

There is More to Internet Invariants Than Meets the Eye

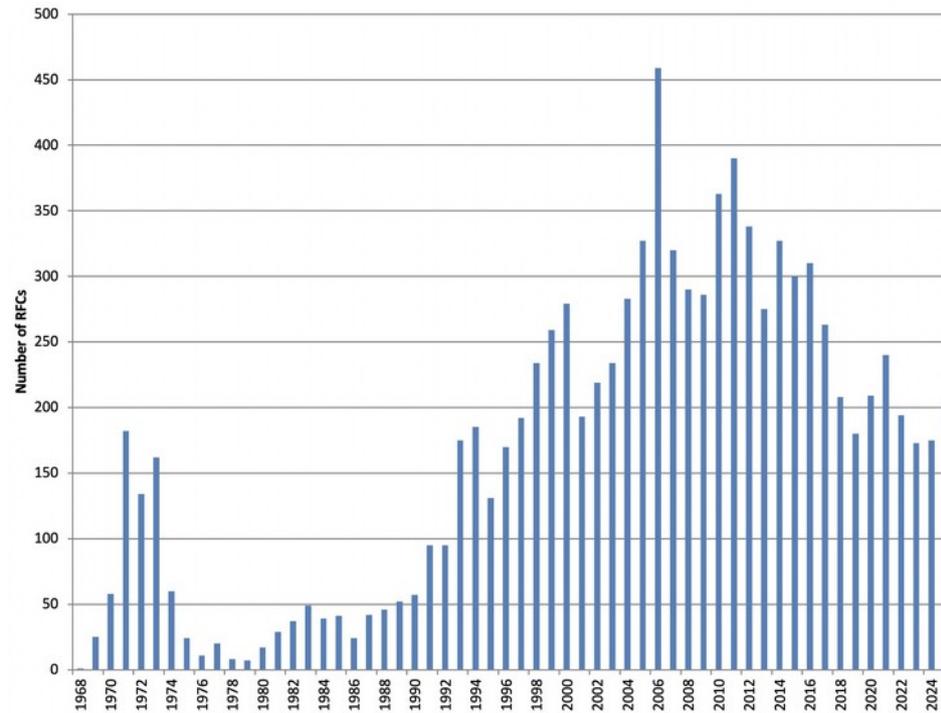
Chris Misa¹, Walter Willinger², Ramakrishnan Durairajan^{1,3}, Reza Rejaie¹

1. University of Oregon 2. NIKSUN, Inc. 3. Link Oregon

The Internet is a moving target

Reference: Difficulties in simulating the Internet, by S. Floyd and V. Paxson, Proc. IEEE/ACM ToN (2001).

The Internet is a moving target

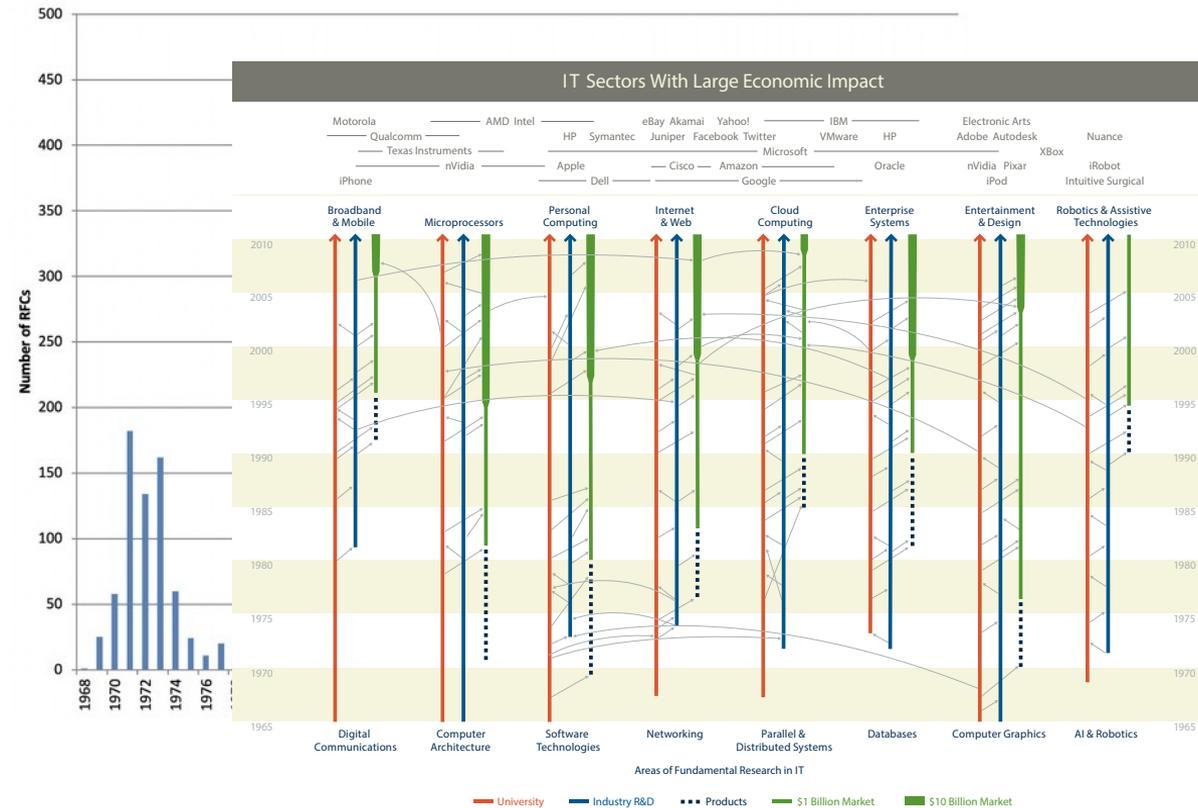


Source: <https://www.rfc-editor.org/rfc-per-year/>

Continuous technological development (e.g., ~200 RFCs / year) ...

- New protocols / applications / services
- Updates to existing protocols / applications / services

The Internet is a moving target

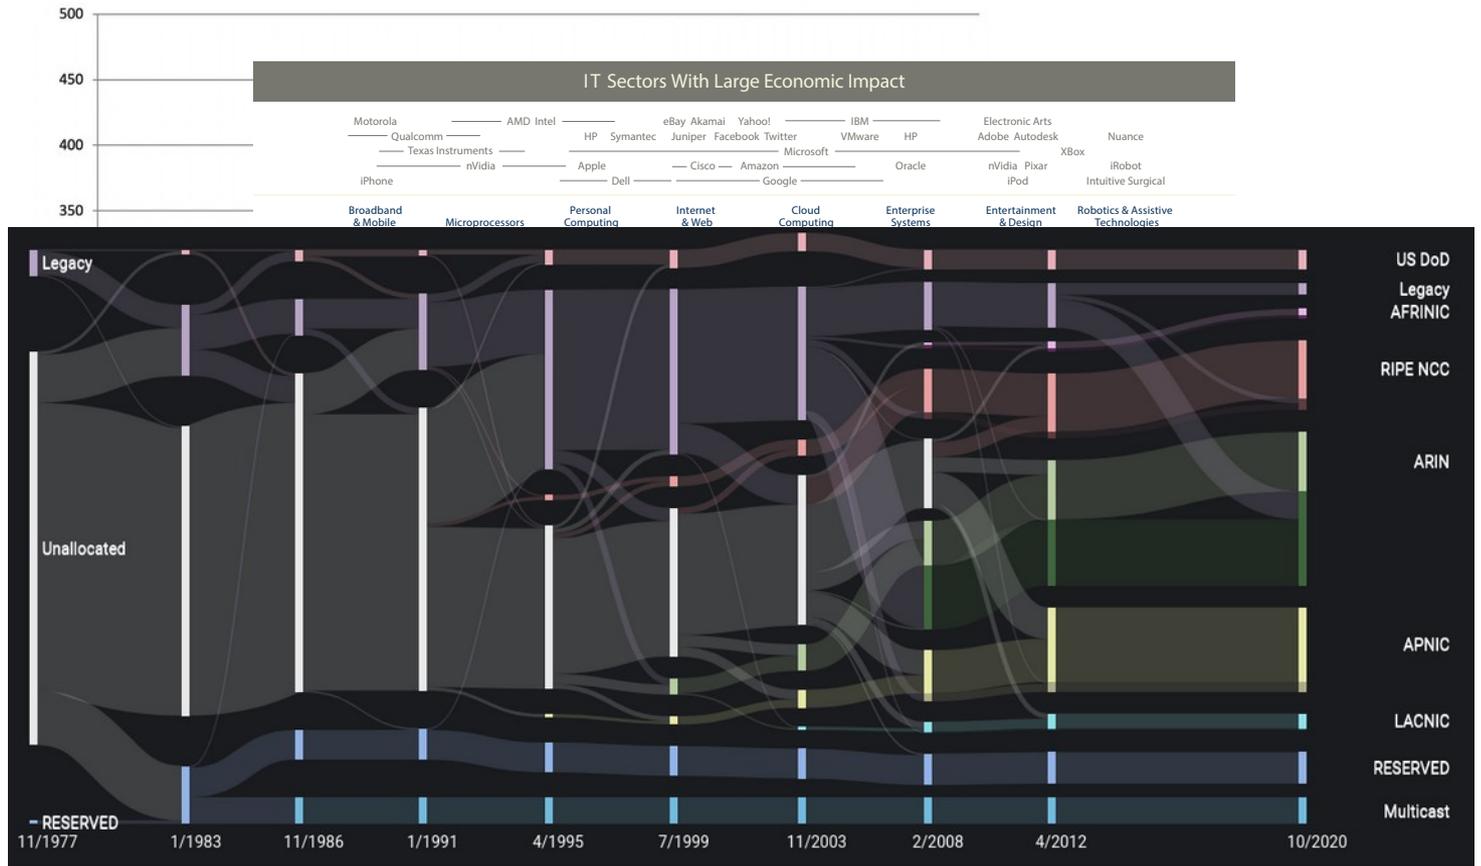


Source: <https://doi.org/10.17226/23393> (page 5)

Enables multiple new >\$10 billion industries ...

- Cloud Computing
- Broadband & Mobile

The Internet is a moving target



Source: <https://www.caida.org/catalog/media/visualizations/ipv4-history/>

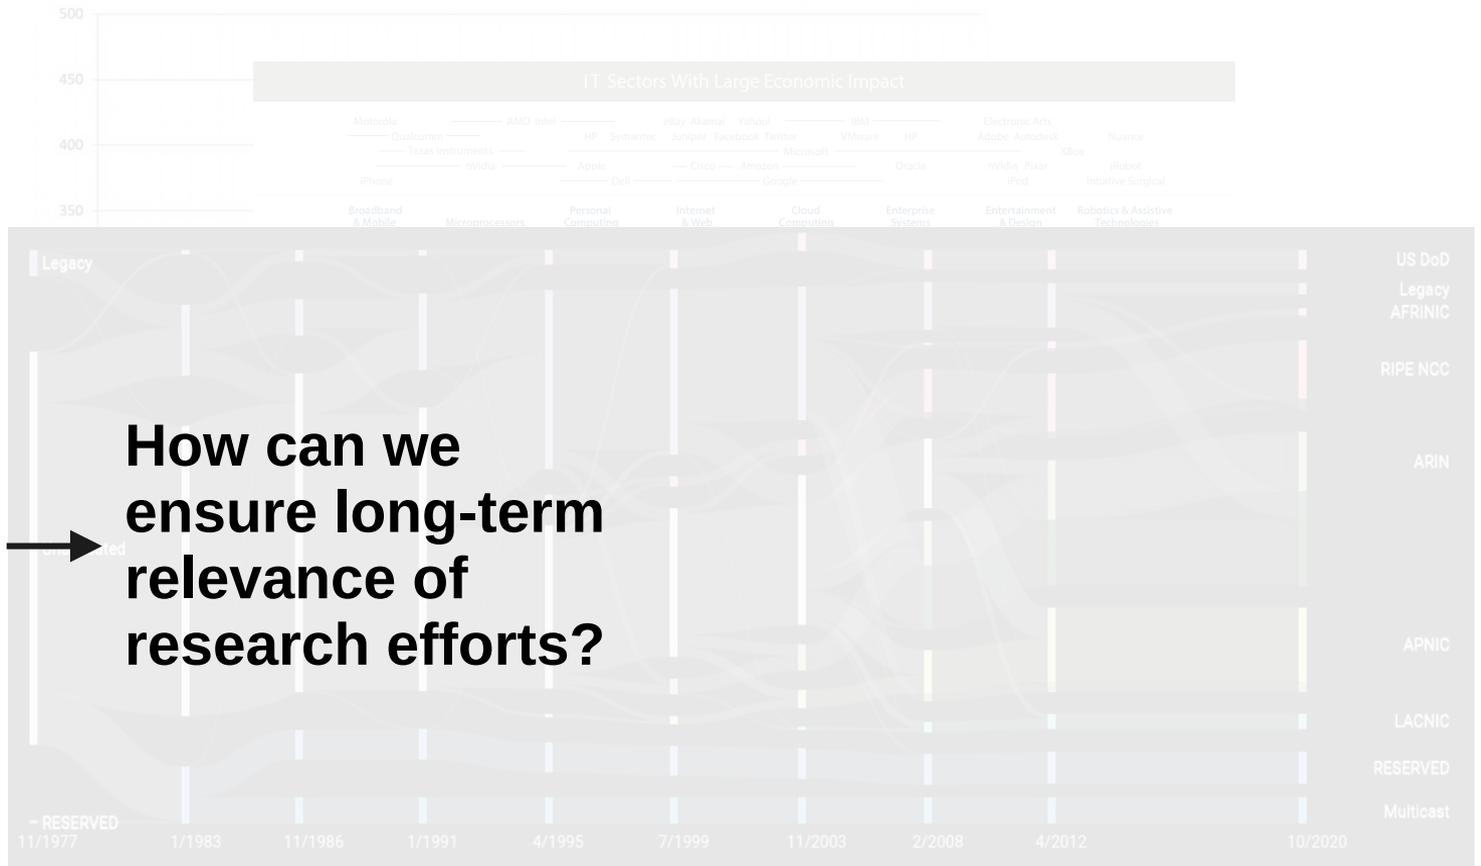
Continuous negotiation of finite resources ...

- Addresses, AS numbers, etc.

The Internet is a moving target

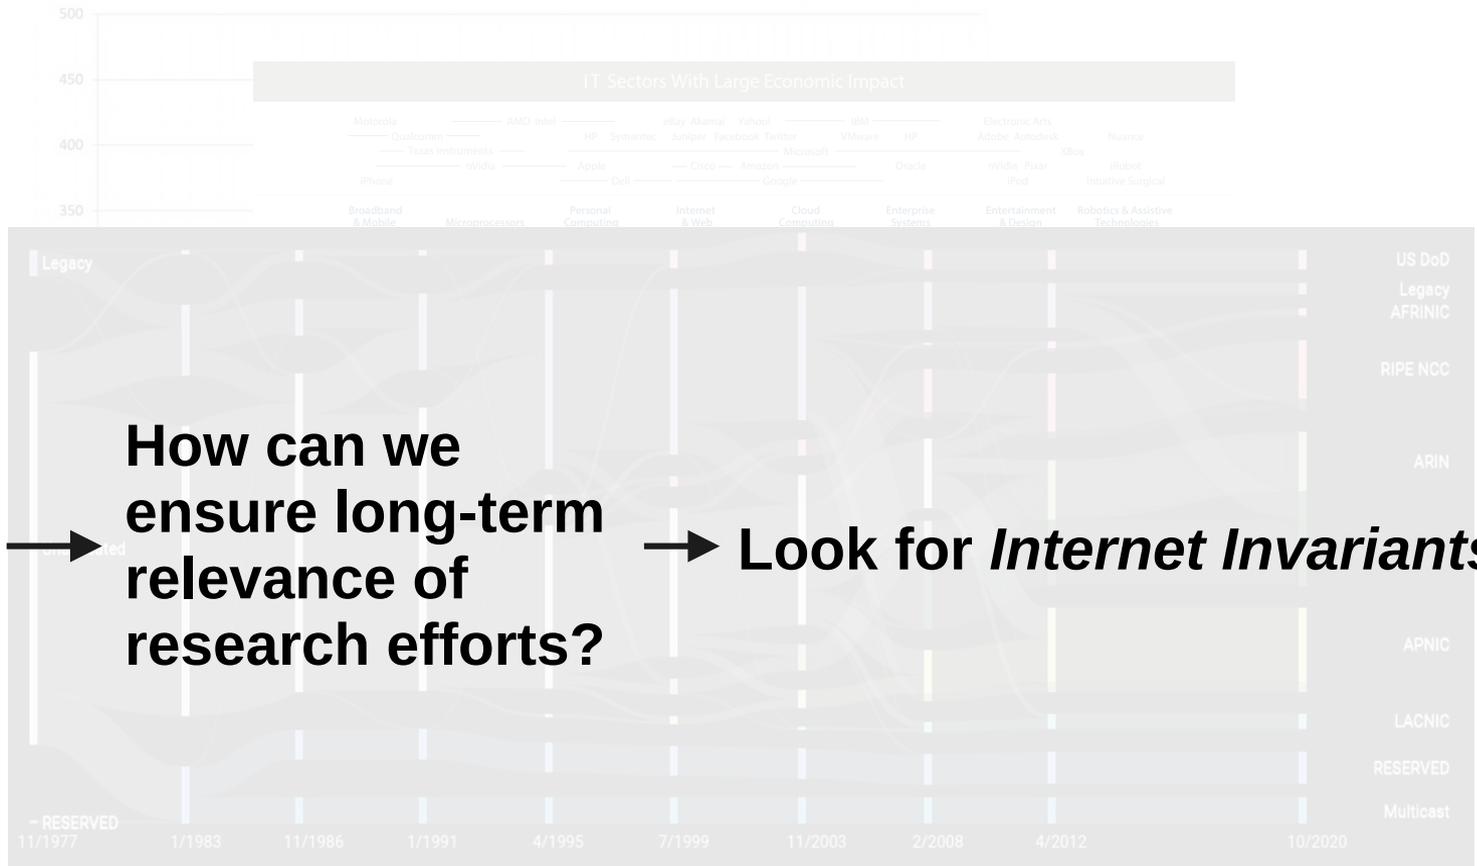


How can we ensure long-term relevance of research efforts?



Source: <https://www.caida.org/catalog/media/visualizations/ipv4-history/>

Reference: Difficulties in simulating the Internet, by S. Floyd and V. Paxson, Proc. IEEE/ACM ToN (2001).



The Internet is a moving target



How can we ensure long-term relevance of research efforts?



Look for *Internet Invariants!*

Source: <https://www.caida.org/catalog/media/visualizations/ipv4-history/>

Reference: Difficulties in simulating the Internet, by S. Floyd and V. Paxson, Proc. IEEE/ACM ToN (2001).

*Empirical
Evidence*



A distinct statistical phenomenon is observed across wide range of scenarios.

*Empirical
Evidence*

Internet Invariant



*Stated/unstated
Design Principles*

A distinct statistical phenomenon is observed across wide range of scenarios.

Is this phenomenon caused by some underlying stated/unstated design principle?

Empirical Evidence

Internet Invariant

Stated/unstated Design Principles

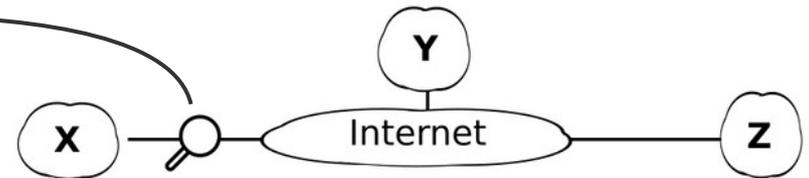


A distinct statistical phenomenon is observed across wide range of scenarios.

Is this phenomenon caused by some underlying stated/unstated design principle?

Internet Traffic Invariants

- **Temporal** (observed traffic dynamics)
- **Spatial** (observed address structure)



*Empirical
Evidence*

Internet Invariant

*Stated/unstated
Design Principles*



Temporal

*Empirical
Evidence*

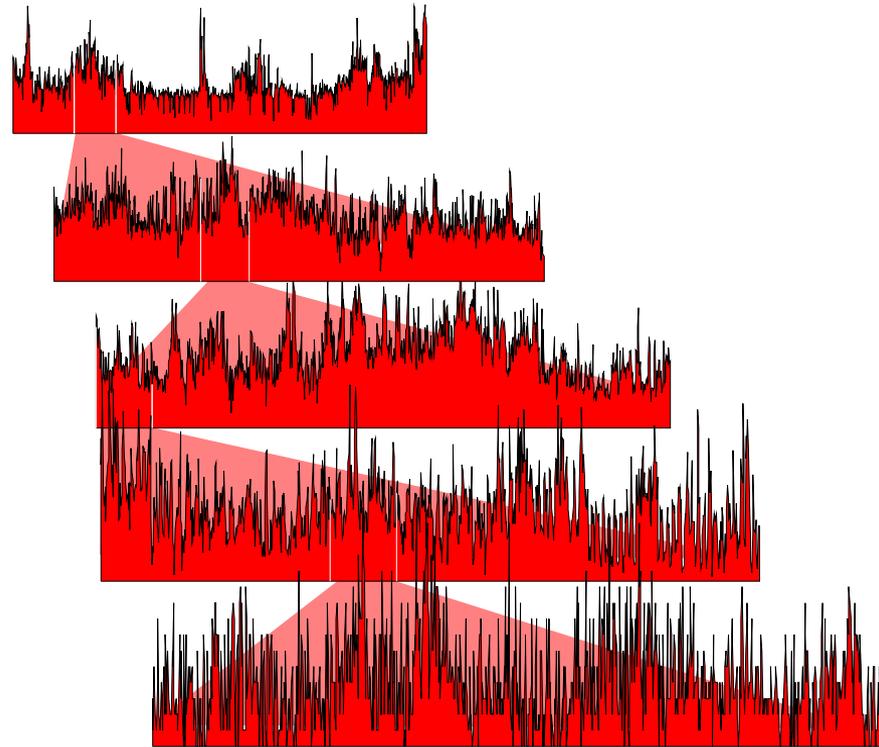
Internet Invariant

*Stated/unstated
Design Principles*



Temporal

Self-similar
scaling



Reference: On the self-similar nature of Ethernet traffic, by W.E. Leland,
M.S. Taqqu, W. Willinger, and D.V. Wilson, Proc. ACM SIGCOMM (1993).

Empirical Evidence

Generating Model



Stated/unstated Design Principles

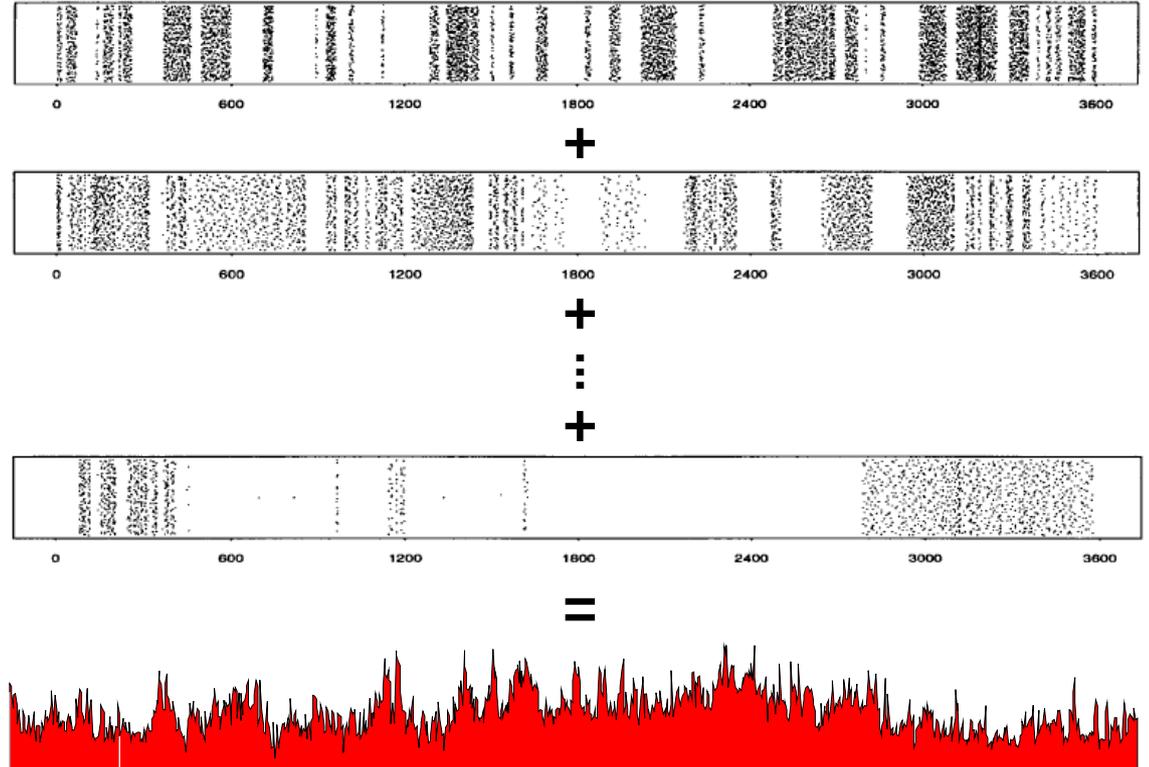
Temporal

Self-similar scaling



ON/OFF burst processes;
Heavy-tailed flow sizes.

Reference: Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level, by W. Willinger, M.S. Taqqu, R. Sherman, and D.V. Wilson, Proc. ACM SIGCOMM (1995).



Empirical Evidence

Generating Model

Reverse Engineering

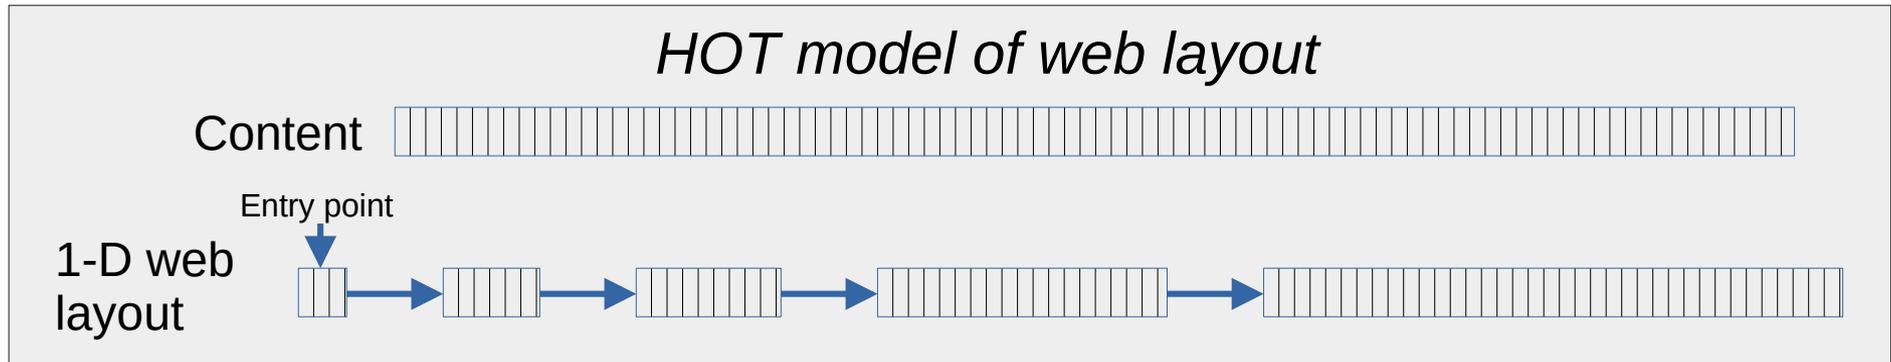
Stated/unstated Design Principles

Temporal

Self-similar scaling

ON/OFF burst processes;
Heavy-tailed flow sizes.

Optimal web layout results in **heavy-tailed flow sizes.**



Reference: Heavy Tails, Generalized Coding, and Optimal Web Layout, by X. Zhu, J. Yu, and J. Doyle, Proc. IEEE INFOCOM (2001).

Empirical Evidence

Generating Model

Reverse Engineering

Stated/unstated Design Principles

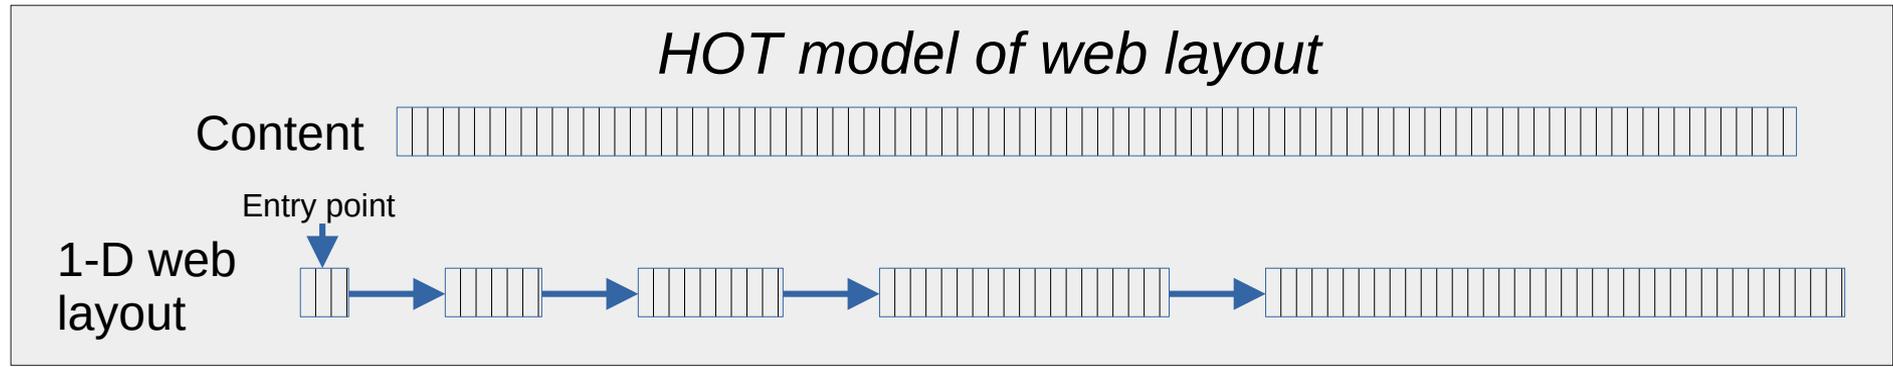
Temporal

Self-similar scaling

ON/OFF burst processes;
Heavy-tailed flow sizes.

Optimal web layout results in **heavy-tailed flow sizes.**

Organizing information for human consumption.



Reference: Heavy Tails, Generalized Coding, and Optimal Web Layout, by X. Zhu, J. Yu, and J. Doyle, Proc. IEEE INFOCOM (2001).

***Empirical
Evidence***

***Generating
Model***

***Reverse
Engineering***

***Stated/unstated
Design Principles***

Temporal

Self-similar
scaling

ON/OFF burst
processes;
**Heavy-tailed
flow sizes.**

Optimal web
layout results
in **heavy-tailed
flow sizes.**

Organizing information
for human consumption.

Spatial

Empirical Evidence

Generating Model

Reverse Engineering

Stated/unstated Design Principles

Temporal

Self-similar scaling

% of addrs



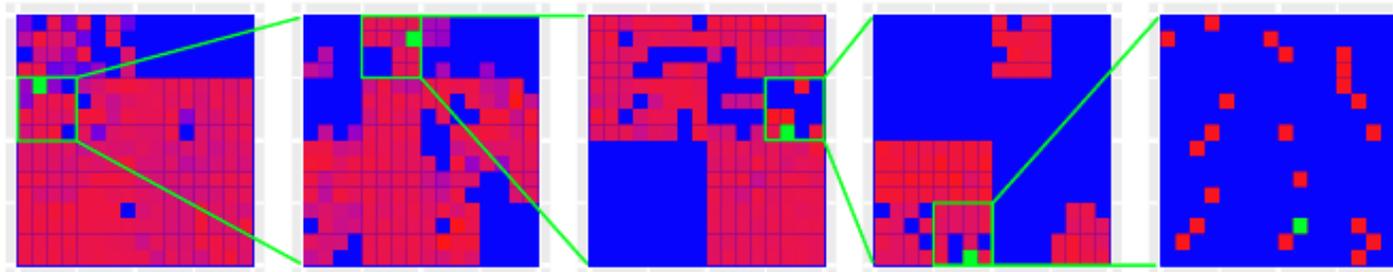
/8

/12

/16

/20

/24



Information
Implication.

Spatial

Multifractal scaling

Reference: Observed structure of addresses in IP traffic, by E. Kohler, J. Li, V. Paxson, and S. Shenker, Proc. SIGCOMM Workshop on Internet measurement (2002).

Empirical Evidence

Generating Model

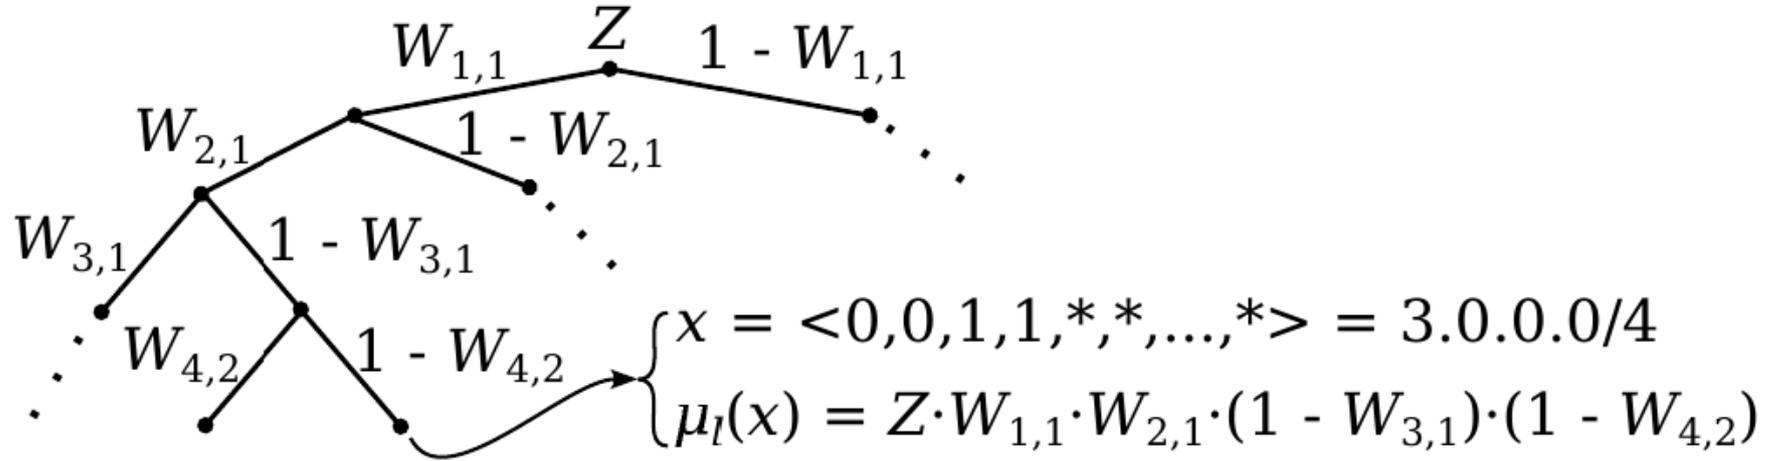
Reverse Engineering

Stated/unstated Design Principles

Temporal

Source

IP



Spatial

Multifractal scaling
Conservative Cascade;
...

Reference: Toward a model for source addresses of Internet background radiation, by P. Barford, R. Nowak, R.Willett, and V. Yegneswaran, Proc. Passive and Active Measurement Conference (2006).

	<i>Empirical Evidence</i>	<i>Generating Model</i>	<i>Reverse Engineering</i>	<i>Stated/unstated Design Principles</i>
Temporal	Self-similar scaling	ON/OFF burst processes; Heavy-tailed flow sizes.	Optimal web layout results in heavy-tailed flow sizes.	Organizing information for human consumption.
Spatial	Multifractal scaling	Conservative Cascade; ...	Known Unknown	

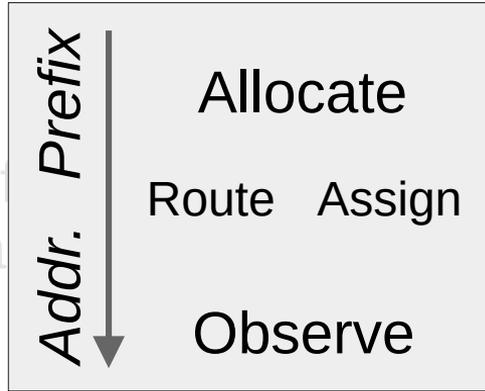
*Empirical
Evidence*

*Generating
Model*

*Reverse
Engineering*

*Stated/unstated
Design Principles*

Temporal



Spatial

Multifractal scaling
Conservative Cascade;
...

burst
es;
ailed
es.

Optimal web layout results in **heavy-tailed flow sizes.**

Organizing information for human consumption.

Empirical Evidence

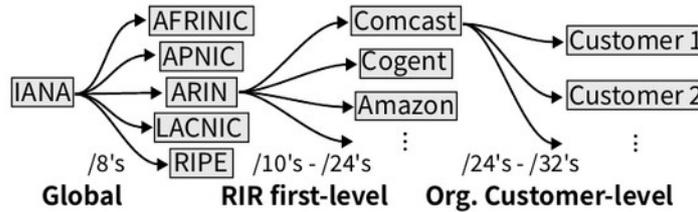
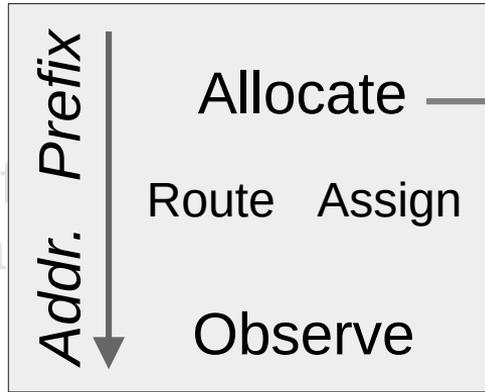
Generating Model

Reverse Engineering

Stated/unstated Design Principles

Temporal

Spatial



Self-similar scaling

bursty flows; heavy-tailed flows.

in heavy-tailed flow sizes.

organizing information for human consumption.

Multifractal scaling

Conservative Cascade;

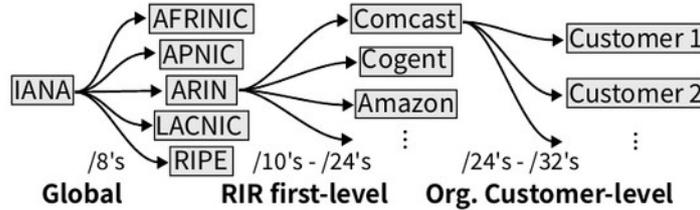
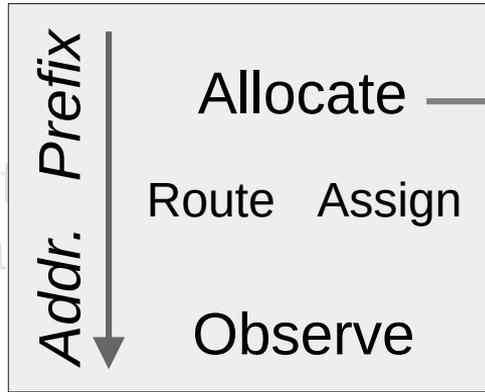
...

Empirical Evidence

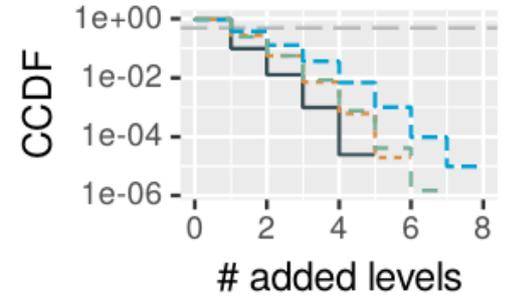
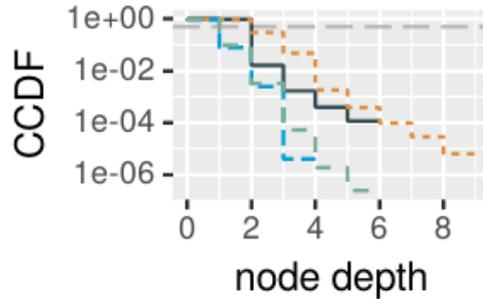
Generating Model

Reverse Engineering

Stated/unstated Design Principles



— AFRINIC - - - APNIC - · - · ARIN - - - RIPE



Temporal

Spatial

Multifractal scaling

Conservative Cascade;

Heavy-tailed allocation sizes.

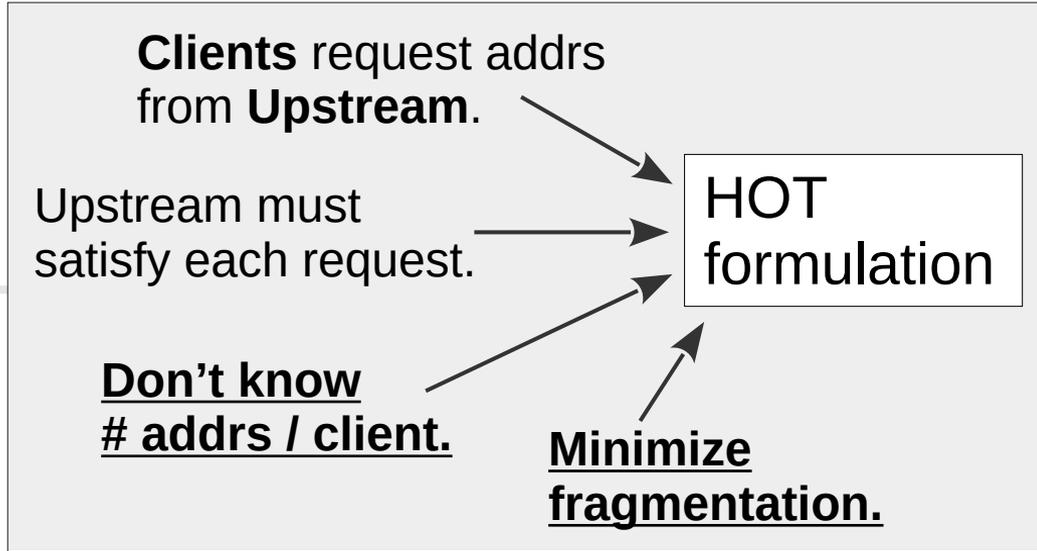
*Empirical
Evidence*

*Generating
Model*

*Reverse
Engineering*

*Stated/unstated
Design Principles*

Temporal



al web
results
vy-tailed
izes.

Organizing information for human consumption.

Spatial

Multifractal scaling ← Conservative Cascade; **Heavy-tailed allocation sizes.**

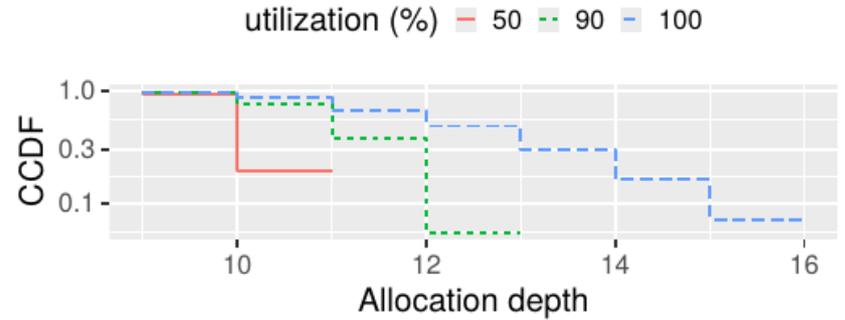
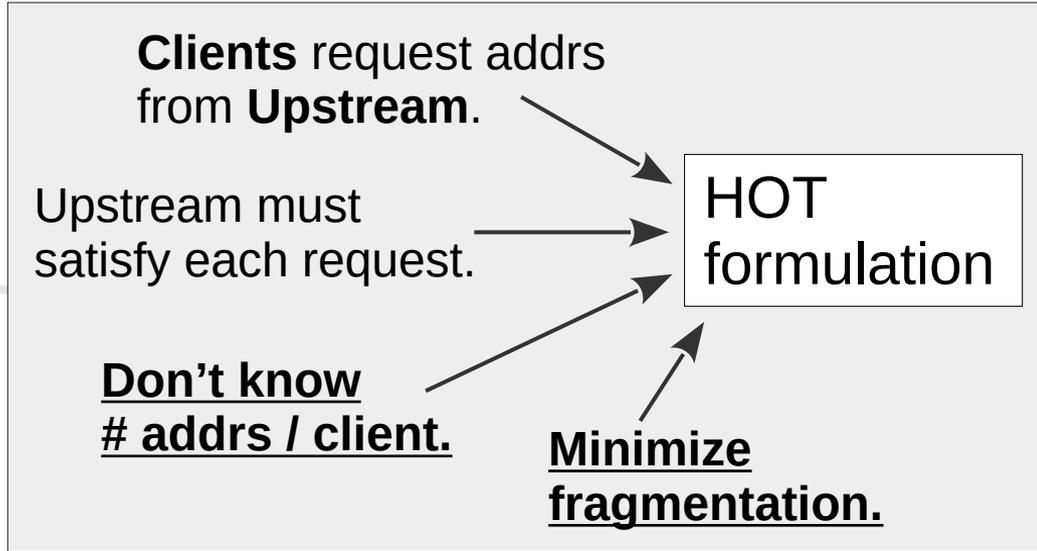
Empirical Evidence

Generating Model

Reverse Engineering

Stated/unstated Design Principles

Temporal



(...regardless of # addrs / client!)

Spatial

Multifractal scaling

Conservative Cascade; **Heavy-tailed allocation sizes.**

Optimal allocation policy results in **heavy-tailed allocation sizes.**

Empirical Evidence

Generating Model

Reverse Engineering

Stated/unstated Design Principles

4.2.3.1. Efficient Utilization

ISPs are required to apply a utilization efficiency criterion in providing address space to their customers. ...

4.2.3.3. Contiguous Blocks

IP addresses are allocated to ISPs in contiguous blocks, which should remain intact. Fragmentation of blocks is discouraged. ...

Temporal

flow sizes.

flow sizes.

Source: <https://www.arin.net/participate/policy/nrpm/>

Principles are actually stated (e.g., in RIR policy).

Spatial

Multifractal scaling

Conservative Cascade;
Heavy-tailed allocation sizes.

Optimal allocation policy results in **heavy-tailed allocation sizes.**

Organizing allocations to maximize utilization and avoid fragmentation.

	<i>Empirical Evidence</i>	<i>Generating Model</i>	<i>Reverse Engineering</i>	<i>Stated/unstated Design Principles</i>
Temporal	Self-similar scaling	ON/OFF burst processes; Heavy-tailed flow sizes.	Optimal web layout results in heavy-tailed flow sizes.	Organizing information for human consumption.
Spatial	Multifractal scaling	Conservative Cascade; Heavy-tailed allocation sizes.	Optimal allocation policy results in heavy-tailed allocation sizes.	Organizing allocations to maximize utilization and avoid fragmentation.

... see our paper for details!

	<i>Empirical Evidence</i>	<i>Generating Model</i>	<i>Reverse Engineering</i>	<i>Stated/unstated Design Principles</i>
Temporal	Self-similar scaling	ON/OFF burst processes; Heavy-tailed flow sizes.	Optimal web layout results in heavy-tailed flow sizes.	Organizing information for human consumption.
Spatial	Multifractal scaling	Conservative Cascade; Heavy-tailed allocation sizes.	Optimal allocation policy results in heavy-tailed allocation sizes.	Organizing allocations to maximize utilization and avoid fragmentation.

But future AI-based applications and services will certainly change everything! Or will they?

	<i>Empirical Evidence</i>	<i>Generating Model</i>	<i>Reverse Engineering</i>	<i>Stated/unstated Design Principles</i>
Temporal	Self-similar scaling	ON/OFF burst processes; Heavy-tailed flow sizes.	Optimal web layout results in heavy-tailed flow sizes.	Organizing information for human consumption.
Spatial	Multifractal scaling	Conservative Cascade; Heavy-tailed allocation sizes.	Optimal allocation policy results in heavy-tailed allocation sizes.	Organizing allocations to maximize utilization and avoid fragmentation.

Why is it important for applications to organize information for human consumption?

	<i>Empirical Evidence</i>	<i>Generating Model</i>	<i>Reverse Engineering</i>	<i>Stated/unstated Design Principles</i>
Temporal	Self-similar scaling	ON/OFF burst processes:	Optimal web layout results	Organizing information
	<div style="border: 2px solid red; padding: 5px; background-color: #f0f0f0;"> <p><i>But future changes in how the Internet is owned and managed will certainly change everything! Or will they?</i></p> </div>			
Spatial	Multifractal scaling	Conservative Cascade; Heavy-tailed allocation sizes.	Optimal allocation policy results in heavy-tailed allocation sizes.	Organizing allocations to maximize utilization and avoid fragmentation.

	<i>Empirical Evidence</i>	<i>Generating Model</i>	<i>Reverse Engineering</i>	<i>Stated/unstated Design Principles</i>
<i>Temporal</i>	Self-similar scaling	ON/OFF burst processes; Heavy-tailed flow	Optimal web layout results in heavy-tailed	Organizing information for human consumption.
<i>Spatial</i>	Multifractal scaling	Conservative Cascade; Heavy-tailed allocation sizes.	Optimal allocation policy results in heavy-tailed allocation sizes.	Organizing allocations to maximize utilization and avoid fragmentation.

Why is it important to optimize for scalability and stability of IP address allocations?

Thanks!

Discord:



Project page:



Chris Misa

cmisa@cs.uoregon.edu

chrismisa.com

