Harnessing Internet Topological Stability in Thorup-Zwick Compact Routing

Stephen D. Strowes
(sds@dcs.gla.ac.uk)
Colin Perkins
(csp@csperkins.org)

26 March 2012

Outline

The contributions of this work are as follows:

- We present a long-term study of k-cores decomposition of AS graphs
- We define an algorithm for TZ compact routing that uses the output of this decomposition
- We demonstrate that the ASes selected by the algorithm perform well as landmarks.

Outline

Graph Model

k-cores Graph Decomposition

- Isolates nodes of increasing importance to graph connectivity
- Deterministic, cheap ($O(n)$) graph decomposition algorithm

k-cores Graph Decomposition

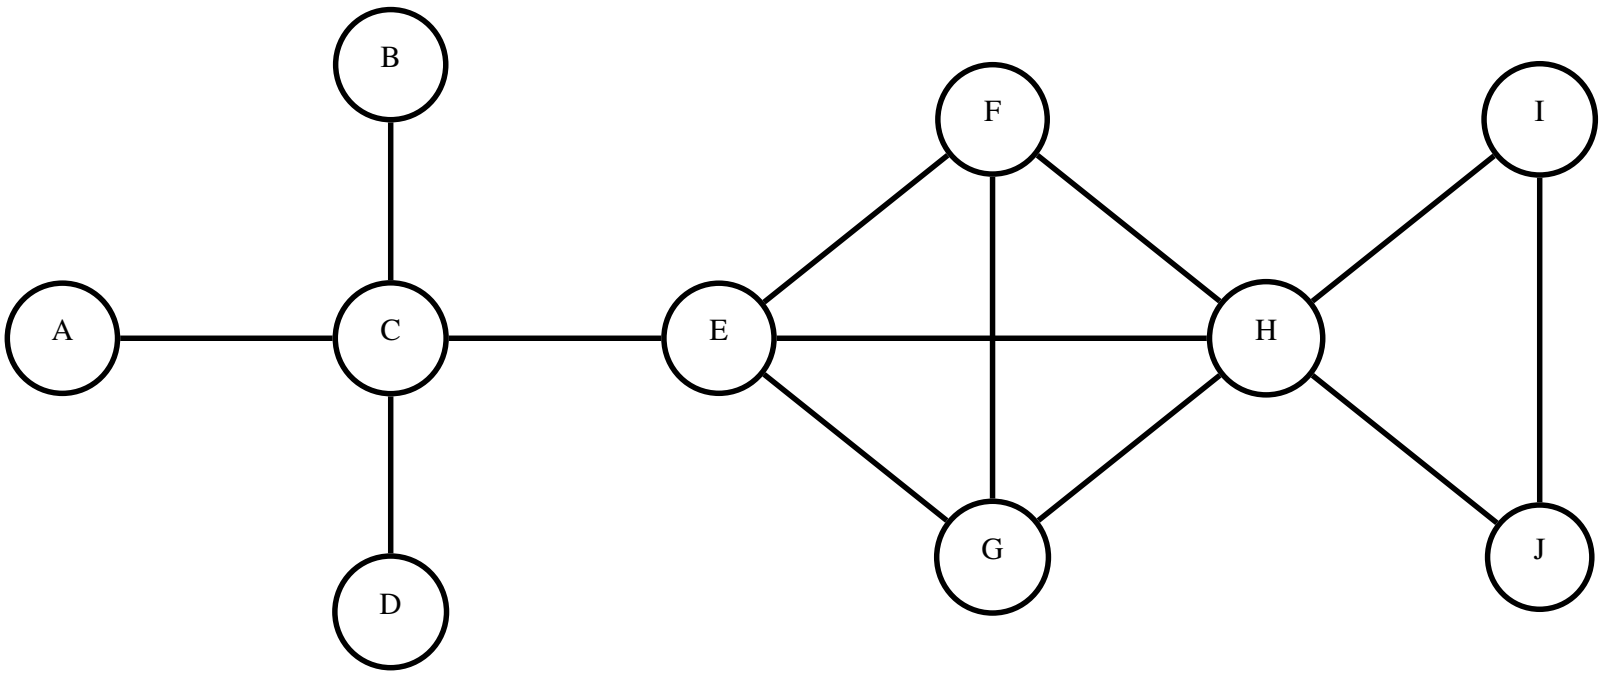

k-cores Graph Decomposition

$k == 1$

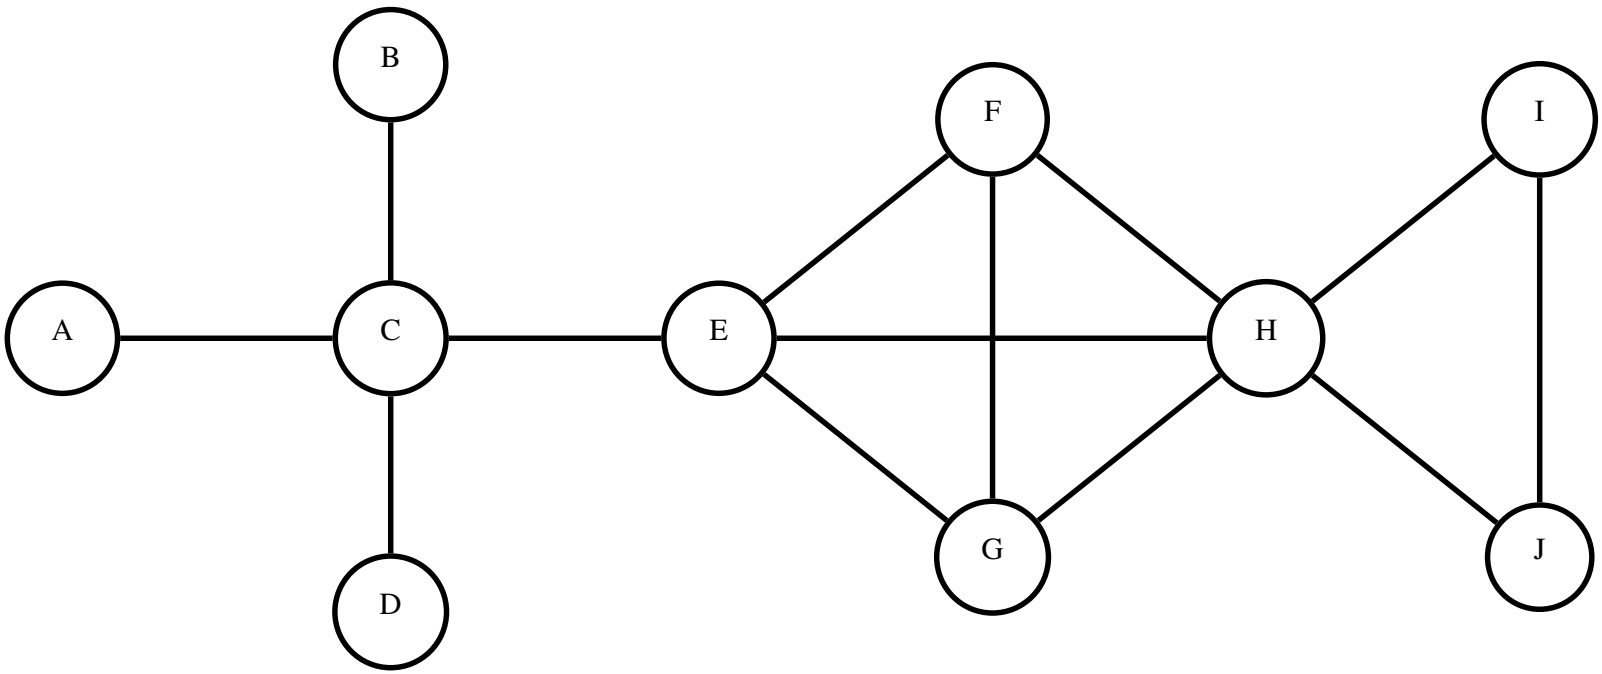

k-cores Graph Decomposition

$k == 1$

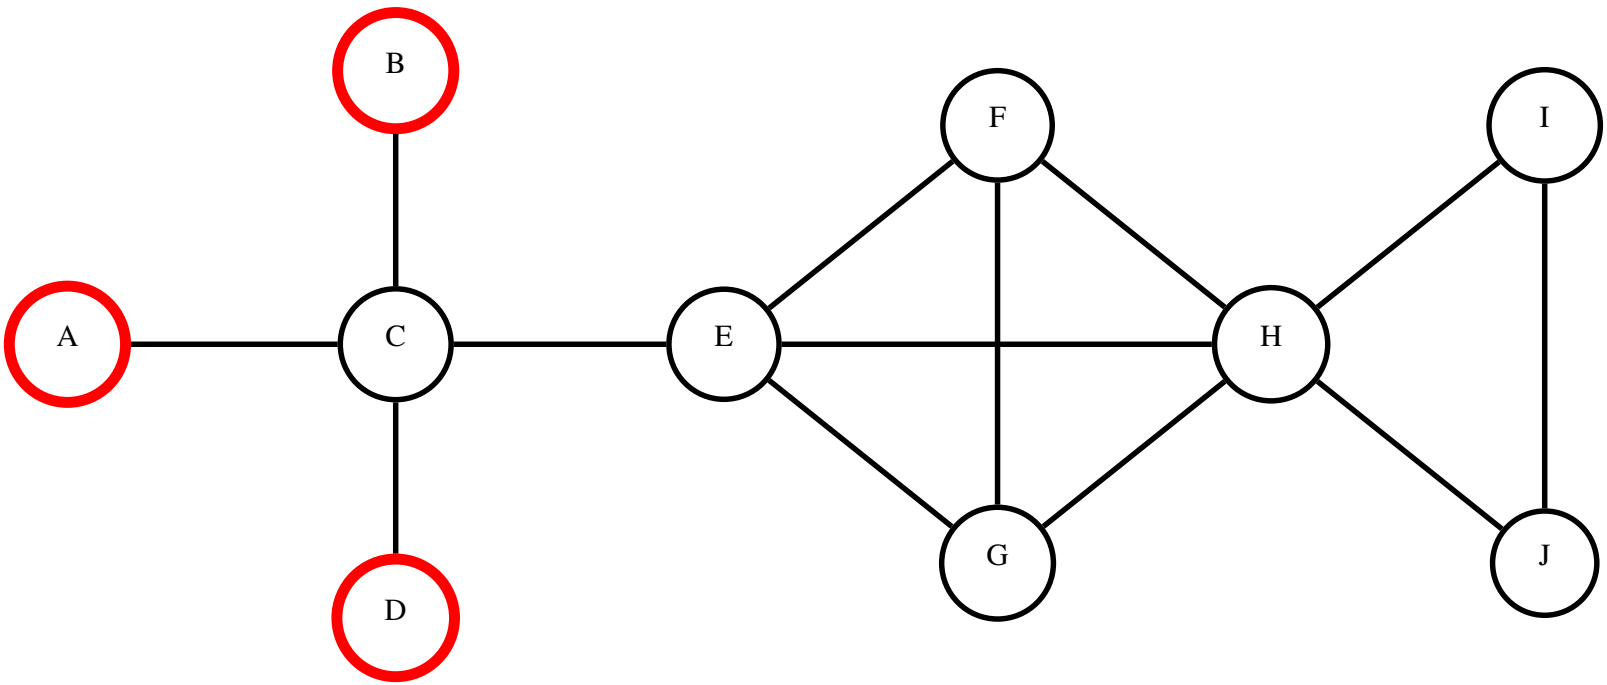

k-cores Graph Decomposition

$k == 1$

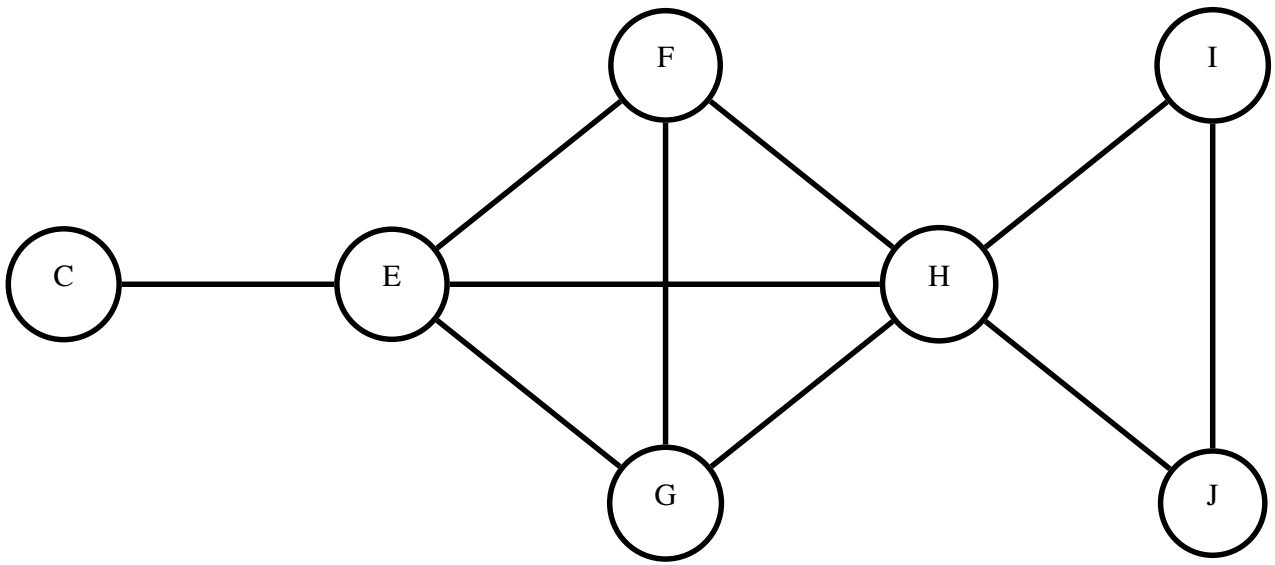

k-cores Graph Decomposition

$k == 1$

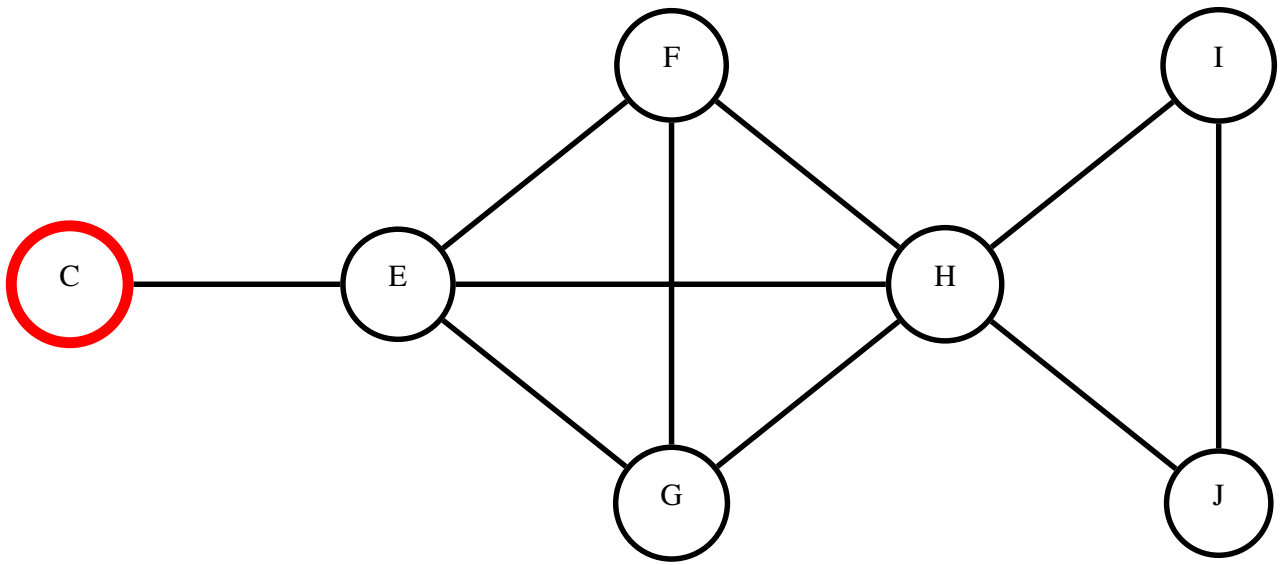

k-cores Graph Decomposition

$k == 2$

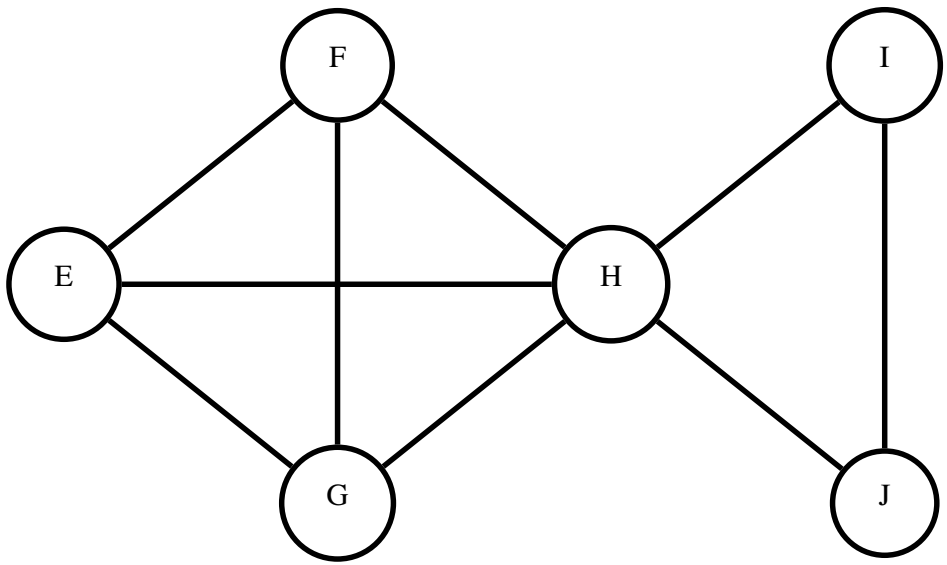

k-cores Graph Decomposition

$k == 2$

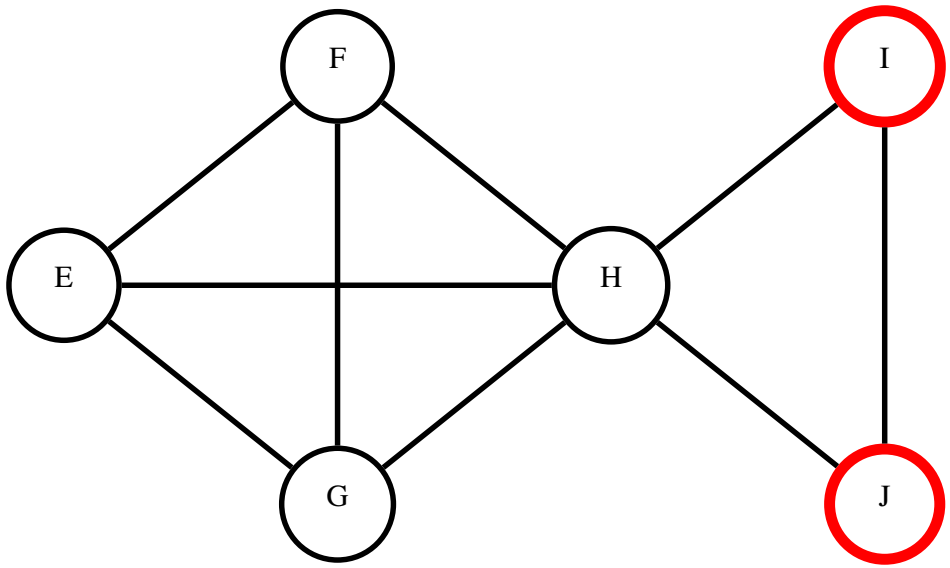

k-cores Graph Decomposition

$k == 3$

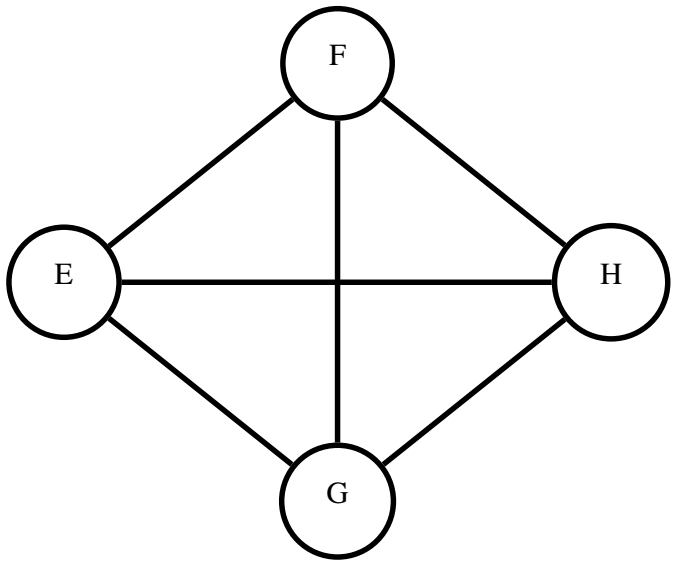

k-cores Graph Decomposition: Nucleus Growth

k-cores Graph Decomposition: Nucleus Growth

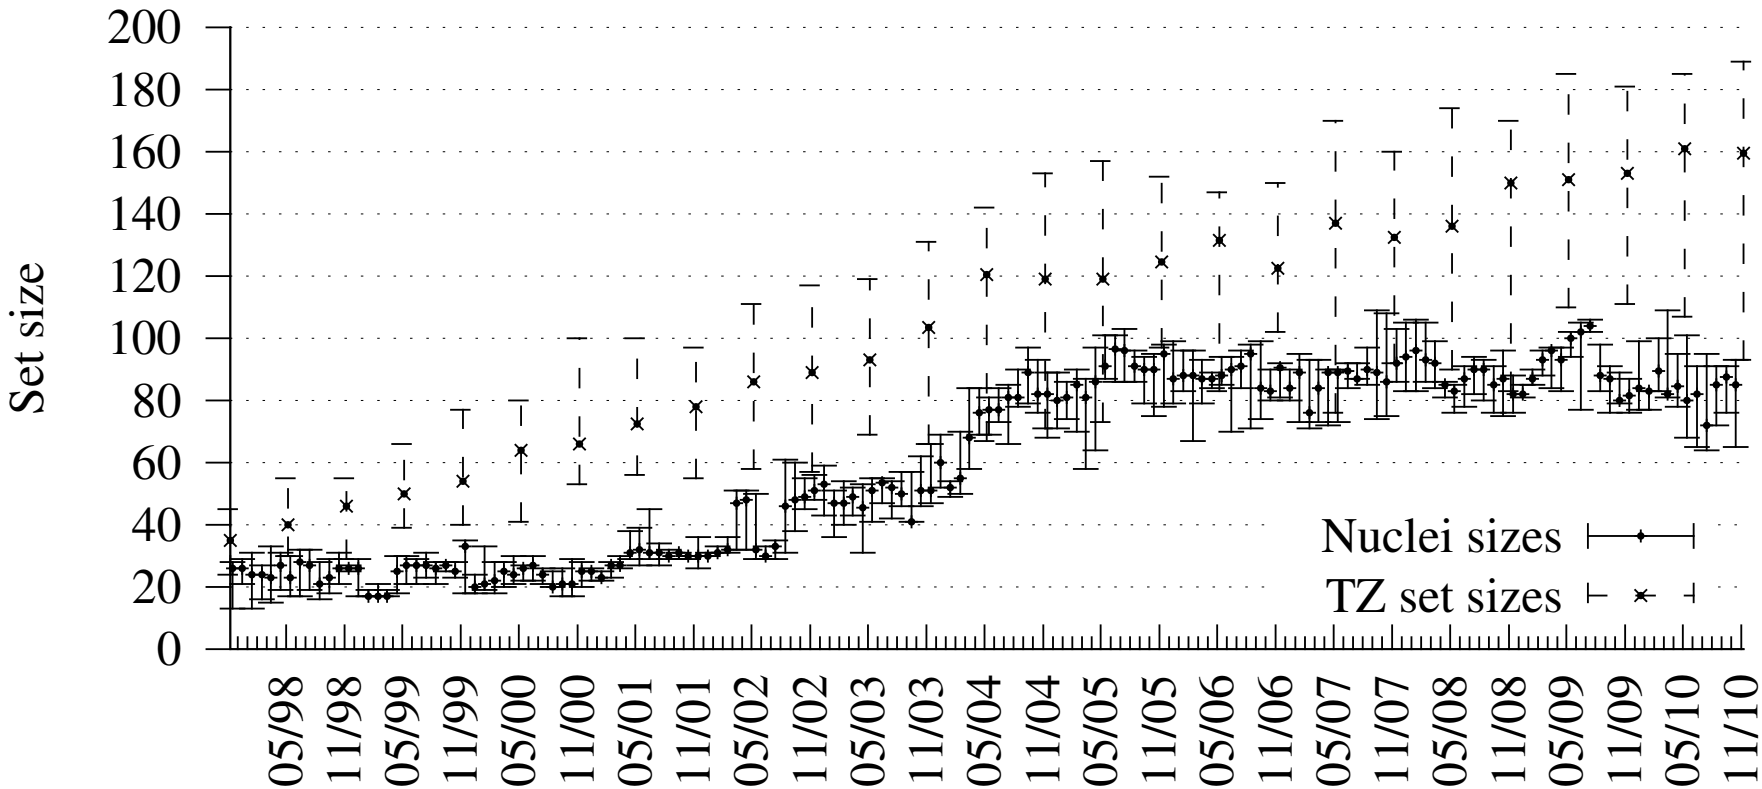

k-cores Graph Decomposition: Nucleus Decay

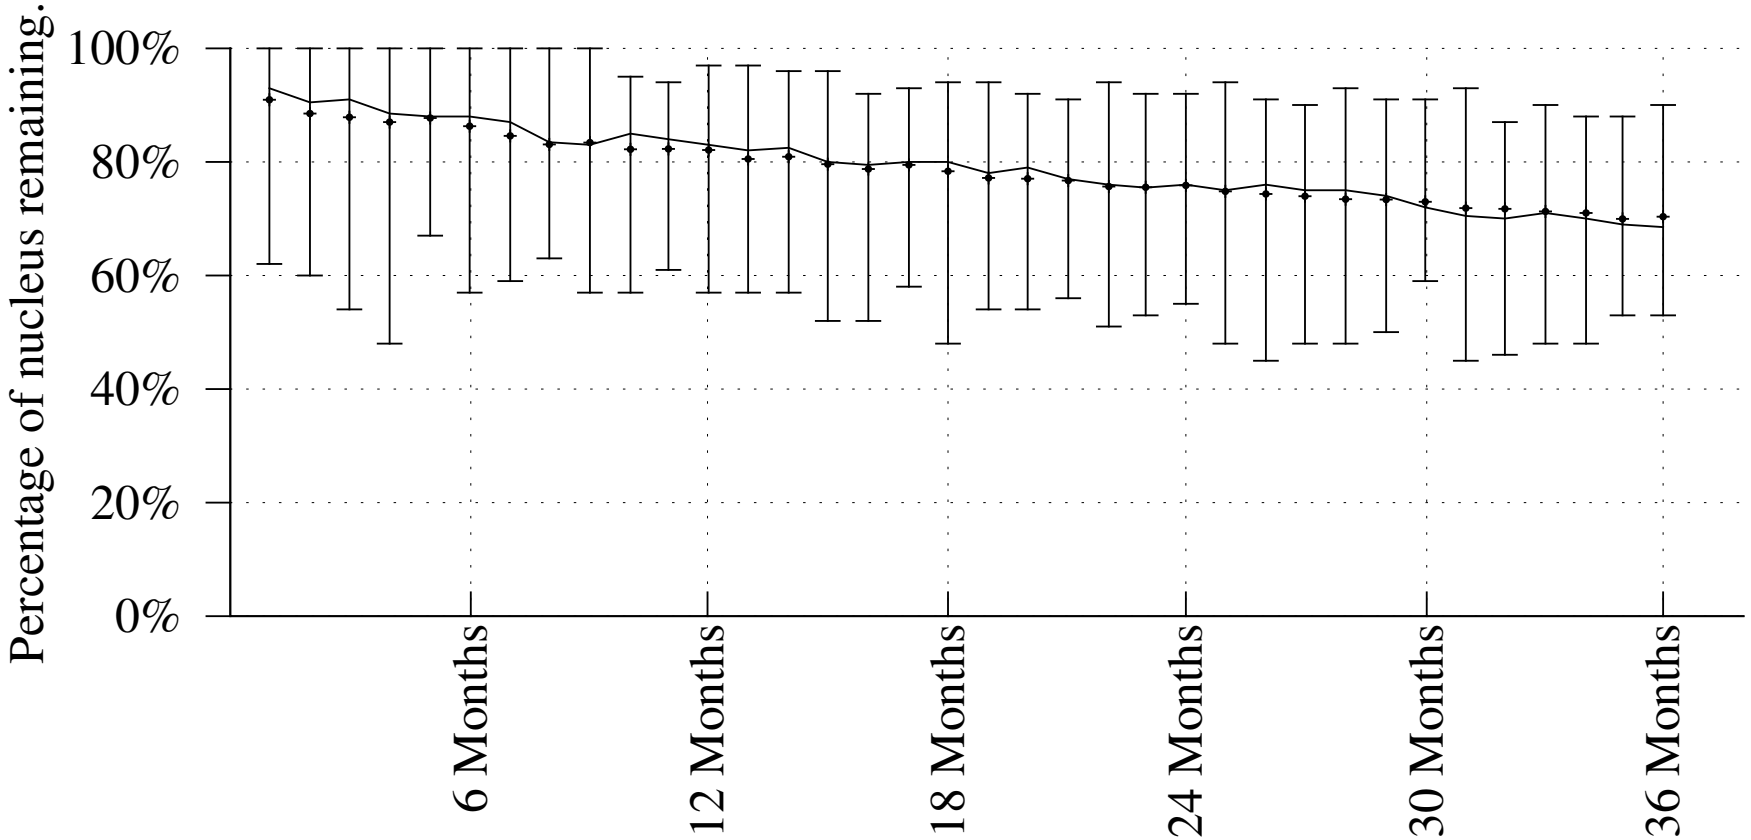

k-cores Graph Decomposition: Nucleus Centrality

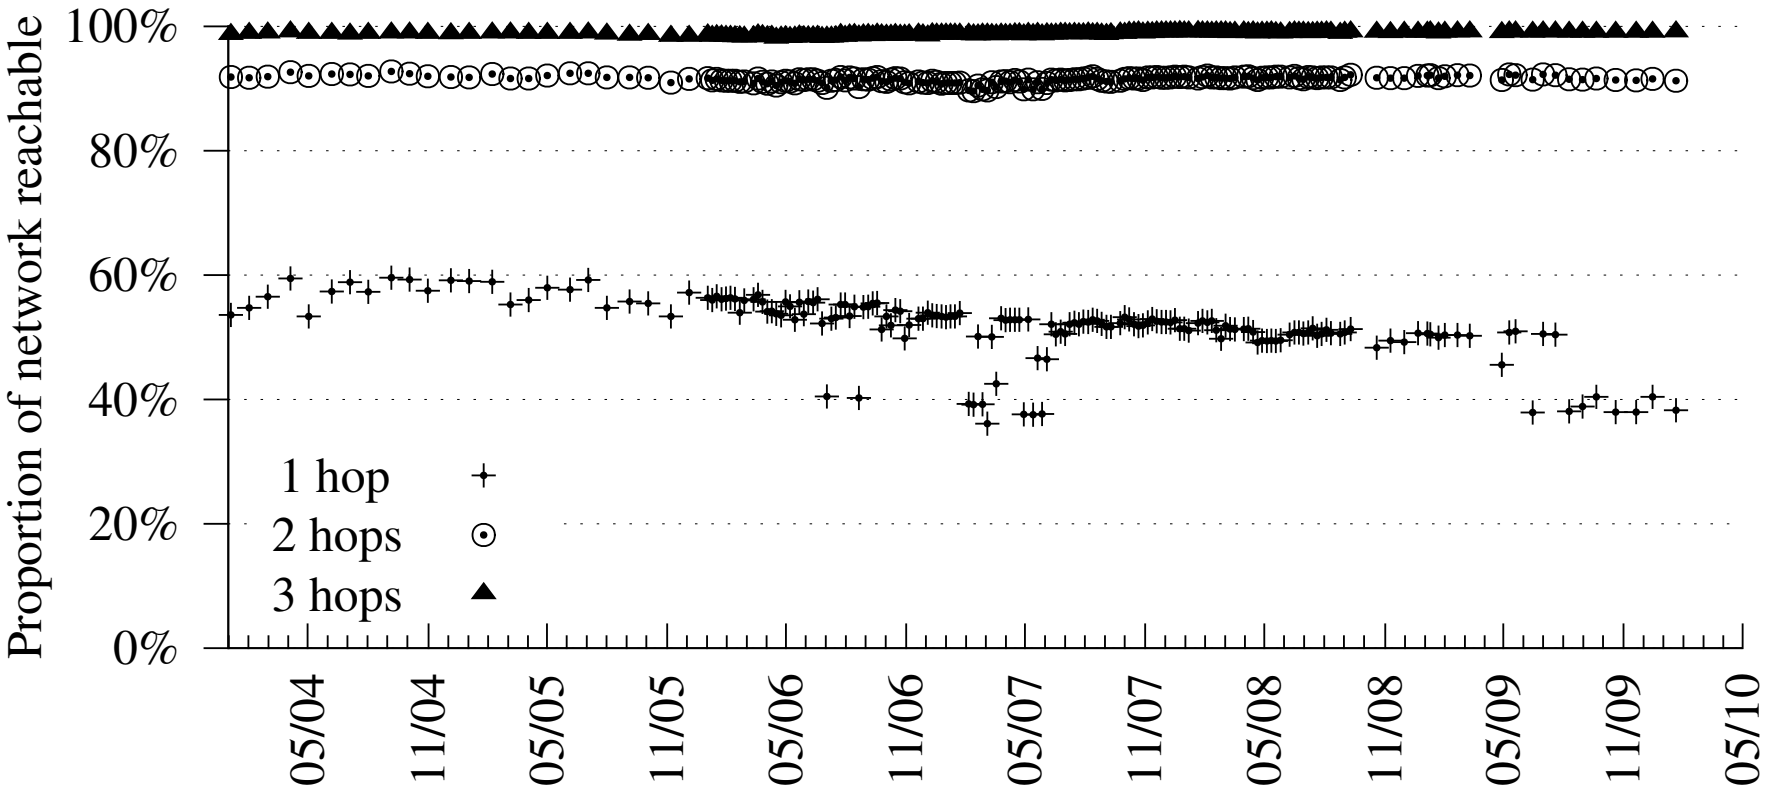

k-cores Graph Decomposition

- 45,000 unique ASes observed in this data
- Only 245 have ever appeared in the nucleus set
- 212 were still in the BGP data on 8 November 2010
- Nucleus sets (now) tend to be on the order of 100 ASes

Outline

Thorup-Zwick Compact Routing

- Guarantees:
  - Forwarding tables of size $O(\sqrt{n}\log(\sqrt{n}))$
  - Path lengths bounded by maximum multiplicative stretch of 3

Thorup-Zwick Compact Routing

But stretch-3 sounds bad!

Thorup-Zwick Compact Routing

But stretch-3 sounds bad!

The averages are much closer to shortest path

Thorup-Zwick Compact Routing

- Relies on:
  - Globally visible landmark nodes
  - Limited visibility around all other nodes

Thorup-Zwick Compact Routing

- Packet forwarding algorithm:
  - Packets forwarded toward landmark if destination is unknown
  - When destination *is* known, packets forwarded directly toward destination

Thorup-Zwick Compact Routing

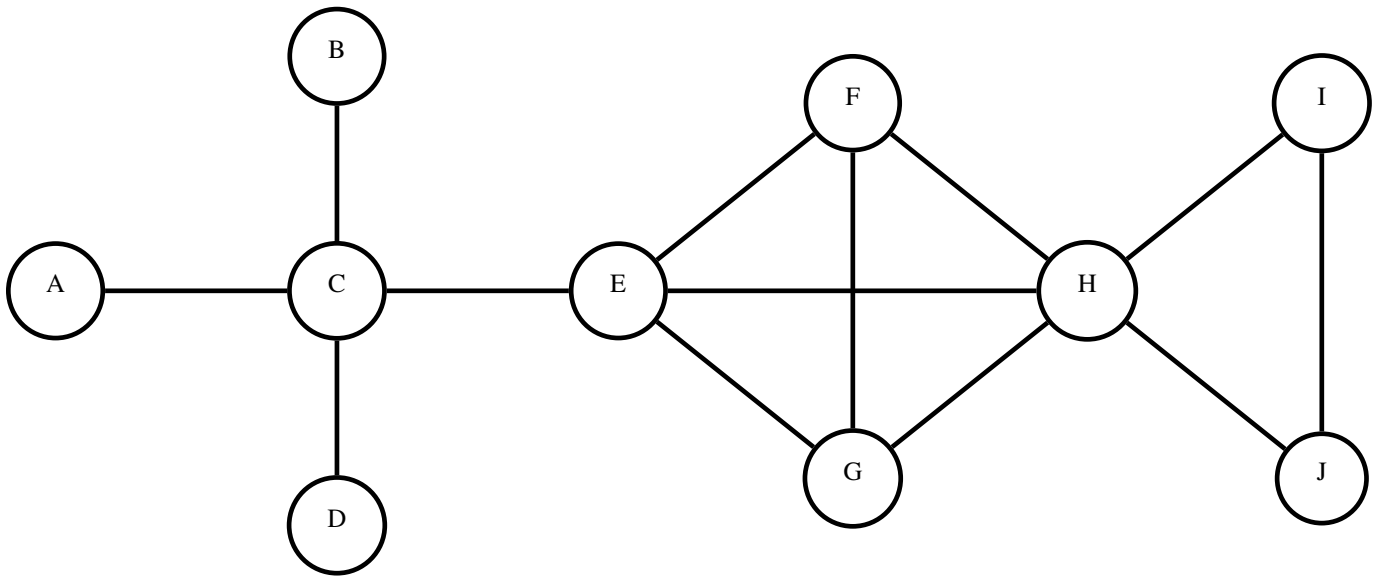

Thorup-Zwick Compact Routing

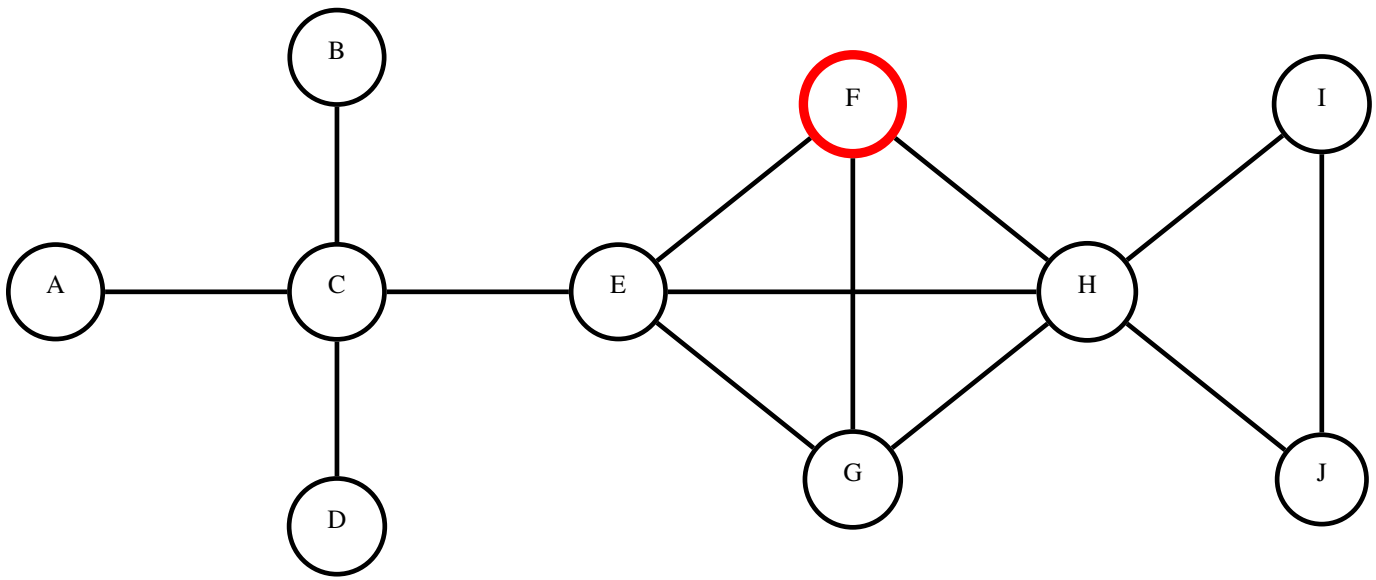

Thorup-Zwick Compact Routing

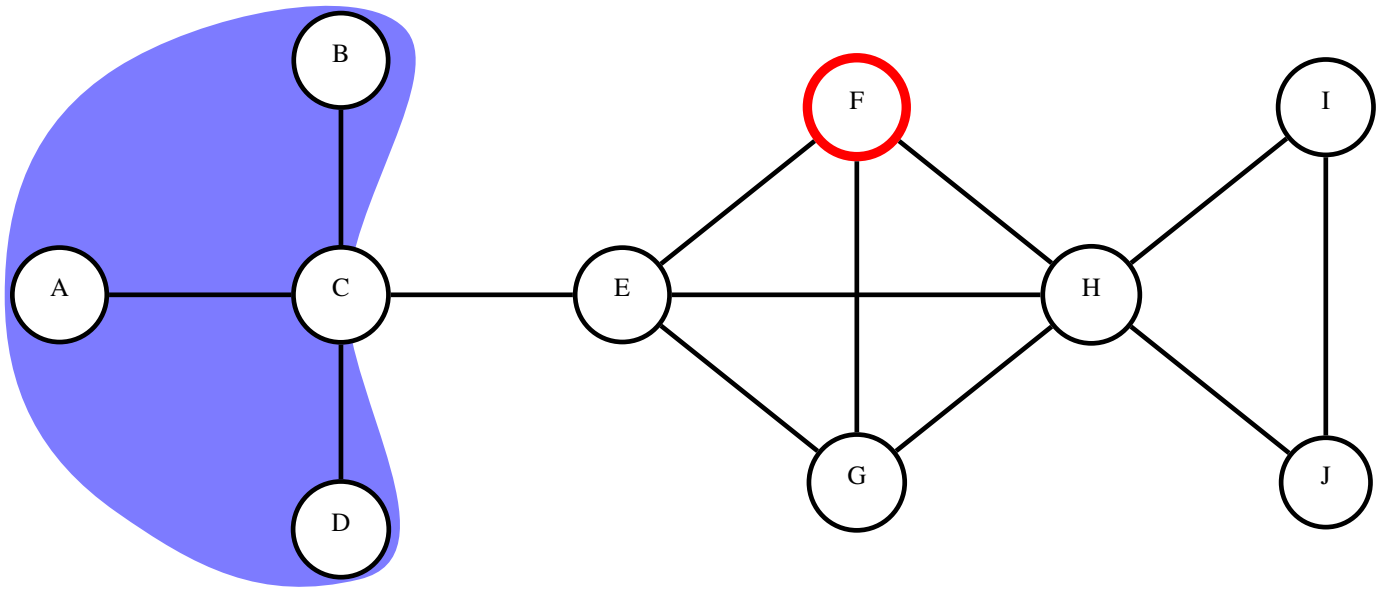

Thorup-Zwick Compact Routing

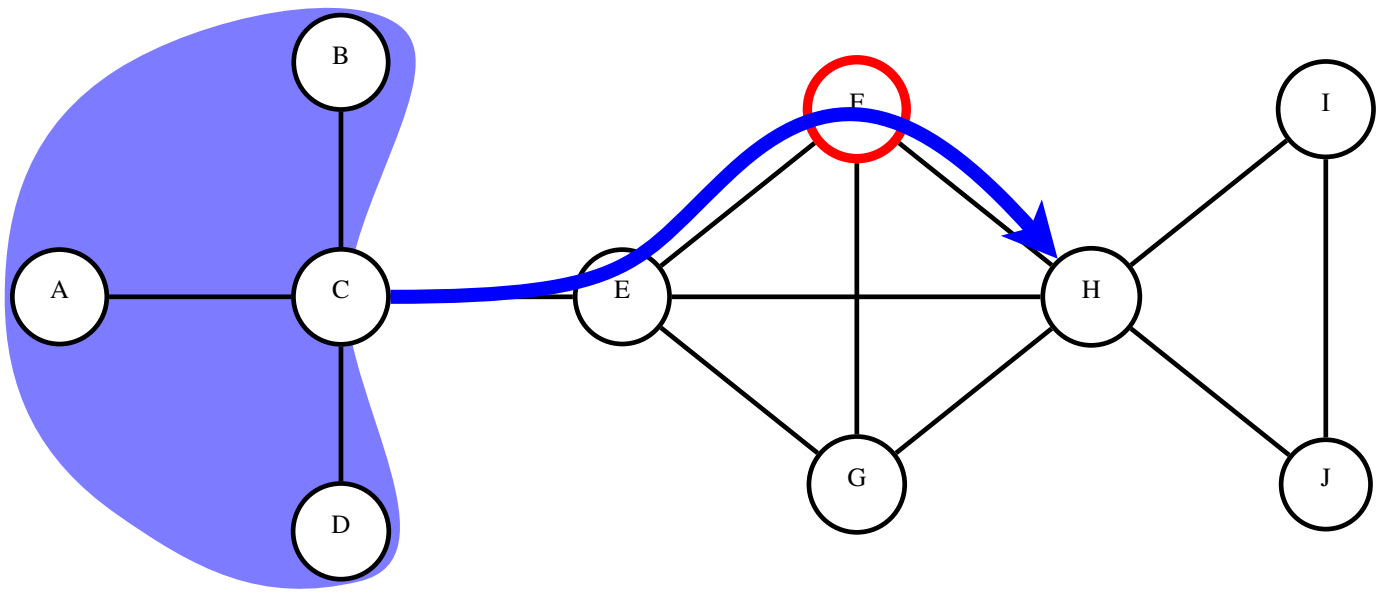

Outline

A better landmark selection

- TZ's landmark selection algorithm is not deterministic, and is inherently unstable
- The nucleus derived by k-cores decomposition is therefore better suits TZ compact routing on these graphs

Generate k-shells $k = 1, 2, ..., max - 1, max$
$A \leftarrow \emptyset$; $W = \emptyset$; $i \leftarrow max$
repeat
  $A \leftarrow A \cup i$-shell
  $i \leftarrow i - 1$
  $C_v = \{w \in V \big| d(v, w) < D(w, A), \text{ for every } v \in V\}$
  $W = \{v \in V \big| |C_v| > 4n/s\}$
until $W \neq \emptyset$
return A

A better landmark selection

- Using this algorithm, on 98.2% of the graphs tested, TZ constraints are met with the nucleus alone

Results: Path Stretch (TZ_k, Multiplicative)

Results: Path Stretch (TZ_k, Multiplicative)

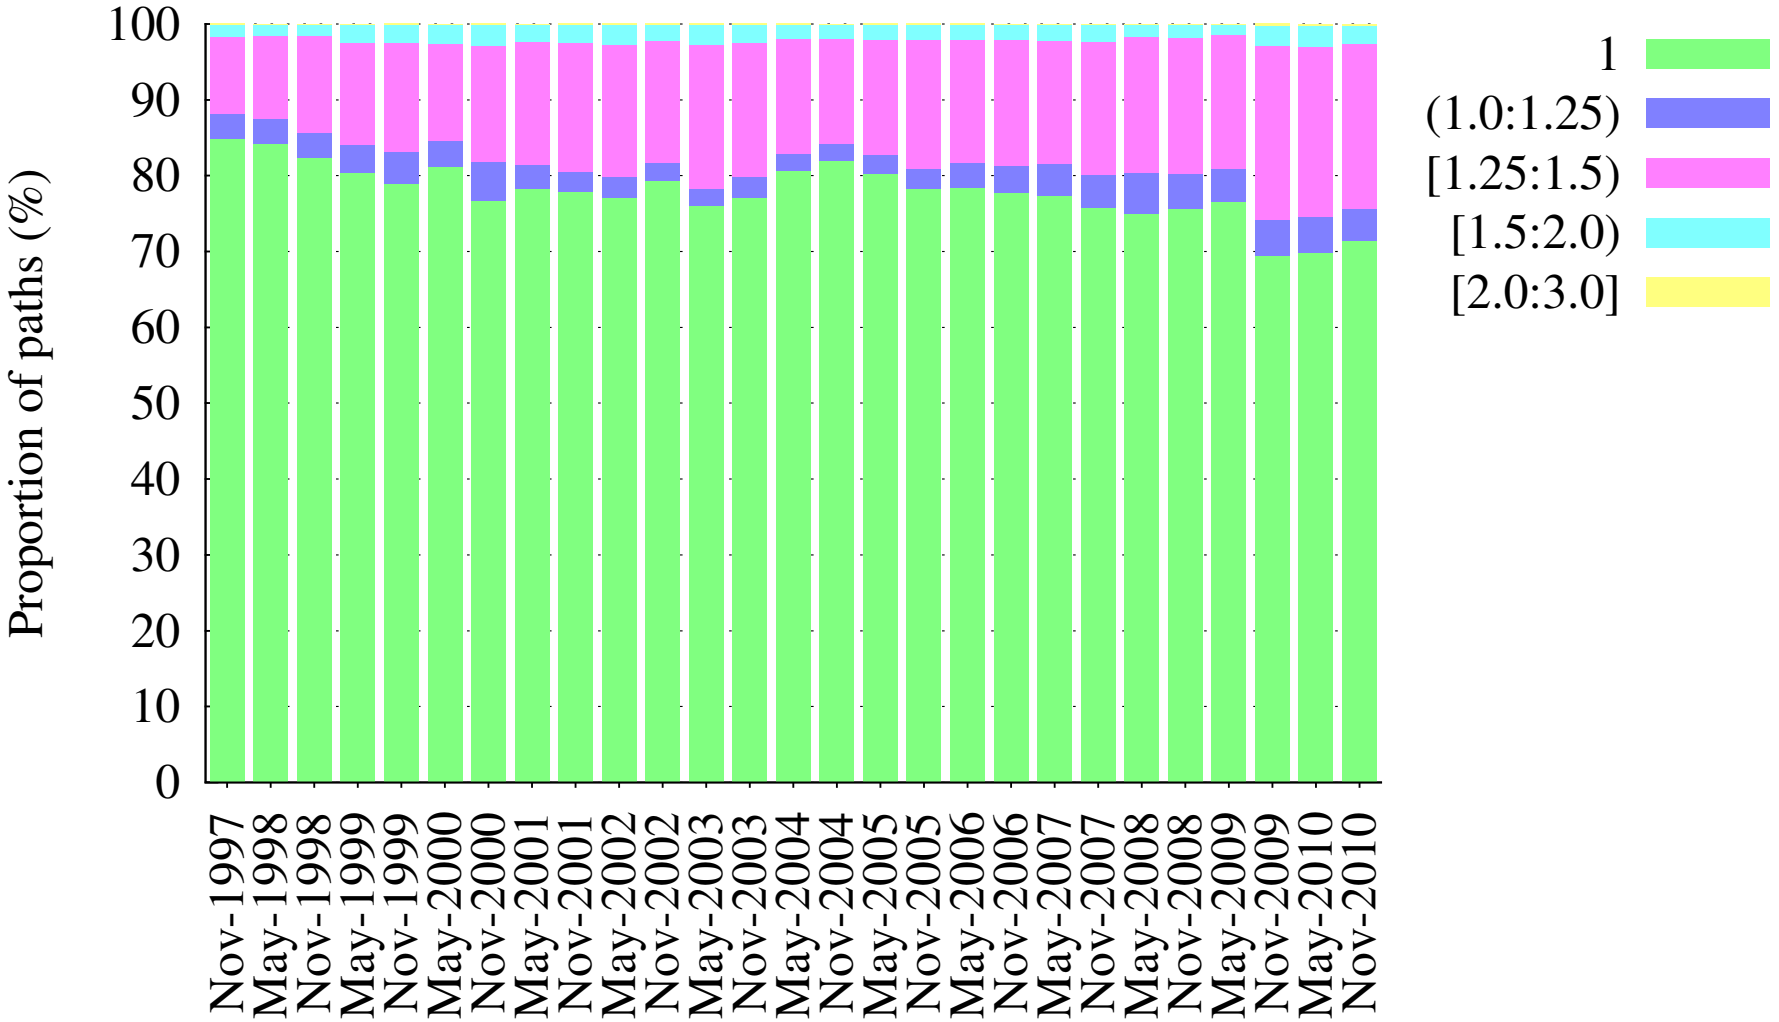

Results: Path Stretch (TZ_k, Additive)

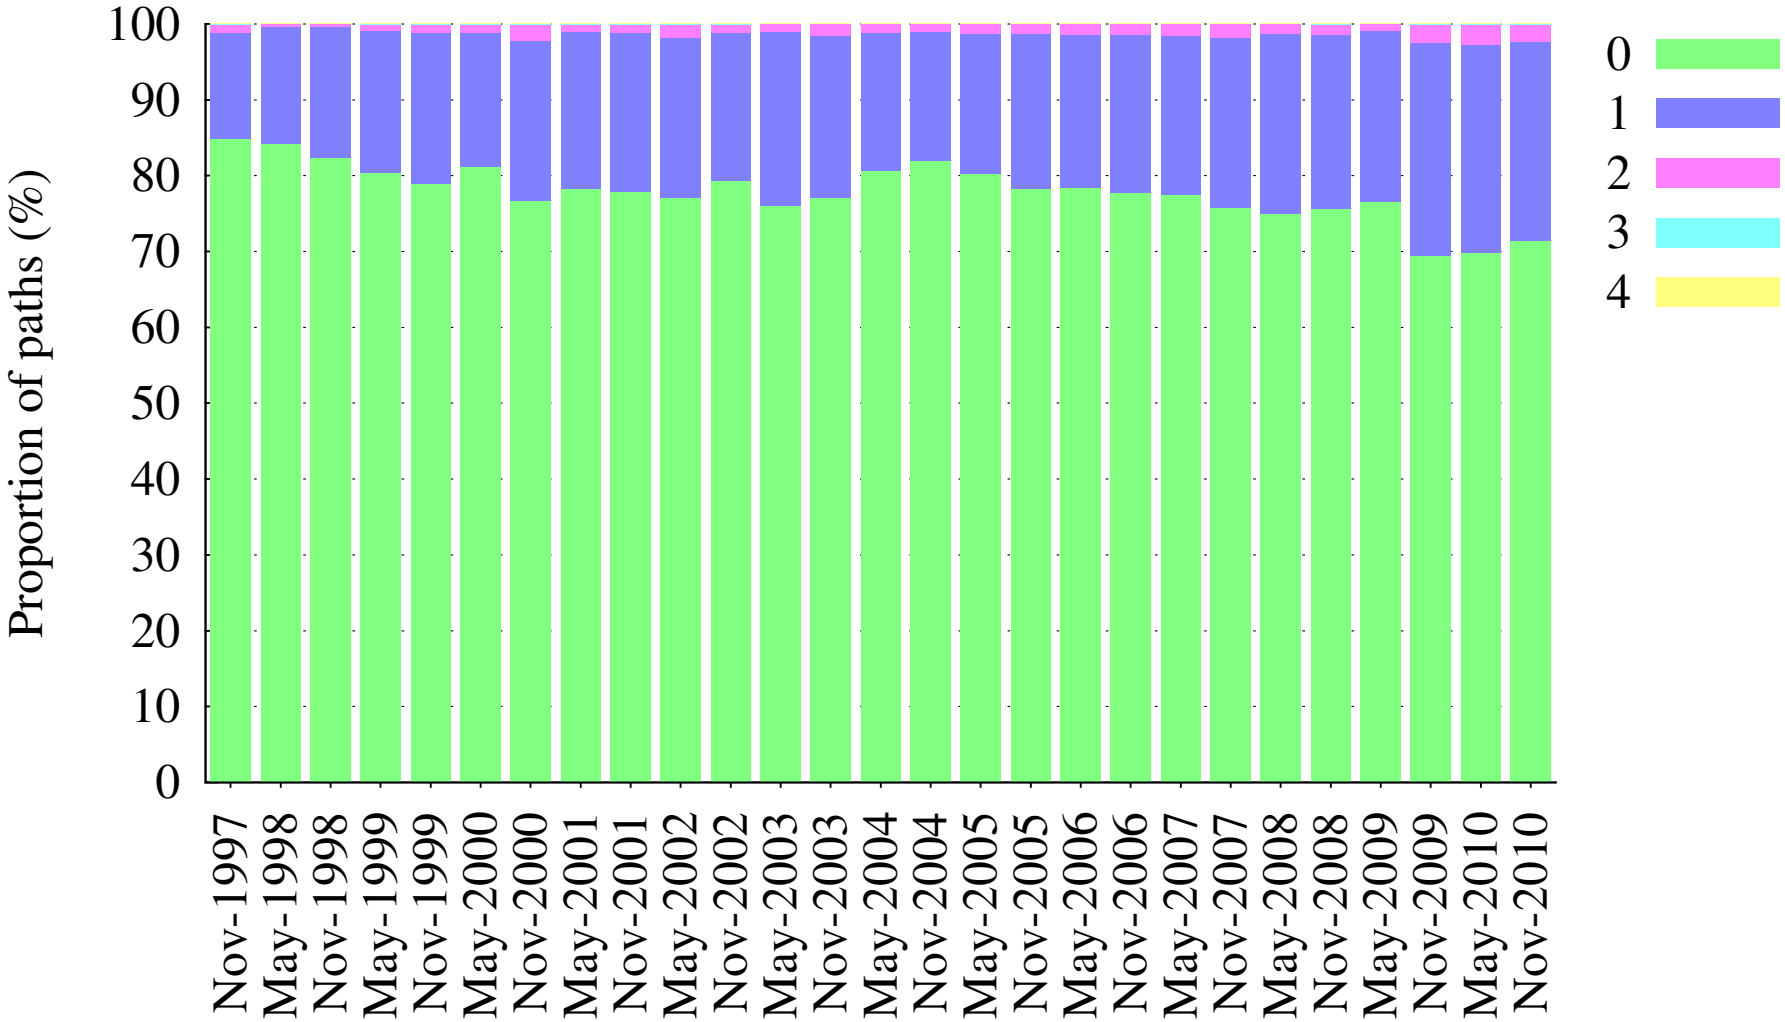

Results: Forwarding Table Sizes

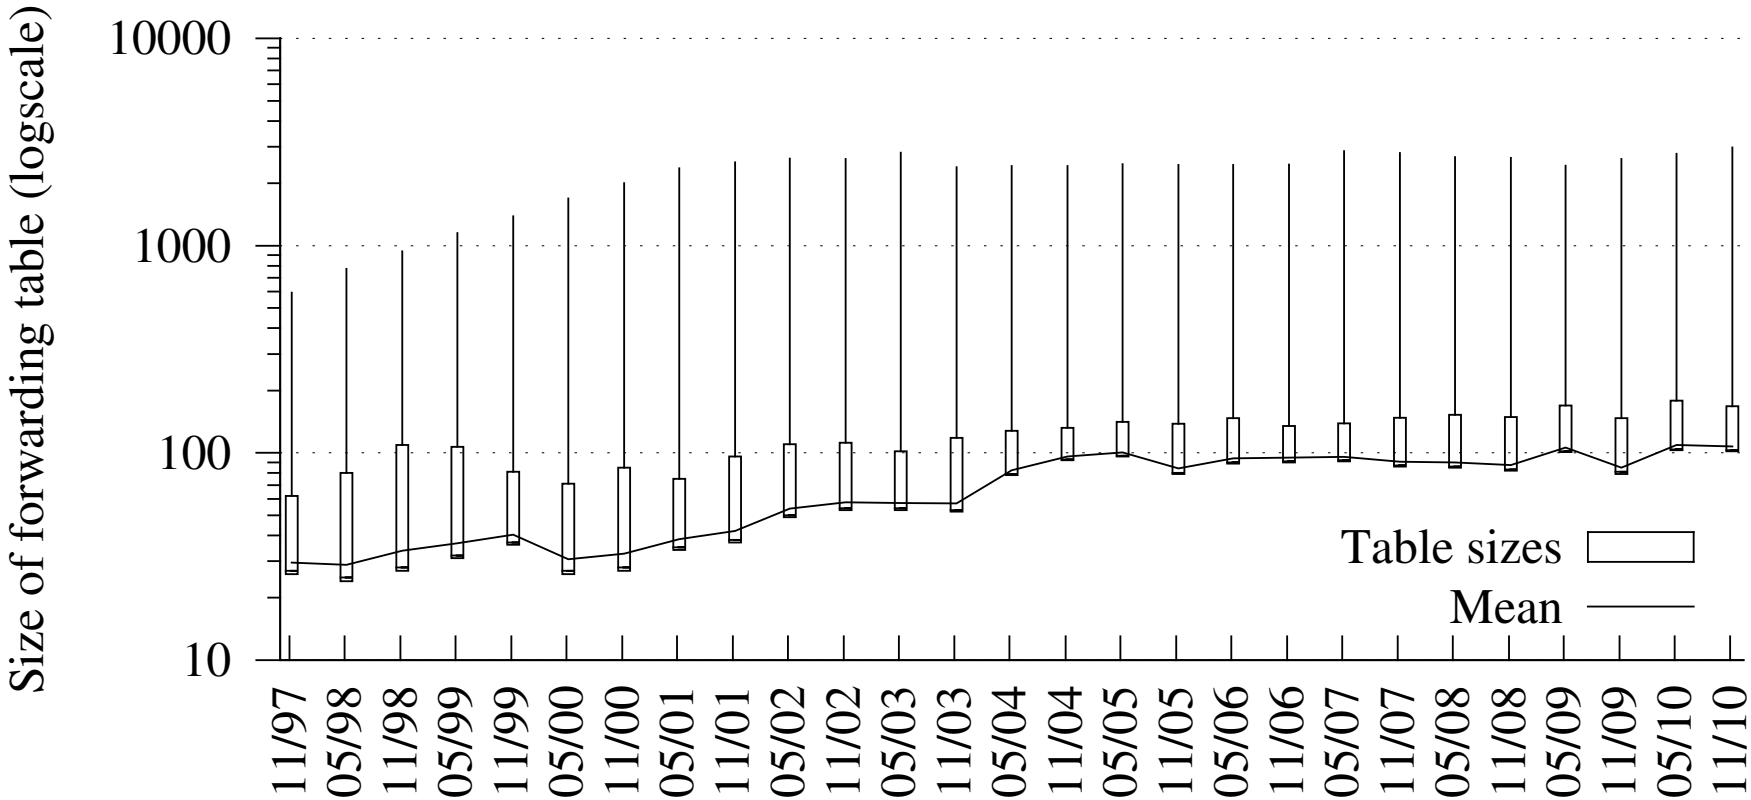

Outline

In Conclusion...

- We have shown the long term stability at the heart of the AS graph can be used by alternative routing algorithms
- The TZ compact routing algorithm on the AS graph generates tiny forwarding tables, which may offer a long-term route toward forwarding place stability for the inter-domain Internet

Future Work

- We are working on policy-compliance for compact routing algorithms
- A distributed k-cores algorithm is feasible
- We can achieve shortest path routing in over 98% of cases just by modifying landmark selection when forwarding packets

Questions?

Results: Path Stretch (TZ, Multiplicative)

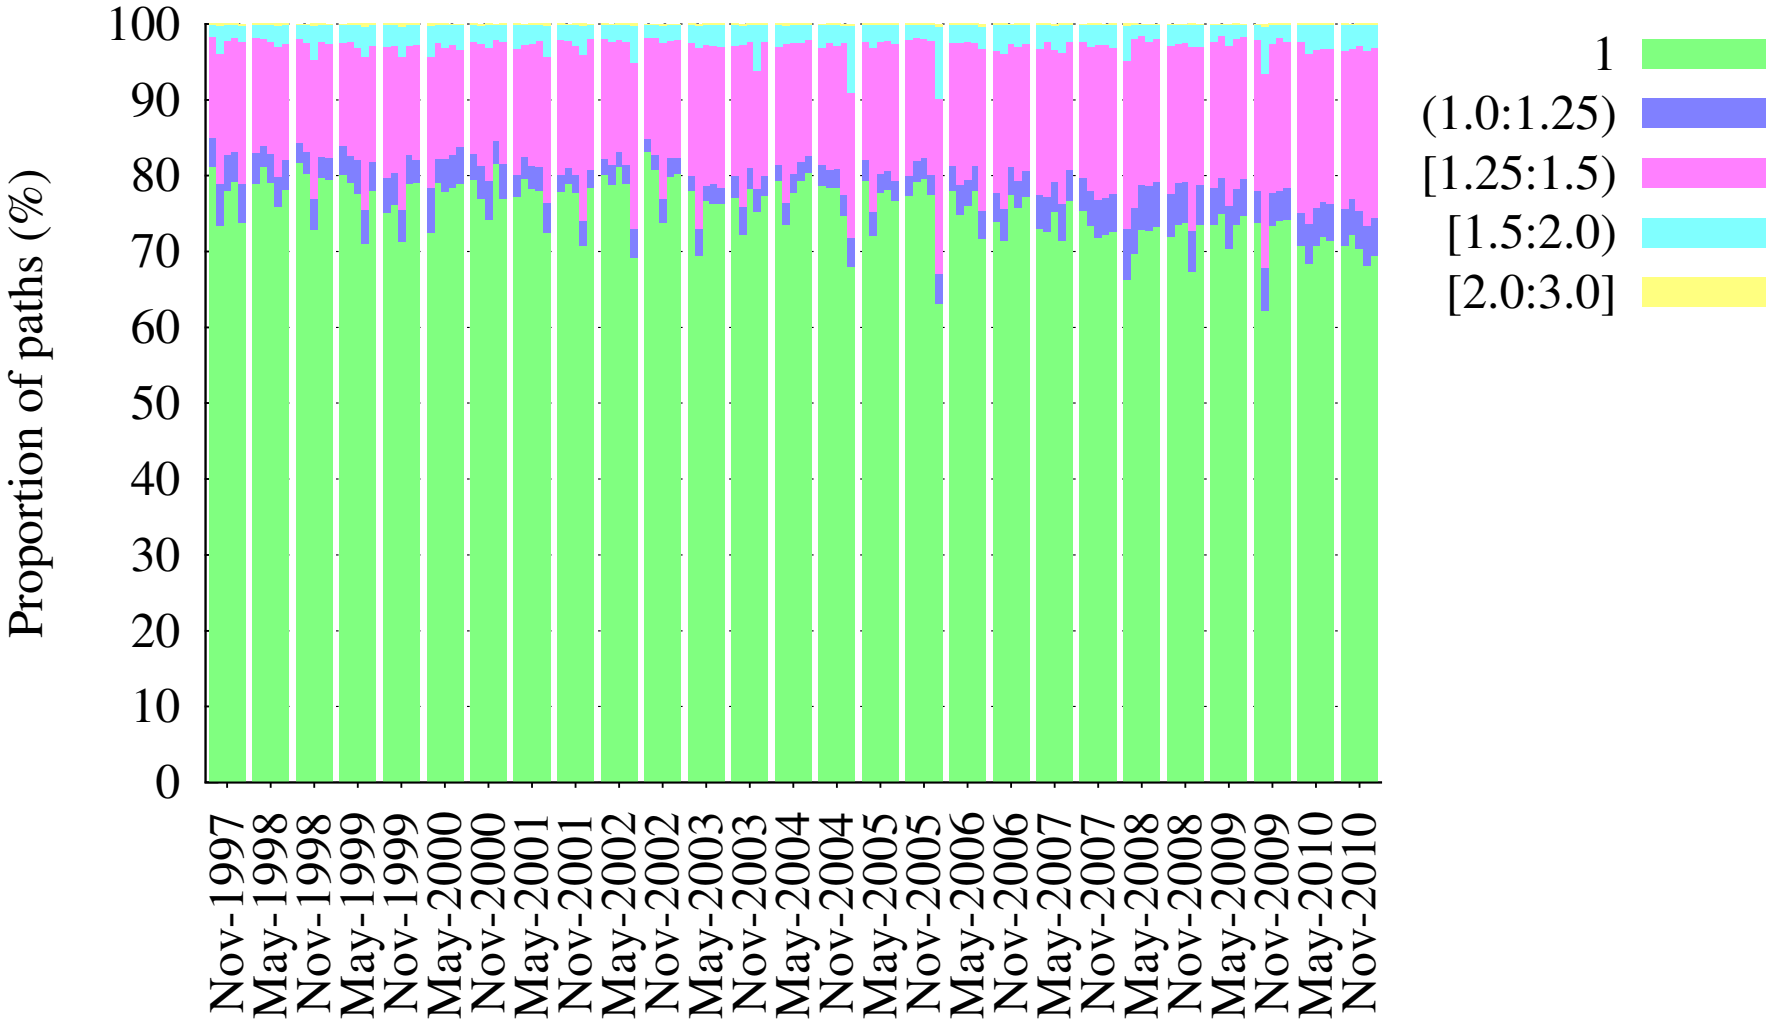

Results: Path Stretch (TZ, Additive)