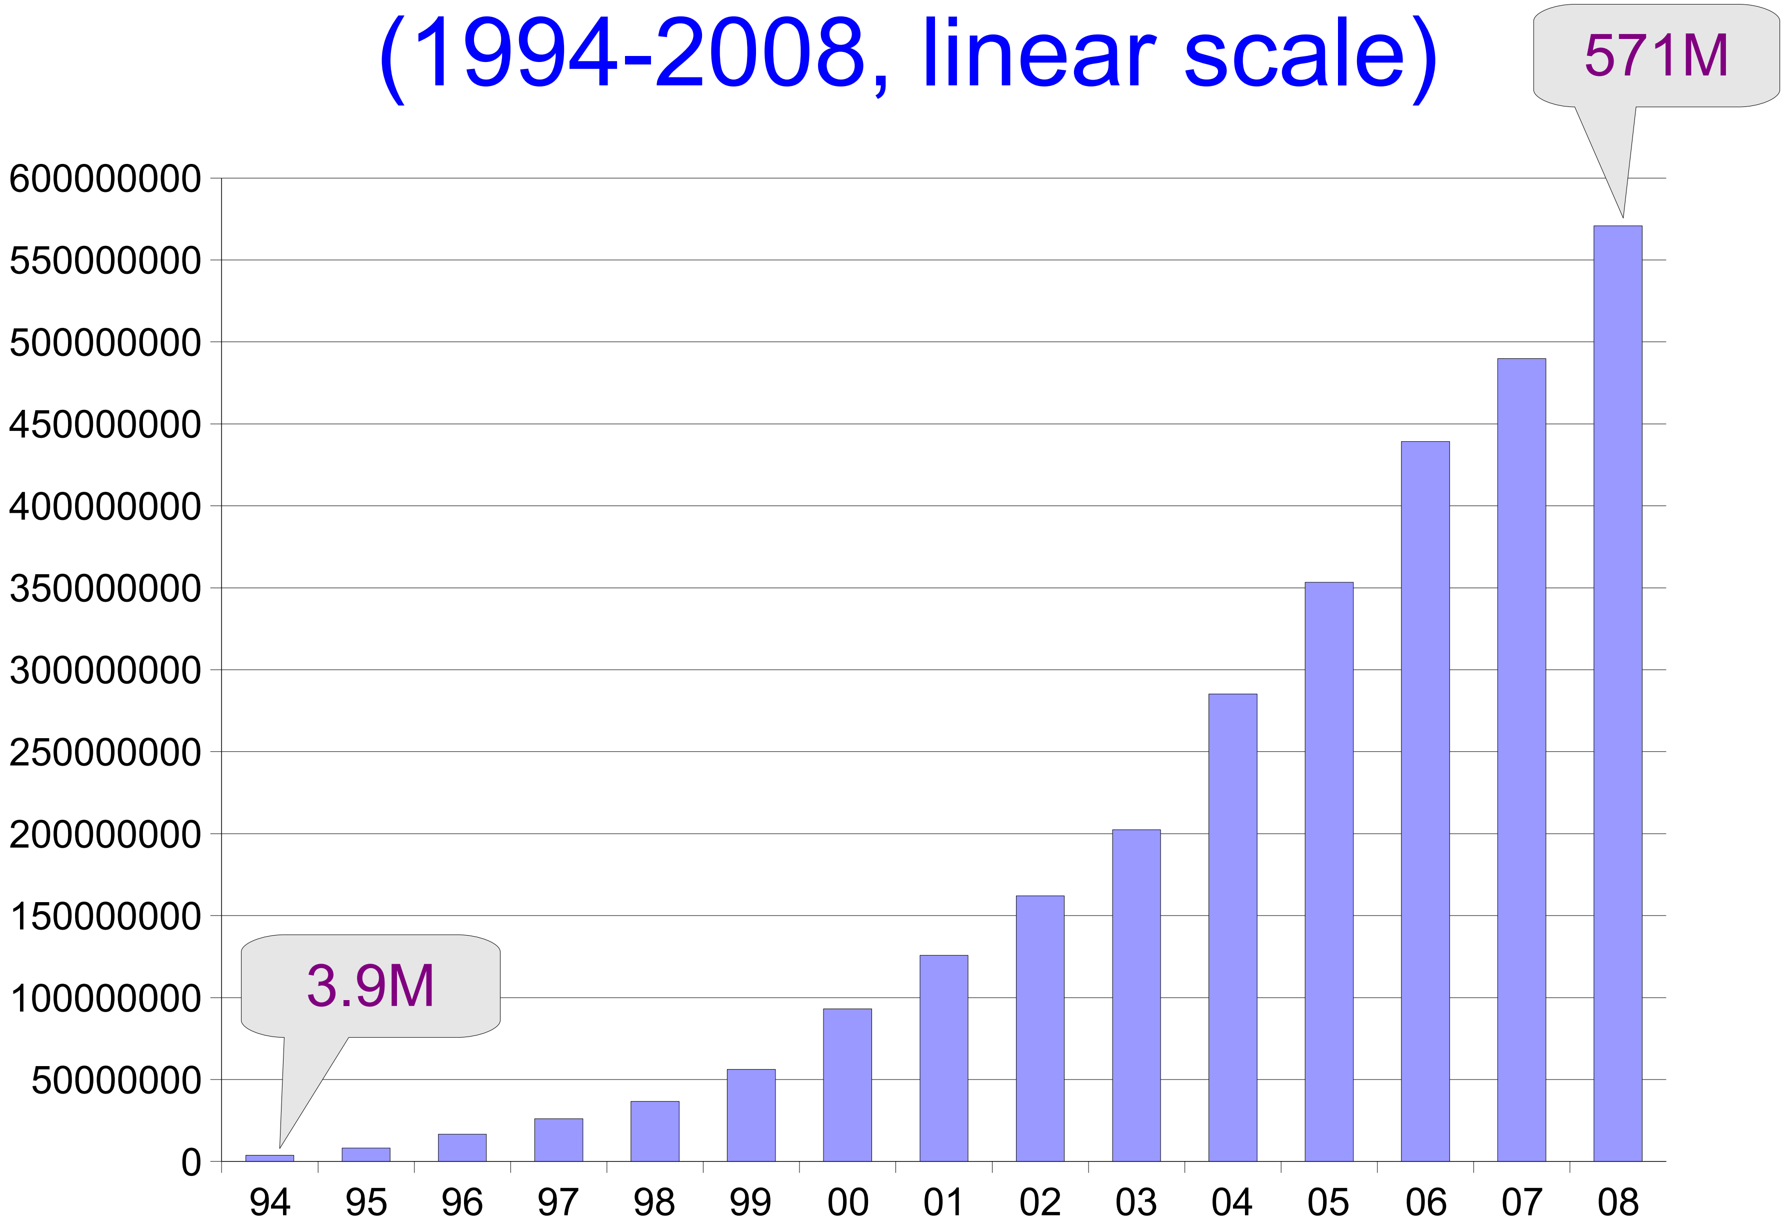

Data from http://www.isc.org/

BGP4 growth
(1989-now, linear scale, APNIC view)

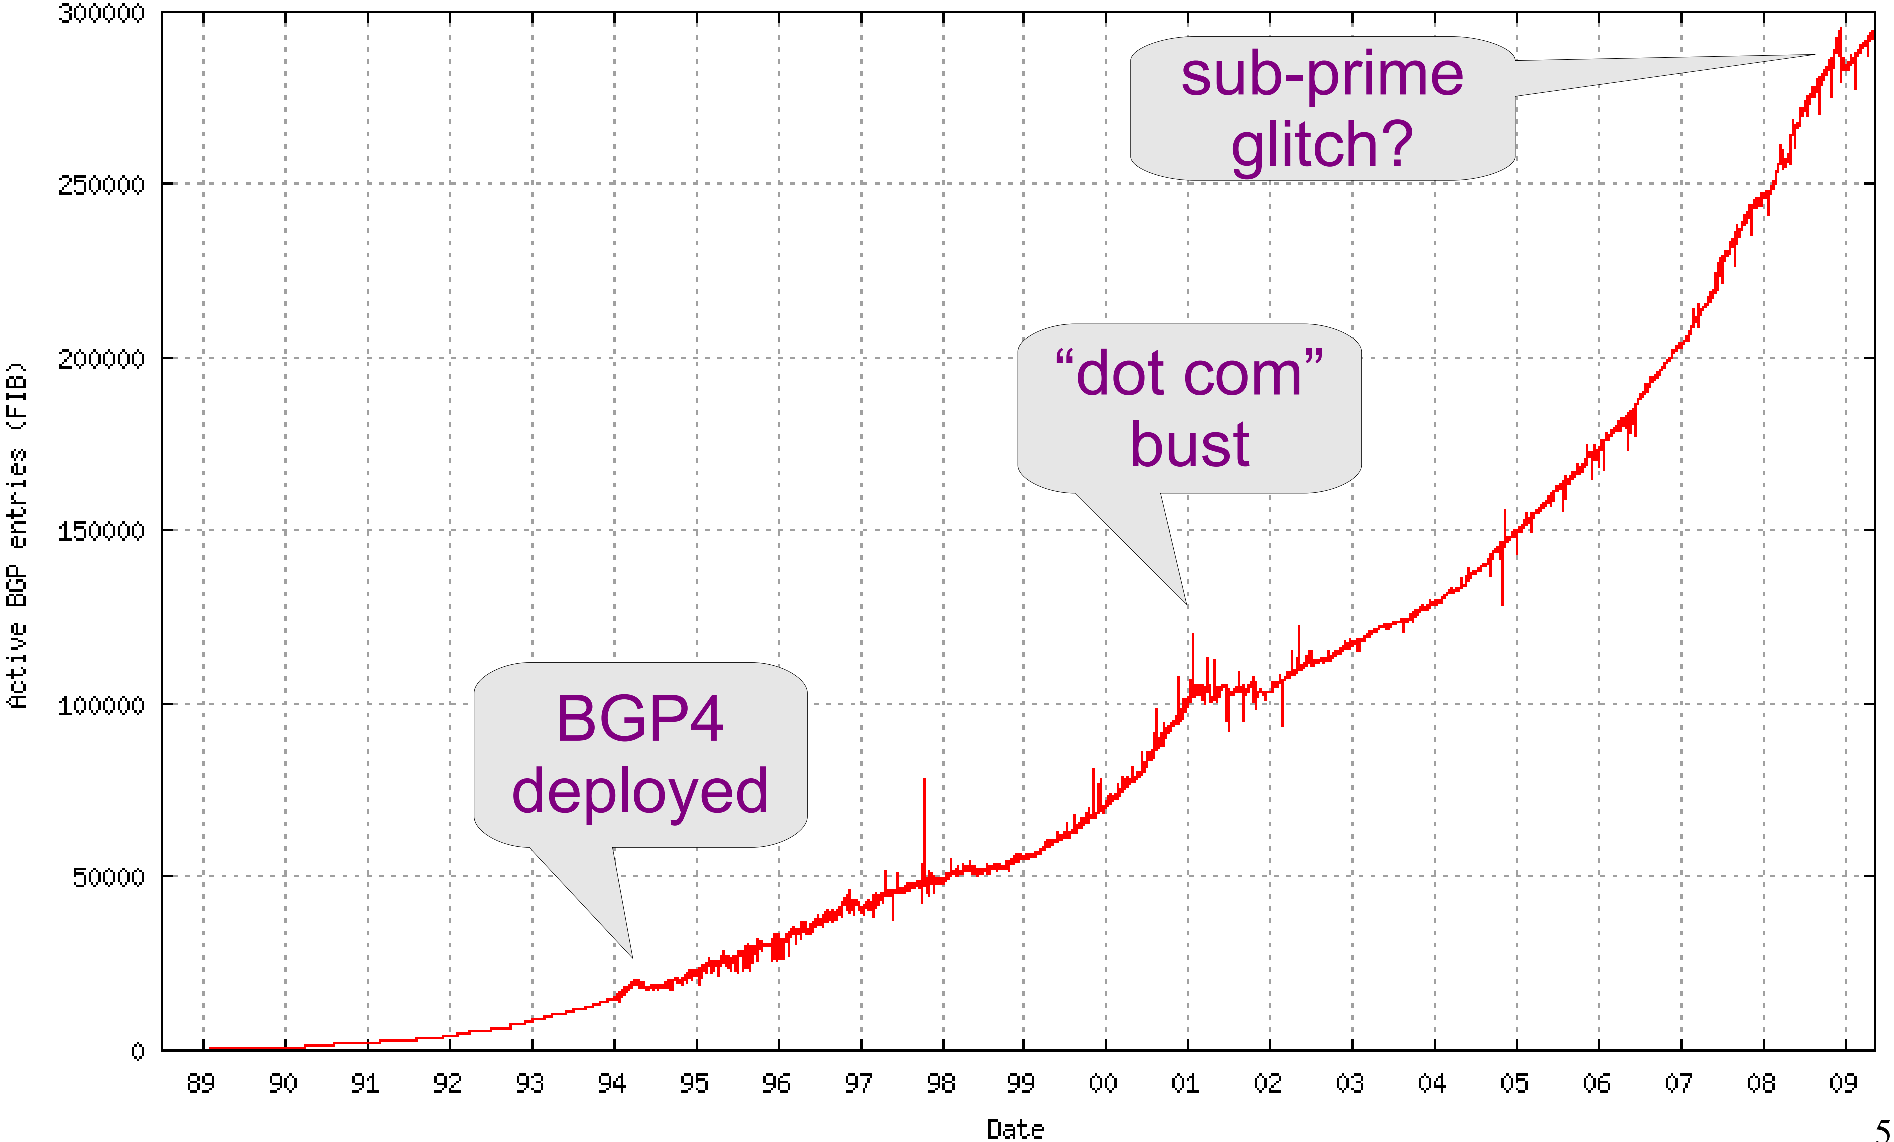

http://www.potaroo.net/

Abstract view of the BGP4 system

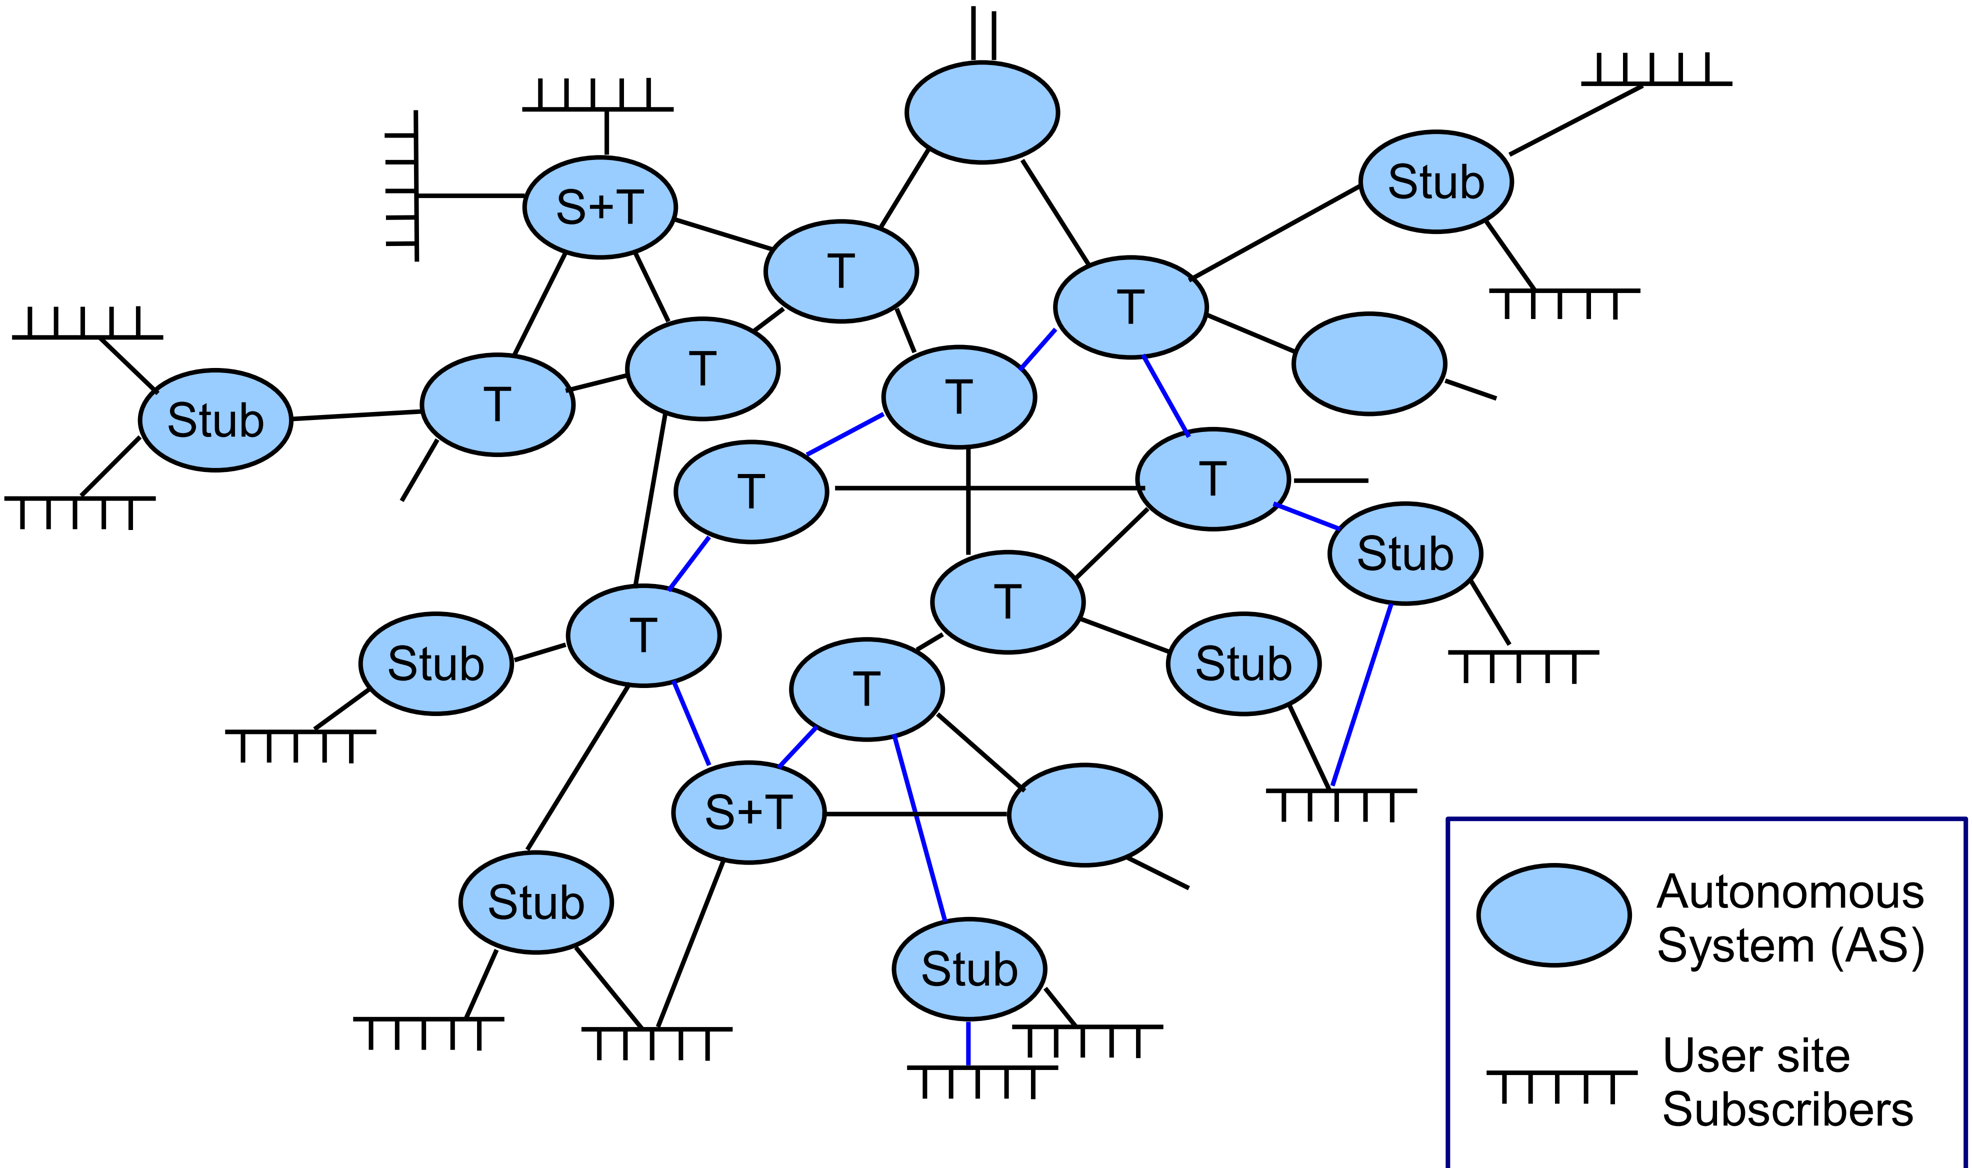

T = Transit AS; S+T = Stub + Transit AS

BGP jargon

- Default-Free Zone (DFZ)
  - The interior part of the BGP4 system, consisting of transit ASes, where all routes are explicit and there is no default (wild card) route.
- Classless Inter-Domain Routing (CIDR)
  - The address allocation scheme adopted in 1994 with BGP4, whereby the old Class A, B and C addresses were abolished, and addresses were allocated in the smallest practical binary blocks.

Gripe: why on earth do some text books still talk about Class A, B and C except as historical remnants?

What questions can we ask about the data?

- Are there any interesting long-term relationships between the size of the globally addressable Internet and the size of the BGP4 system?
- Things that are (relatively) easy to count:
  - Total number of addressable devices on the Internet (~600 million today)
  - Total number of active Autonomous Systems in BGP4 (~30,000 today)
  - Total number of BGP4 routes (~300,000 today)

What data do we have?

- We have BGP4 data back to 1994 and active AS data back to 1997 (thanks potaroo.net)
- We have DNS domain count data back to 1994 (thanks ISC.org)
  - The domain count is a reasonable lower bound on the number of directly accessible IPv4 interfaces with global addresses (± addresses with no DNS name and names with unactivated address).
  - Note that this estimates *sometimes connected* hosts, not *simultaneously connected* hosts.
  - Heidemann et al IMC'08 paper: 2003-2007 ICMP censuses show about 25% as many *sometimes pingable* hosts.

BGP4 table size vs Active AS count

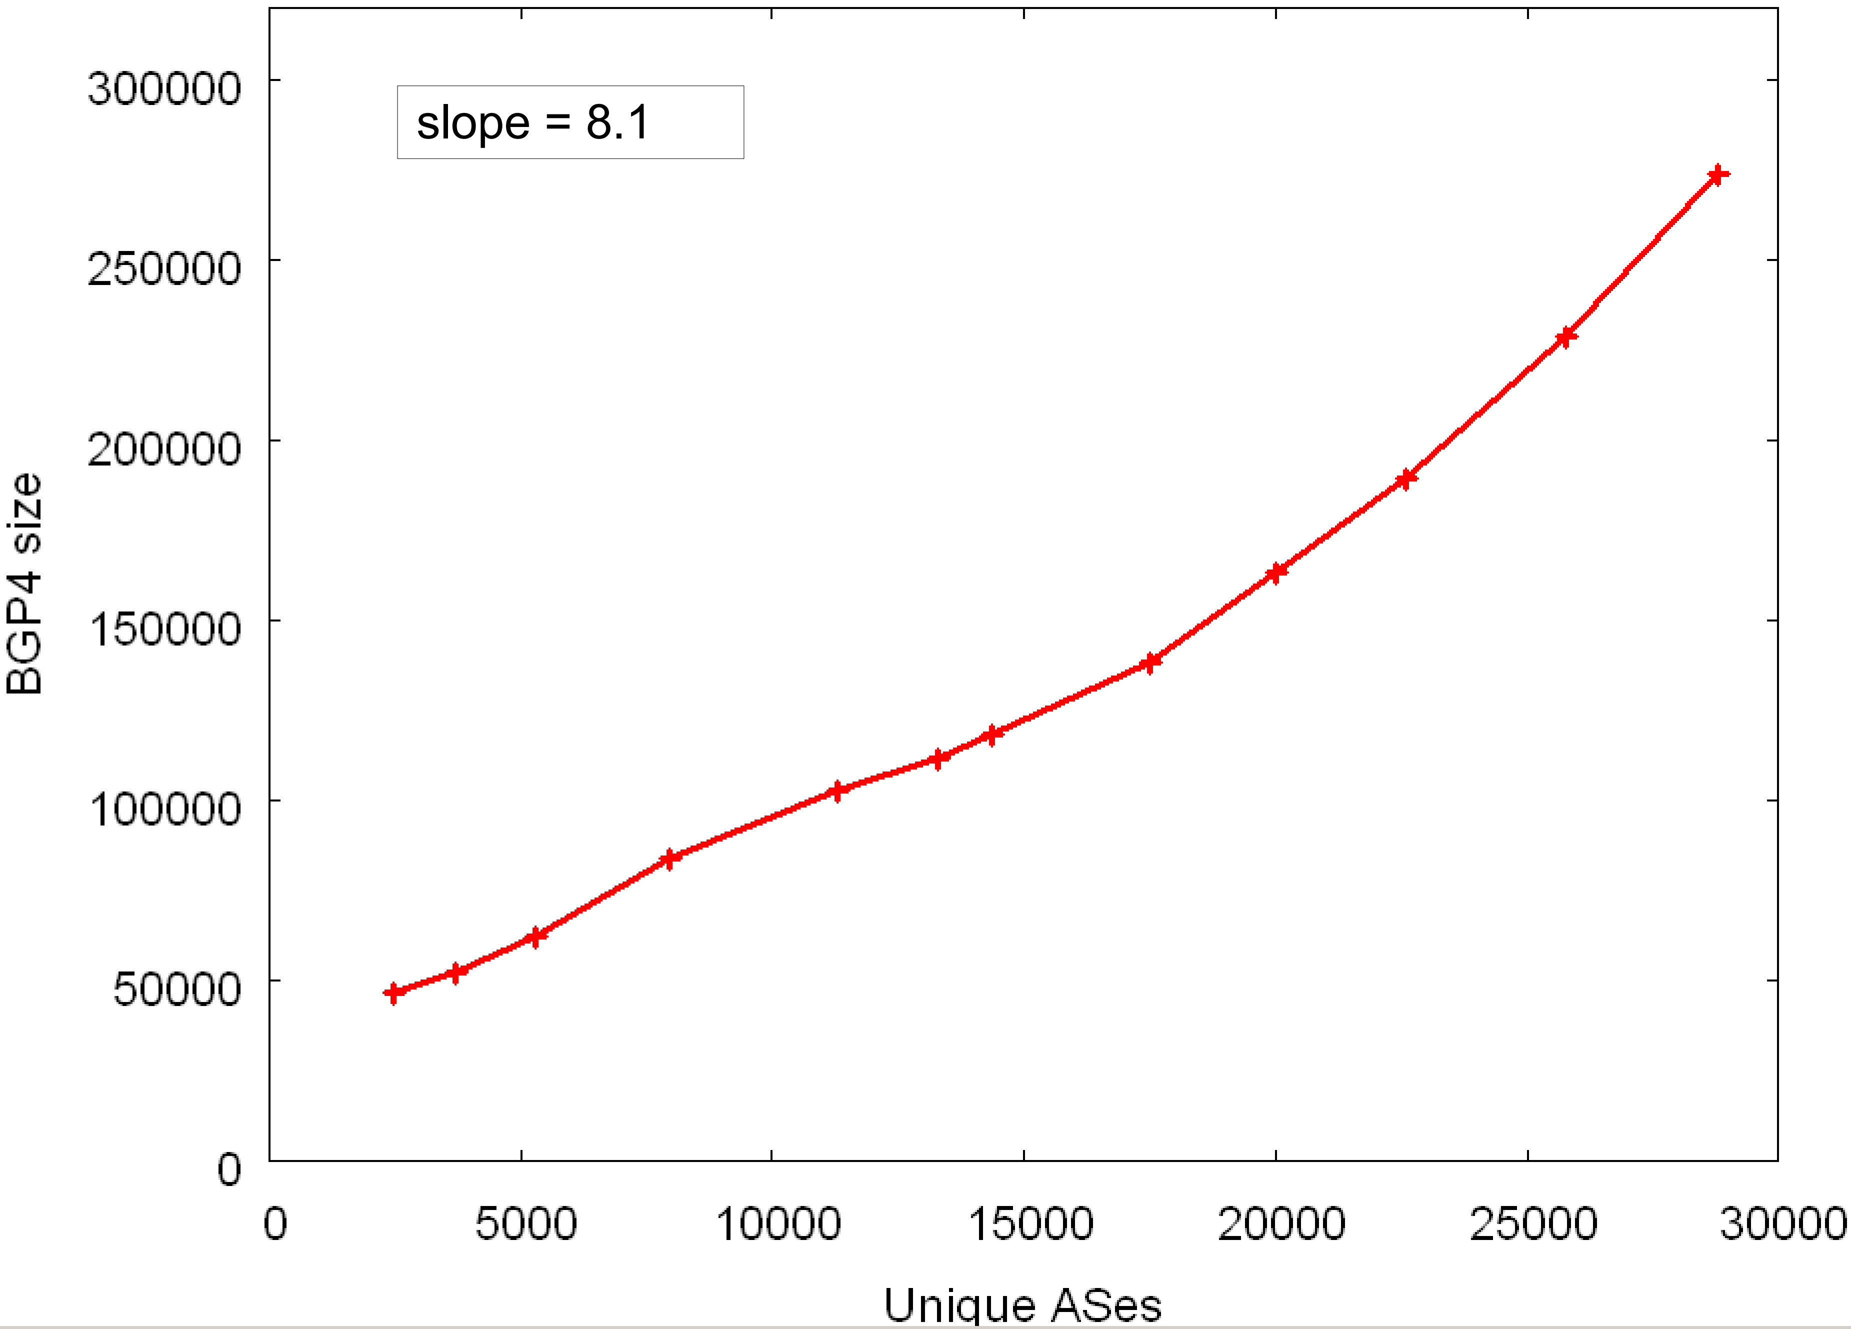

BGP4 table size vs Active AS count

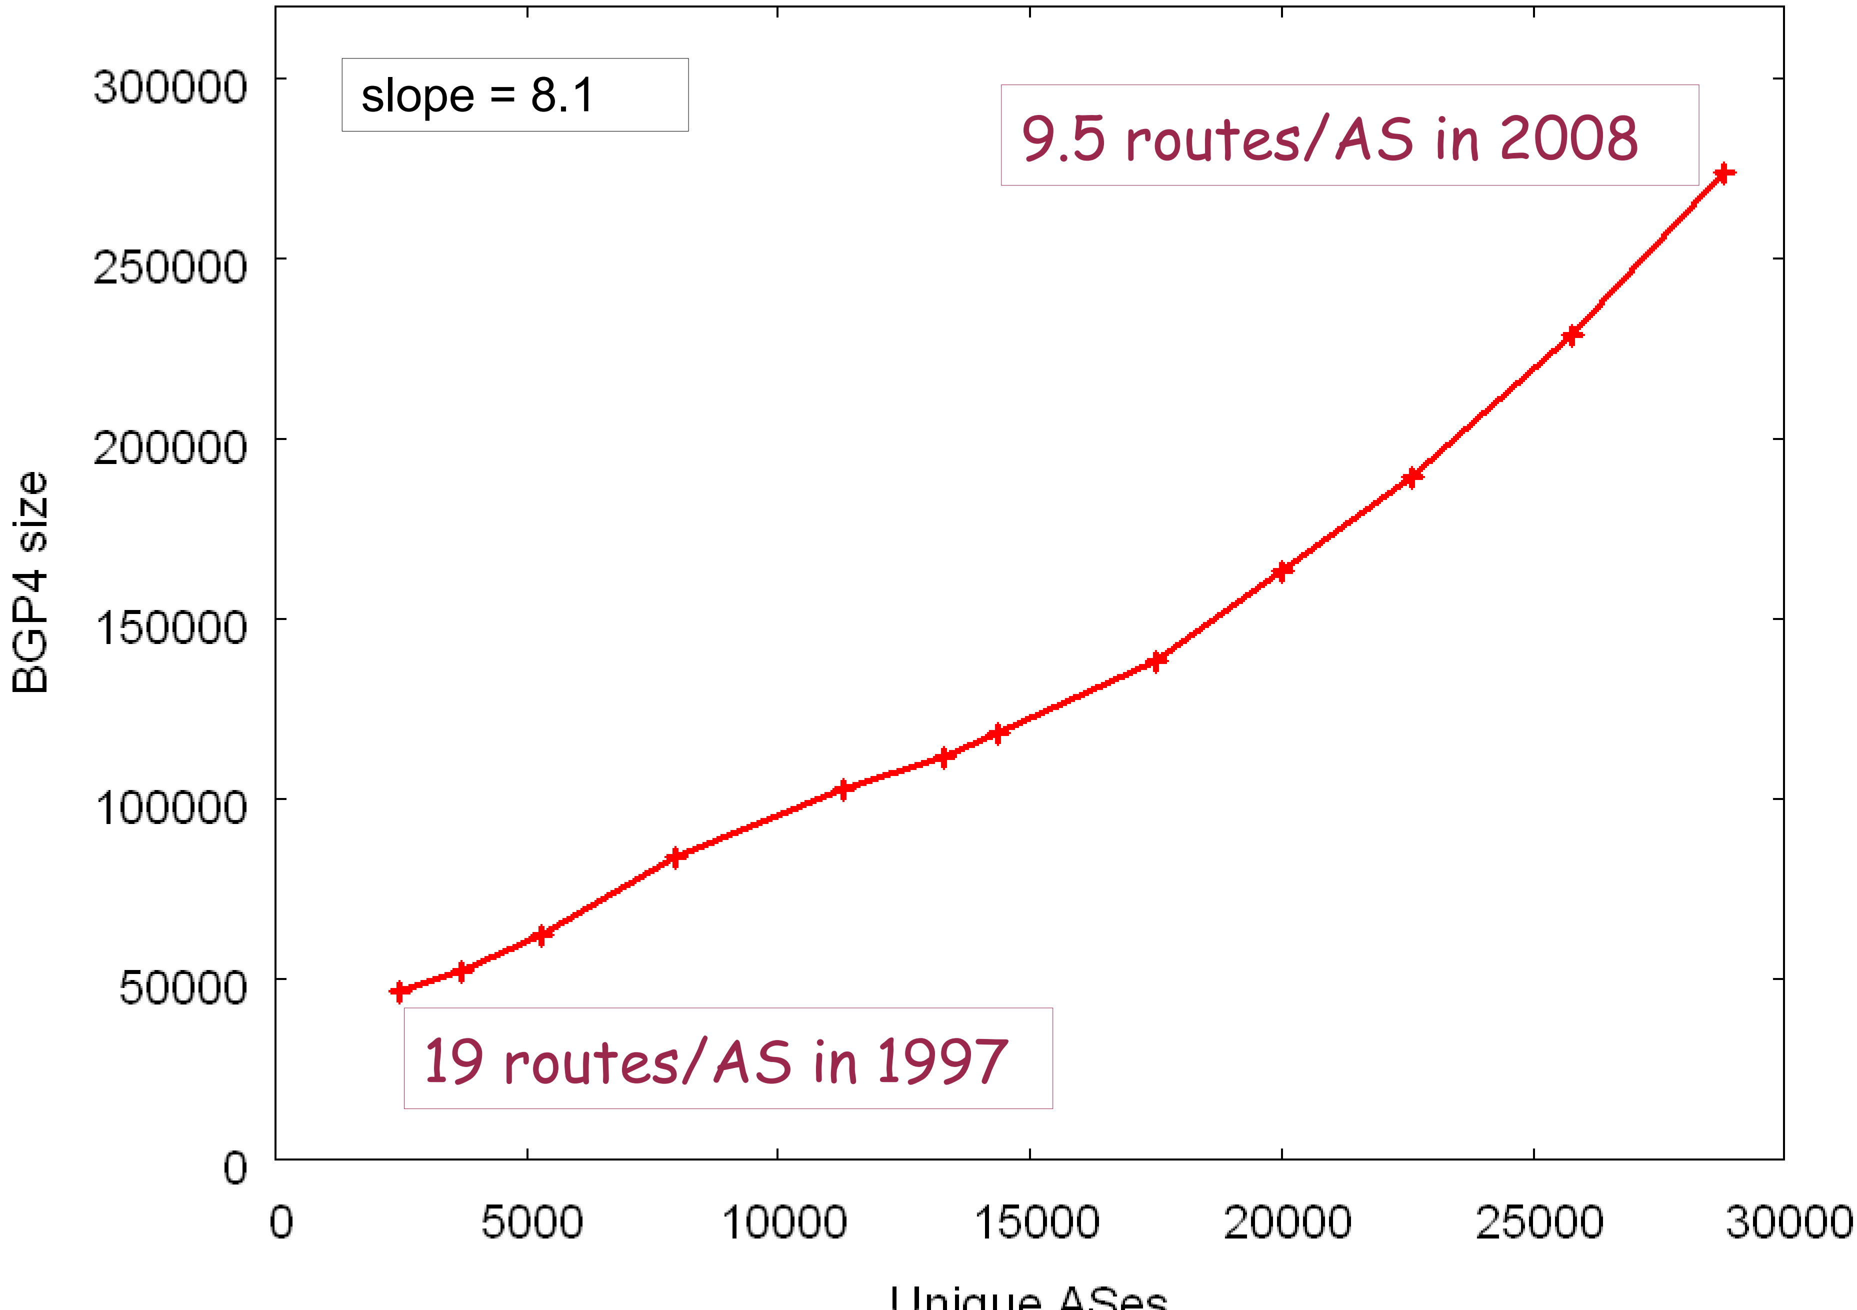

BGP4 table size vs Active AS count

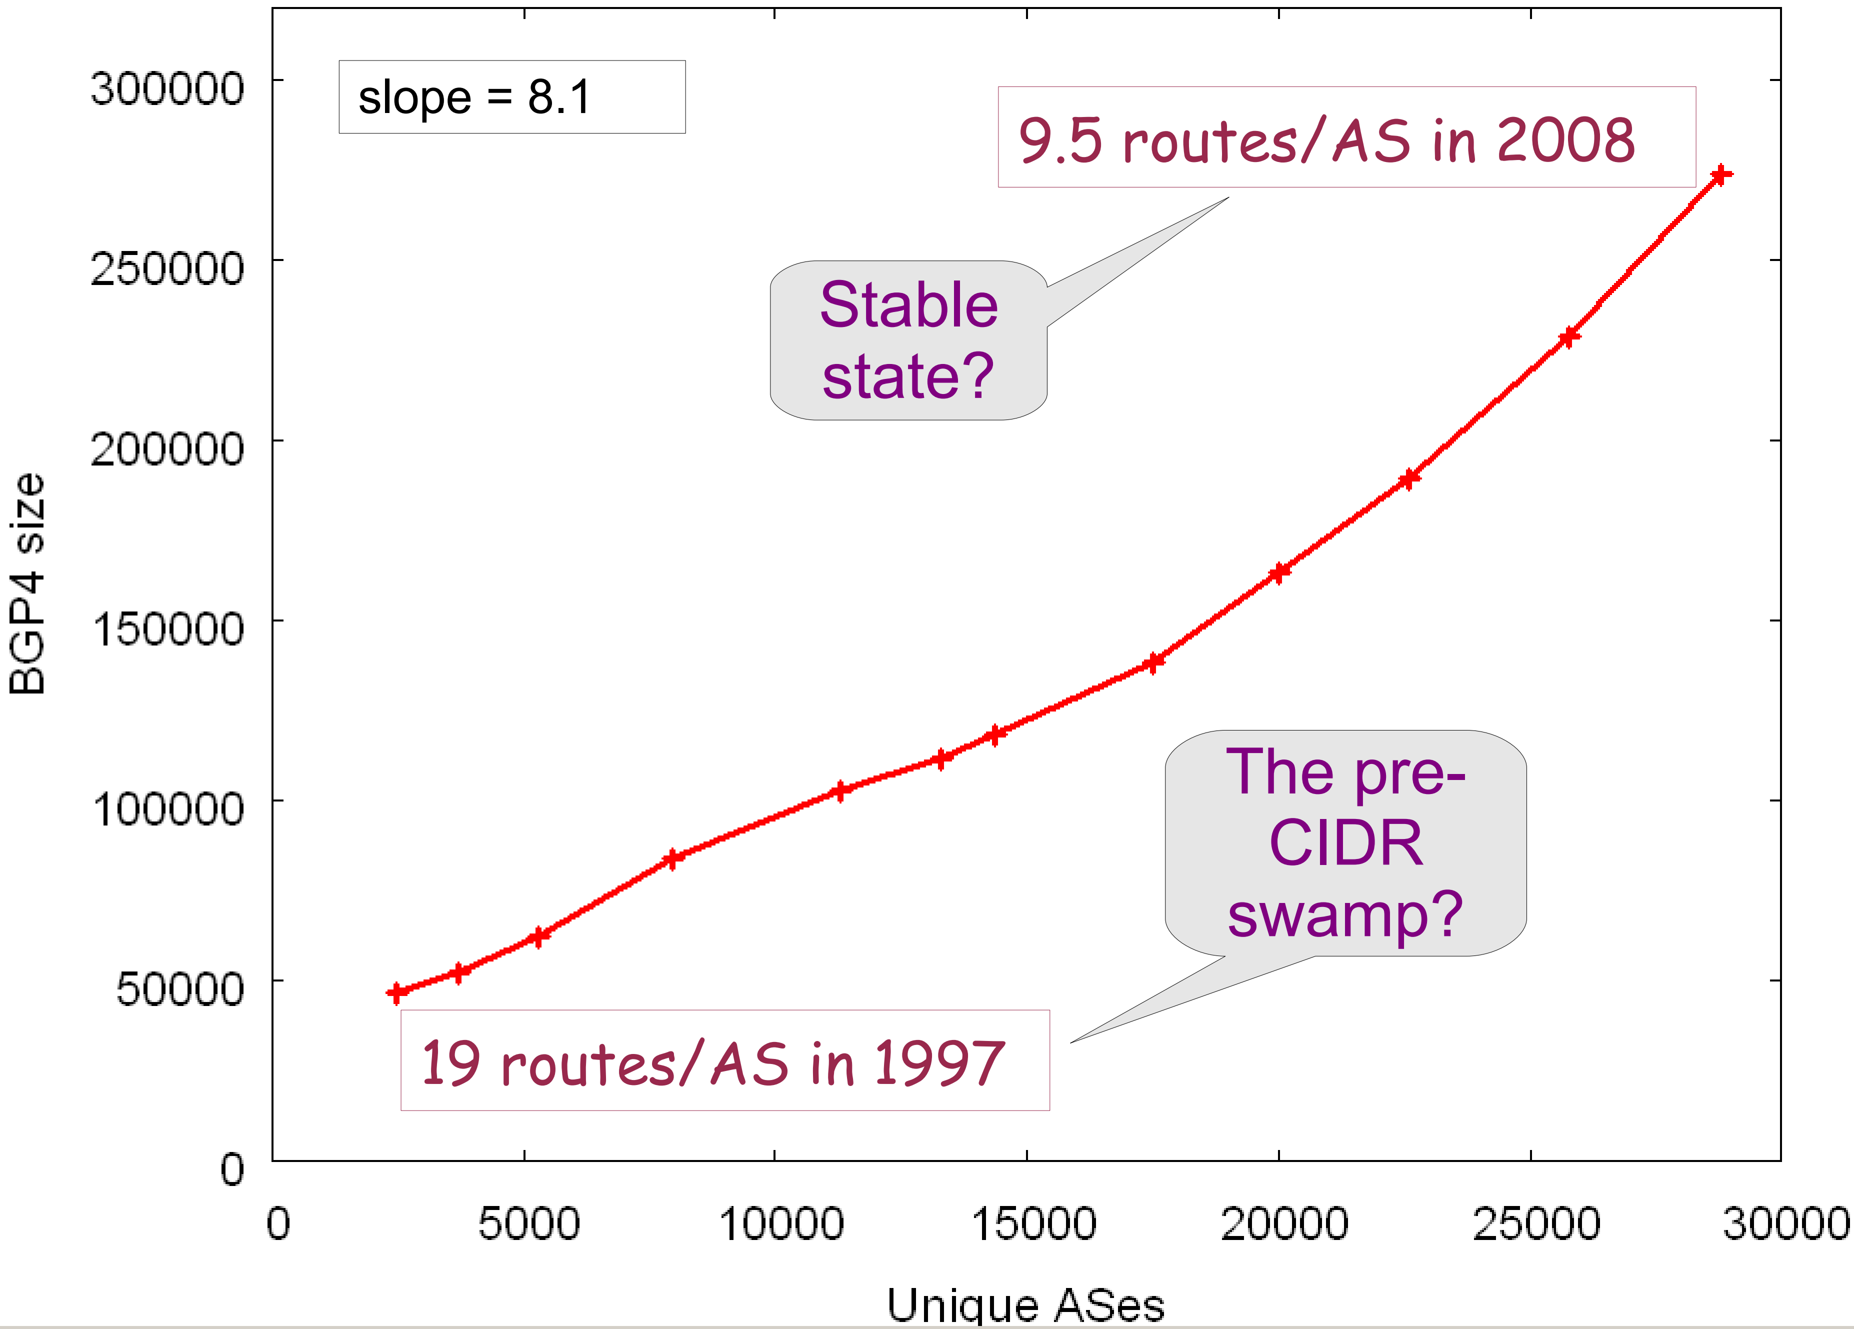

BGP4 table size vs Sqrt of domain count

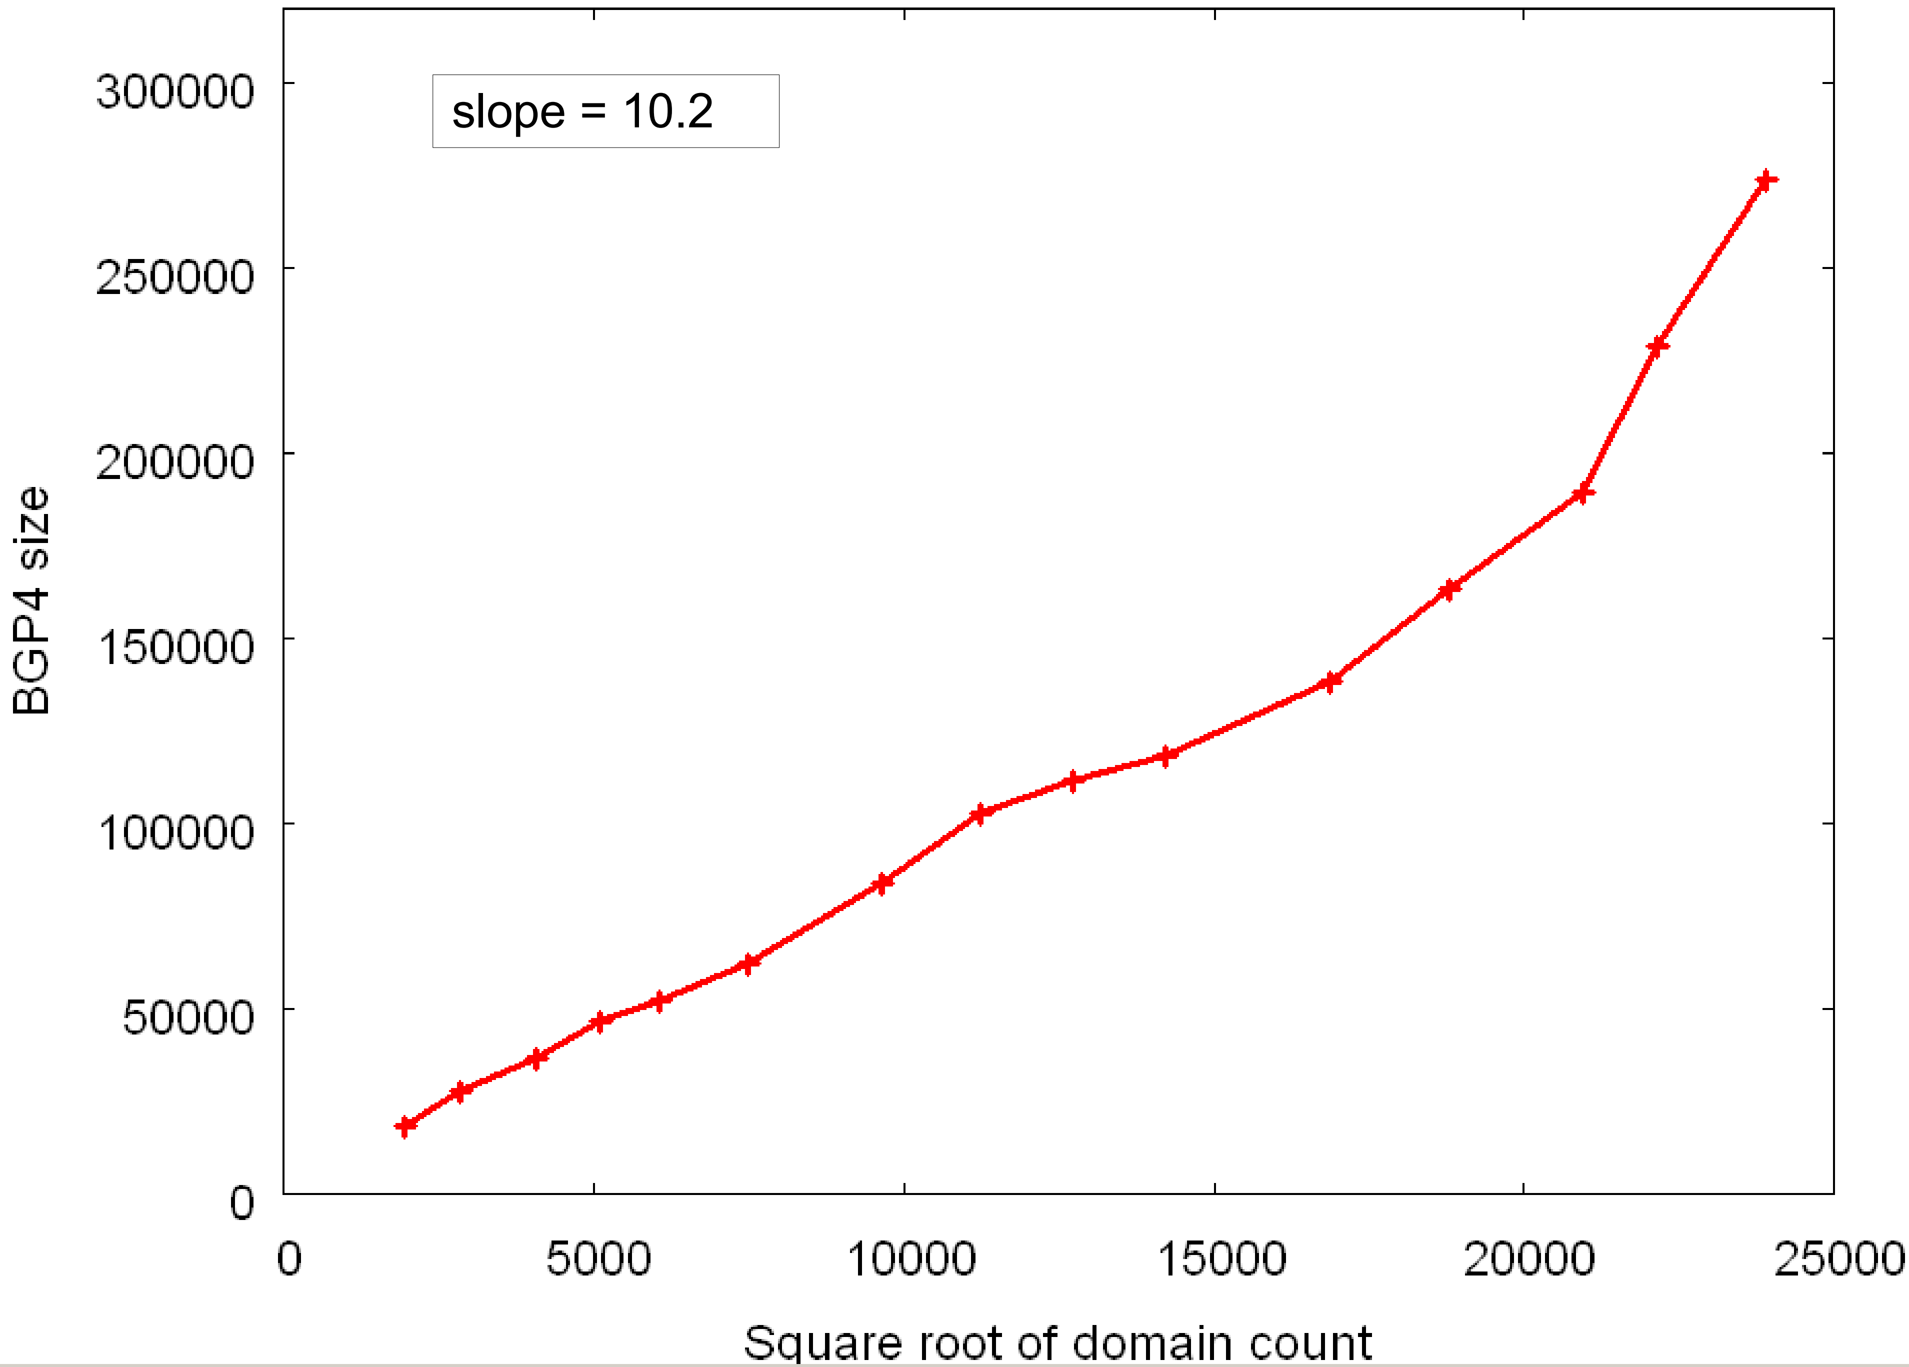

BGP4 table size vs Sqrt of domain count

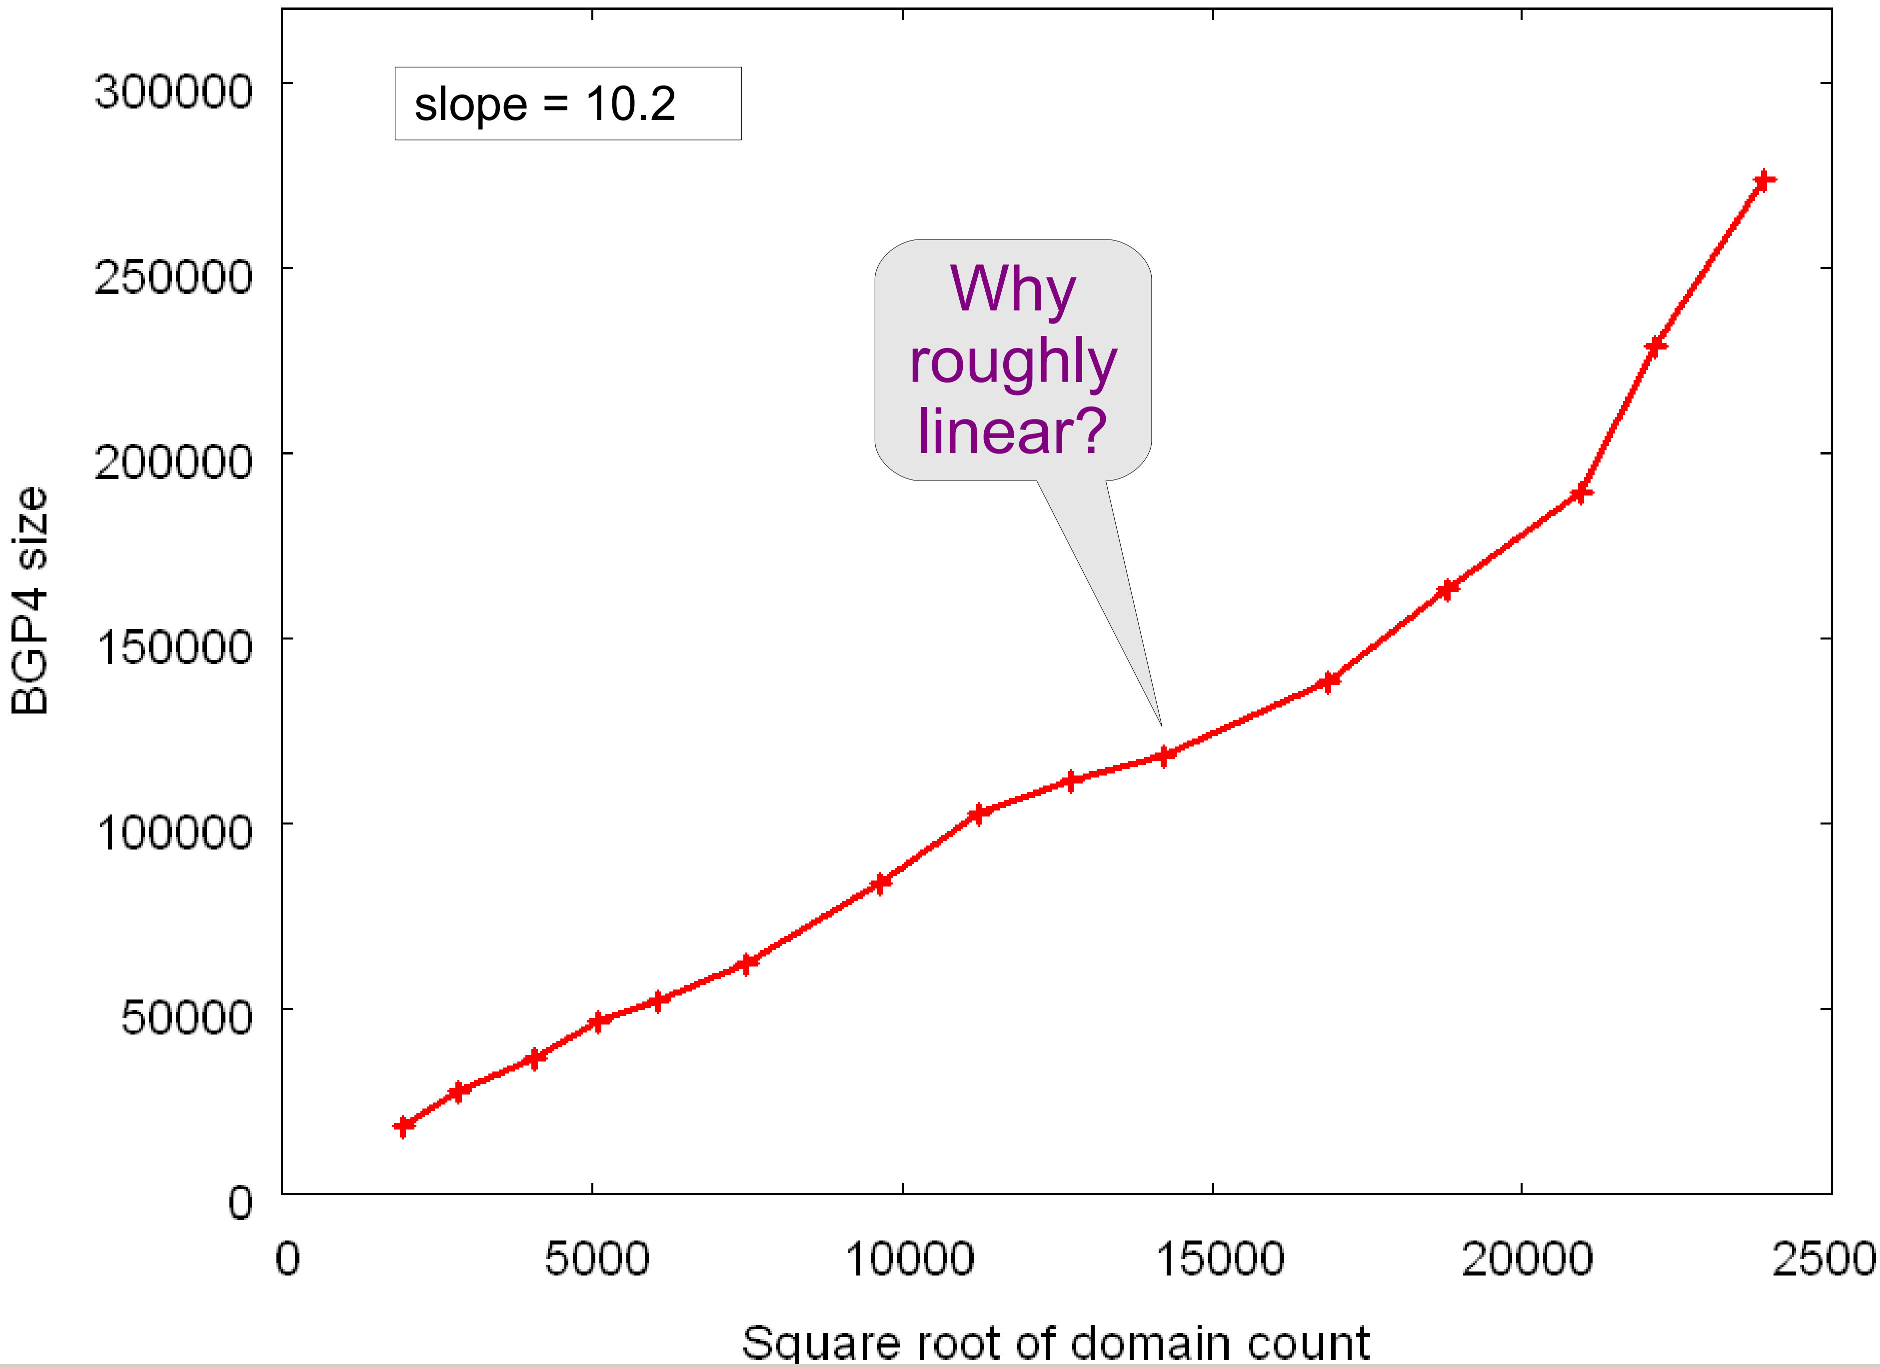

Active AS count vs Sqrt of domain count

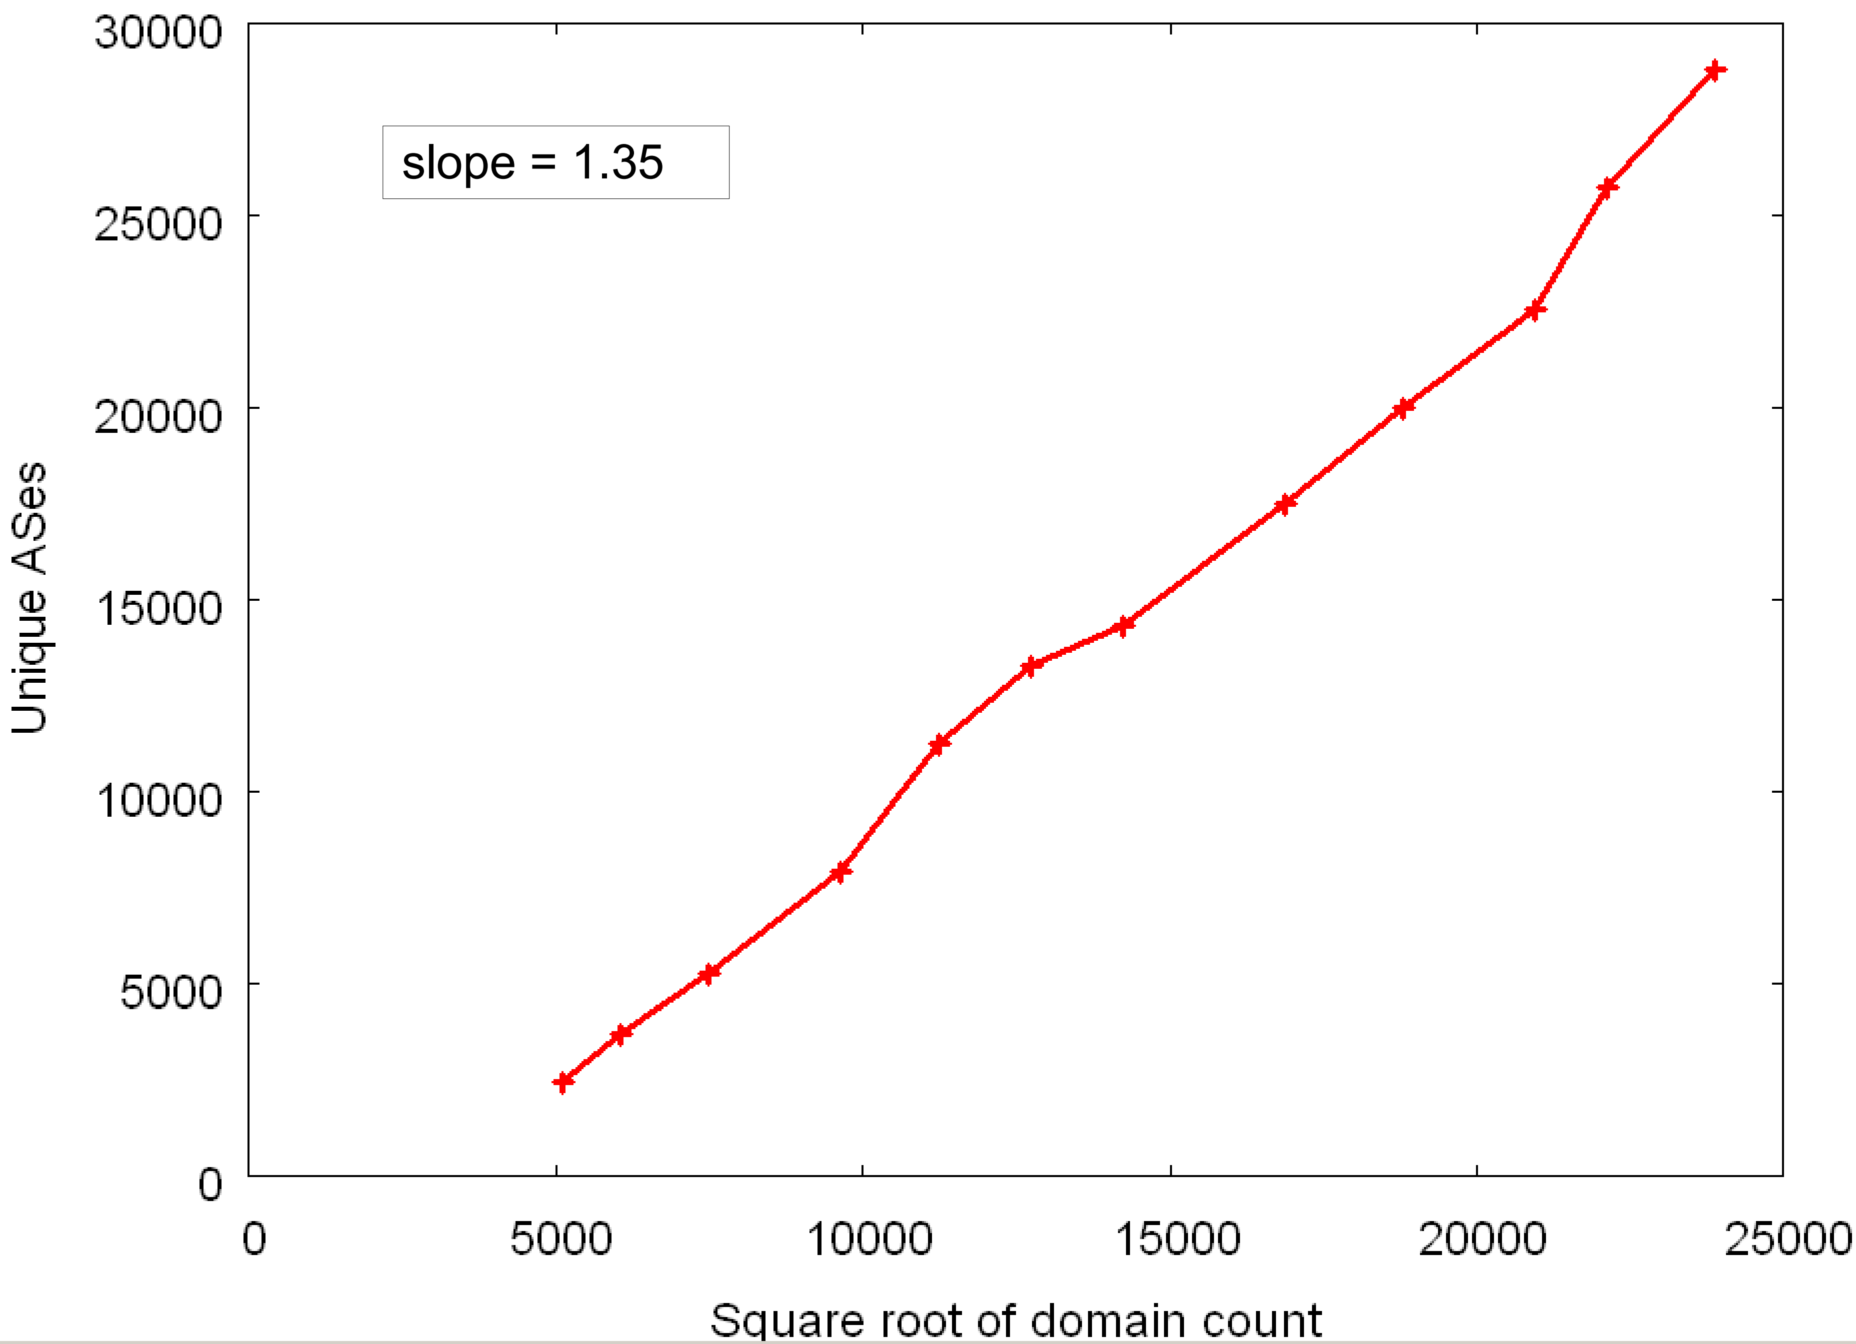

Active AS count vs Sqrt of domain count

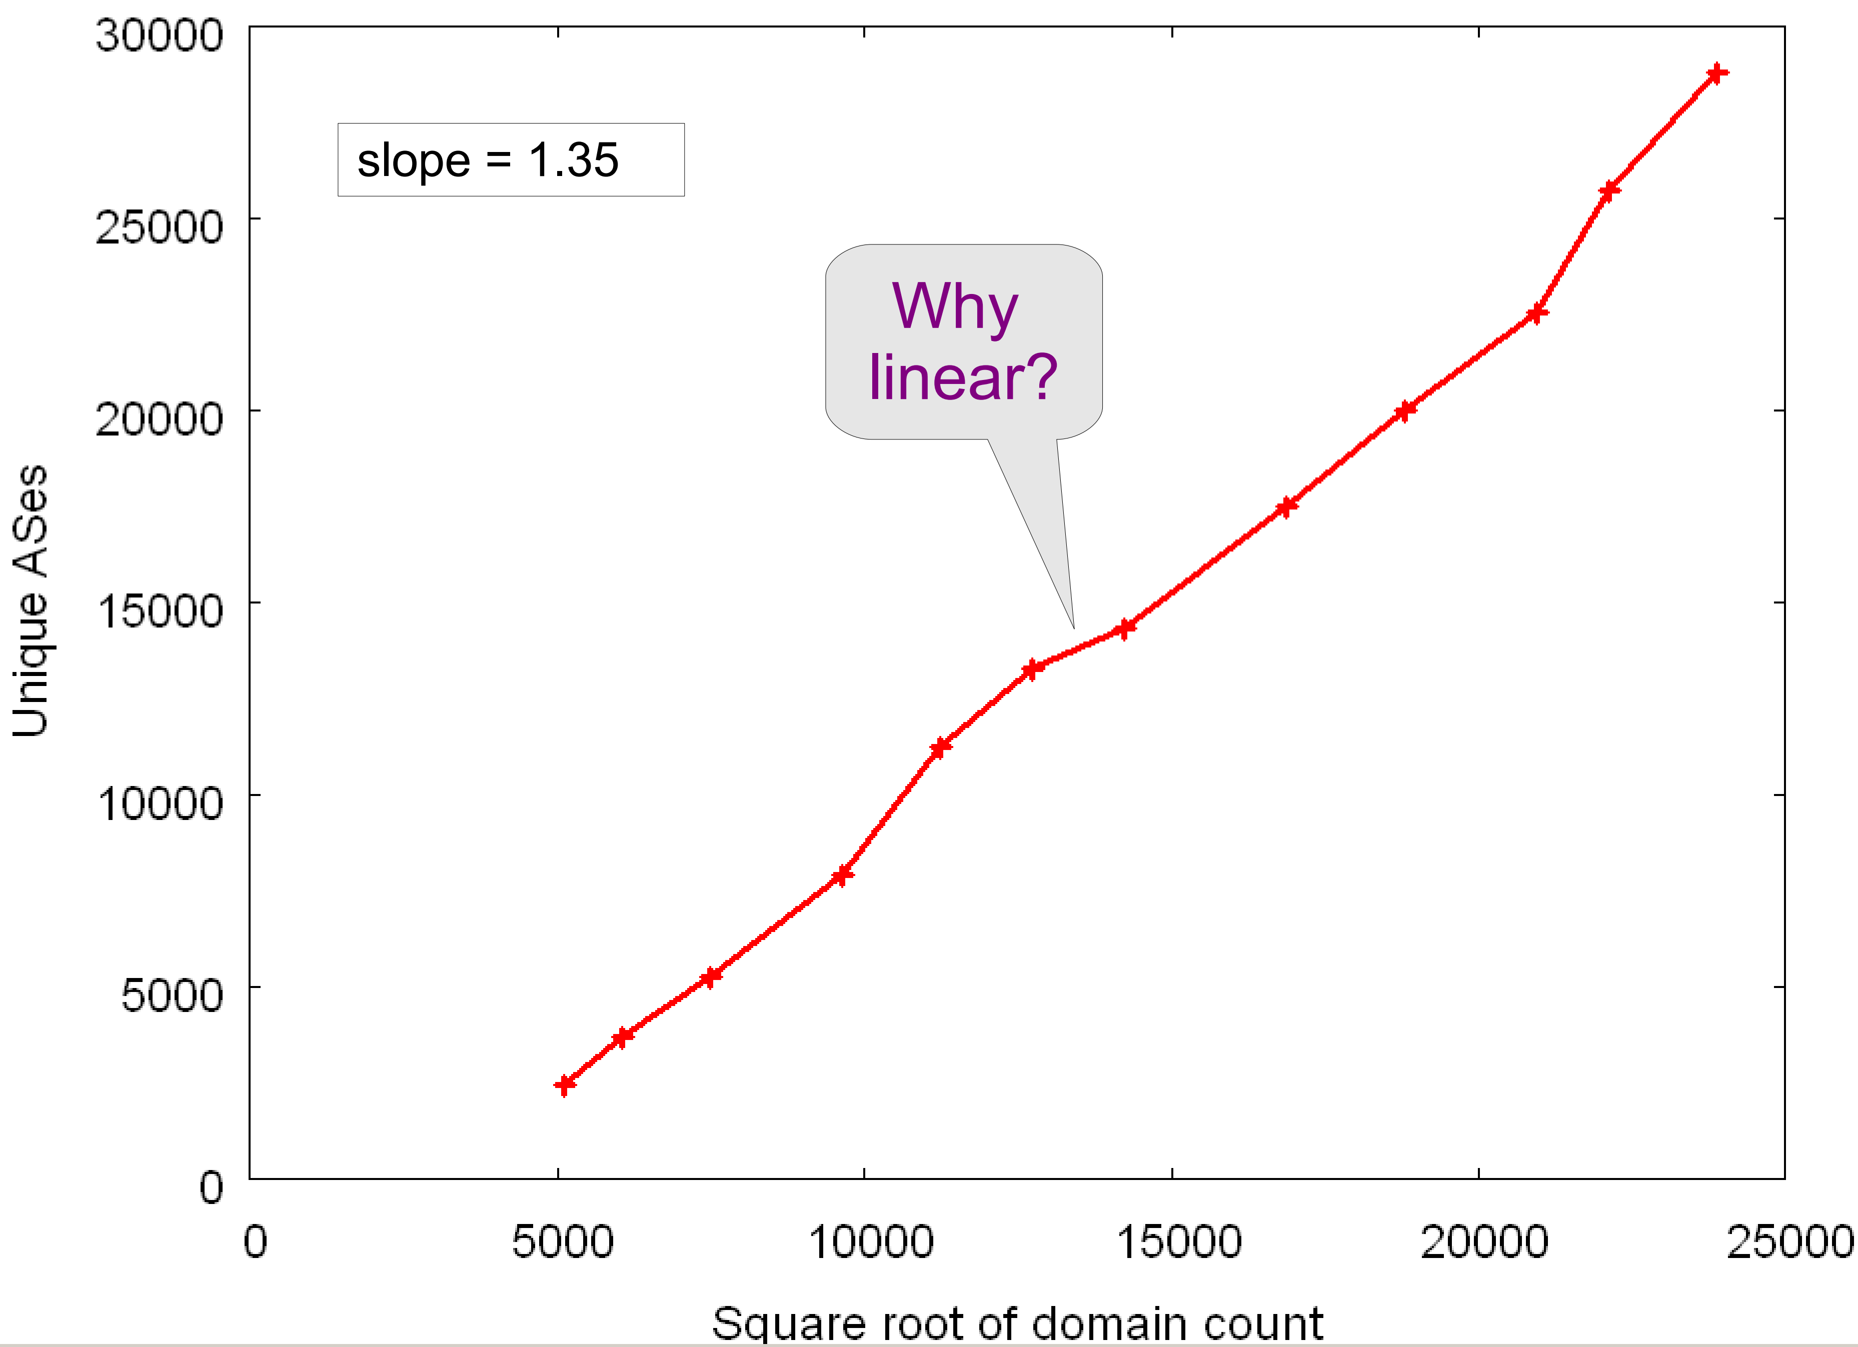

Discussion (1)

- For the record: using data from other BGP4 viewpoints, or from the recent ICMP census, would not change things much. Everything still looks ~ linear.

- The BGP4/AS relationship shows that CIDR worked: the number of routes per AS significantly declined during BGP4's lifetime, and seems to be in an equilibrium.
- 15000 residual pre-CIDR routes + 8 routes per AS

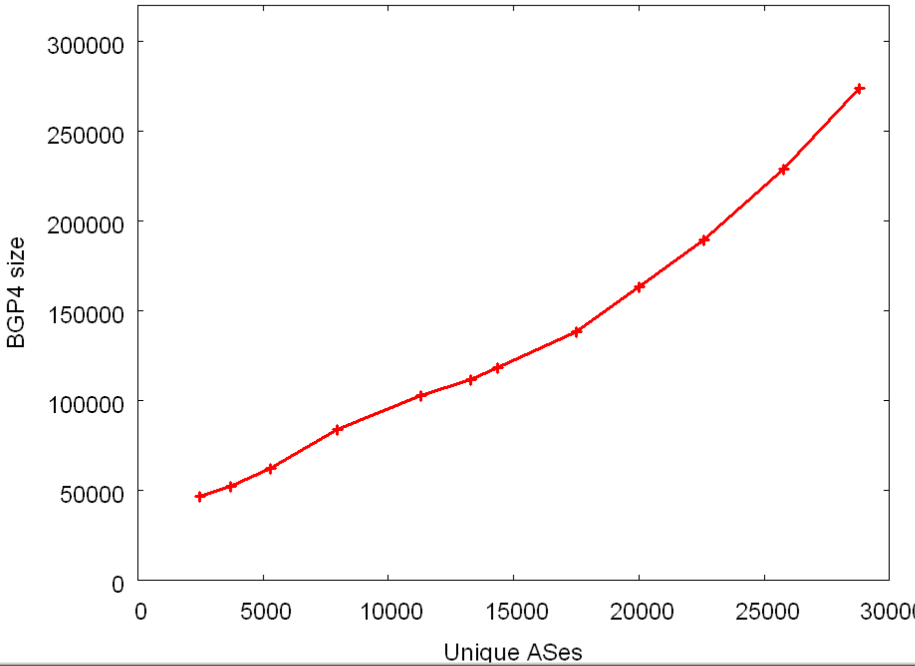

Discussion (2)

- Why has the square law relationship between BGP size and host count been sustained over 14 years?
- If B = size of BGP4 routing table (DFZ),

  A = number of active ASes,

  D = domain count,

  then $B = 10.19\sqrt{D} - 10563$

  $A = 1.35\sqrt{D} - 4615$

- Imagine the Internet as a star...

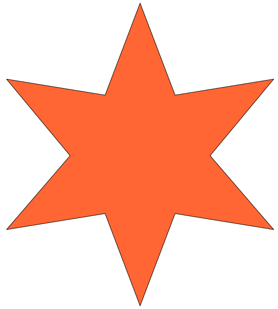

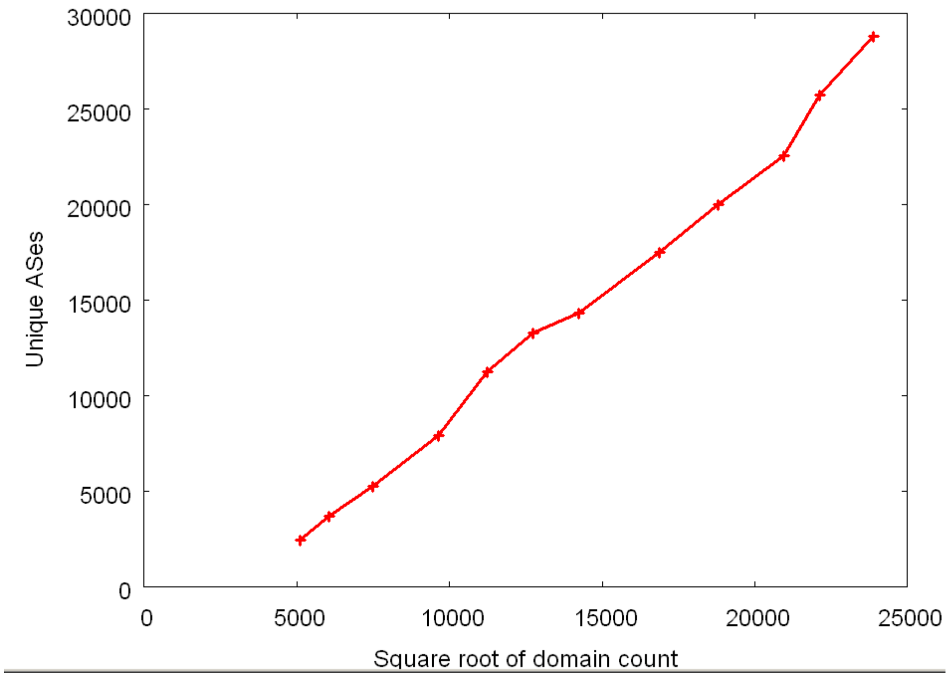

Imagine

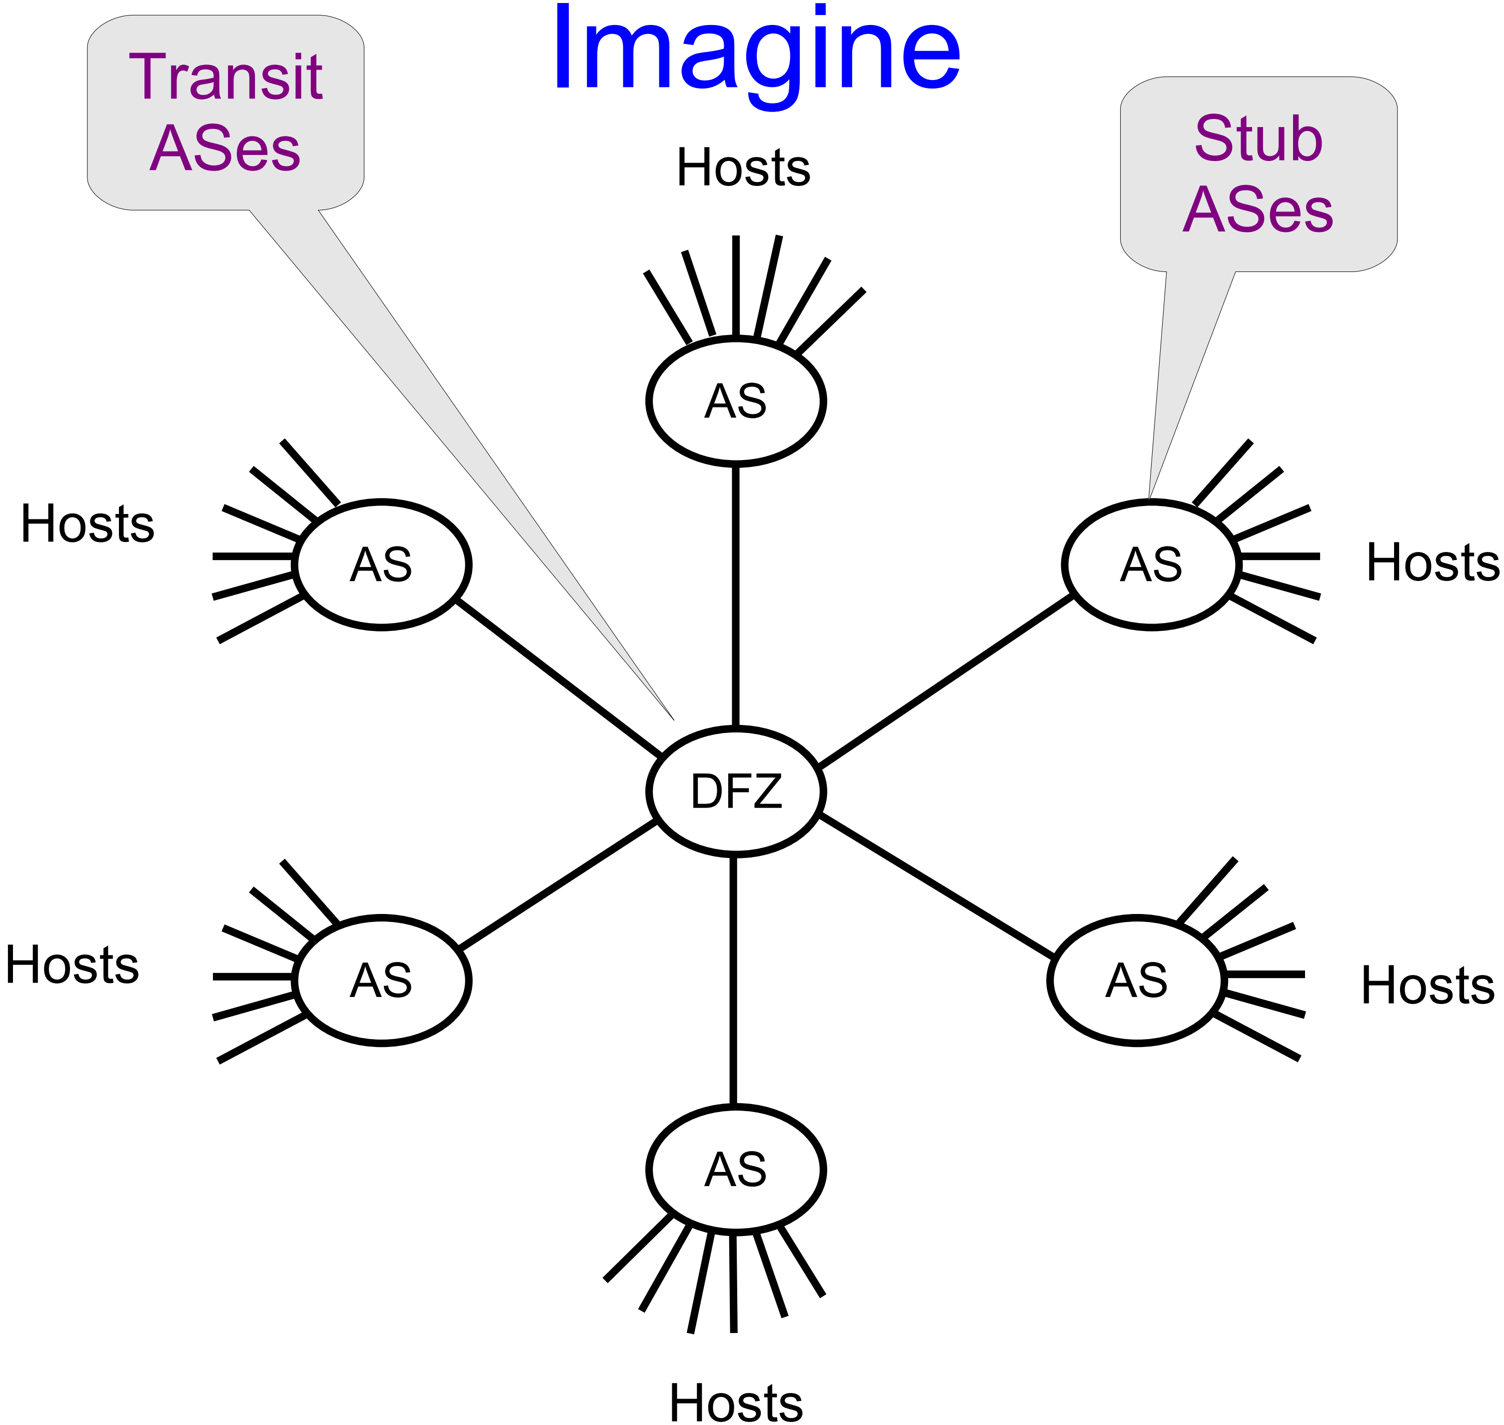
Transit
ASes
Stub
ASes
Hosts
AS
Hosts
AS
AS
Hosts
DFZ
Hosts
AS
AS
Hosts
AS
Hosts

Discussion (3)

- If N stubs each support N hosts, the total host count would be N^2.
- If each stub supports kN hosts, the total host count would be kN^2.
- If we set $D = kN^2$ and $A = N$, then

  $$A = 1.35\sqrt{D} - 4615$$

  becomes

  $$N = 1.35\sqrt{kN^2} - 4615$$

  giving $k = 0.55$ for large N.

Discussion (4)

- We can argue that the Internet has grown like a star topology, with 0.55 X slower growth in the hosts per AS than in the number of ASes.
  - at the same time, the theoretical address span of each BGP4 route has decreased by 1.6 X
  - symptomatic of increasingly efficient address usage
- The centre of the star is the mesh of transit-ASes
  - Note that in reality 86% of ASes are pure stubs, 14% are stub+transit, and only 0.35% are pure transit. The star model is only an approximate explanation.
- However, this star-like history tends to explain why the growth of the BGP4 system has been *significantly* slower than the growth of the Internet.

Observations aren't predictions

- Note that we have no grounds for extrapolating any of the graphs.
  - They reflect past practice by ISPs and do not predict future practice.
  - Changes in technical or business pressures could change the slopes.
- Continuing to plot these graphs will allow us to monitor the scaling of the BGP4 system.

Summary

- We’re fortunate that the relationships shown above have been linear. Otherwise, the core routers would have melted some years ago.
- CIDR worked - routes per AS are stable.
- Address conservation worked - address utilisation per AS became much more efficient.
- The (unplanned) star topology worked - caused $\sqrt{N}$ growth instead of linear growth.
  – therefore, it’s a Good Thing for the large majority of ASes to be origin-only stubs.
- Observations aren’t predictions.

References

- Carpenter, Brian E.,
  Observed Relationships between Size Measures of the Internet,
  ACM SIGCOMM CCR, 39(2) (April 2009) 6-12
  http://www.cs.auckland.ac.nz/~brian/sqlaw.pdf
- Heidemann, John, et al,
  Census and Survey of the Visible Internet,
  ACM IMC’08 (October 2008) 169-182
  http://doi.acm.org/10.1145/1452520.1452542

Domain count vs ICMP census

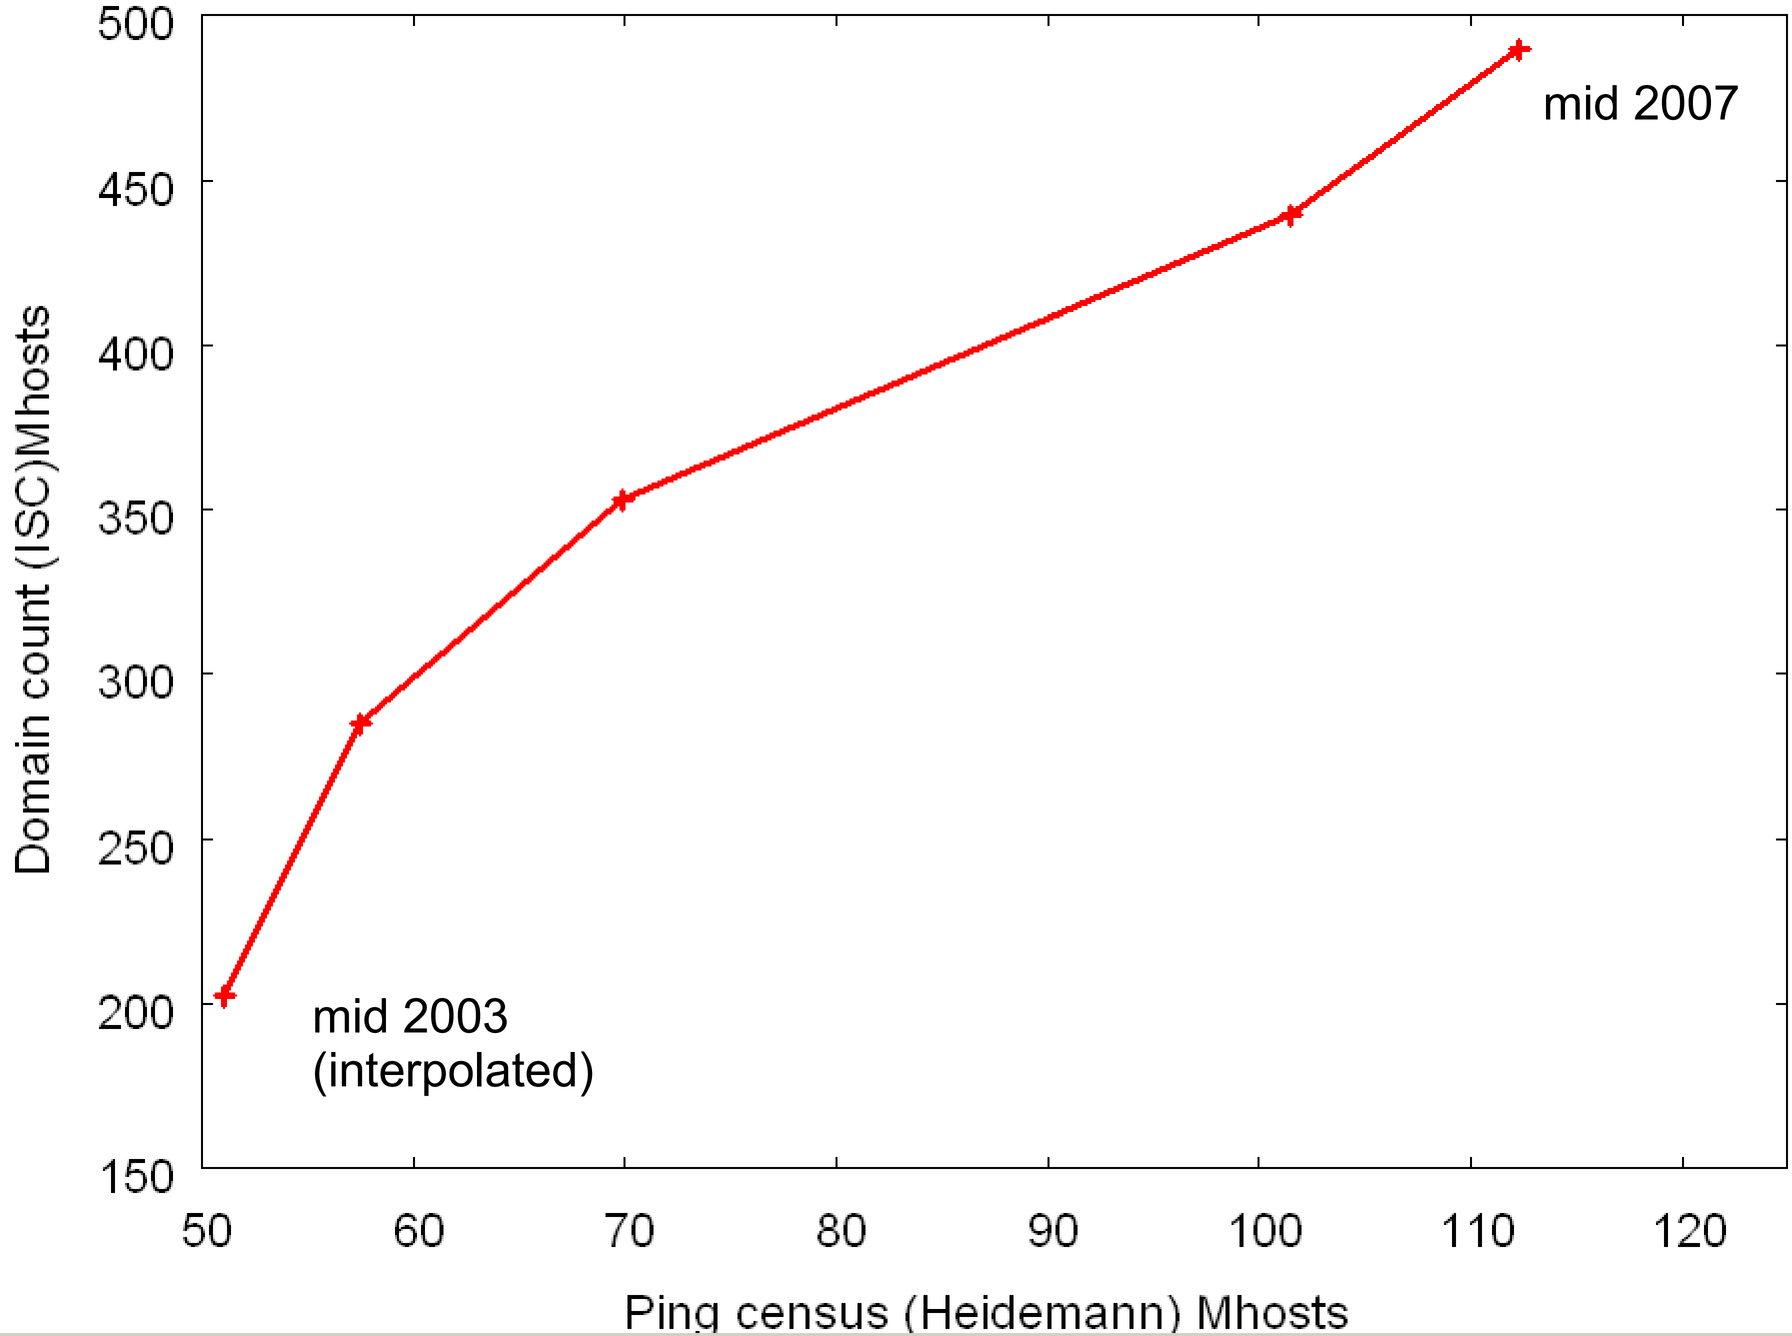

BGP4 table size vs Sqrt of ICMP census

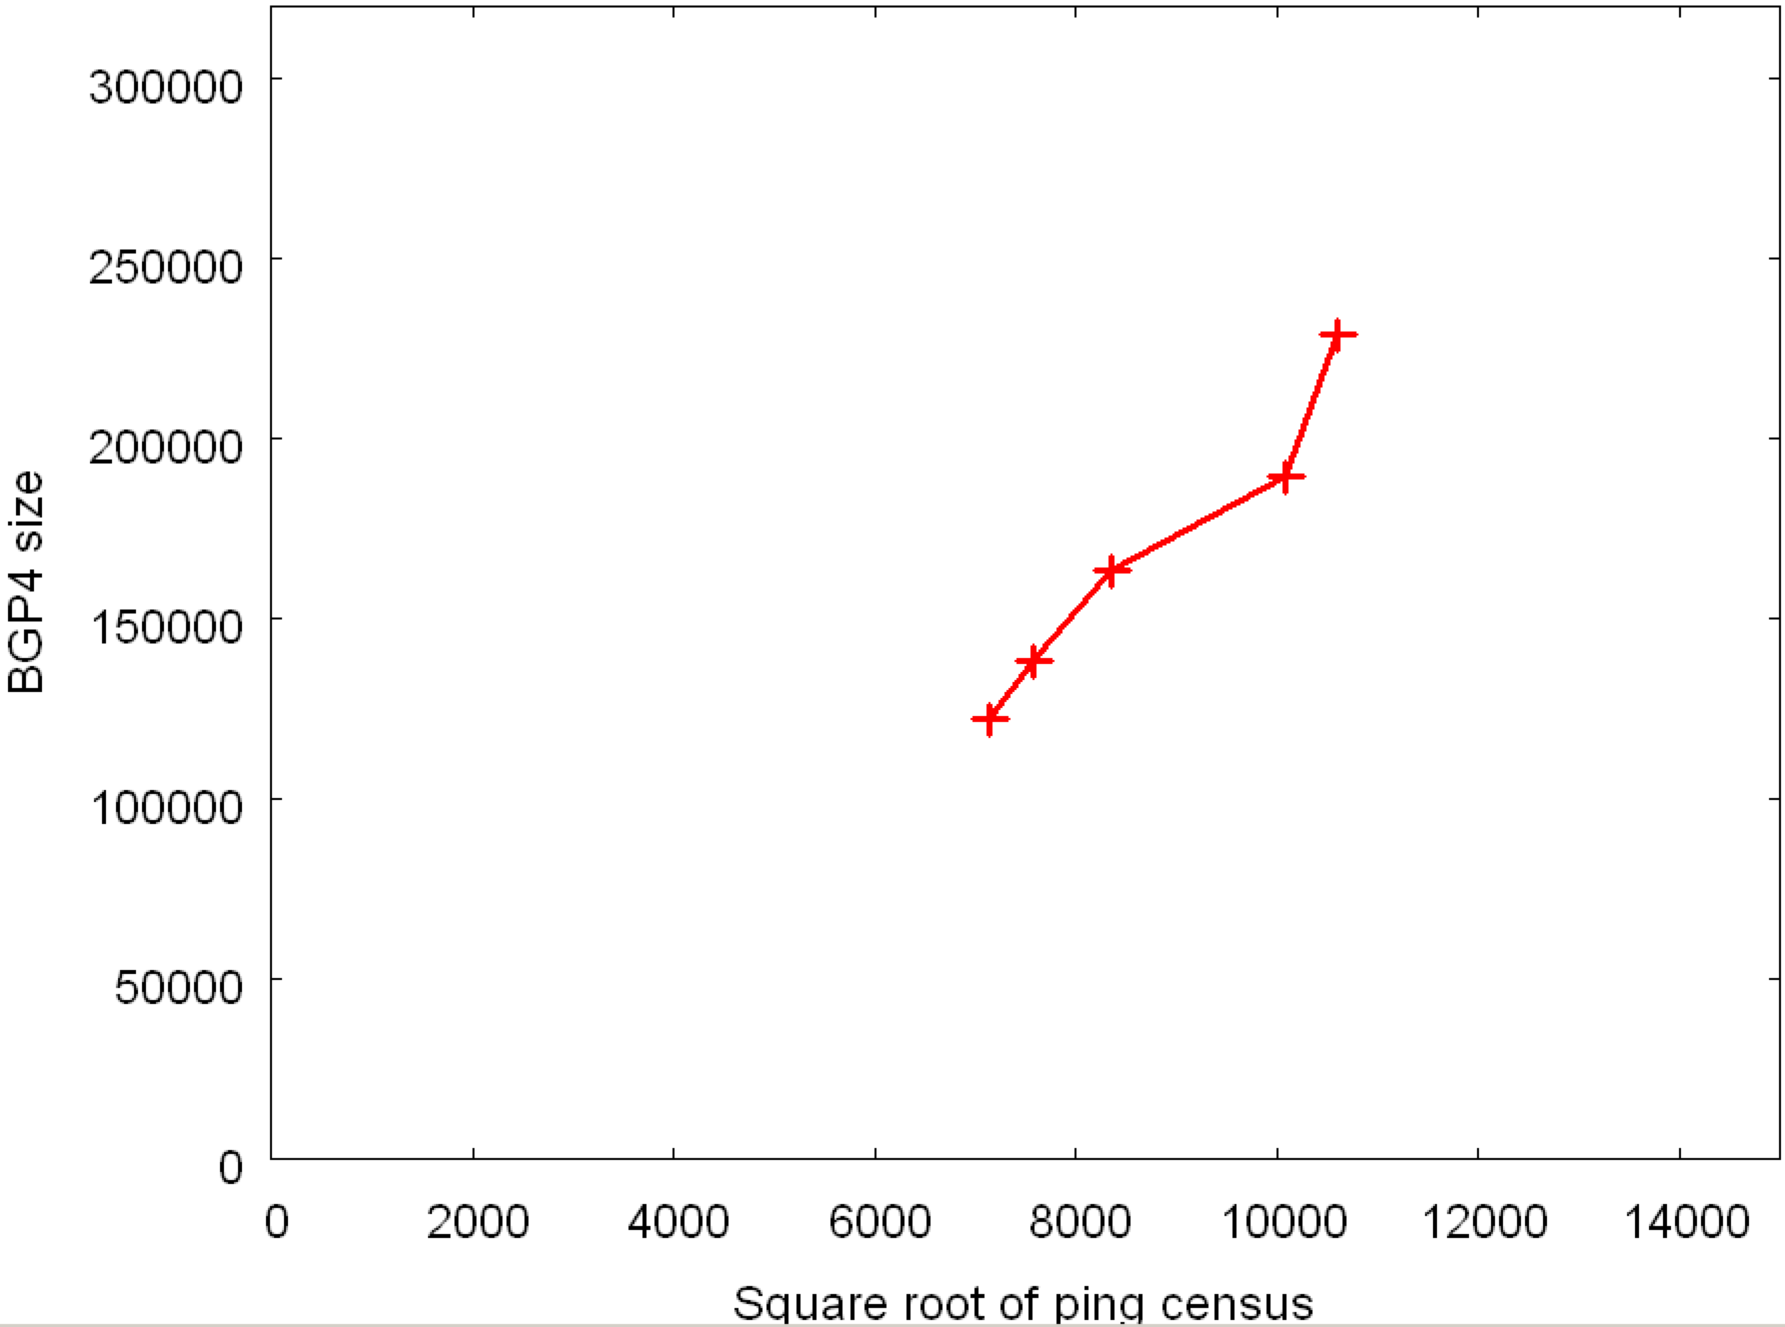

Active AS count vs Sqrt of ICMP census