Observed structure of addresses in IP traffic

Eddie Kohler, Jinyang Li, Vern Paxson, Scott Shenker
ICSI Center for Internet Research

Thanks to David Donoho and Dick Karp

Problem

- How can we model the set of destination IP addresses visible on some link? (And does it matter?)

  Example from a 4-hour trace at a university access link:

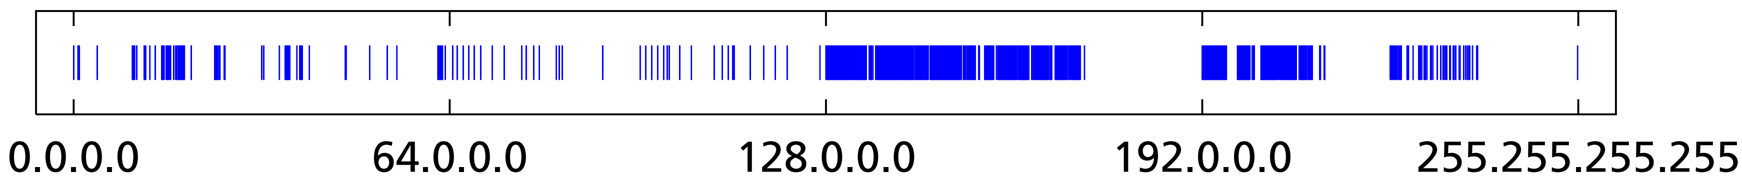

  In particular, can we model *how the addresses aggregate*?

  We call this *address structure.*

- Applications might include average-case route lookup, analysis of aggregate-based congestion control, realistic sets of addresses for simulations, . . .

Results

- Address structure dominates the characteristics of medium-scale prefix aggregates, such as /16s.
- The medium-scale aggregation behavior of real addresses is well modeled by a multifractal Cantor set construction with two parameters.

  The model captures both fractal metrics and metrics we developed for address structures.

- Address structure can serve as a site "fingerprint".

  Structural metrics differ between sites.

  At a given site, these metrics are stable over short time scales.

  New communication dynamics, such as worm propagation, show up in the metrics.

Outline

- Terminology
- Address structure and aggregate packet counts
- Model
- Metrics
- Fingerprints

Terminology

- **Active address**: an IP address visible in the trace as a destination
- N: the number of active addresses in a trace

  $N \leq 2^{32}$ by definition; $N \ll 2^{32}$ for all our traces
- p-**aggregate**: a set of addresses that share the same p-bit address prefix ($0 \leq p \leq 32$)

  Also called a $/p$

  1.0.0.0 and 1.99.130.14 are in the same /8, but different /10s
- **Active** p-**aggregate**: a $/p$ containing at least one active address

Traces

Name	Description	ΔT	# pkts	N	
U1	large university access link	$\sim$ 4 h	62M	69,196	
U2	large university access link	$\sim$ 1 h	101M	144,244	
A1	ISP	$\sim$ 0.6 h	34M	82,678	
A2	ISP	1 h	29M	154,921	
R1	link from regional ISP	1 h	1.5M	168,318	§
R2	link from regional ISP	2 h	1M	110,783	§
W1	large Web site access link	$\sim$ 2 h	5M	124,454	

- Collected between 1998 and 2001

  Most anonymized while preserving prefix and class relationships

  § means sampled (1 in 256)

Does address structure matter?

- Assume that aggregate packet counts matter.

  Accounting, fairness, congestion control . . .

- What factors affect aggregate packet counts?

  Packet counts per address: probably a heavy-tailed distribution

  Addresses per aggregate = *address structure*

  Correlation

- Analyze the contributions of these factors to an observed packet count distribution

  Medium scales are most interesting (/16s and thereabouts)

R1 packet count distributions

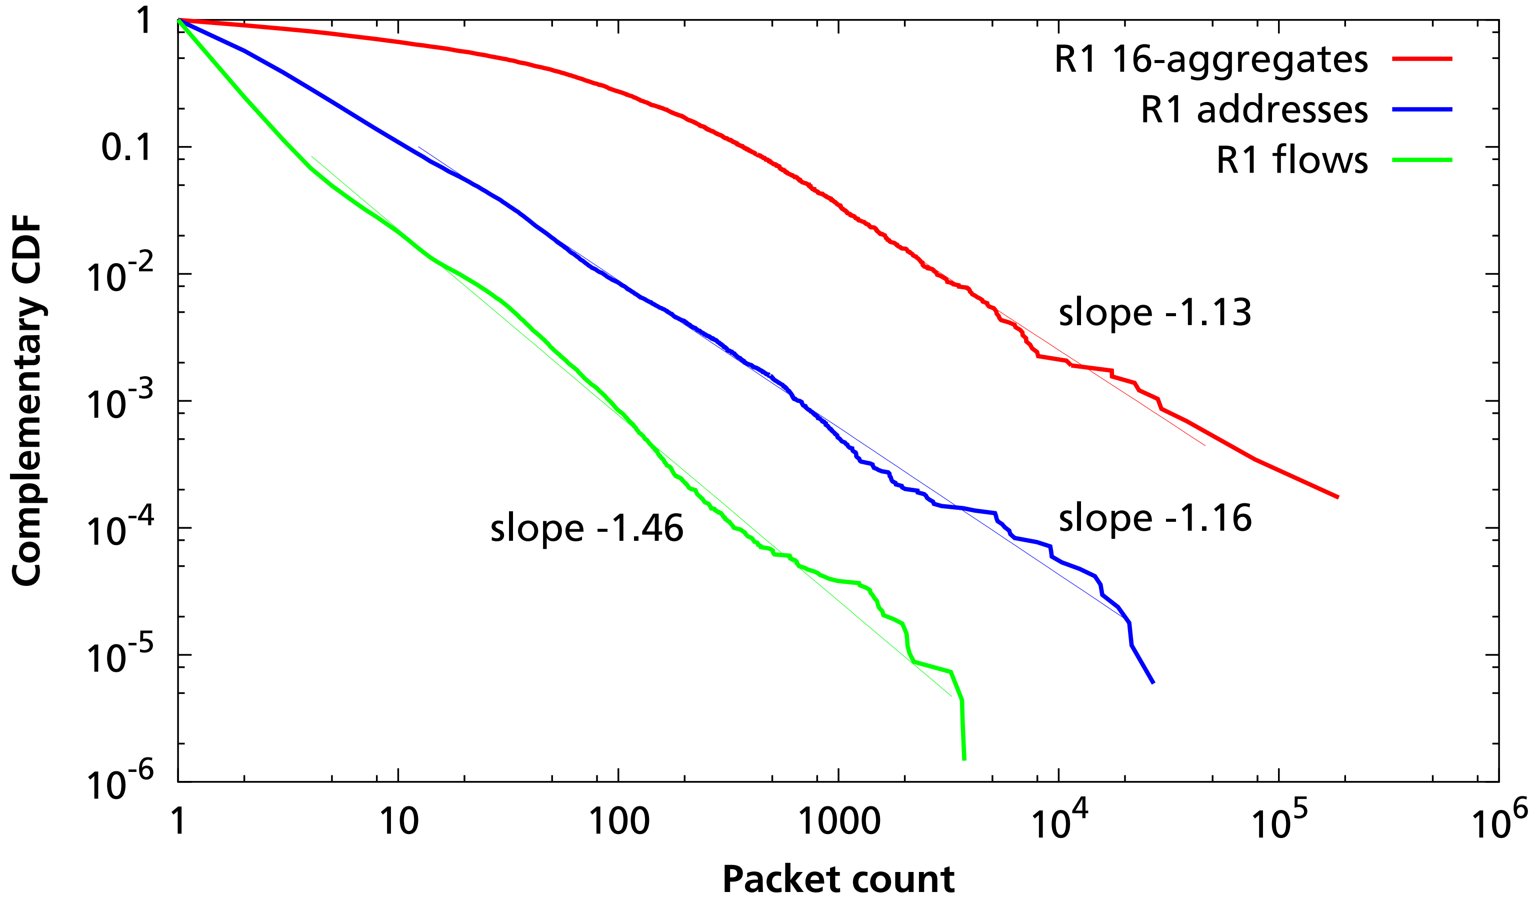

Semi-experiments

- Manipulate the data, destroying one factor at a time; see which factors impact aggregate packet counts
- "Random counts": destroy per-address packet counts

  Replace the (heavy-tailed) per-address packet count distribution with a uniform distribution over $[0, 17.54]$
- "Random addresses": destroy address structure

  Replace address structure with a uniform random distribution over the entire IP address space
- "Permuted counts": destroy correlation

  Permute per-address packet counts among the active addresses

Address structure matters most

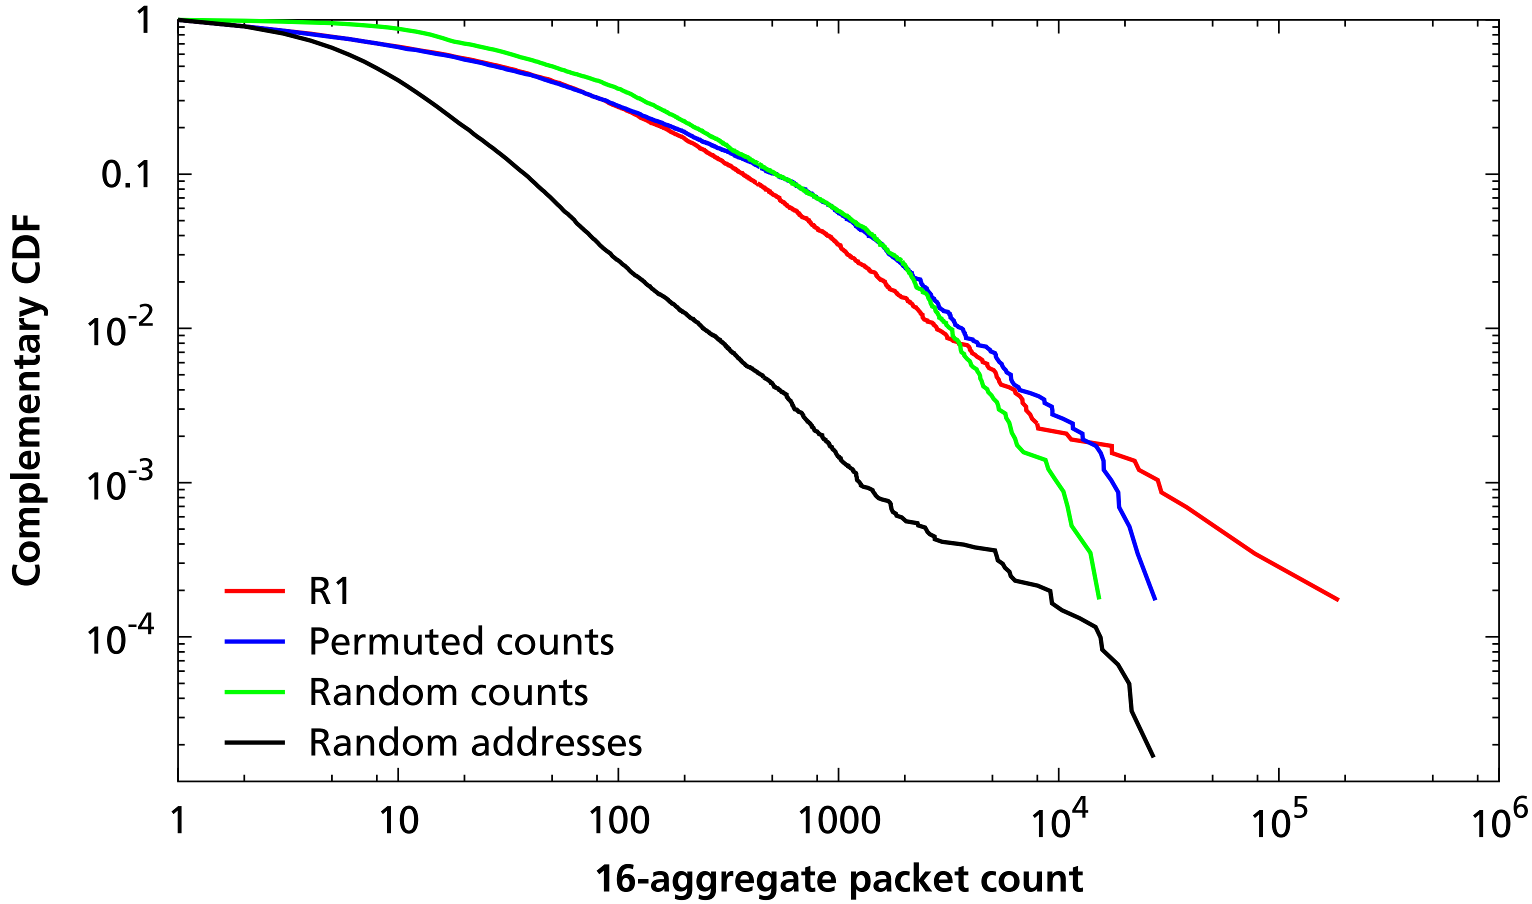

Tour of U1's address structure

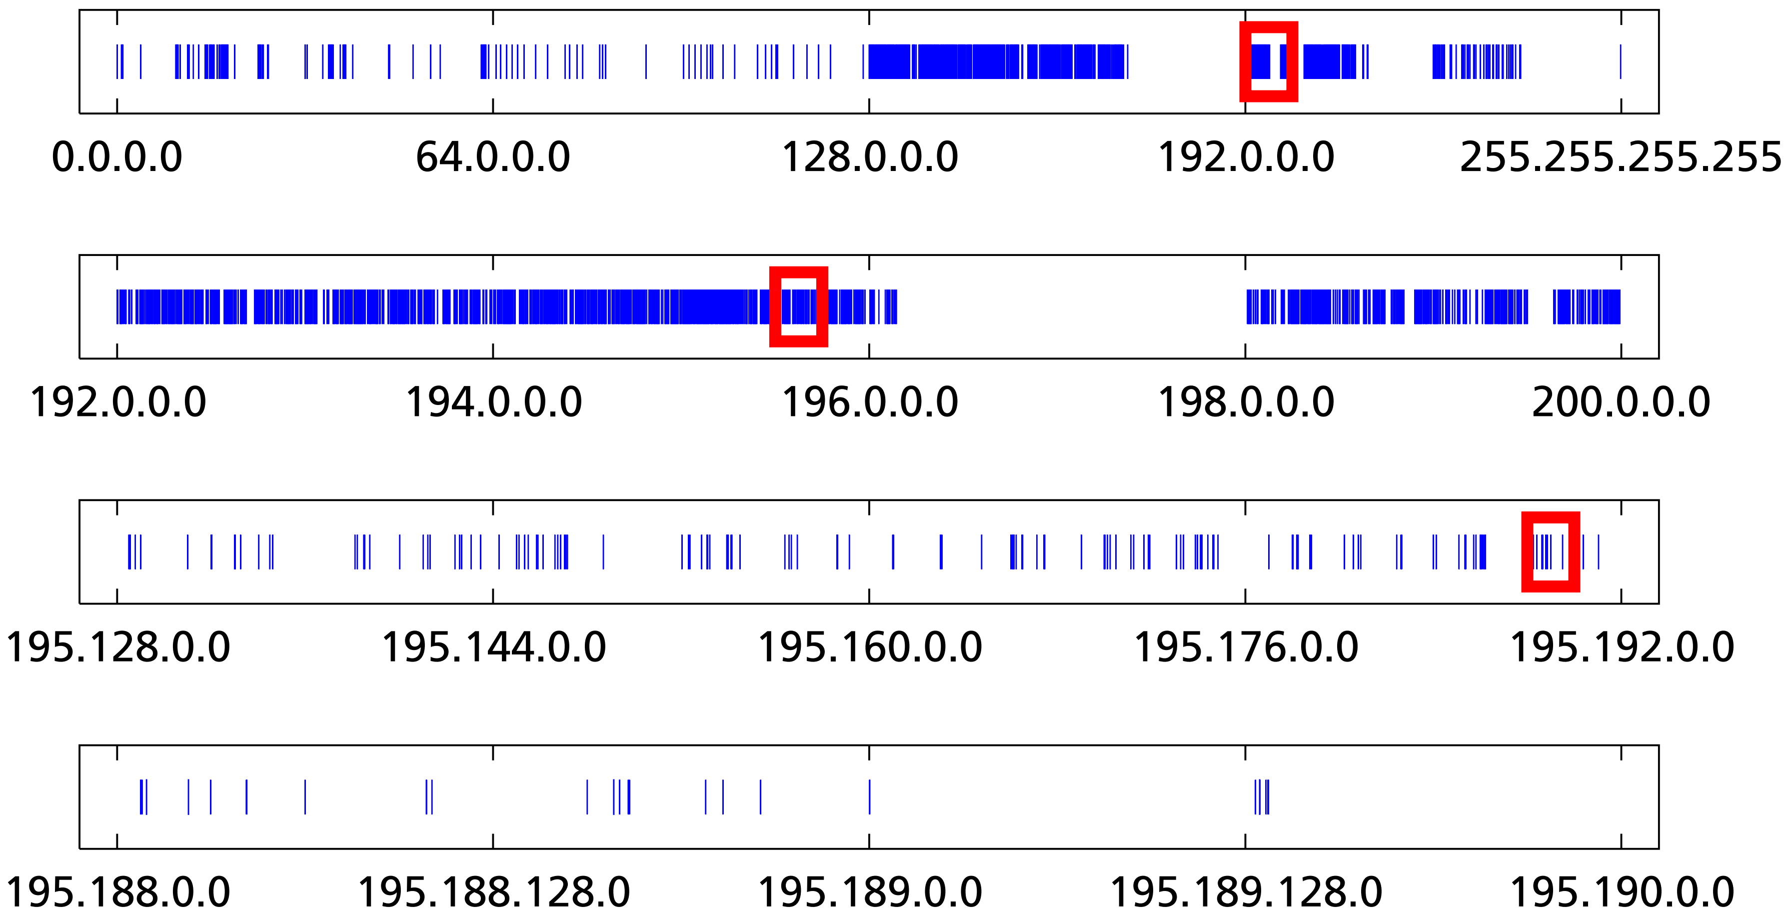

Self-similarity?

- **Interesting structure all the way down**

    Visually “self-similar” characteristics

- **Might address structure be usefully modeled by a fractal?**

    Treat an address structure as a subset of the unit interval

    Fractal dimension $D \in [0, 1]$?

Fractal dimension for address structure

- Use *lattice box-counting dimension*

  Corresponds nicely to prefix aggregation

- Let n_p equal the number of active $/p$s in a trace

  $n_{32} = N$

  $n_p \leq n_{p+1} \leq 2n_p$

  each $/p$ contains and is covered by 2 disjoint $/(p+1)$s

- Then $D = \lim_{p \to \infty} \frac{\log n_p}{p \log 2}$

  But $p \leq 32$ here, and expect sampling effects for high p

  Examine medium p to see if the limit exists

$\log n_p$ is linearly related to p at medium scales

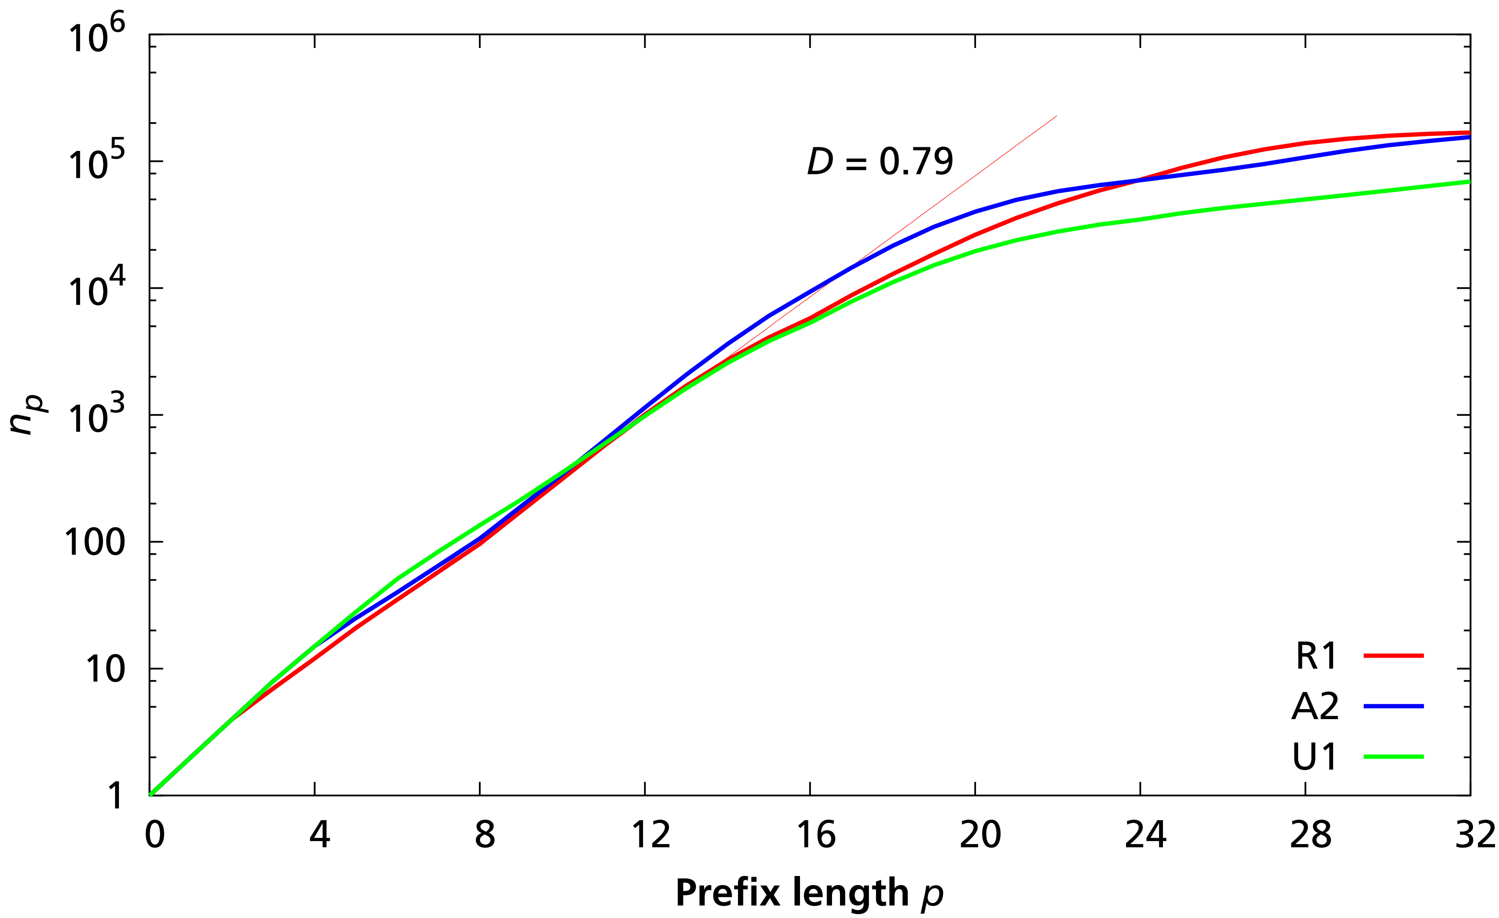

Multifractality

- **Monofractal may not be sufficient**
  - Same scaling behavior everywhere
  - Not what we saw in the tour
- **Examine the *multifractal spectrum* to test for multifractality (different local scaling behavior)**
  - Binned approximation (Histogram Method)
  - If multifractal, spectrum will cover a wide range of scaling exponents

Address structure is multifractal at /16

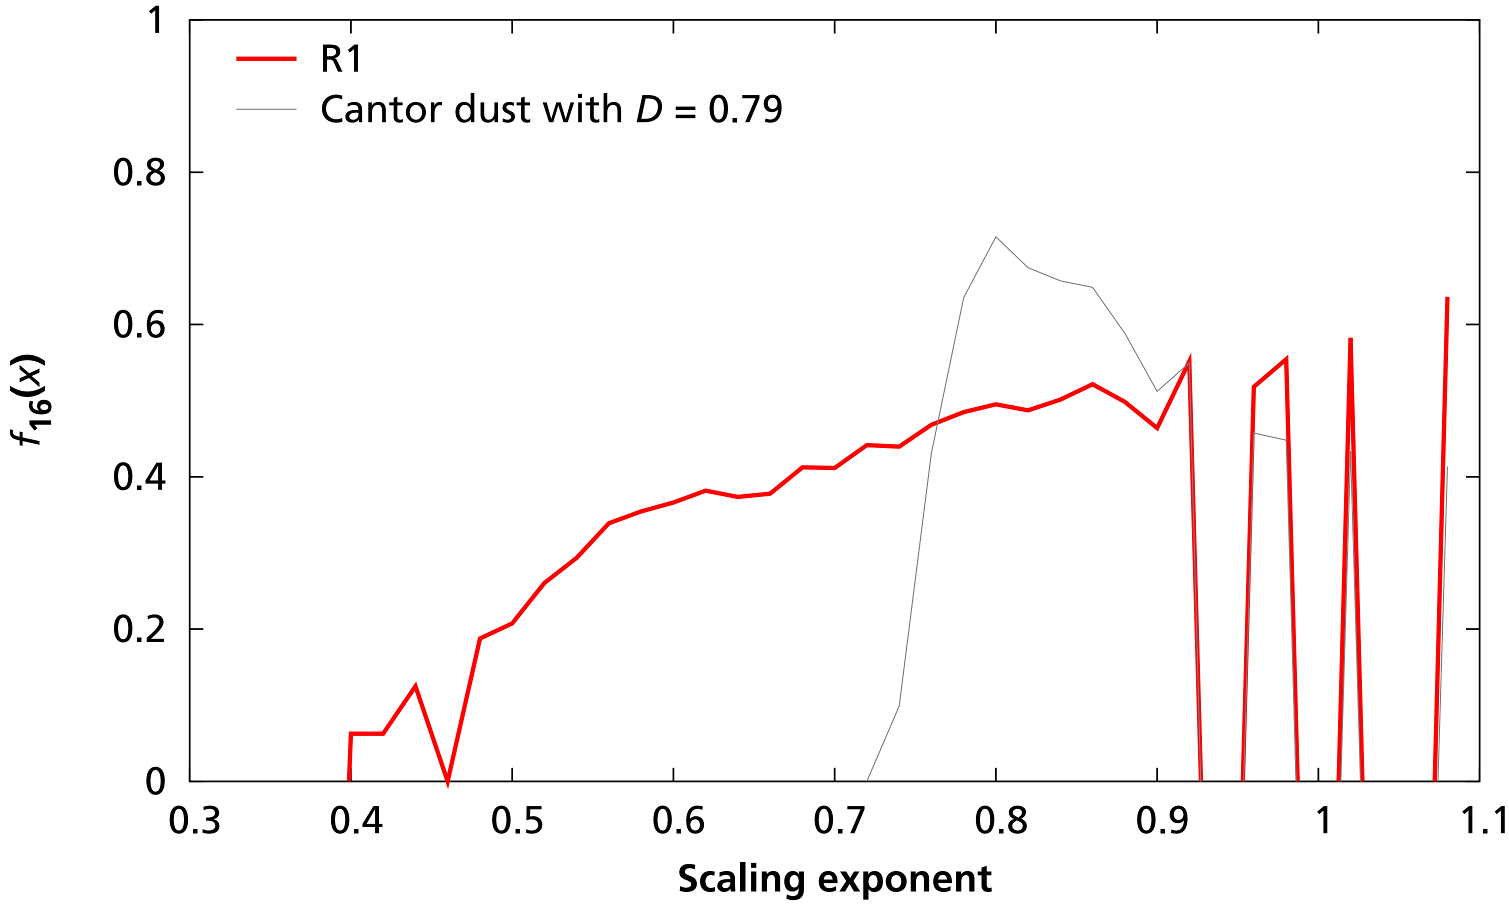

Multifractal model

- Make a multifractal Cantor measure matching this spectrum
- Start with a Cantor dust with dimension D

  Repeatedly remove middle subinterval with proportion $h = 1 - 2^{1-1/D}$

- Sample unequally from left and right subintervals

  Distribute a unit of "mass" between subintervals; left gets m_0, middle gets $\mathbf{0}$ (removed), right gets $m_2 = 1 - m_0$

  Produces a sequence of measures μ_k that weakly converge to μ

  Sample an address with probability equal to its measure

  Result: different local scaling behavior

The model fits well

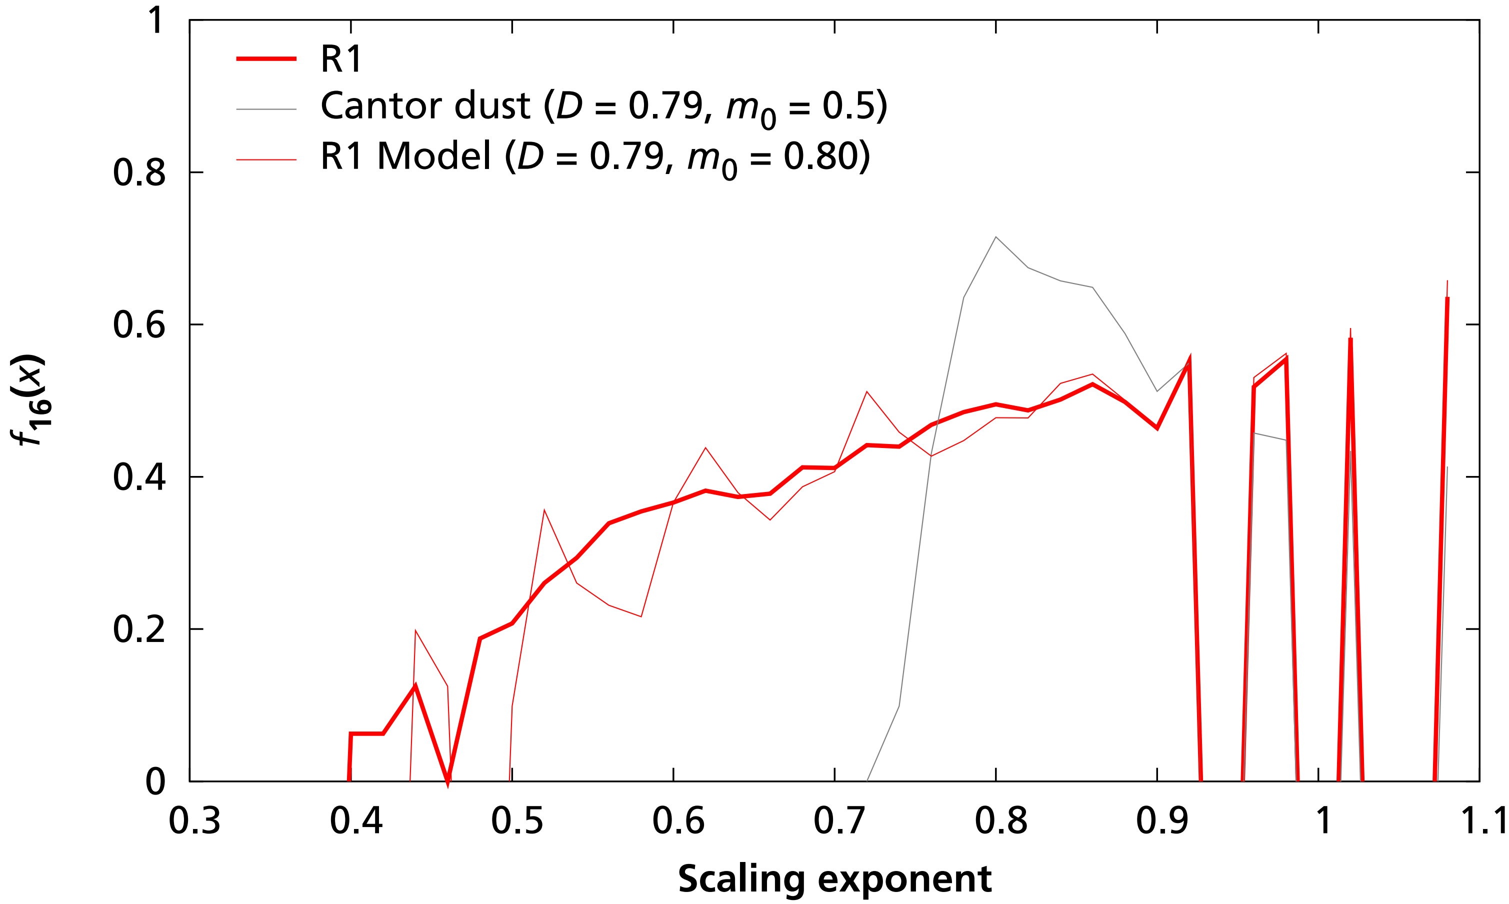

The model fits well

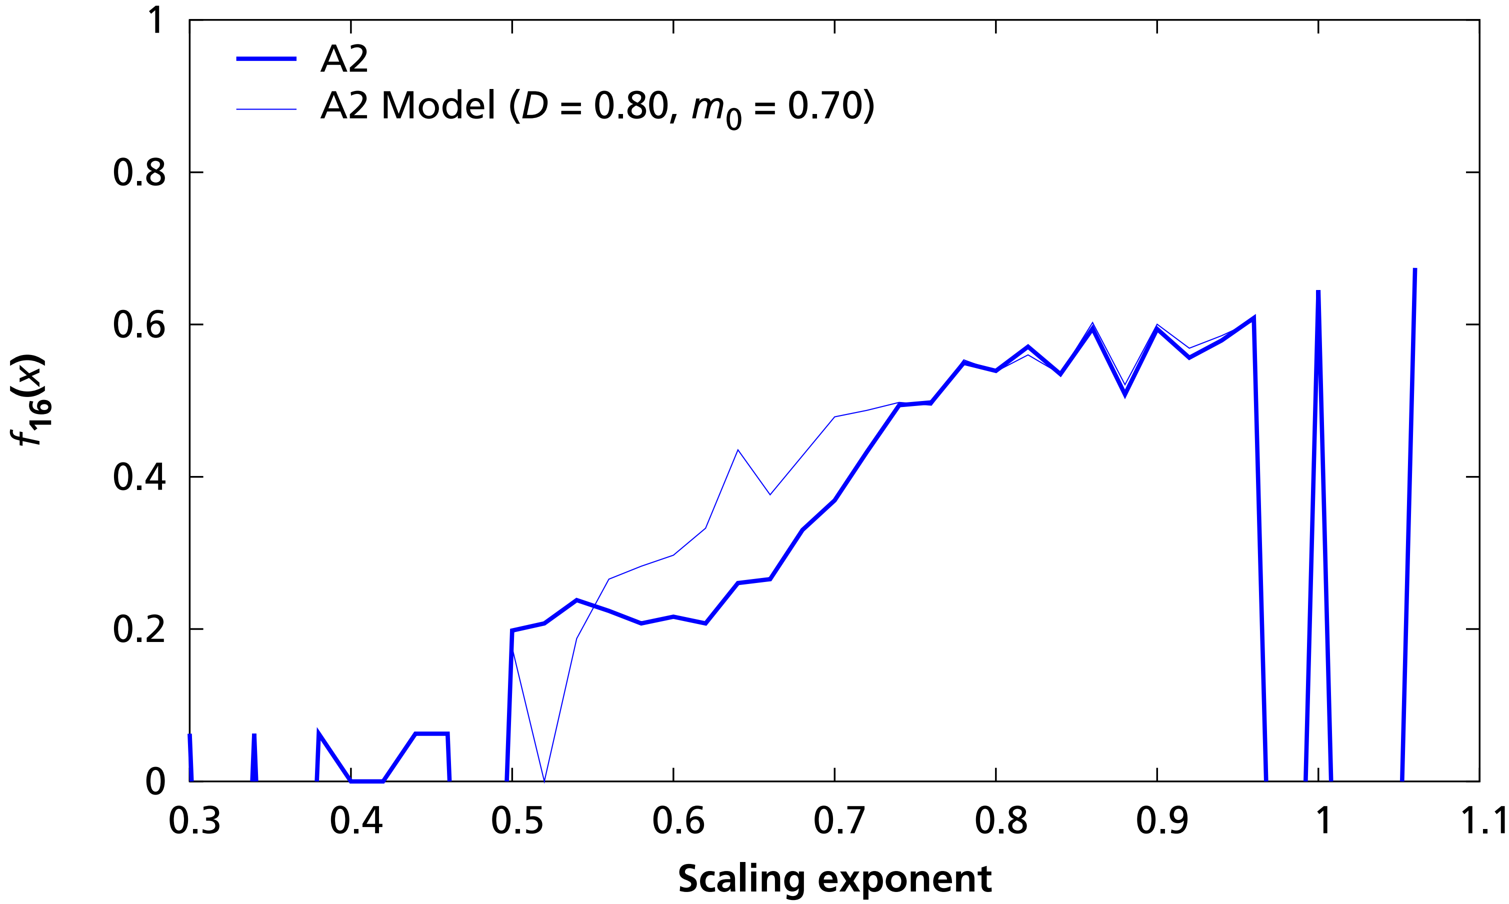

Why multifractal?

- Perhaps it's due to a cascade

  Recursive subdivision plus a rule for distributing mass

- For example, address allocation

  Pure speculation!

  ICANN allocates short prefixes to providers

  Providers allocate longer prefixes to their customers

  All parties might allocate basically from left to right

Does the multifractal spectrum matter?

- Certainly the model doesn't *look* like real data:

U1

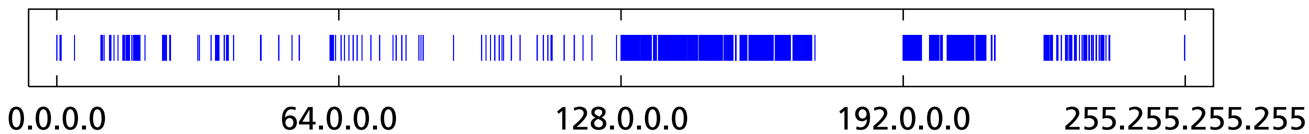

U1 Model

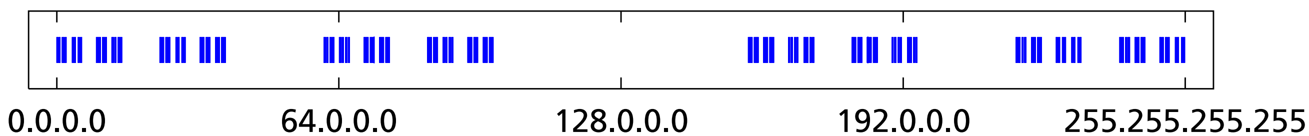

  How do we know whether we've captured relevant properties?

- Develop application metrics for address structures

  Contrast metrics among traces

  Compare with model

Active aggregate counts: n_p and γ_p

- n_p again equals the number of active $/p$s in a trace
- n_p measures how densely addresses are packed

  If $N = 2^{16}$ and $n_{16} = 1$, addresses are closely packed

  If $N = 2^{16}$ and $n_{16} = 2^{16}$, addresses are well spread out

  Useful for algorithms keeping track of aggregates—shows how many aggregates there tend to be

- $\gamma_p = n_{p+1}/n_p$ more convenient for graphs

  $N = \prod_{1 \leq p < 32} \gamma_p$

γ_p

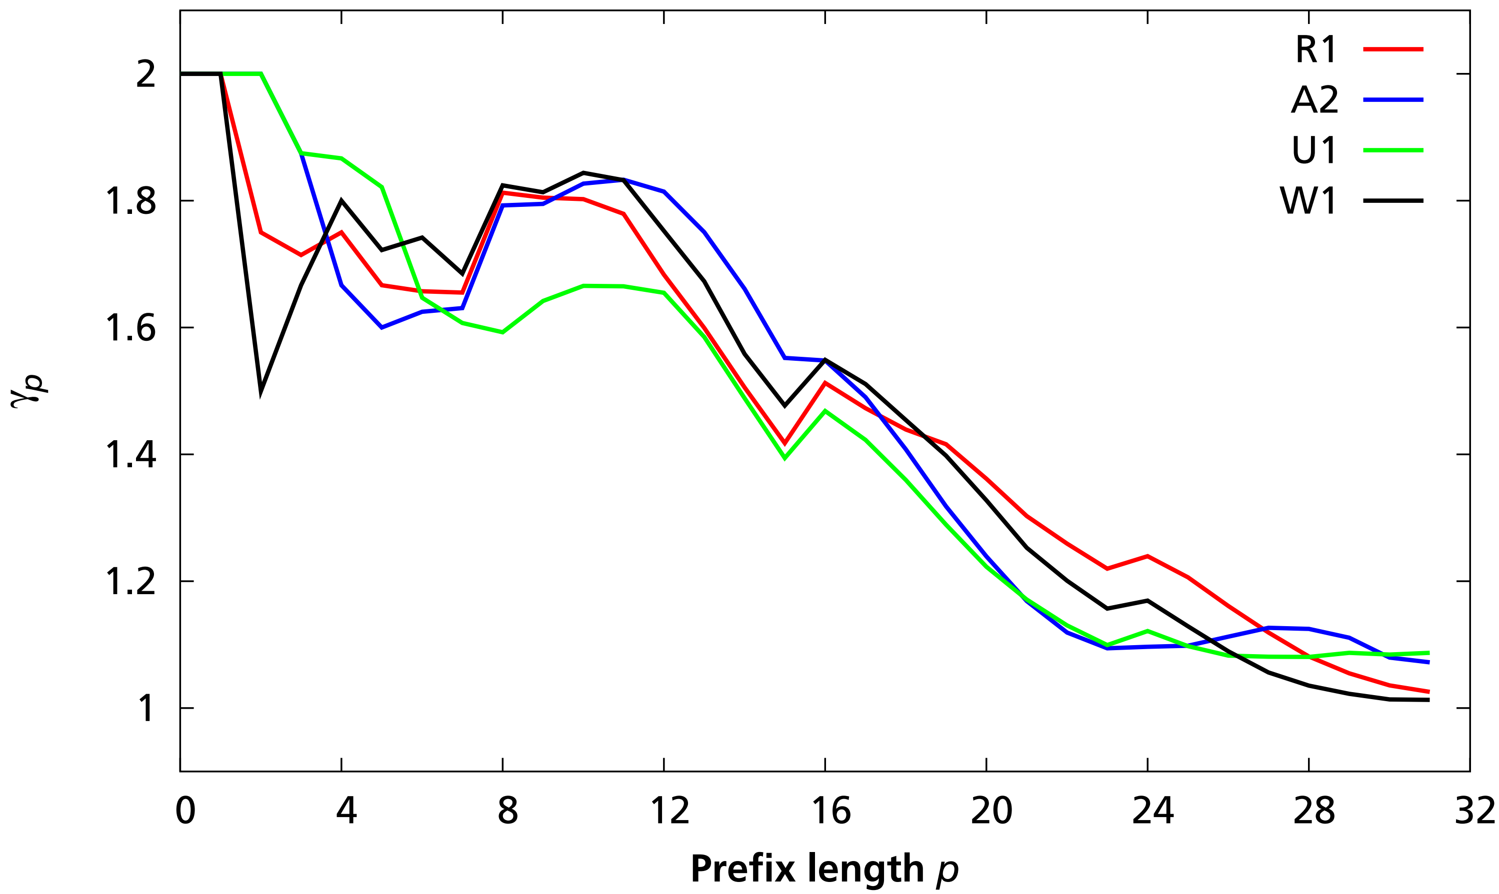

Models' γ_p

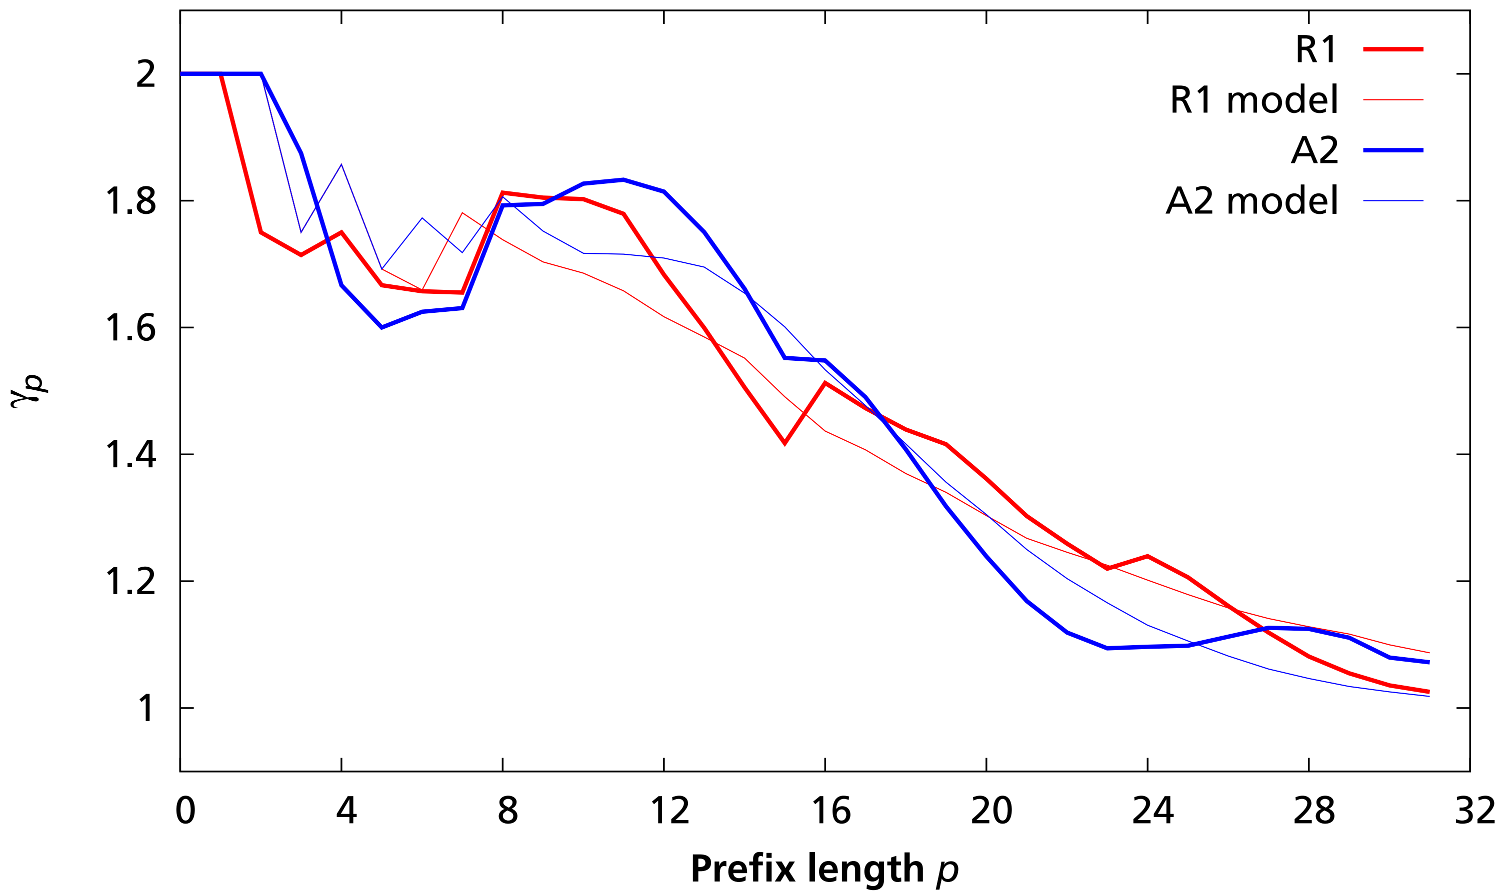

Discriminating prefixes

- The **discriminating prefix** of an active address, *a*, is the prefix length of the *largest* aggregate that contains *only one* active address, namely *a*.

  Example with 4-bit addresses:

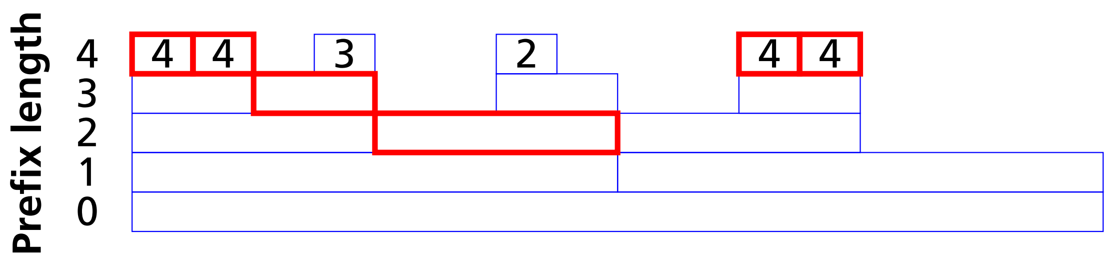

- Measures address *separation*

  If many addresses have d.p. < **20**, say, then addresses are well separated

  How depopulated do aggregates become?

Discriminating prefixes: π_p

- Let π_p equal the number of addresses with d.p. p

  $\sum \pi_p = N$

  Turns discriminating prefixes into a metric

π_p

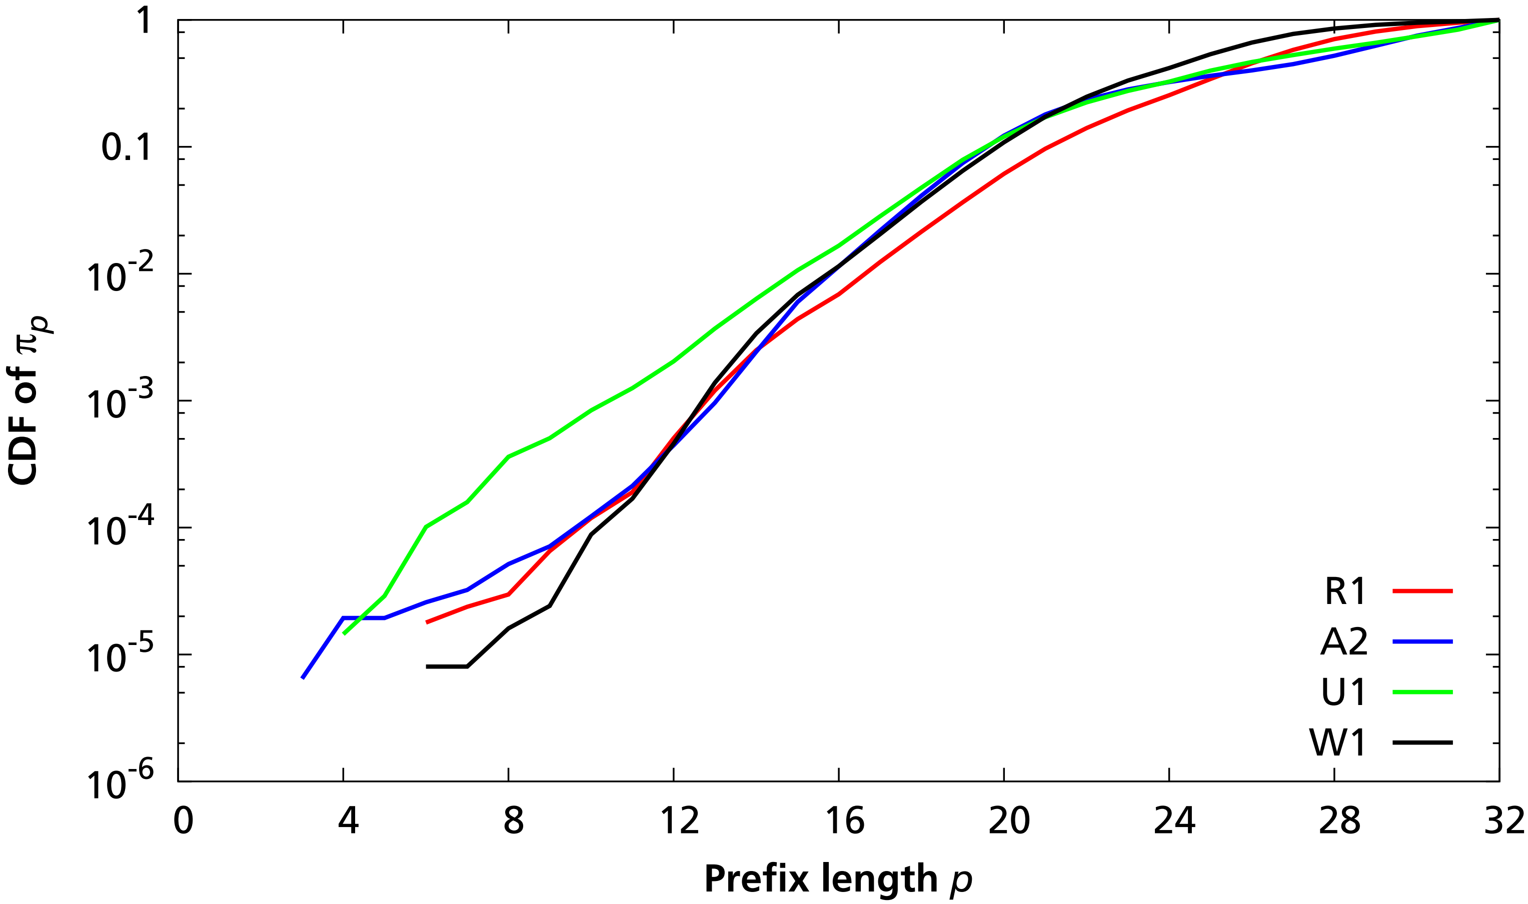

1
0.1
10-2
10-3
10-4
10-5
10-6
CDF of πp
0
4
8
12
16
20
24
28
32
Prefix length p
R1
A2
U1
W1

Models' π_p

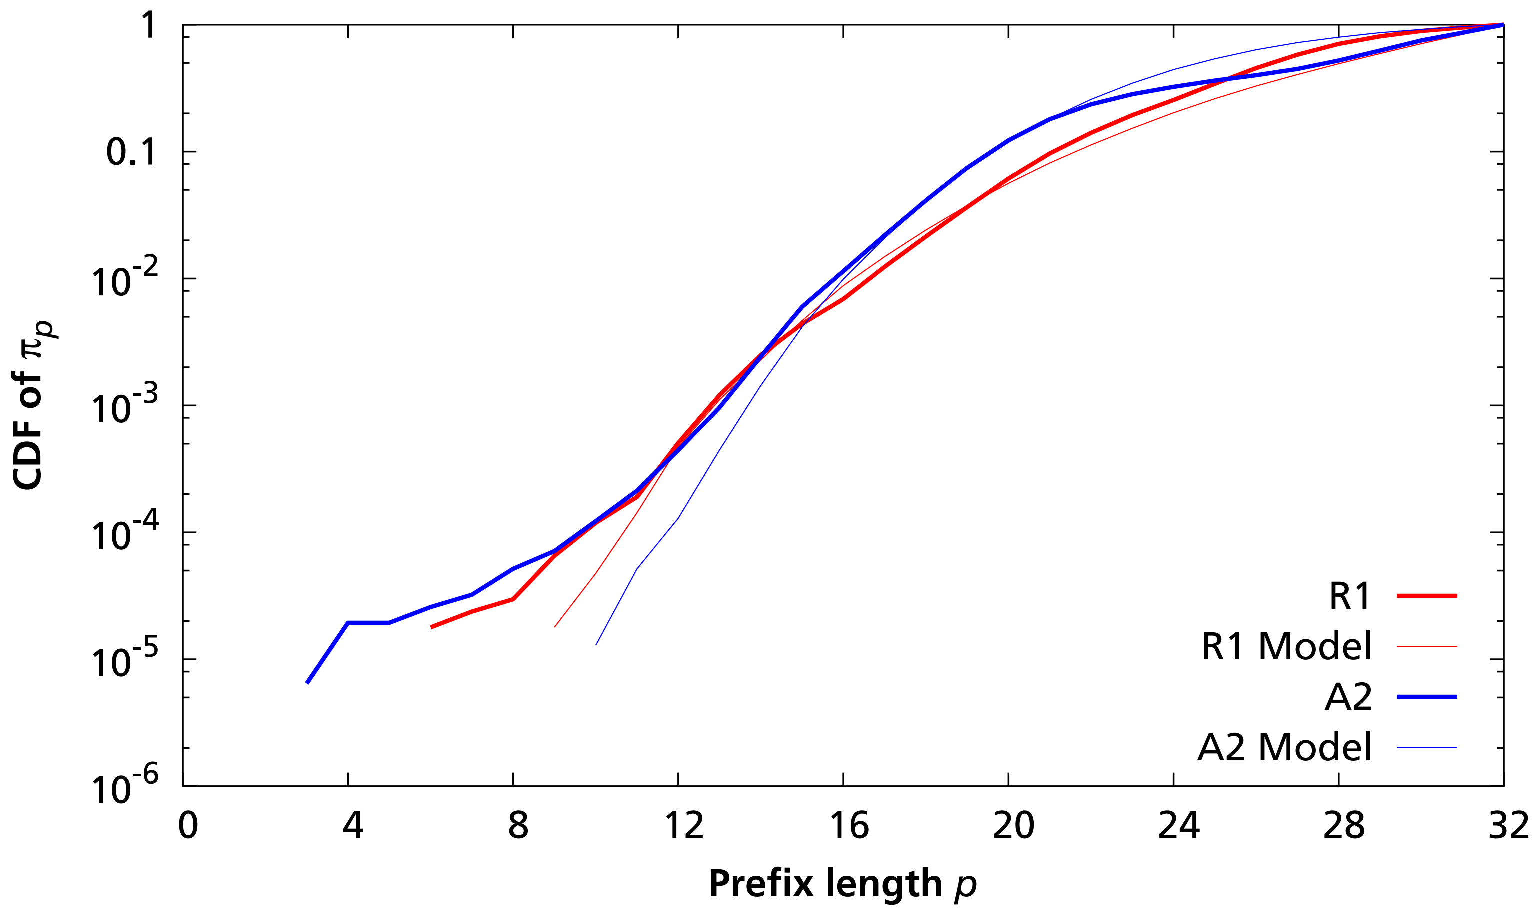

CDF of π_p
1
0.1
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
0
4
8
12
16
20
24
28
32
Prefix length p
R1
R1 Model
A2
A2 Model

Aggregate population distribution

- Like aggregate packet count distribution, but count the number of *active addresses* per aggregate
    - Expect a wide range of variation, just as with the other metrics

Aggregate population distribution

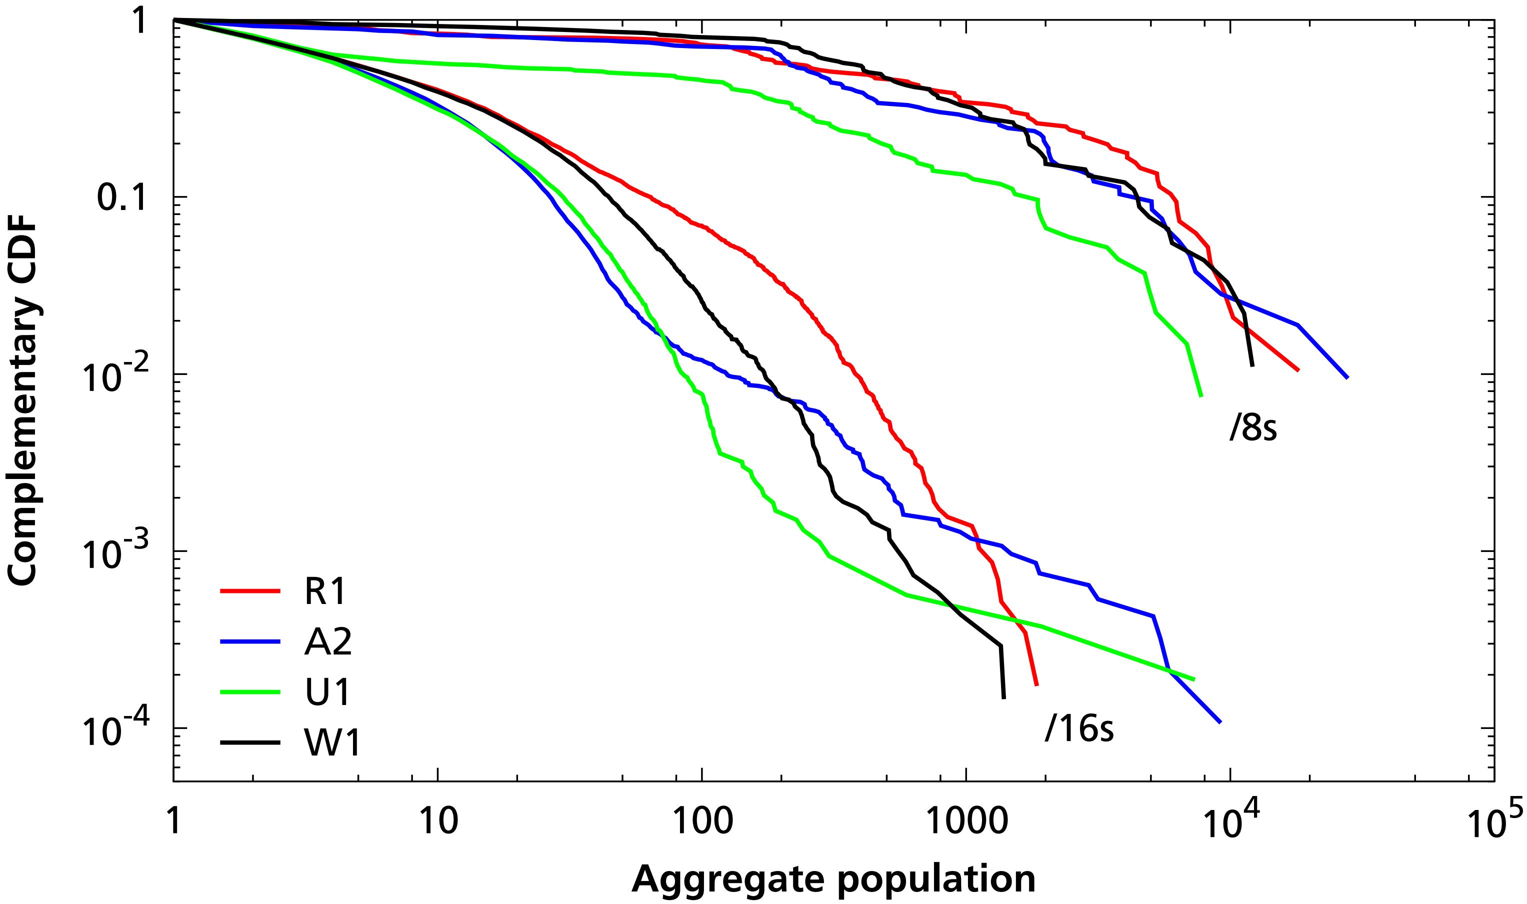

Models’ aggregate population distribution

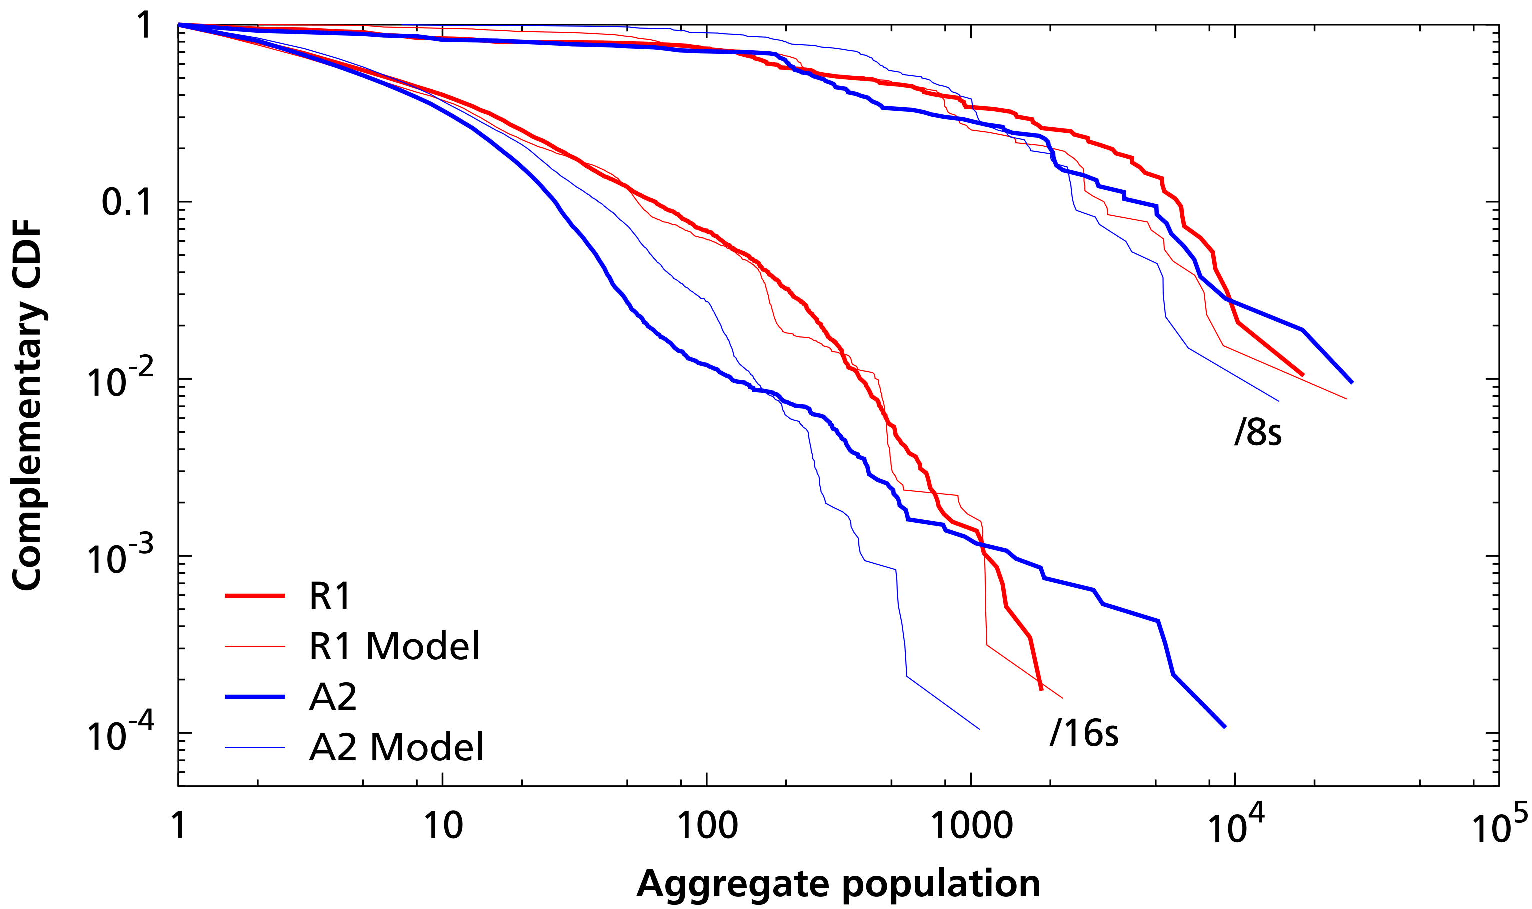

A tough metric

- The model for A2 doesn't match A2's aggregate populations

  R1, W1 match well, A2, U1 do not

  Significant aggregation in A2, U1 at long prefixes . . . ?

- Aggregate population distribution is difficult to match

- Consider random allocation constrained to match γ_p and π_p exactly

  Heck, match "generalized discriminating prefixes"—d.p.s for aggregates—as well

  Call this the "Match-DP" model

  How well does this do?

Match-DP fails aggregate population distribution

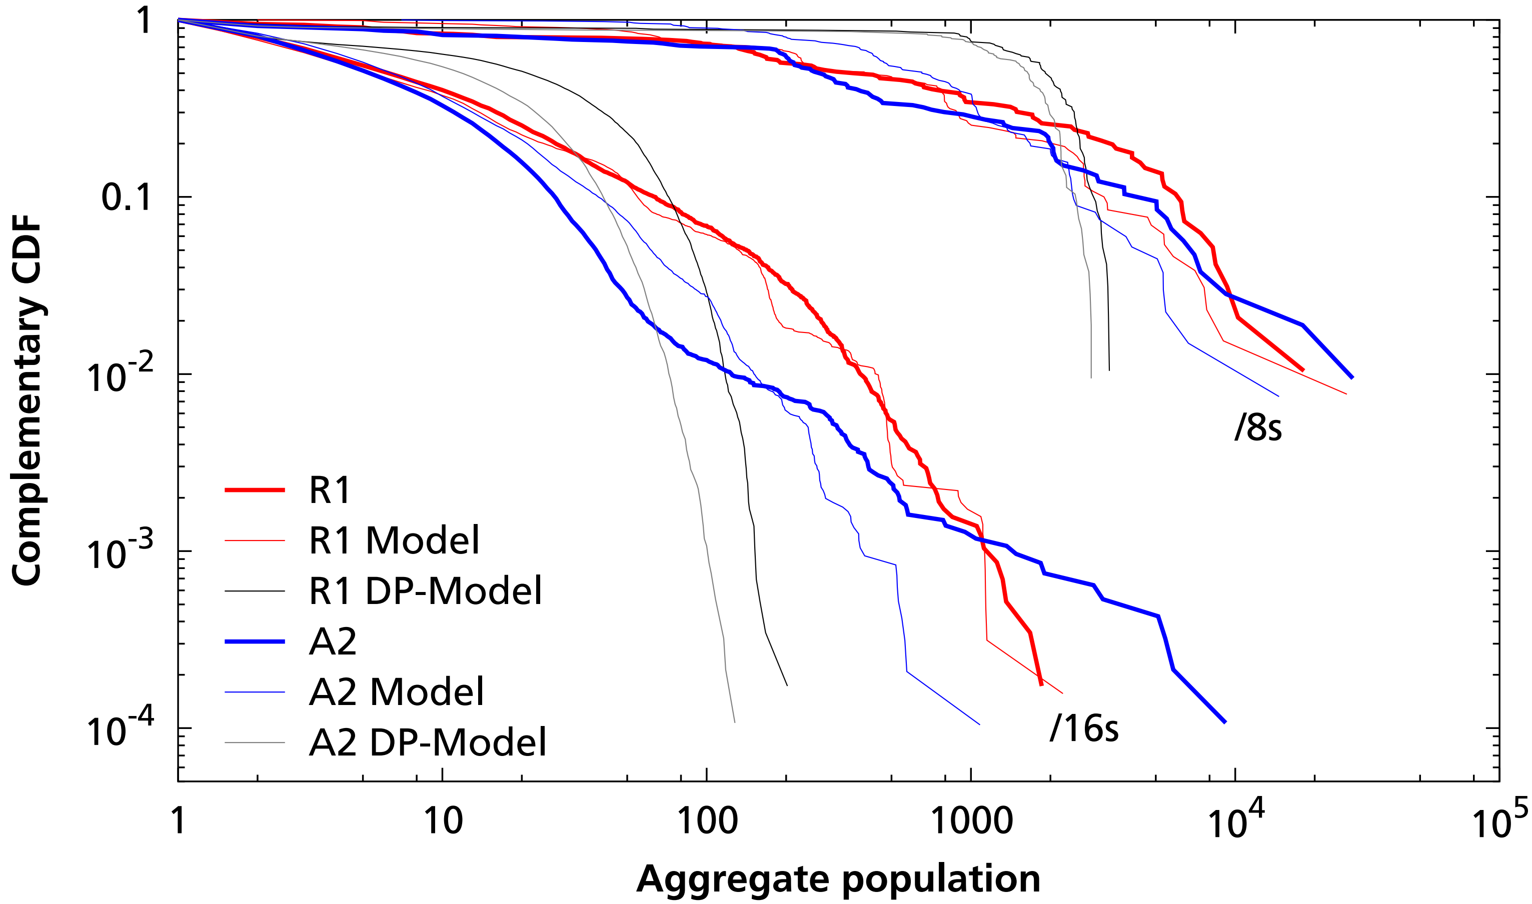

Another tough metric: The multifractal spectrum

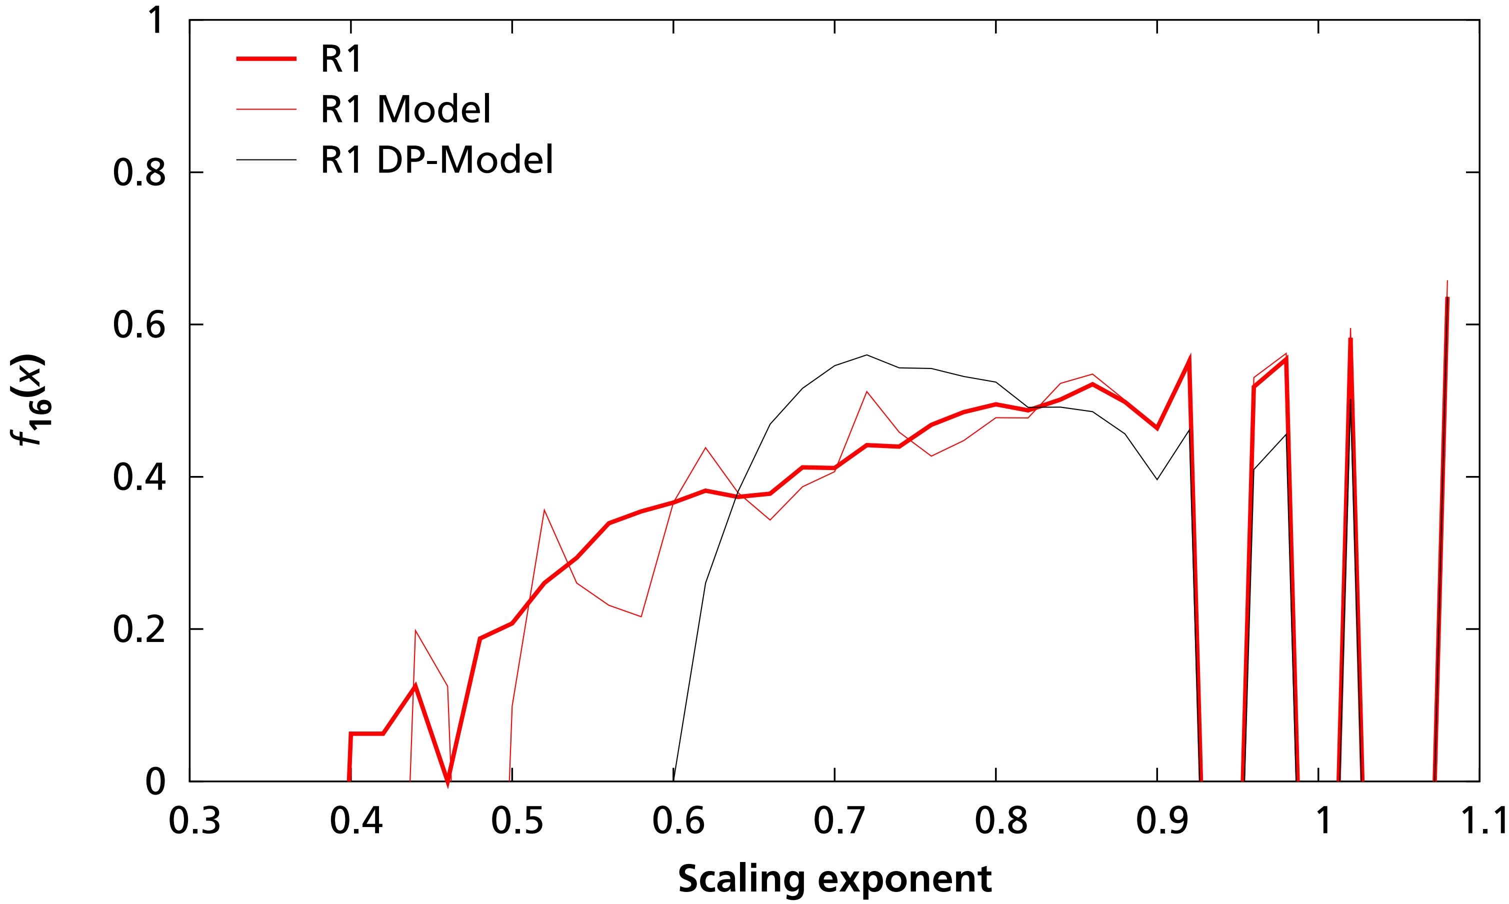

Properties of γ_p: Sampling effects?

- Turn from the multifractal model to properties of our γ_p metric
- First: Is γ_p dominated by sampling effects?

  N is effectively a sample size

  How does the shape of the γ_p curve depend on N?
- Plot γ_p for longer and shorter sections of trace U1

  24 hours $\rightarrow$ 6 minutes; $N = 161{,}560 \rightarrow 11{,}838$

Shape of γ_p similar for wide range of sample sizes

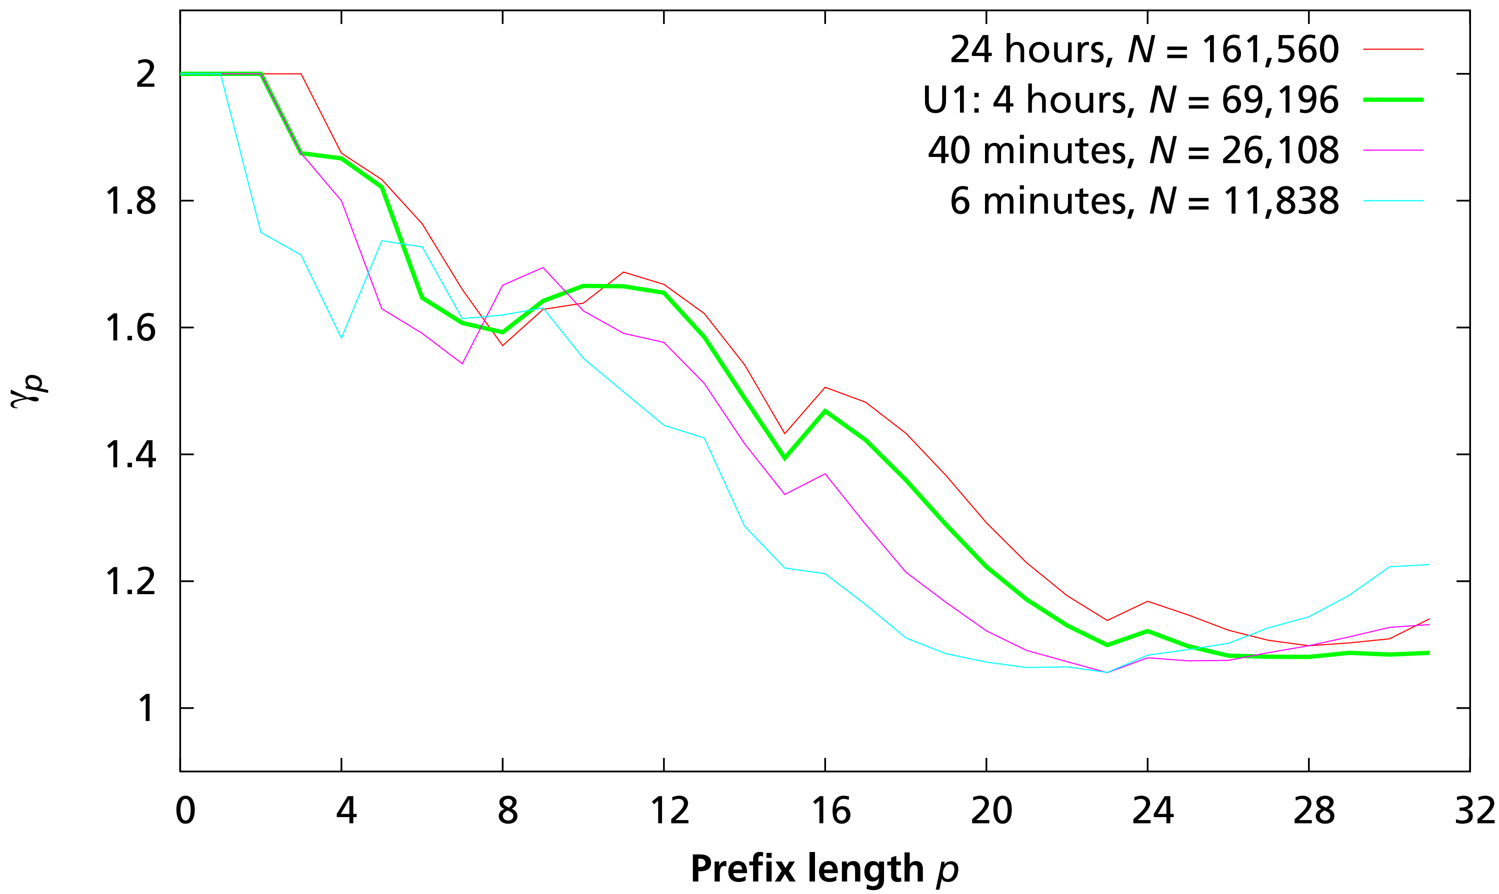

Shape of γ_p similar for wide range of sample sizes

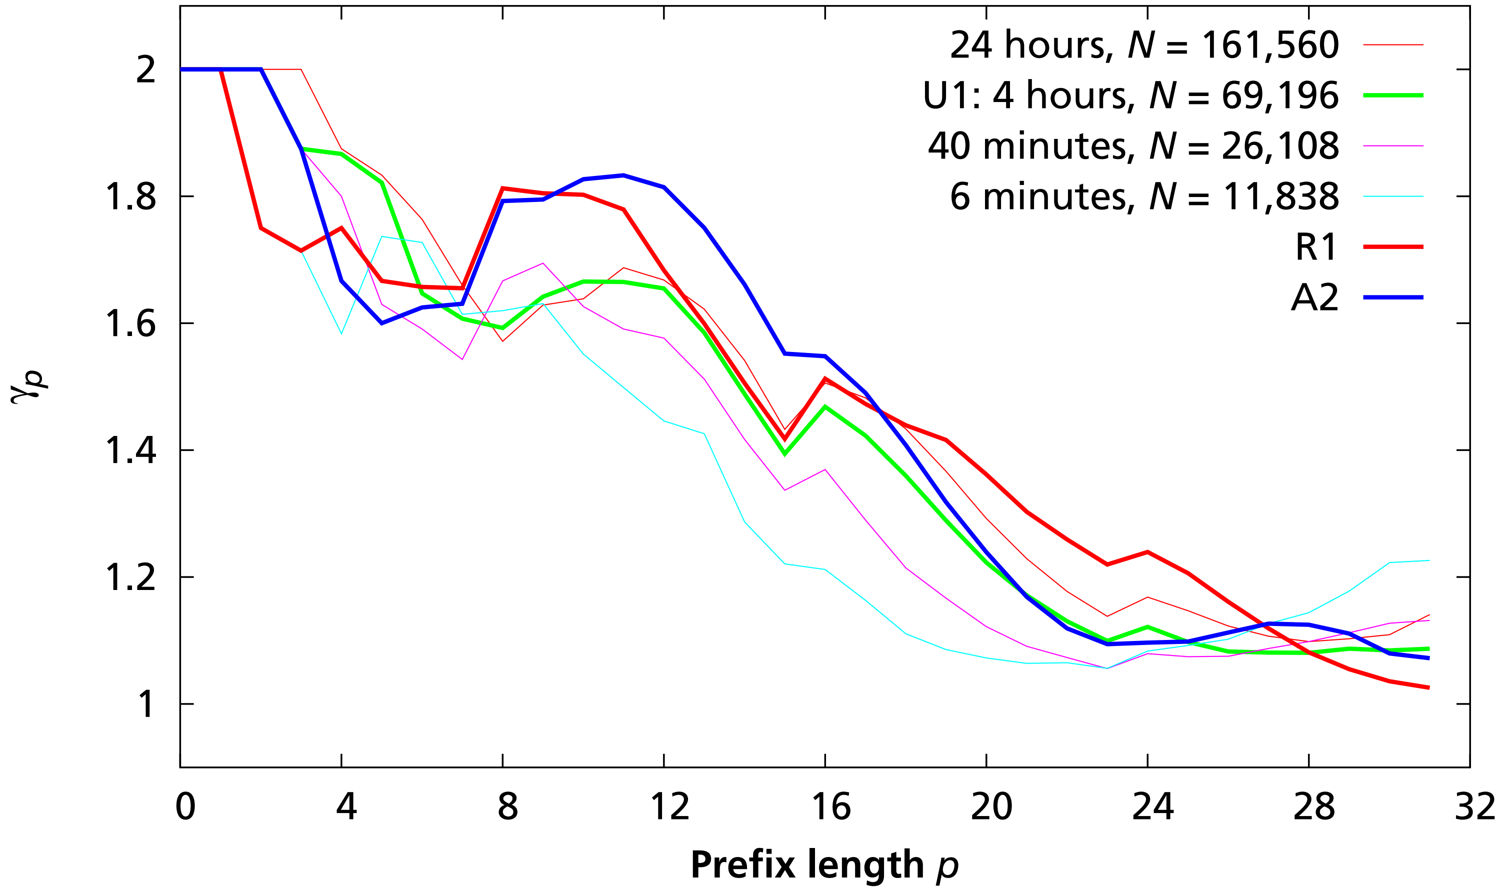

Short-term stability?

- Is γ_p stable over short time scales?
- Divide traces into short sections, each with $N = 32{,}768$

  Plot maximum, minimum, and mean γ_p over all sections

  R1, A2, and U2; sections last about 6–7 minutes each

Shape of γ_p relatively stable over short time scales

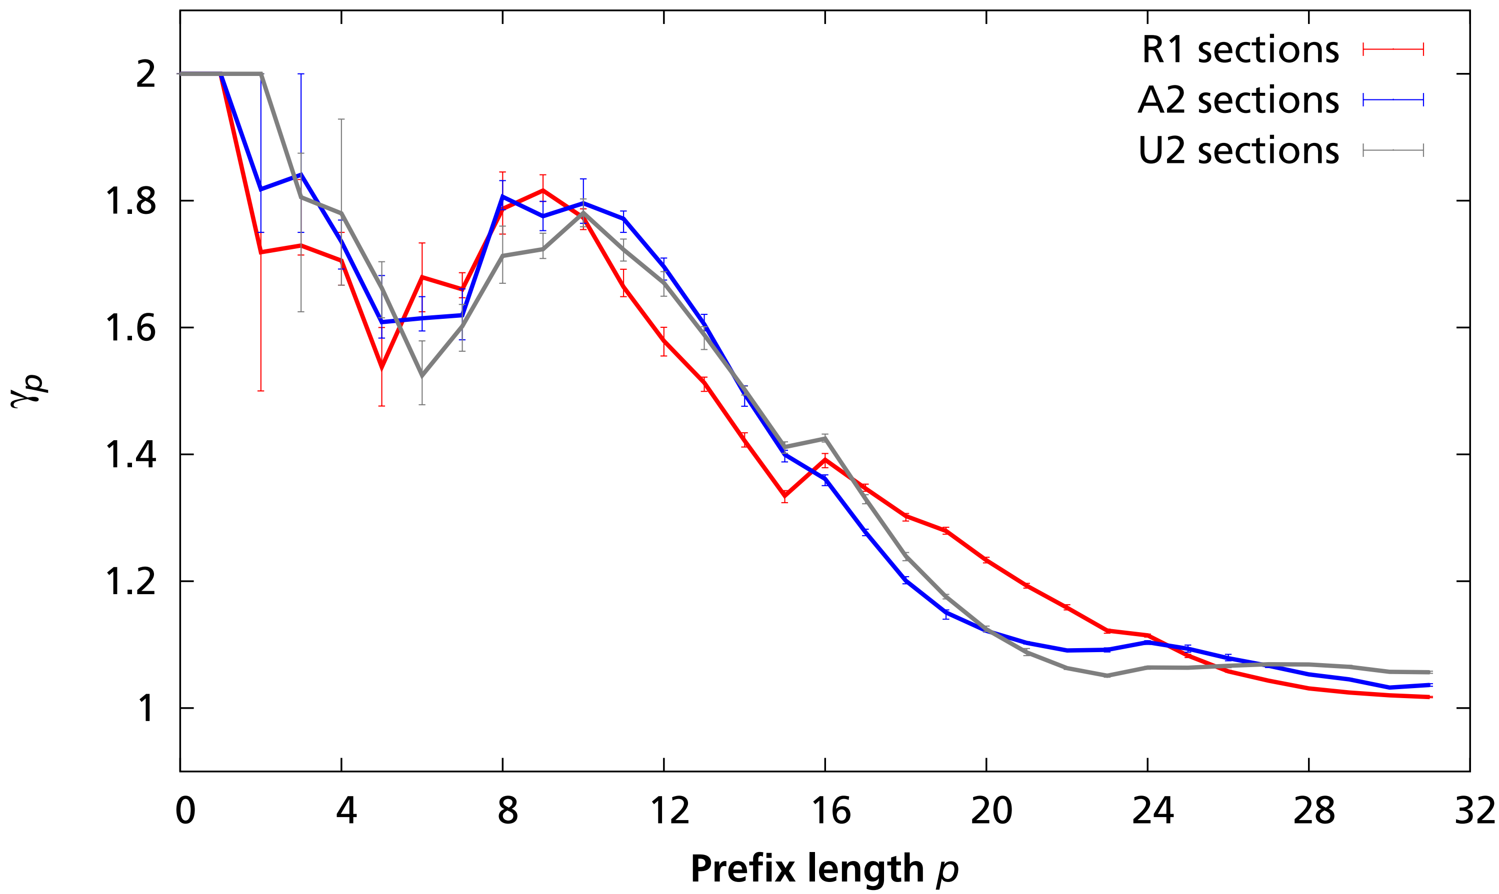

New communication dynamics?

- **How does γ_p change given a different communication pattern, such as worm propagation?**

  Expect worm propagation to significantly change the destination addresses visible at an access link, since every possible internal address will be contacted.

  Not the best detection metric . . .

- **Take a new data set, collected at a national laboratory, before and after Code Reds 1 and 2**

  Consider γ_p and aggregate population distribution

Shape of γ_p changes during worm propagation

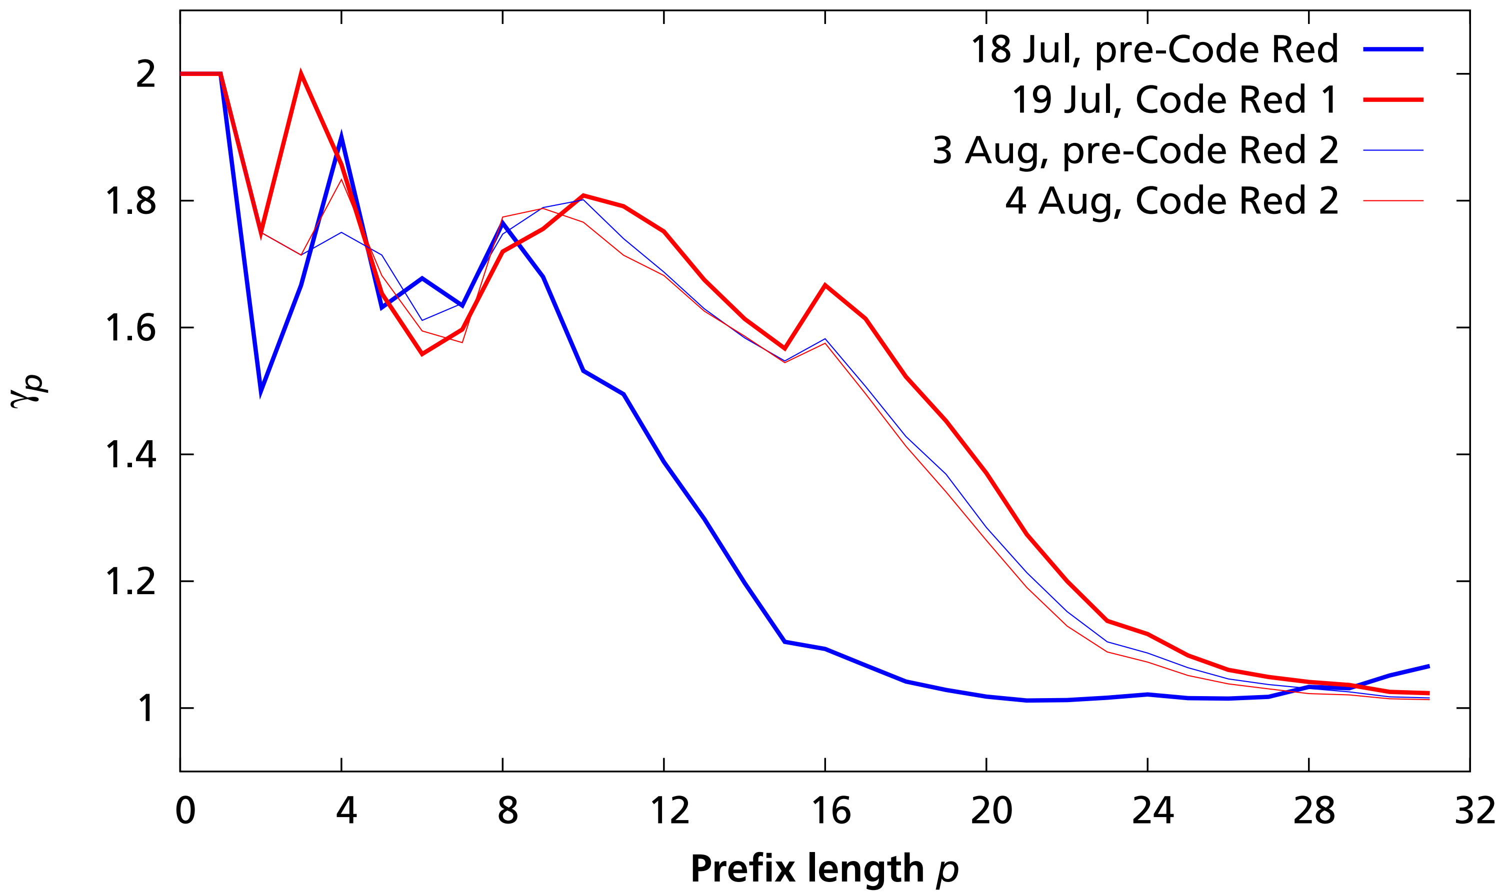

Agg. packet counts change during worm propagation

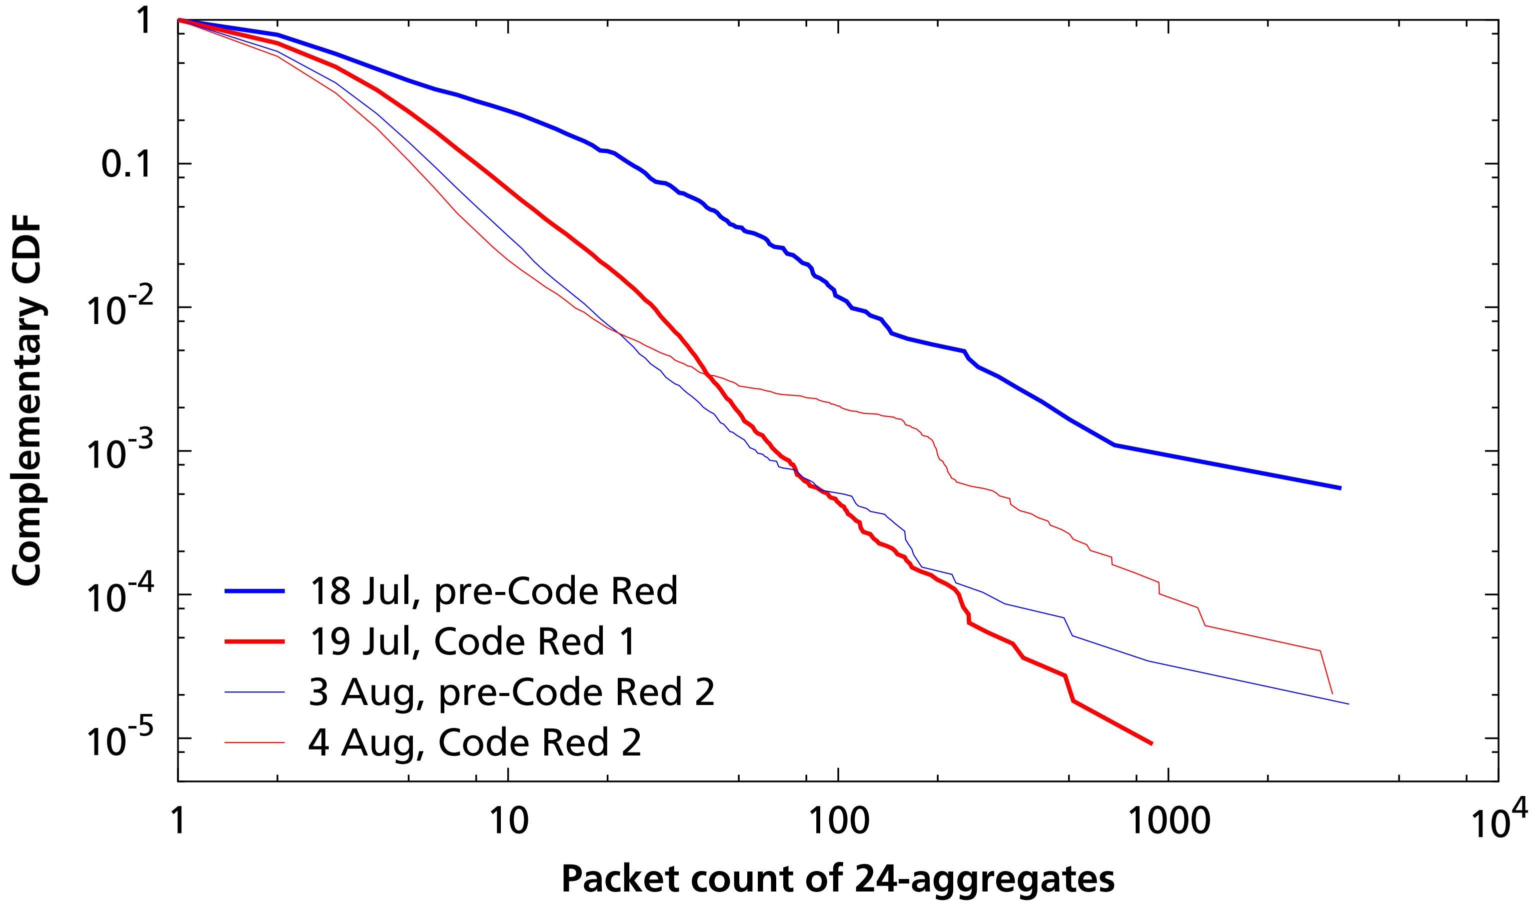

Address stability

- Divide a trace into sections, each lasting t seconds.
- How many addresses in section 1 recur in section 2?

  . . . in sections 1, 2, and 3? and so forth

  Indicates how quickly address sets change
- Model: there are long-lived addresses and short-lived addresses

  Every section contains n_S short-lived and n_L long-lived

  Addresses survive into the next section with probabilities p_S and p_L (where $p_L > p_S$)

  How well does this model match?

U2, 6-minute sections

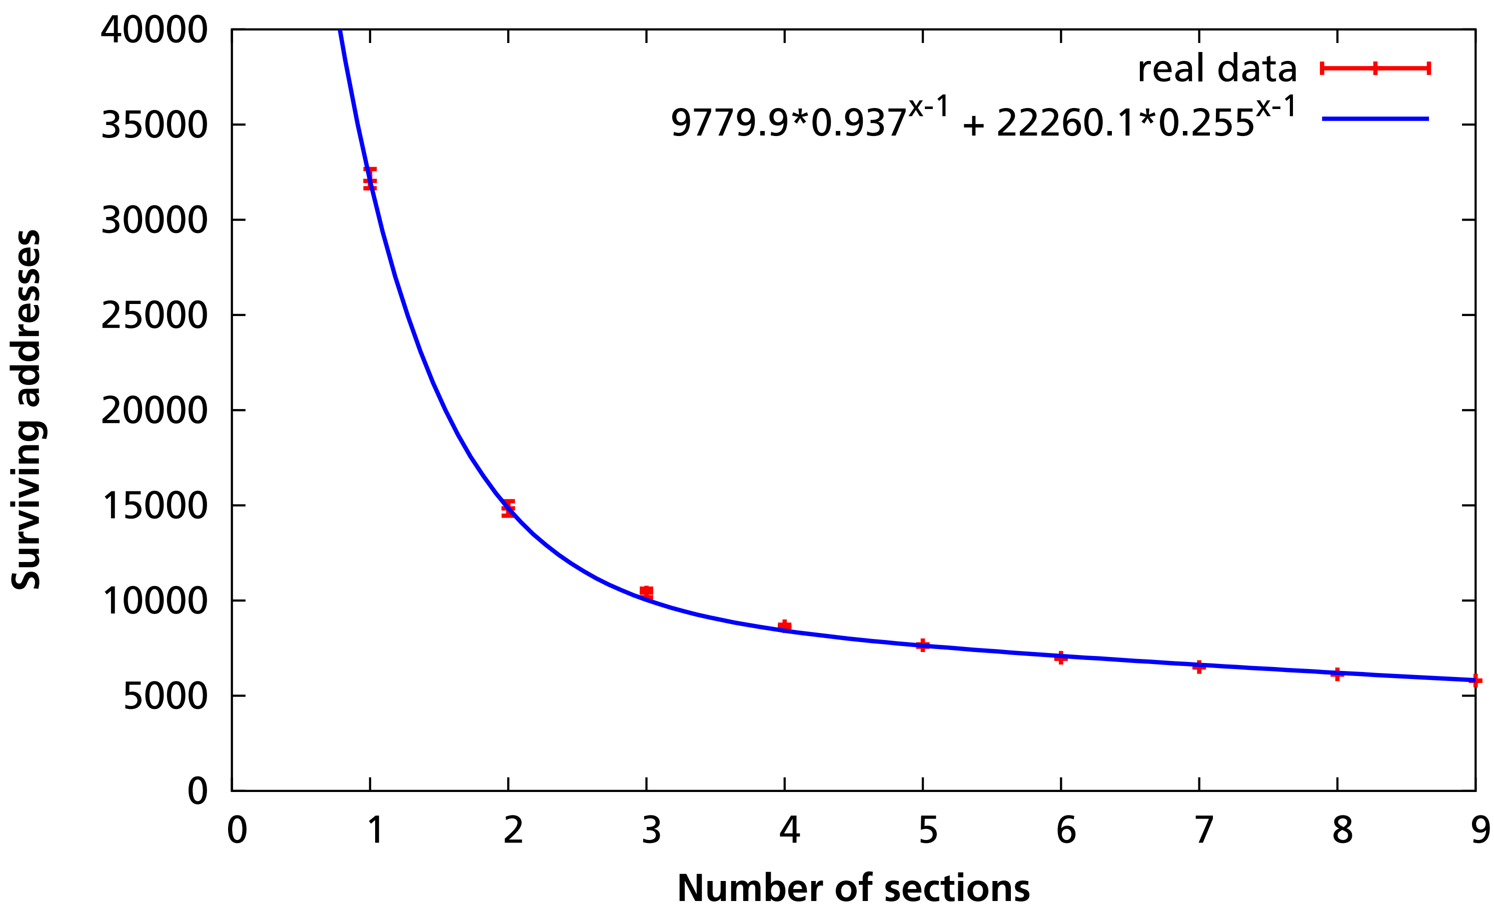

Other time scales

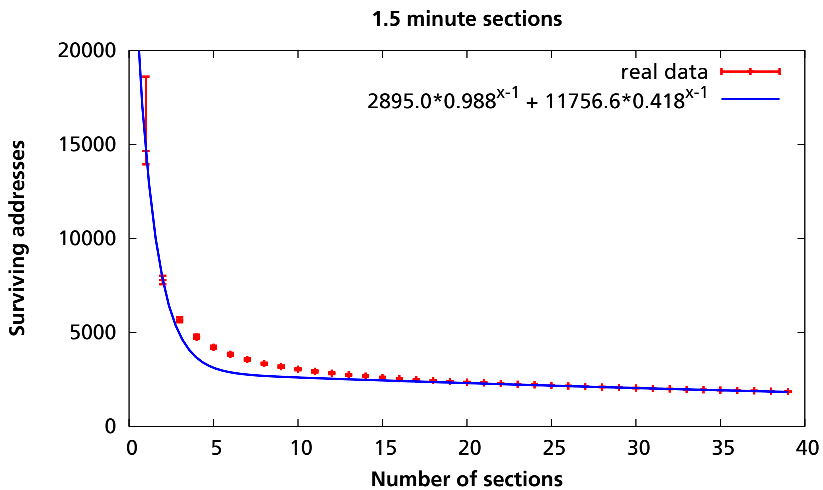

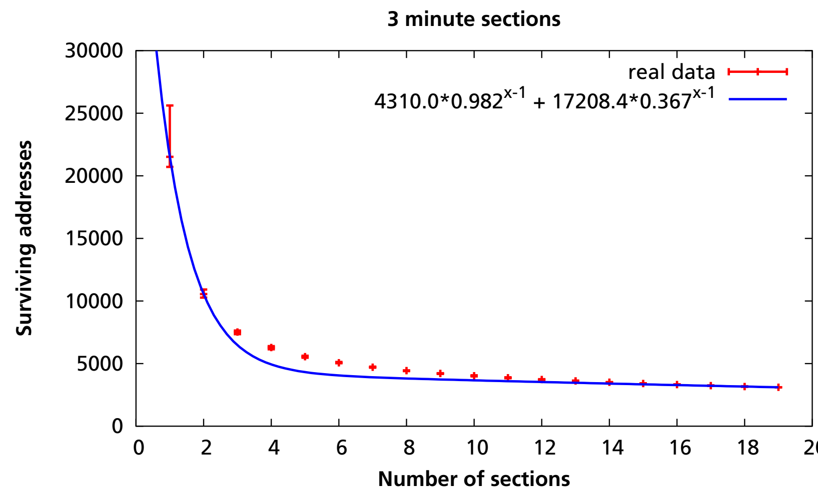

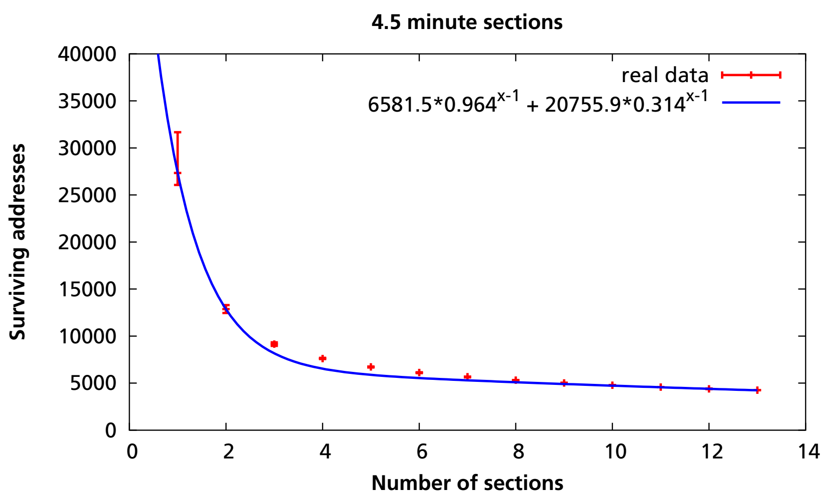

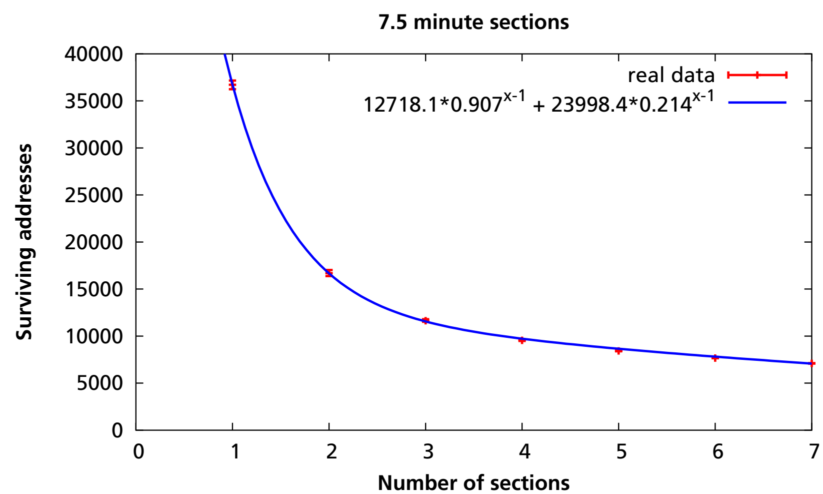

Conclusions

- Demonstrated importance of address structure
- Real address structure well modeled by a two-parameter multifractal

  Captures some aggregation behavior better than models built using metrics from real data

- Use of structural metrics as site fingerprints

  Metrics differ between sites, are stable over short time scales

Future work

???

Analysis details

- Sections are numbered $1 \ldots k$.

  $n[A]$ is number of active addresses in intersection of sections A.

- n_L long-lived addresses per section, n_S short-lived addresses.

- p_L long-lived survival probability, p_S short-lived.

- $p_L \sim n[1 \ldots k]/n[1 \ldots k-1]$.

- $n_L = n[1 \ldots k]/{p_L}^k$.

- $n_S = n[1] - n_L$.

- $p_S = (n[1,2] - n_L p_L)/n_S$.