Compact Routing on the Internet AS-Graph

Stephen Strowes, University of Glasgow
Graham Mooney, Cisco Systems Ltd.
Colin Perkins, University of Glasgow

Context

Previous work:

- Compact routing has shown promise for reducing forwarding state
- Previous work has evaluated synthetic “Internet-like” graphs

We use Internet topologies spanning 14 years

Compact routing?

Shortest-path routing:

Space: linear space

Stretch: 1

Compact routing:

Space: sublinear space

Stretch: 3

Stretch-3 sounds bad!

Previous work showed actual performance to be much closer to stretch-1

(on synthetic graphs at least)

We perform systematic experimental analysis of two algorithms on Internet topologies

Overview

Two compact routing algorithms:

- Thorup-Zwick (TZ)
- Brady-Cowen (BC)

Overview: Thorup-Zwick

Routes via landmark nodes

Defines forwarding entries using proximity to landmarks

Nodes forward packets destination's landmark if destination not in local forwarding table

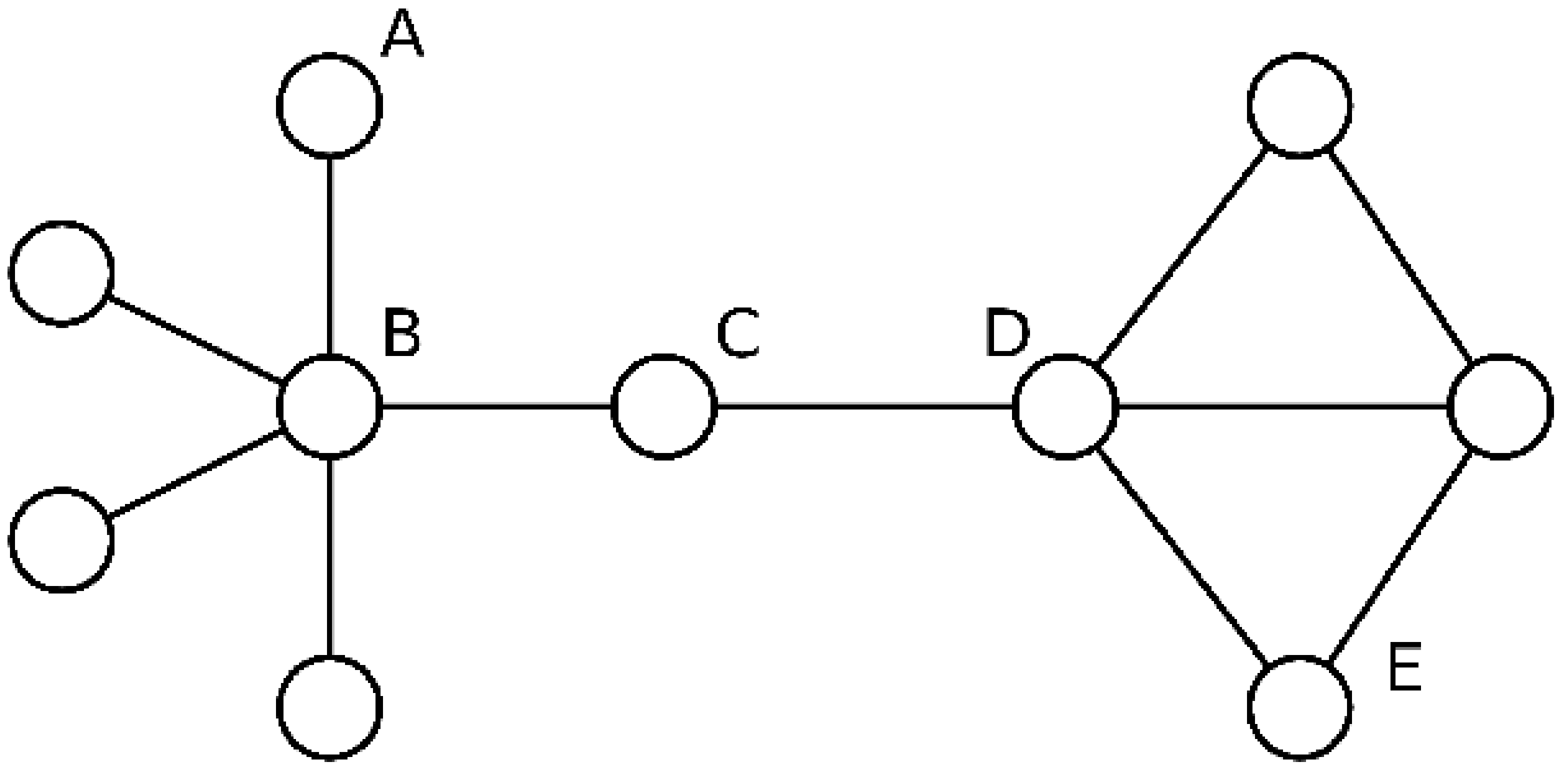
A
B
C
D
E

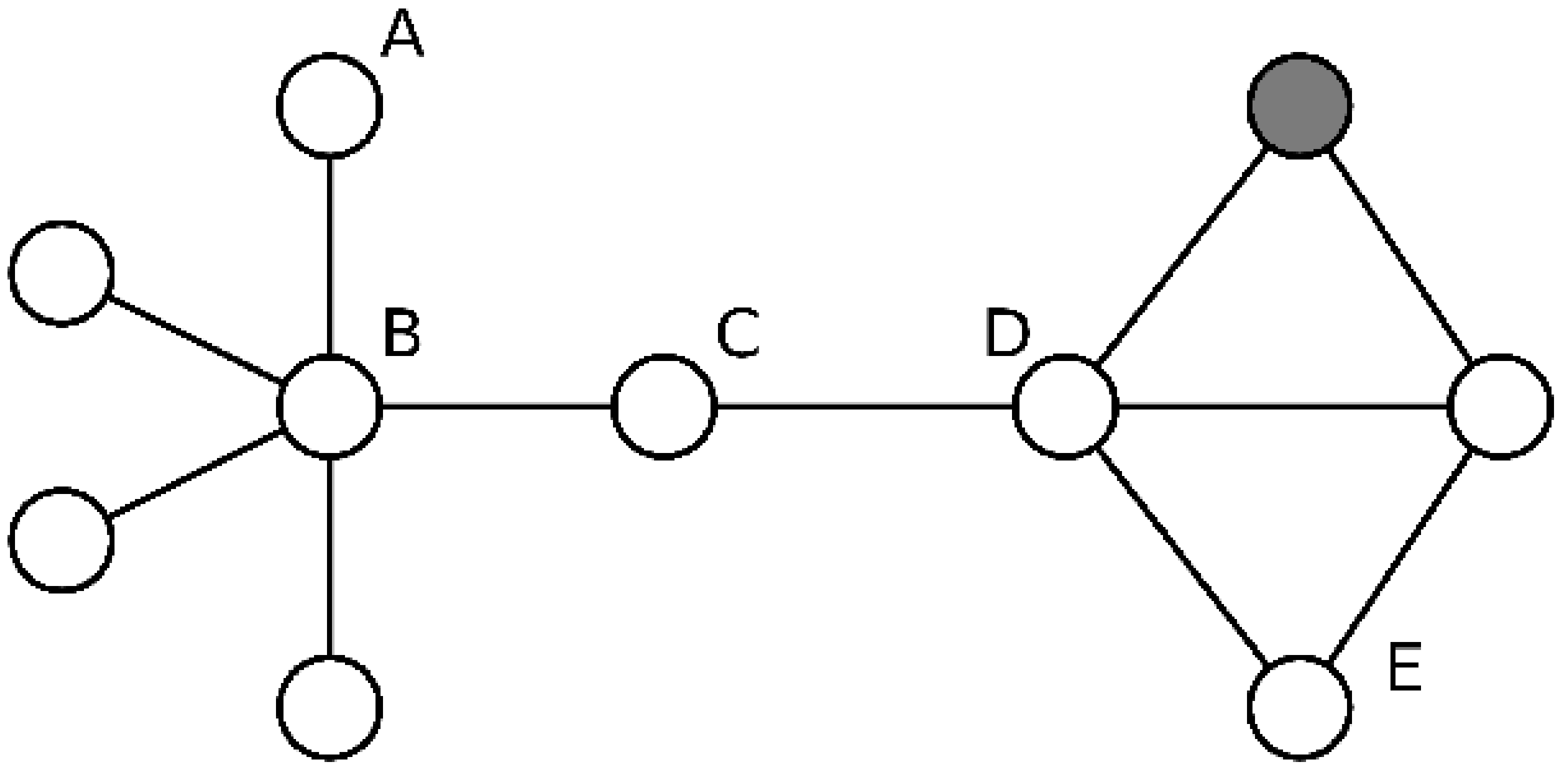
A
B
C
D
E

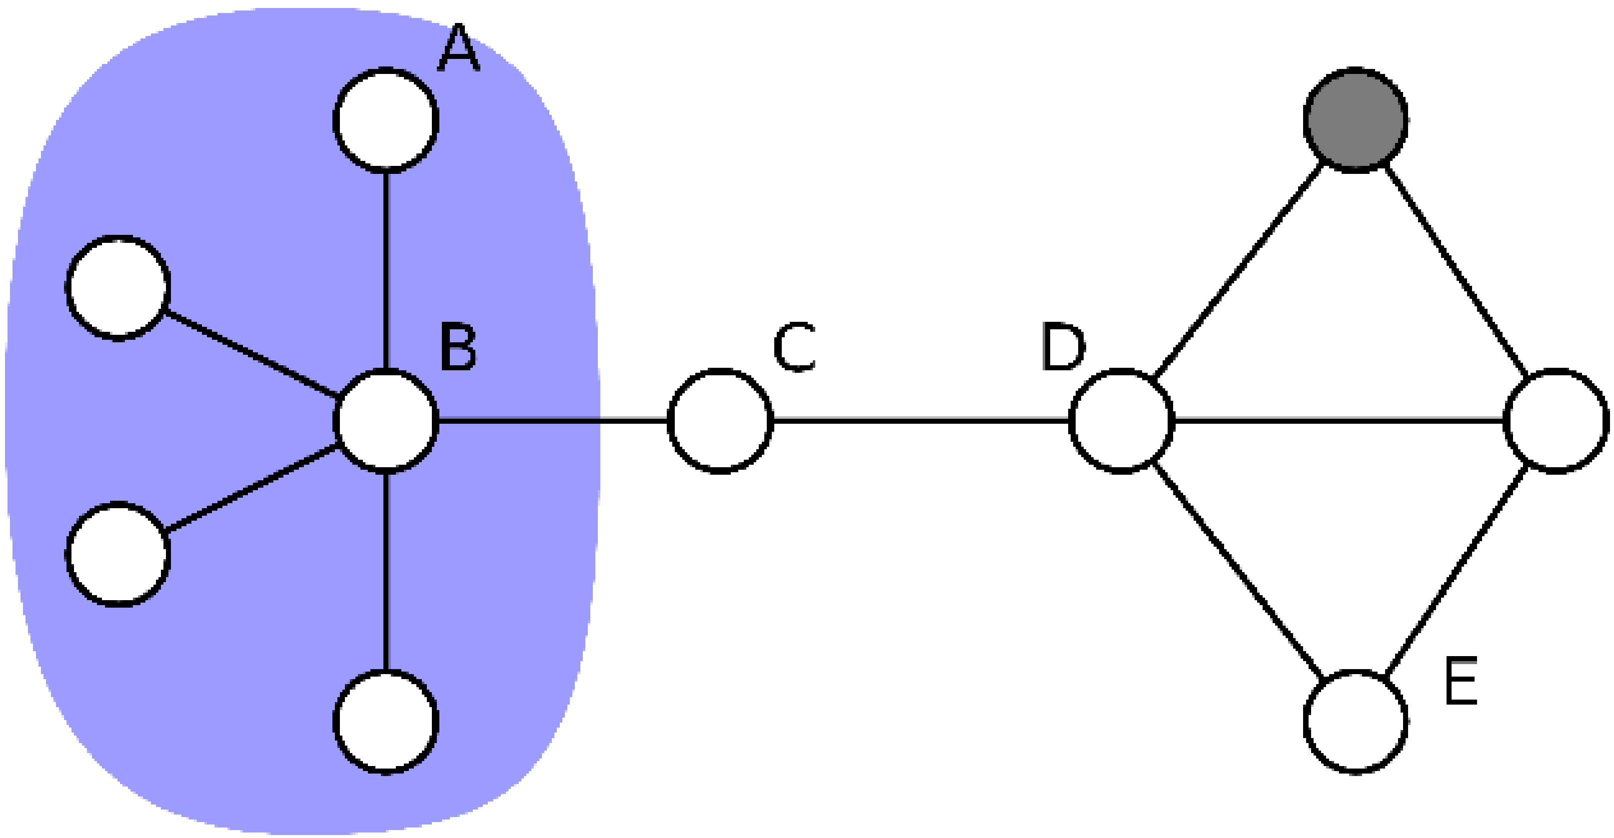
A
B
C
D
E

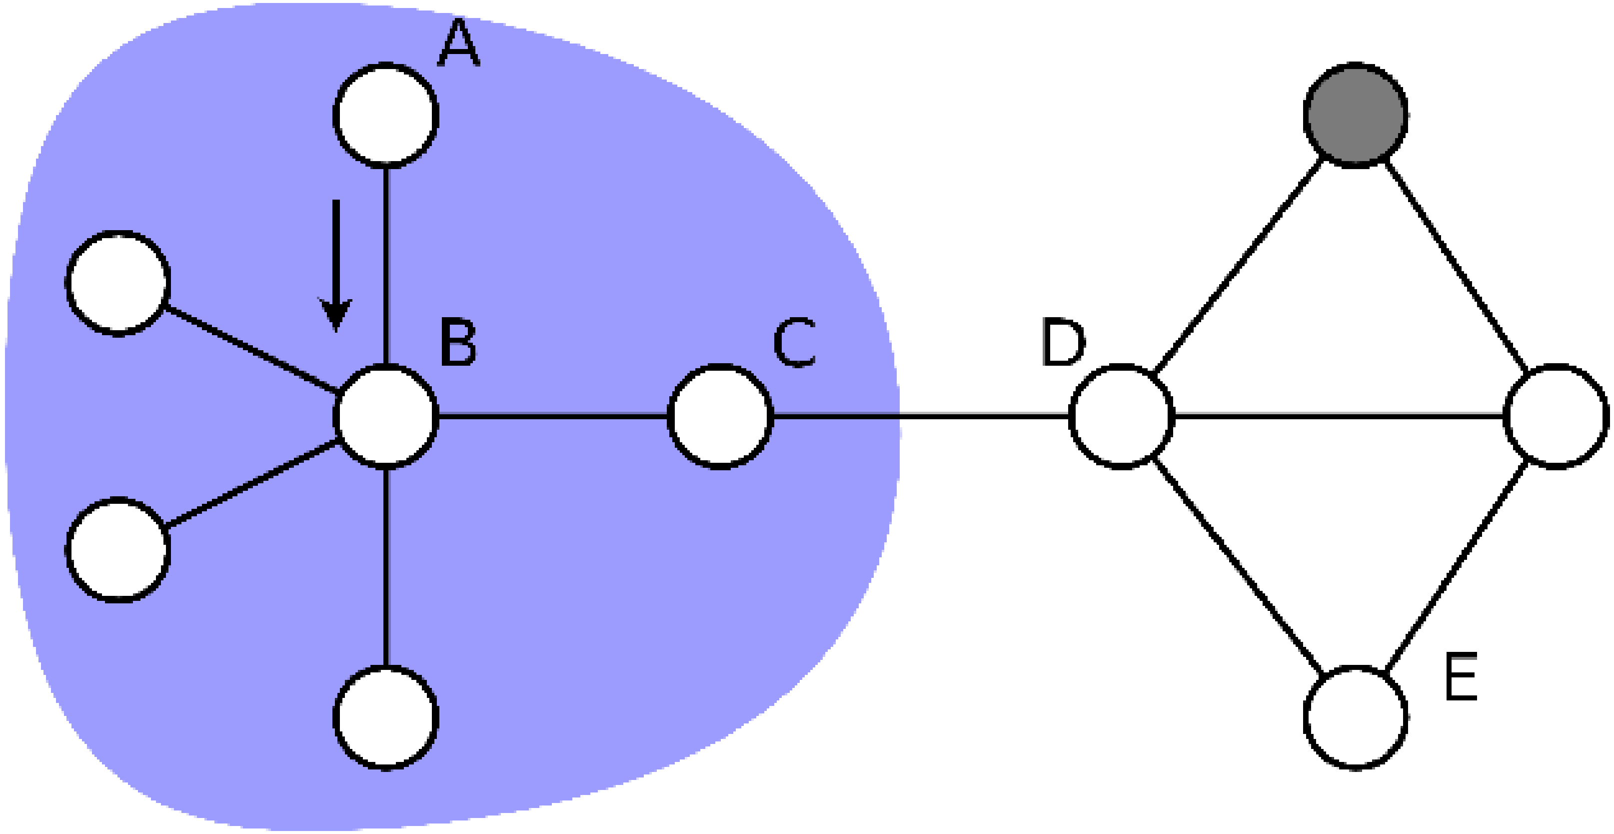
A
B
C
D
E

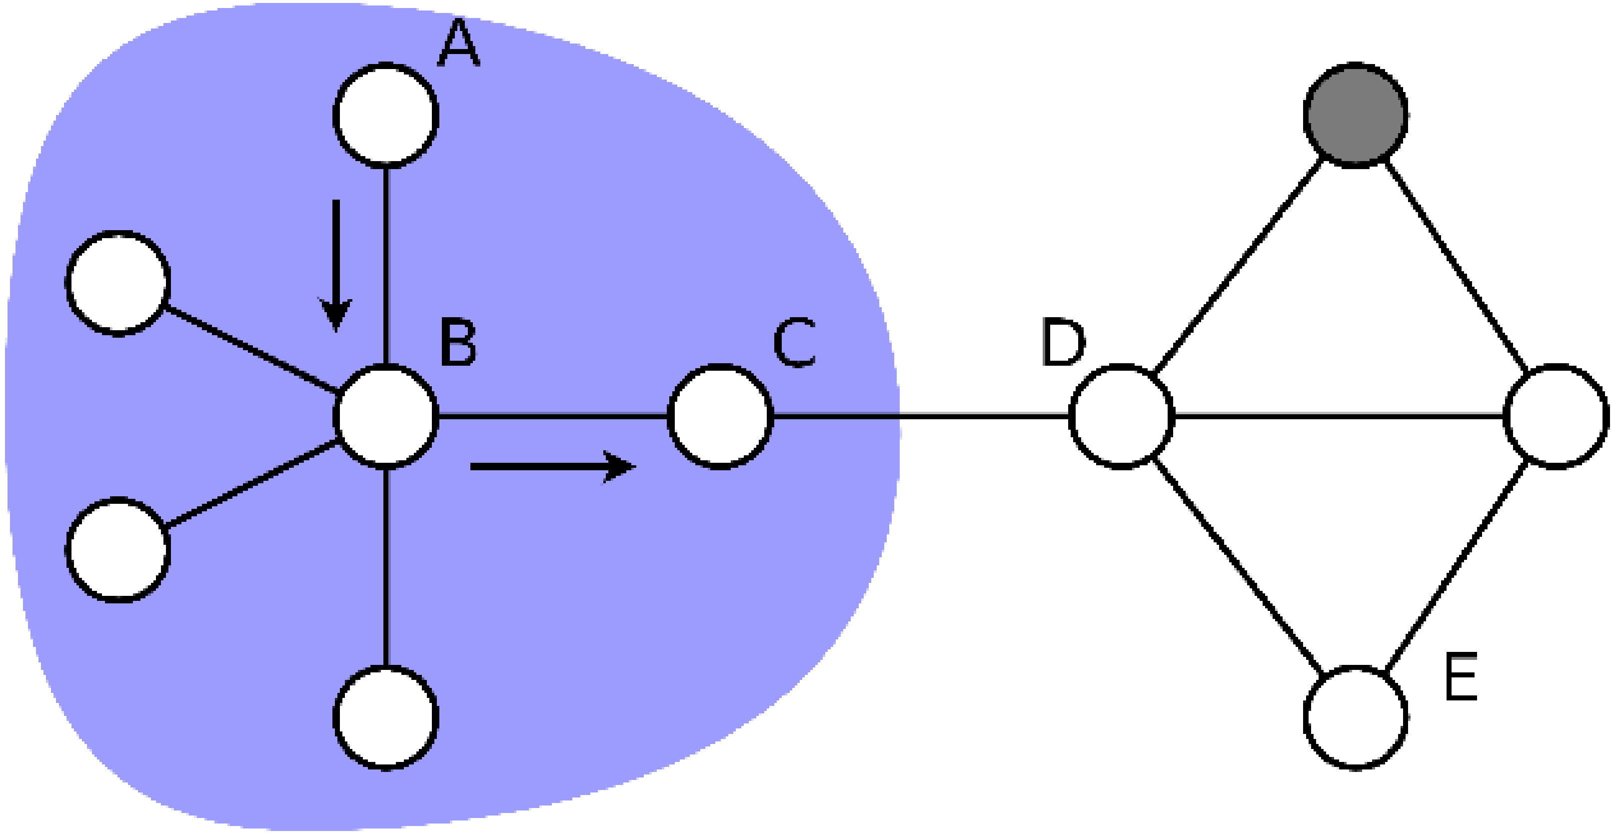
A
B
C
D
E

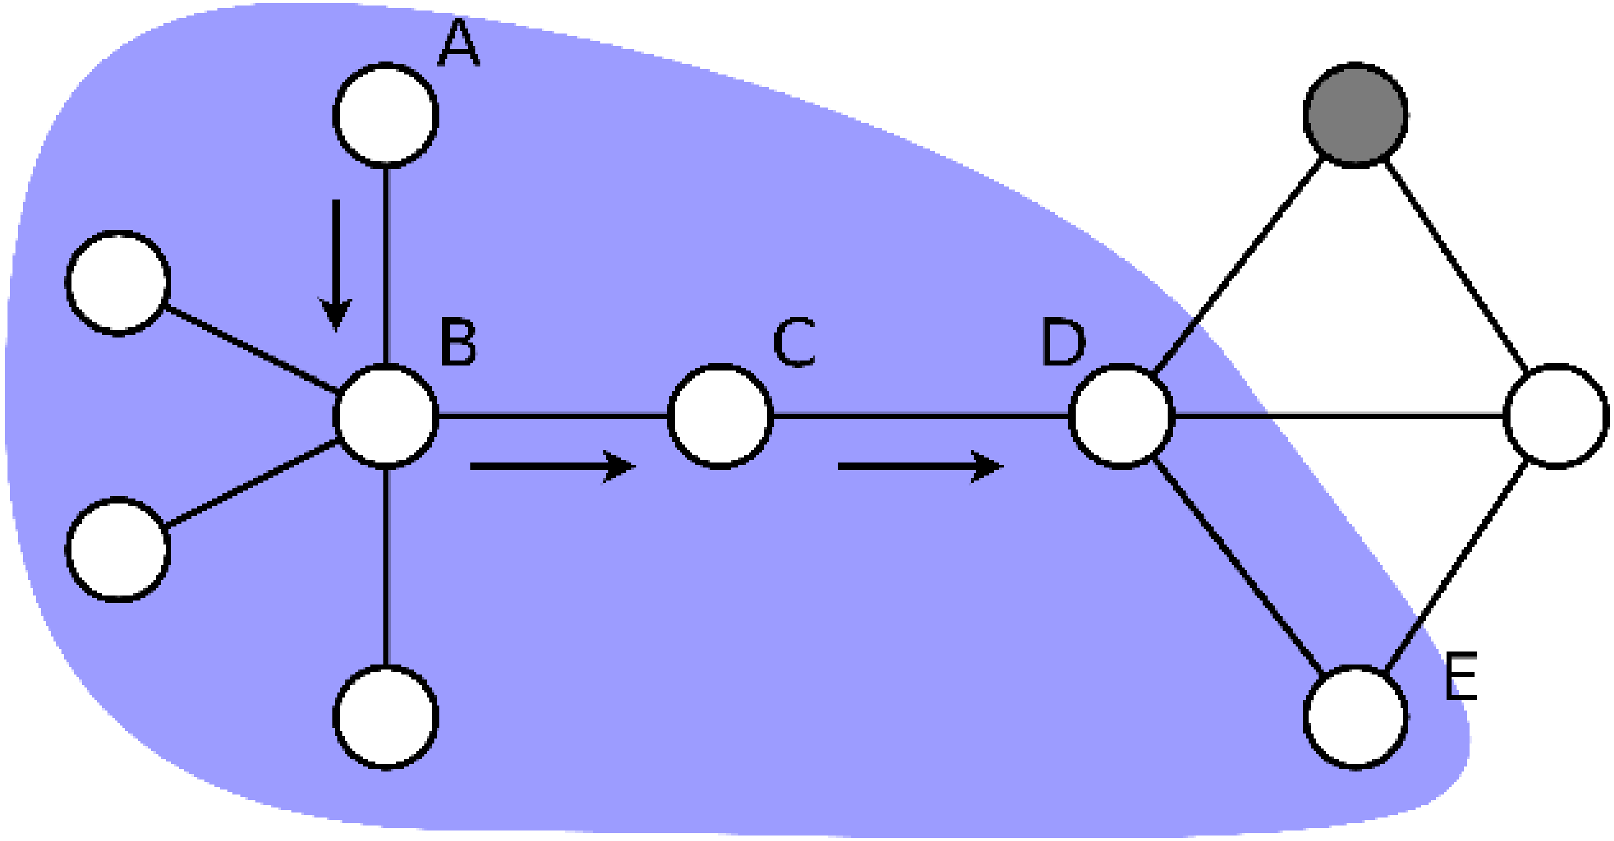
A
B
C
D
E

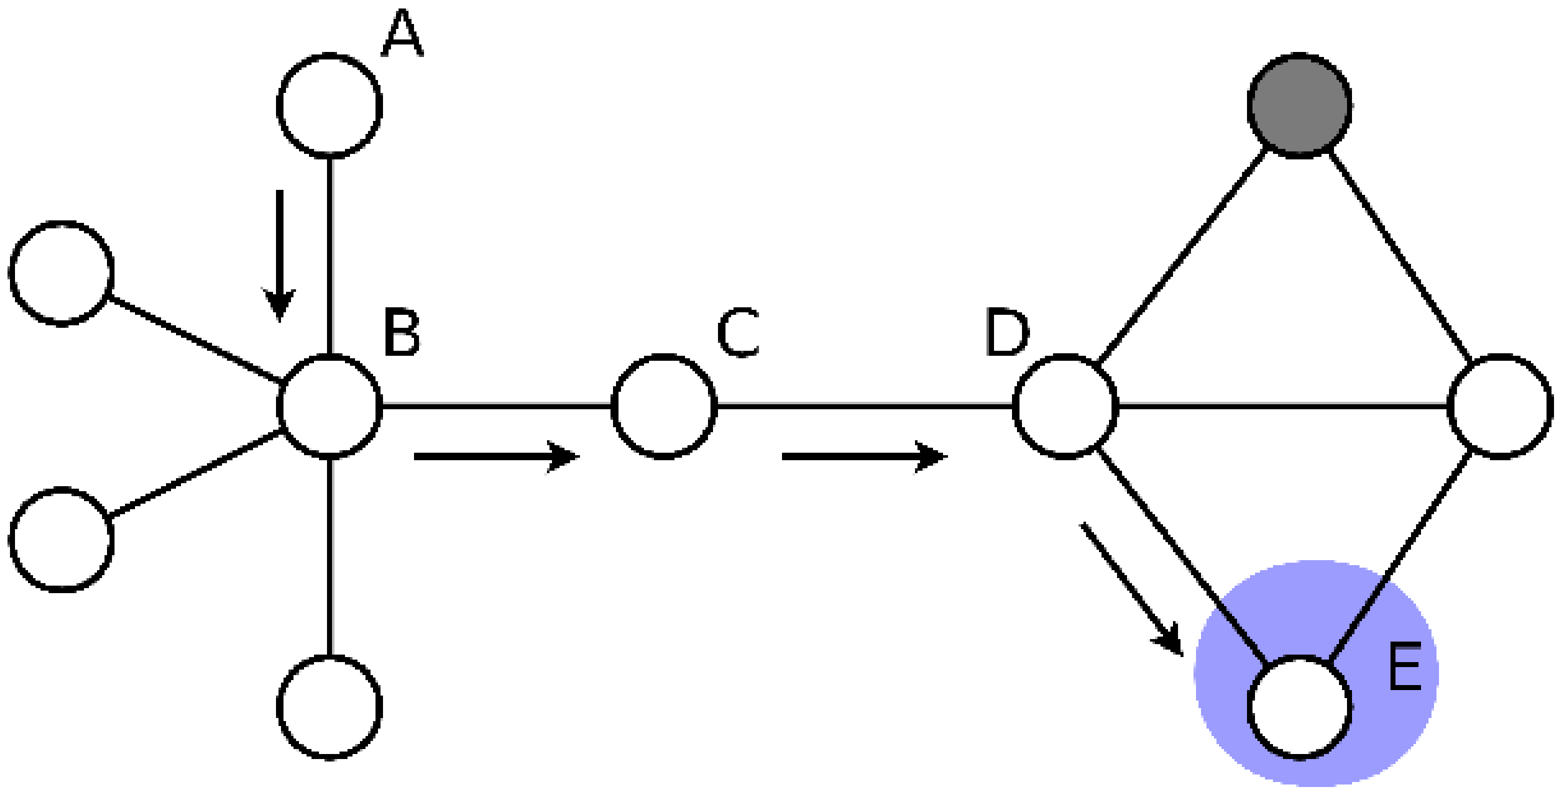
A
B
C
D
E

TZ Asymmetry

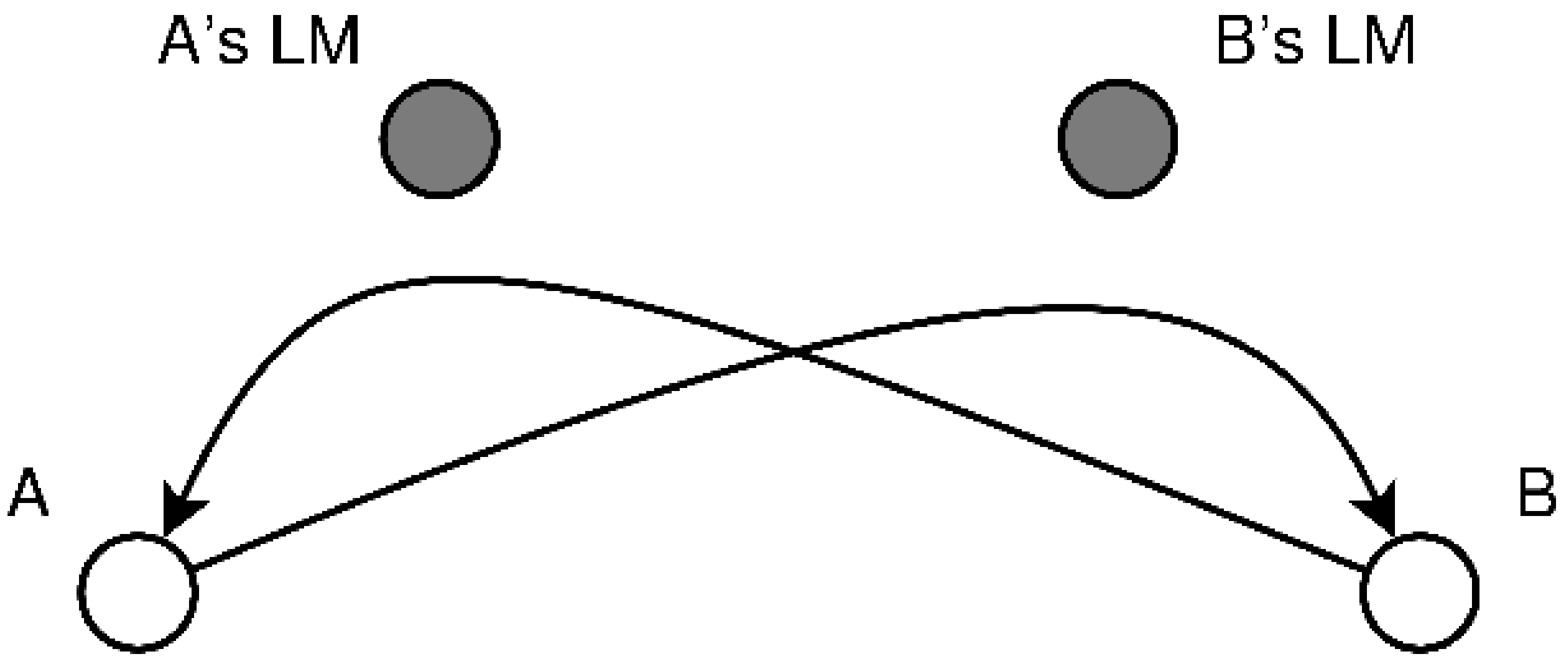

Overview: Brady-Cowen

Builds a primary spanning tree

Then, builds additional, smaller spanning trees on the periphery

Uses a distance labelling to select, at source, best tree to forward packets to desination

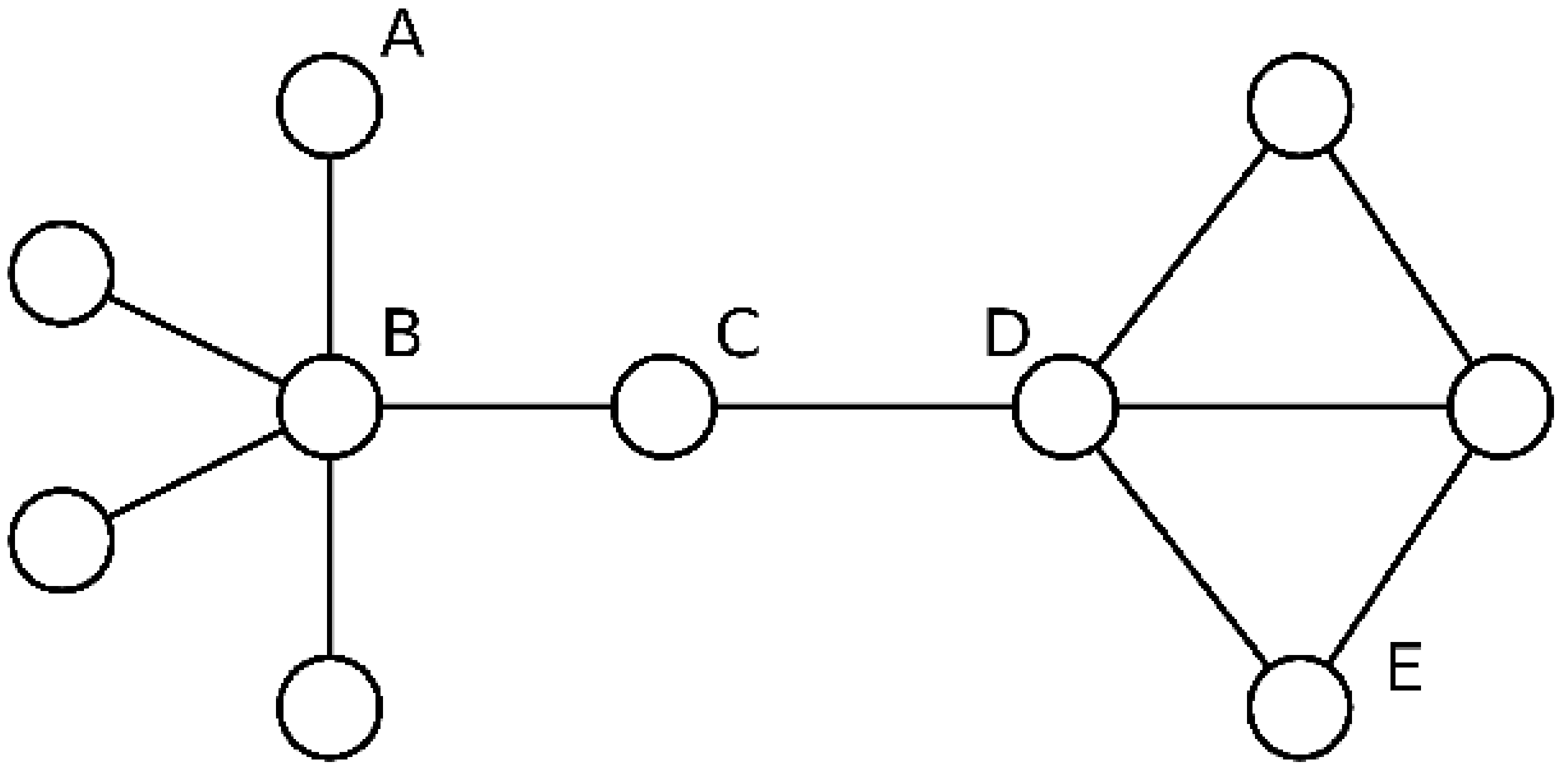
A
B
C
D
E

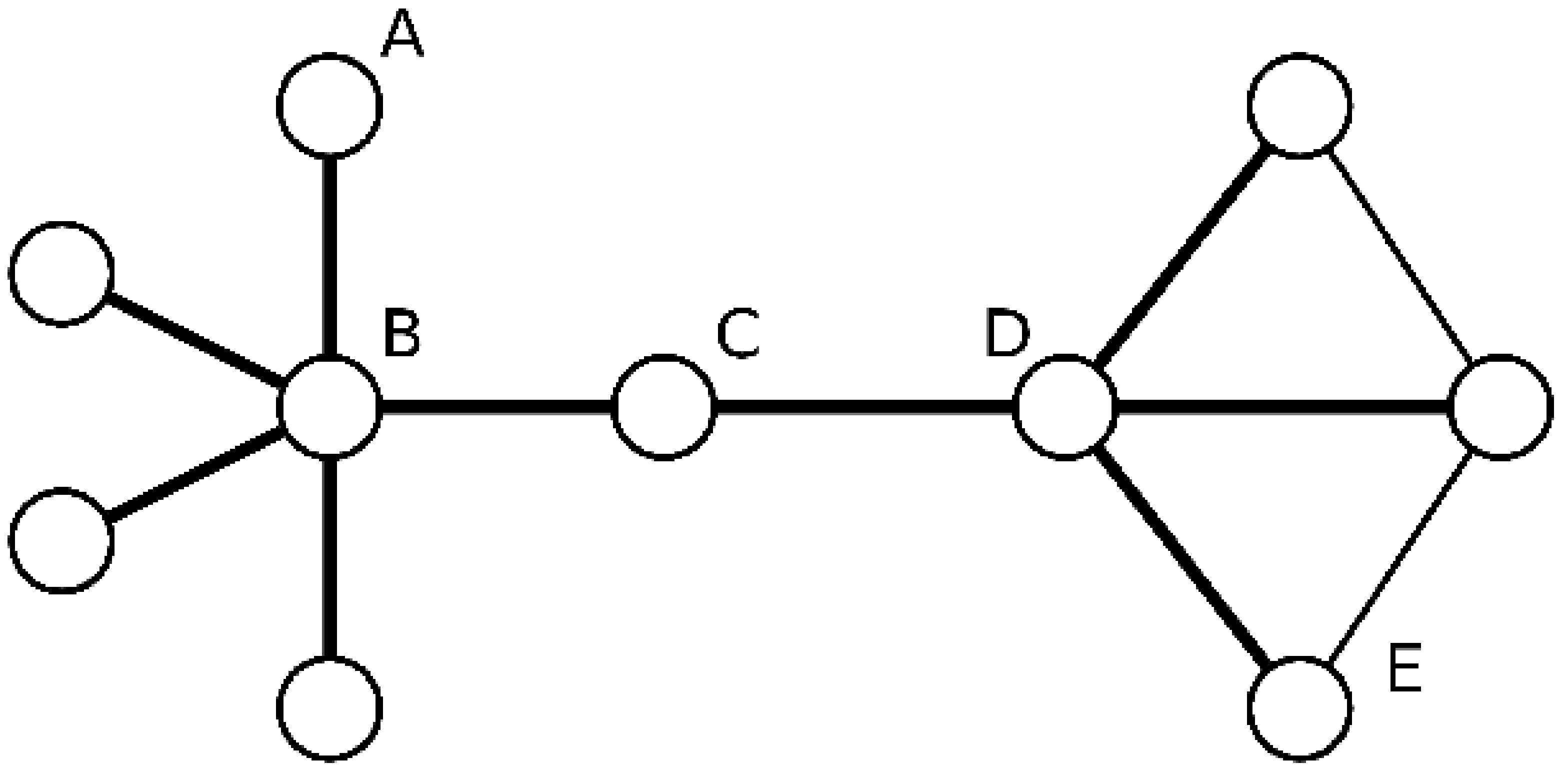
A
B
C
D
E

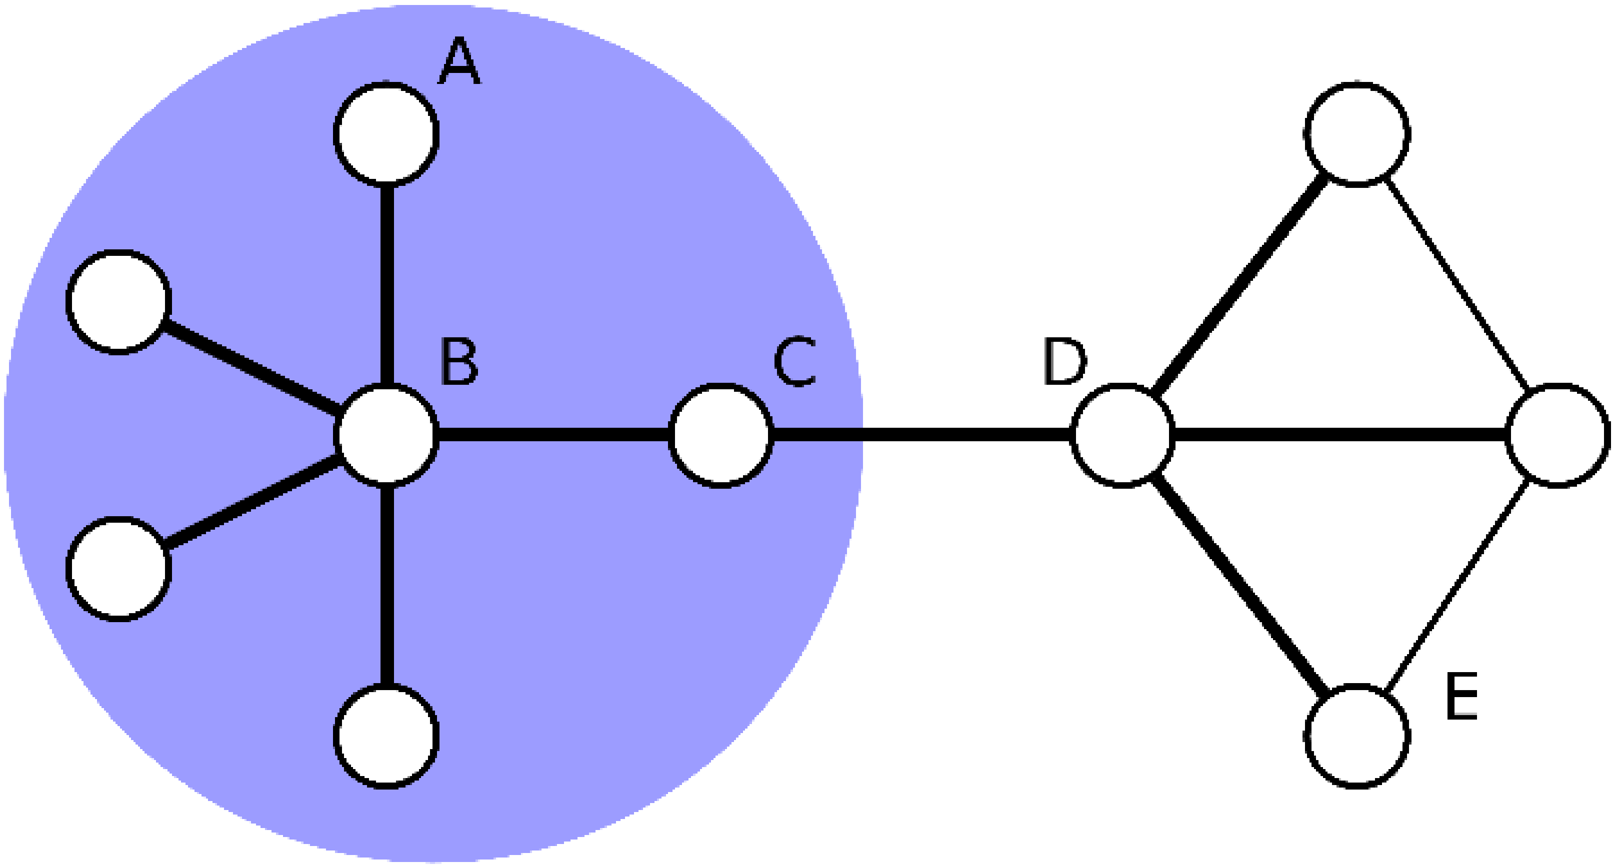
A
B
C
D
E

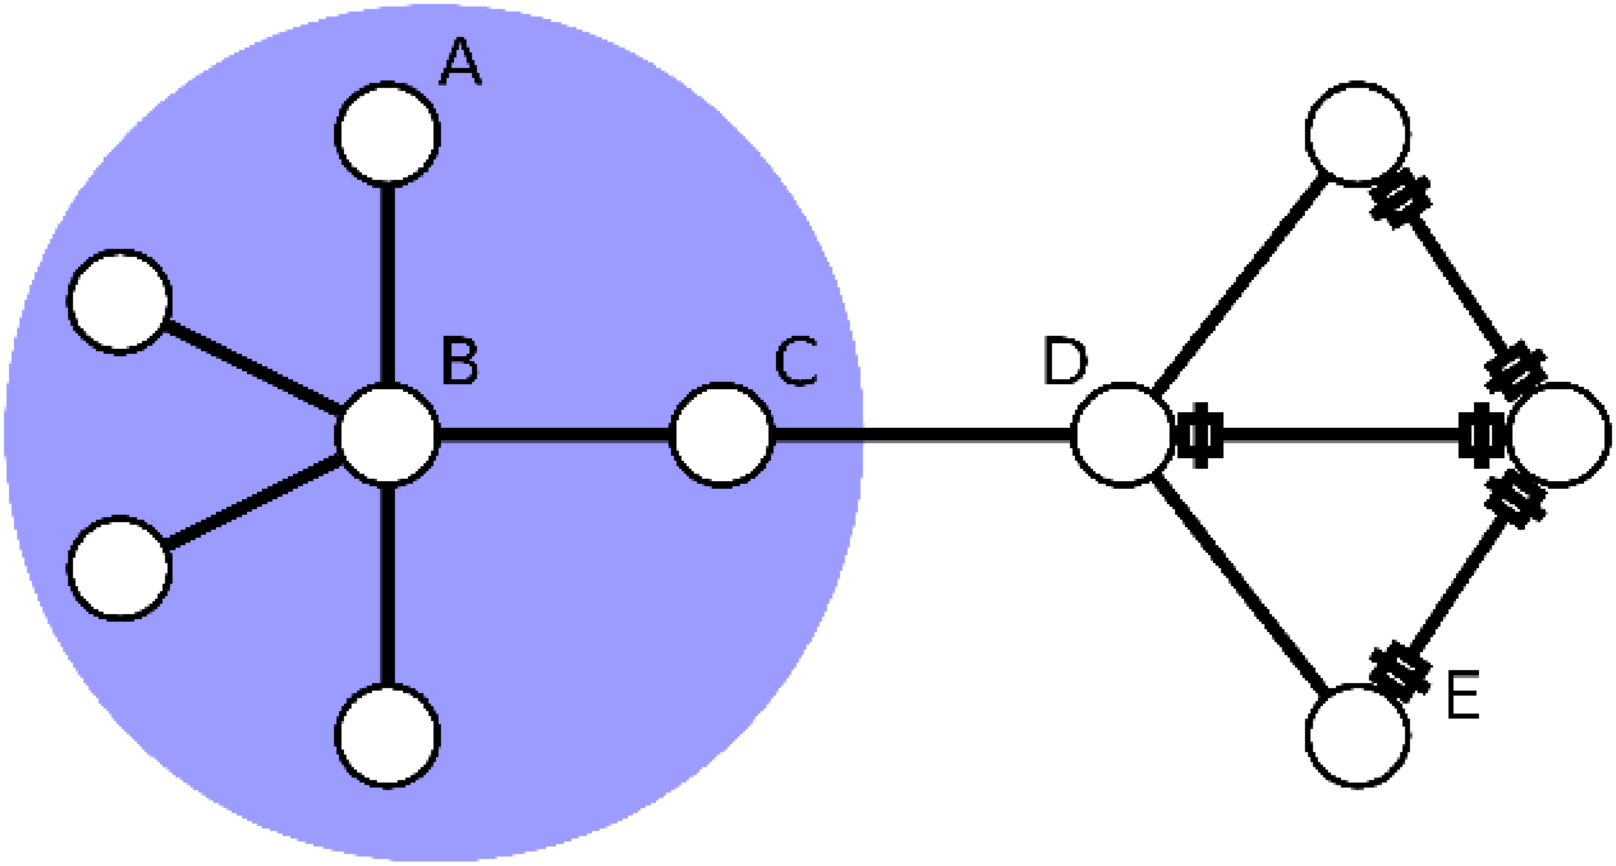
A
B
C
D
E

Experimental Methodology

- Generate AS topologies
  - Route Views BGP tables from 1997 through 2010
- Run TZ & BC algorithms on all graphs to generate appropriate node labels/landmarks
- Determine path lengths
  - TZ: Simulate forwarding from all nodes to 1% (random) of rest of network, and back
  - BC: Deterministic, and known after pre-computation

Results

In the paper, we evaluate:

- Tweaks/parameters for each algorithm
- Multiplicative path stretch
- *Additive* path stretch
- Forwarding table sizes
- TZ landmark selection frequency and distribution of landmark set sizes

Results: TZ, multiplicative stretch

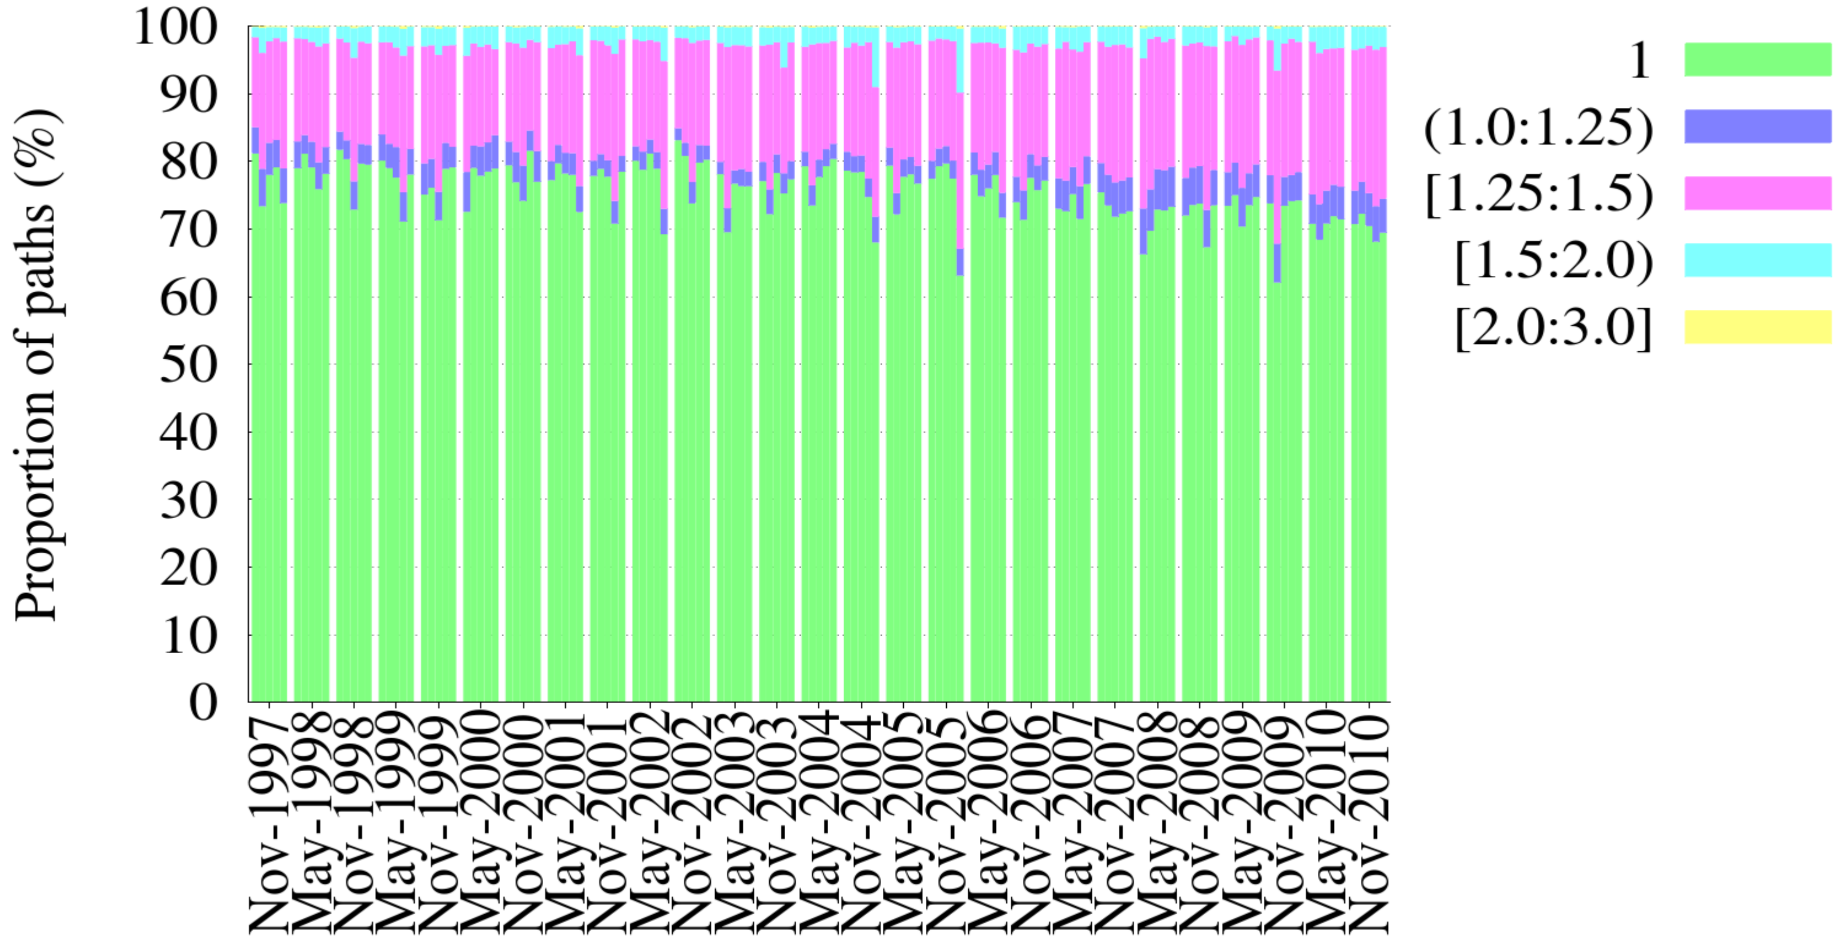

Results: TZ, additive stretch

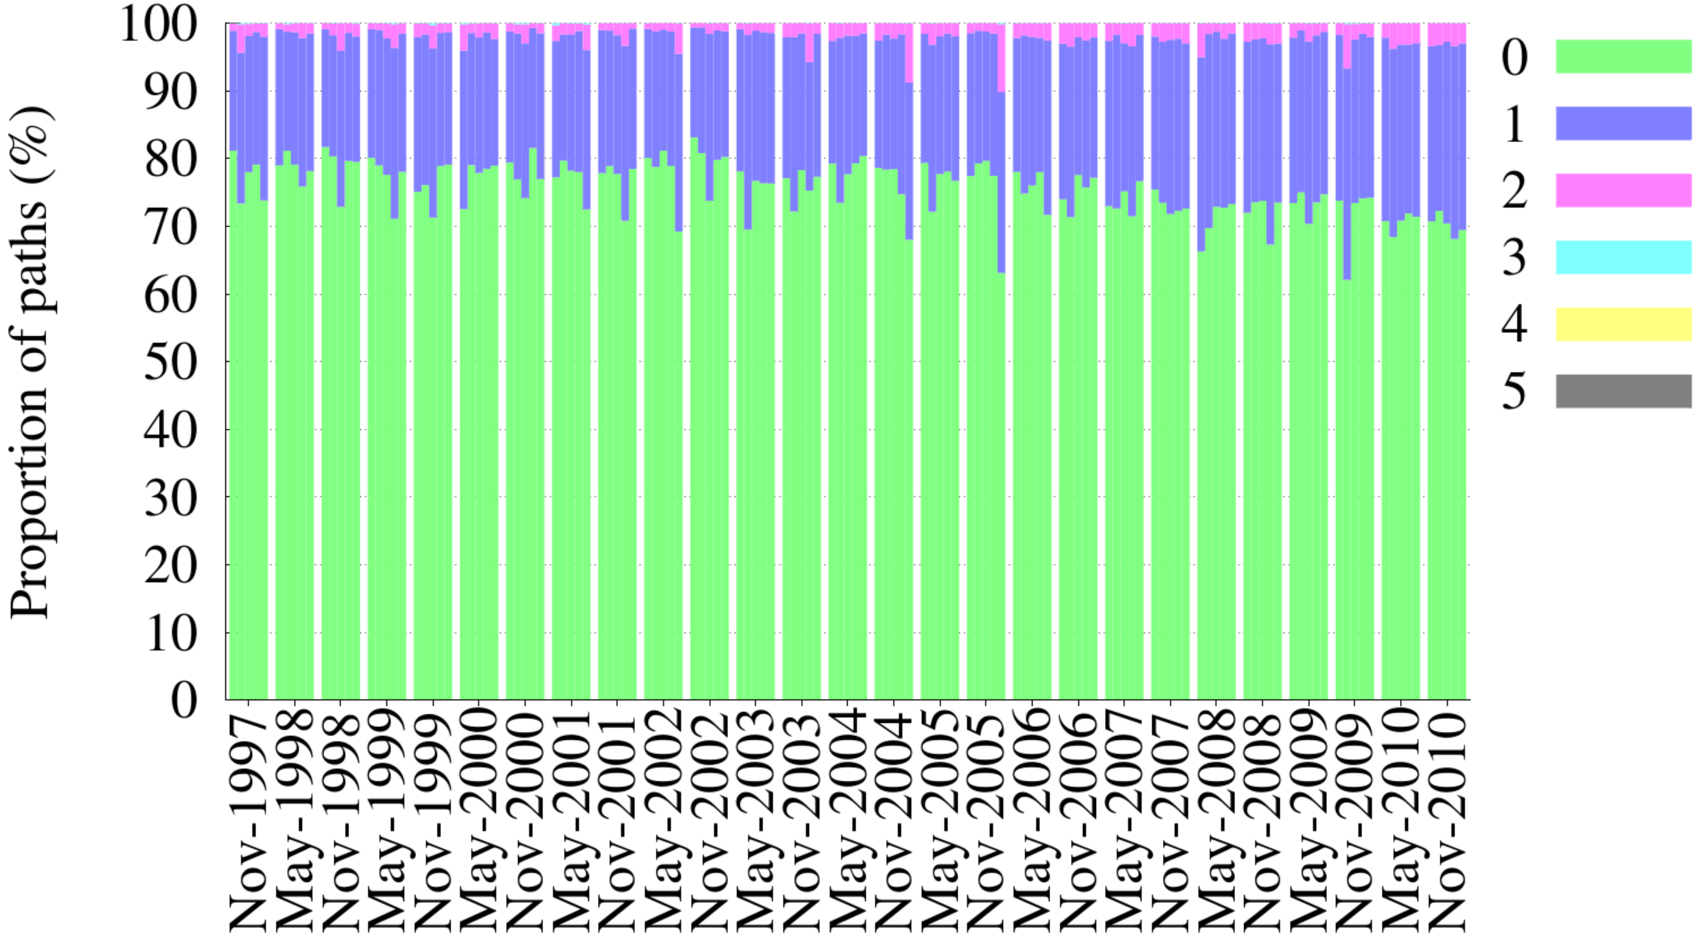

Results: BC multiplicative stretch

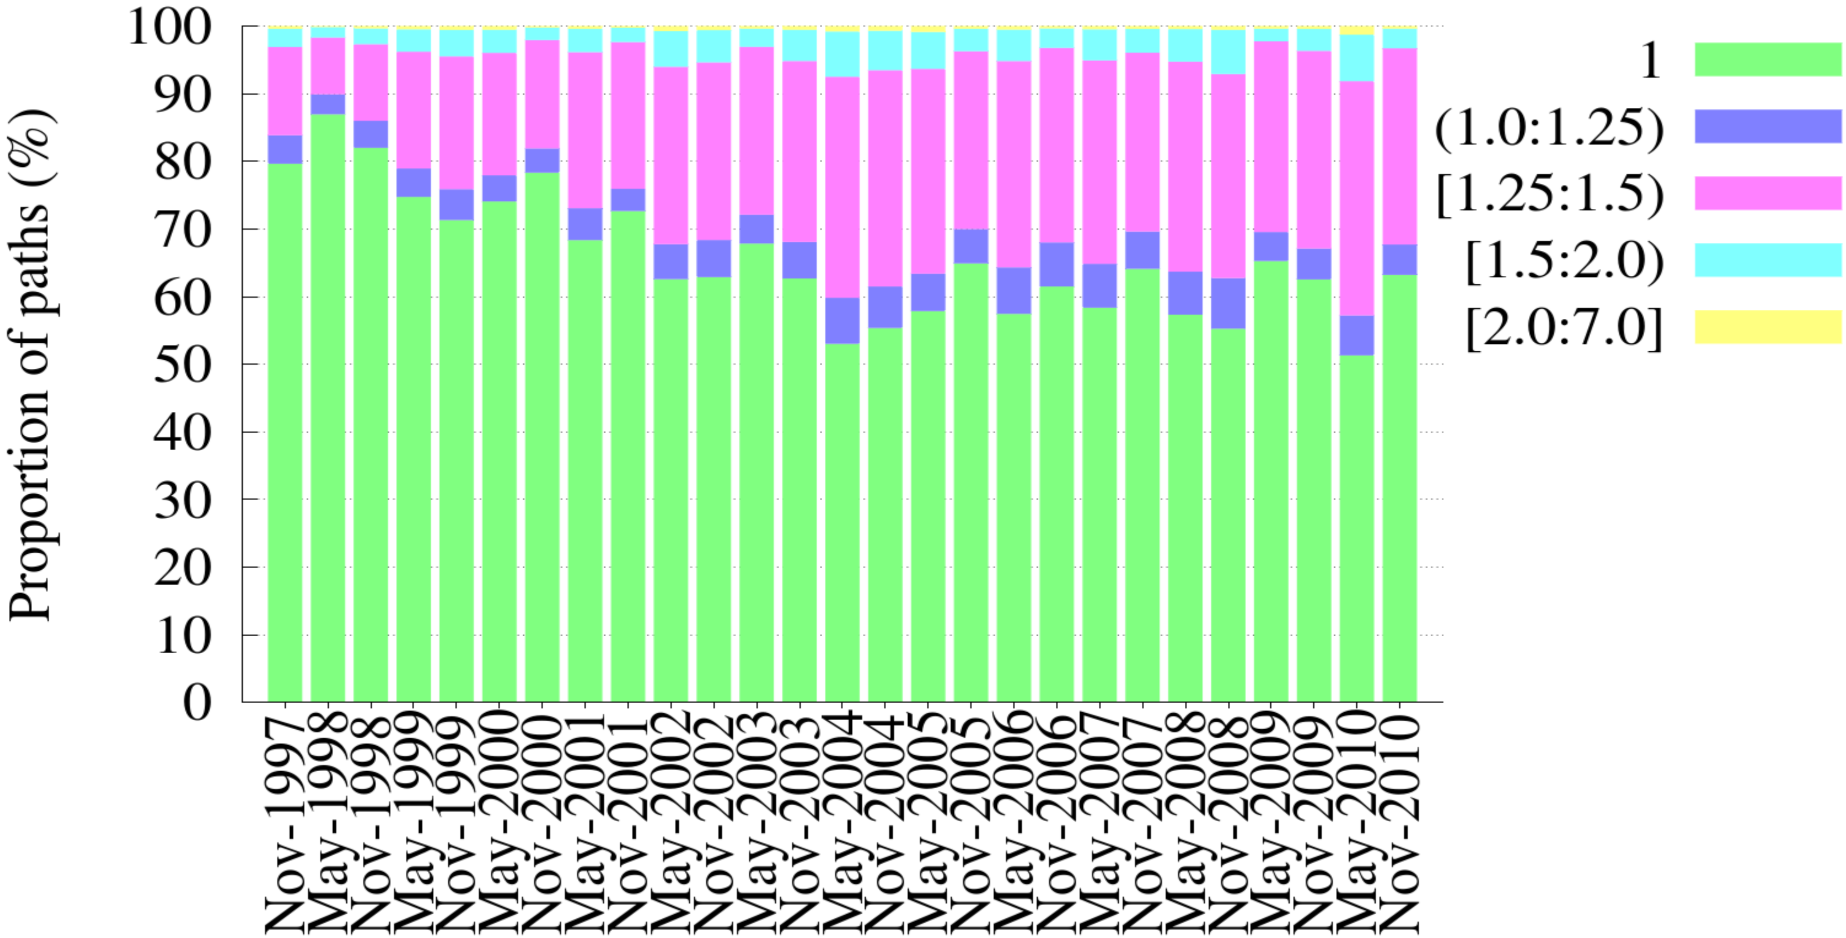

Results: BC, additive stretch

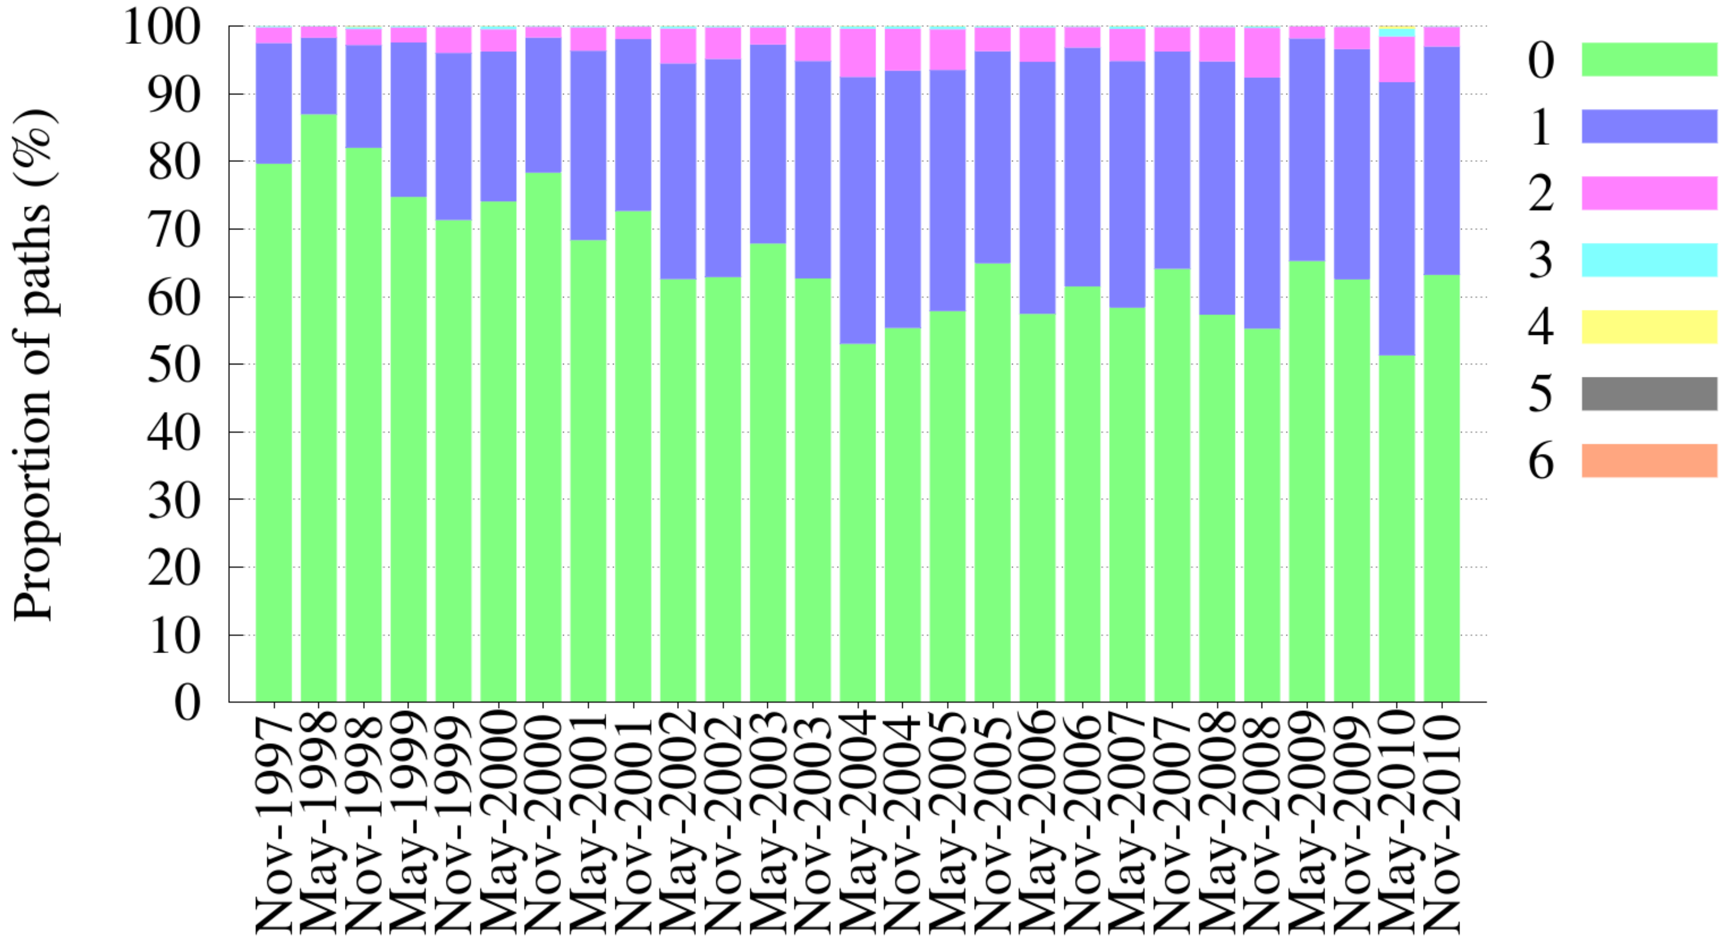

Table sizes

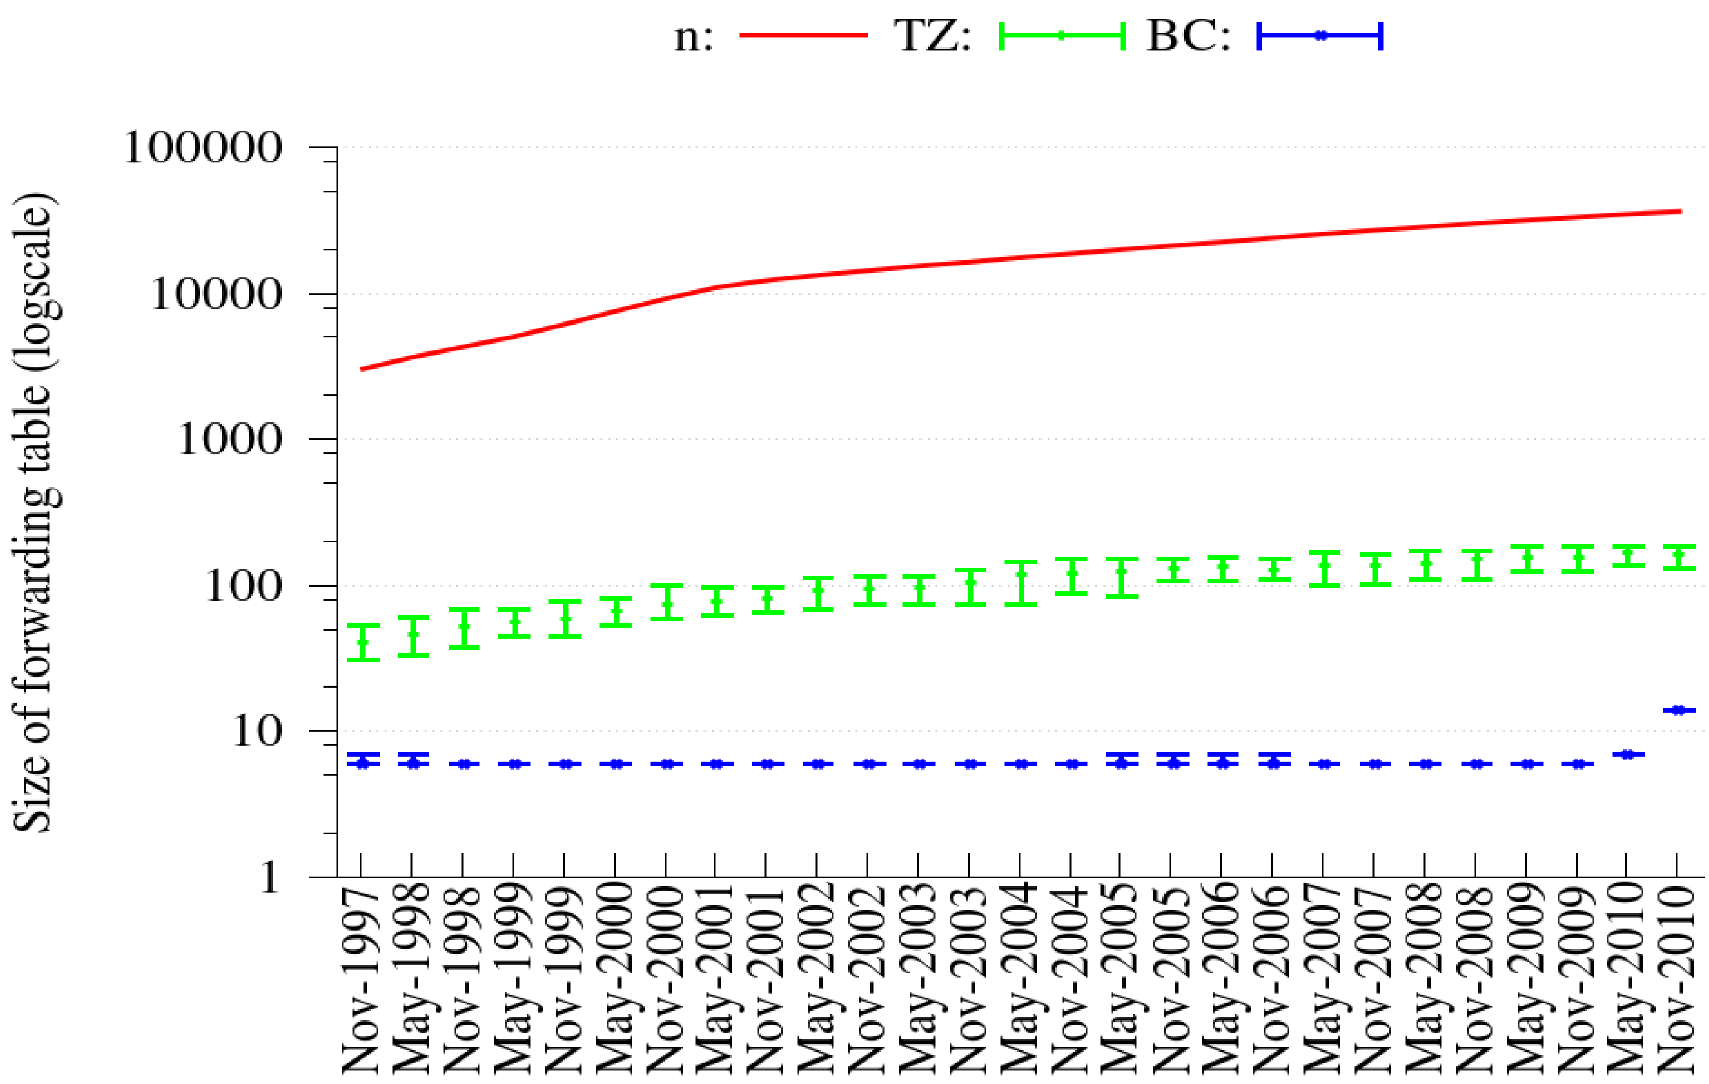

n:
TZ:
BC:
Size of forwarding table (logscale)
100000
10000
1000
100
10
1
Nov-1997
May-1998
Nov-1998
May-1999
Nov-1999
May-2000
Nov-2000
May-2001
Nov-2001
May-2002
Nov-2002
May-2003
Nov-2003
May-2004
Nov-2004
May-2005
Nov-2005
May-2006
Nov-2006
May-2007
Nov-2007
May-2008
Nov-2008
May-2009
Nov-2009
May-2010
Nov-2010

TZ Table sizes

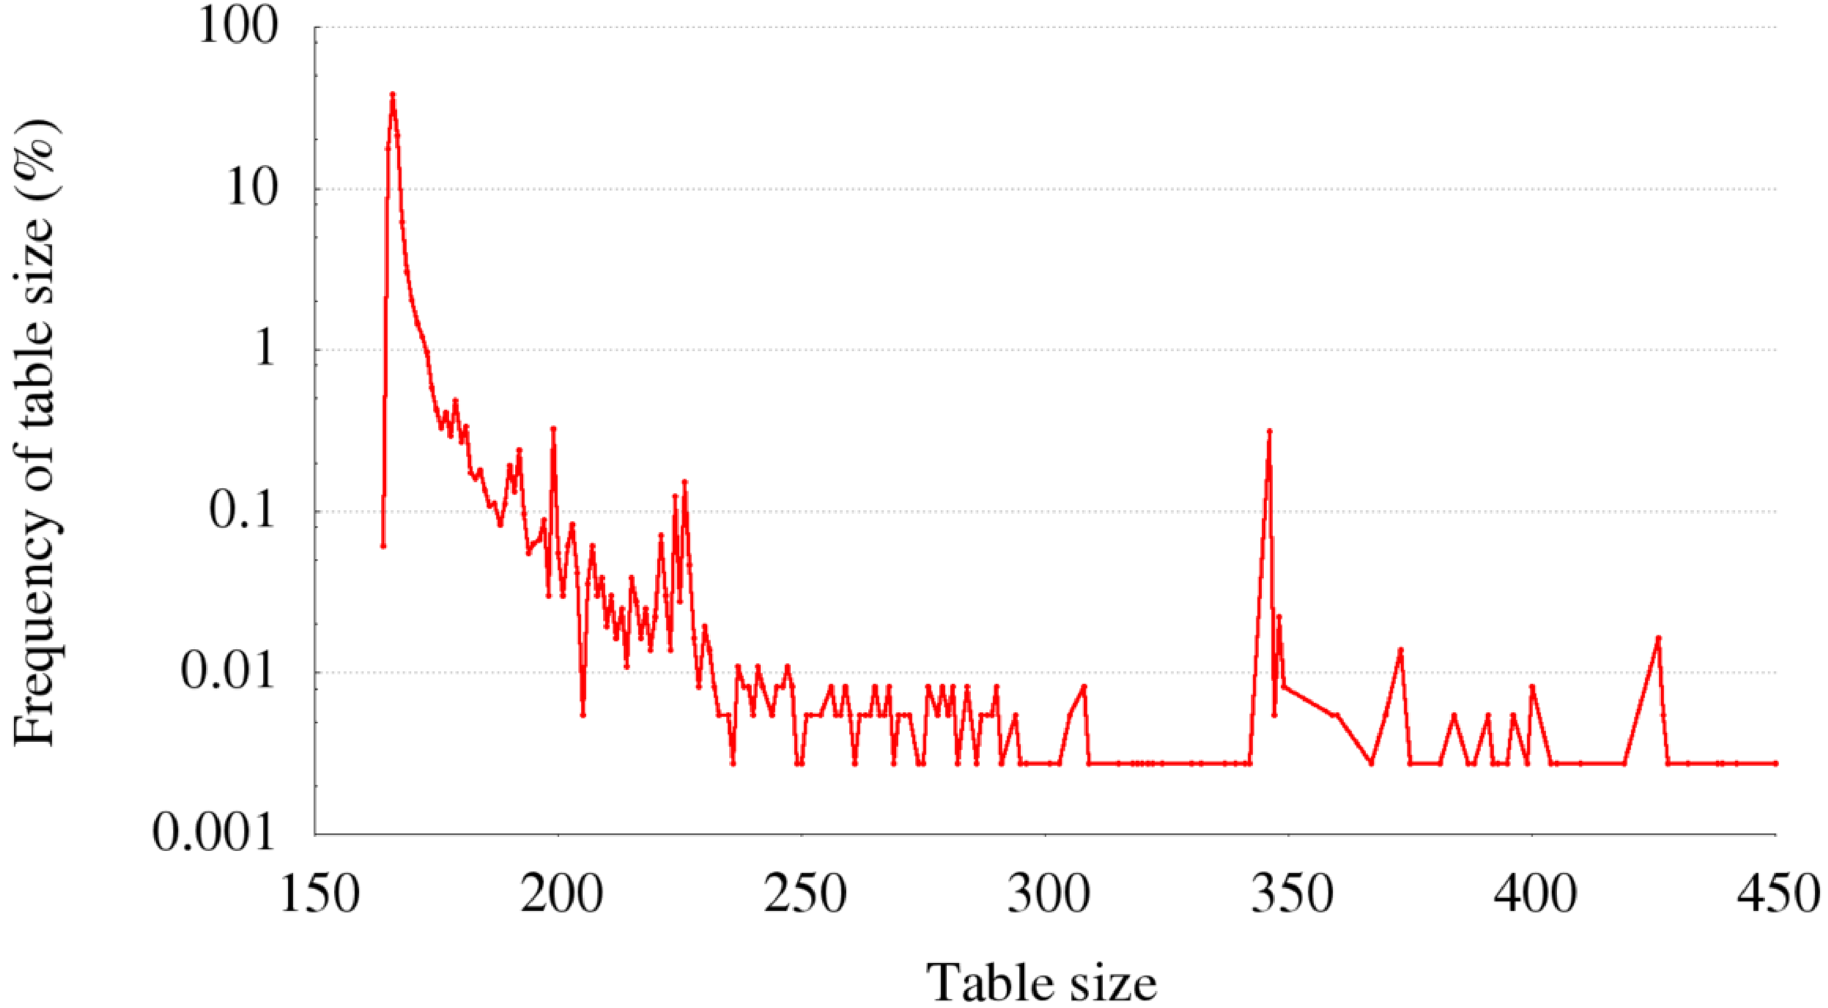

TZ landmark selection

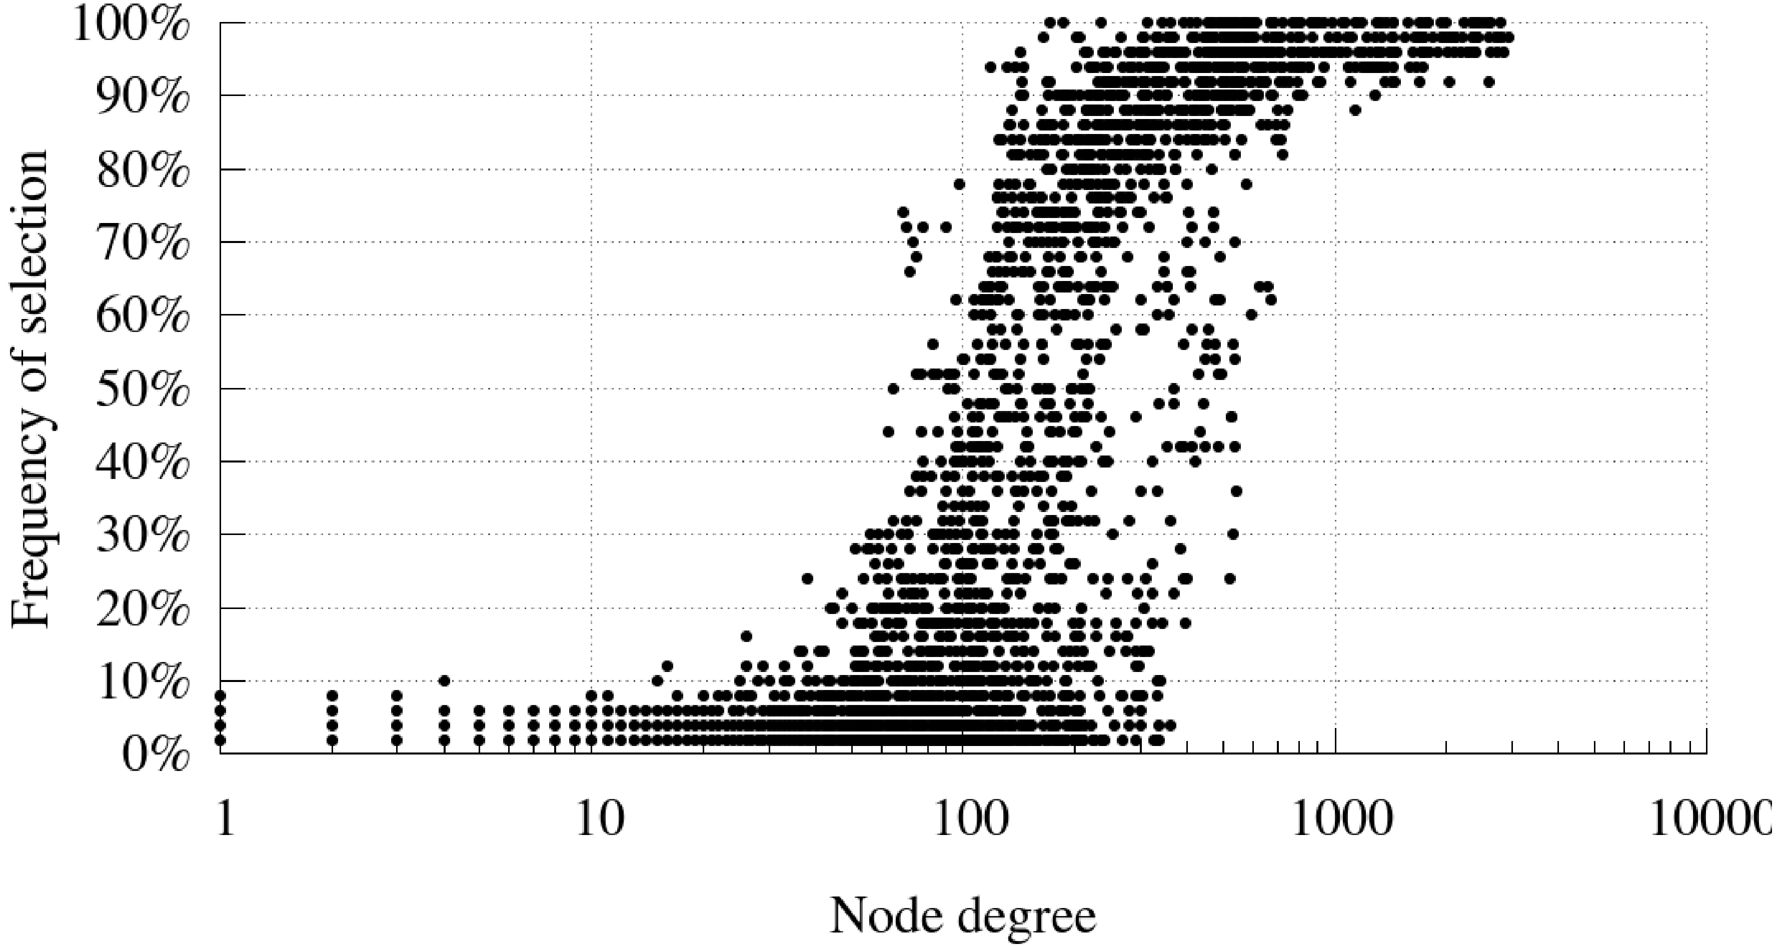

Conclusions

Experimental validation of:

… strong performance on Internet graphs

… consistent performance over long periods of time

Insight into why the TZ algorithm performs well on this type of graph

Future Work

There is scope for further analysis how a decentralised protocol based on the TZ algorithm behaves in a dynamic network

… and how this works at the *router* level rather than the *AS* level

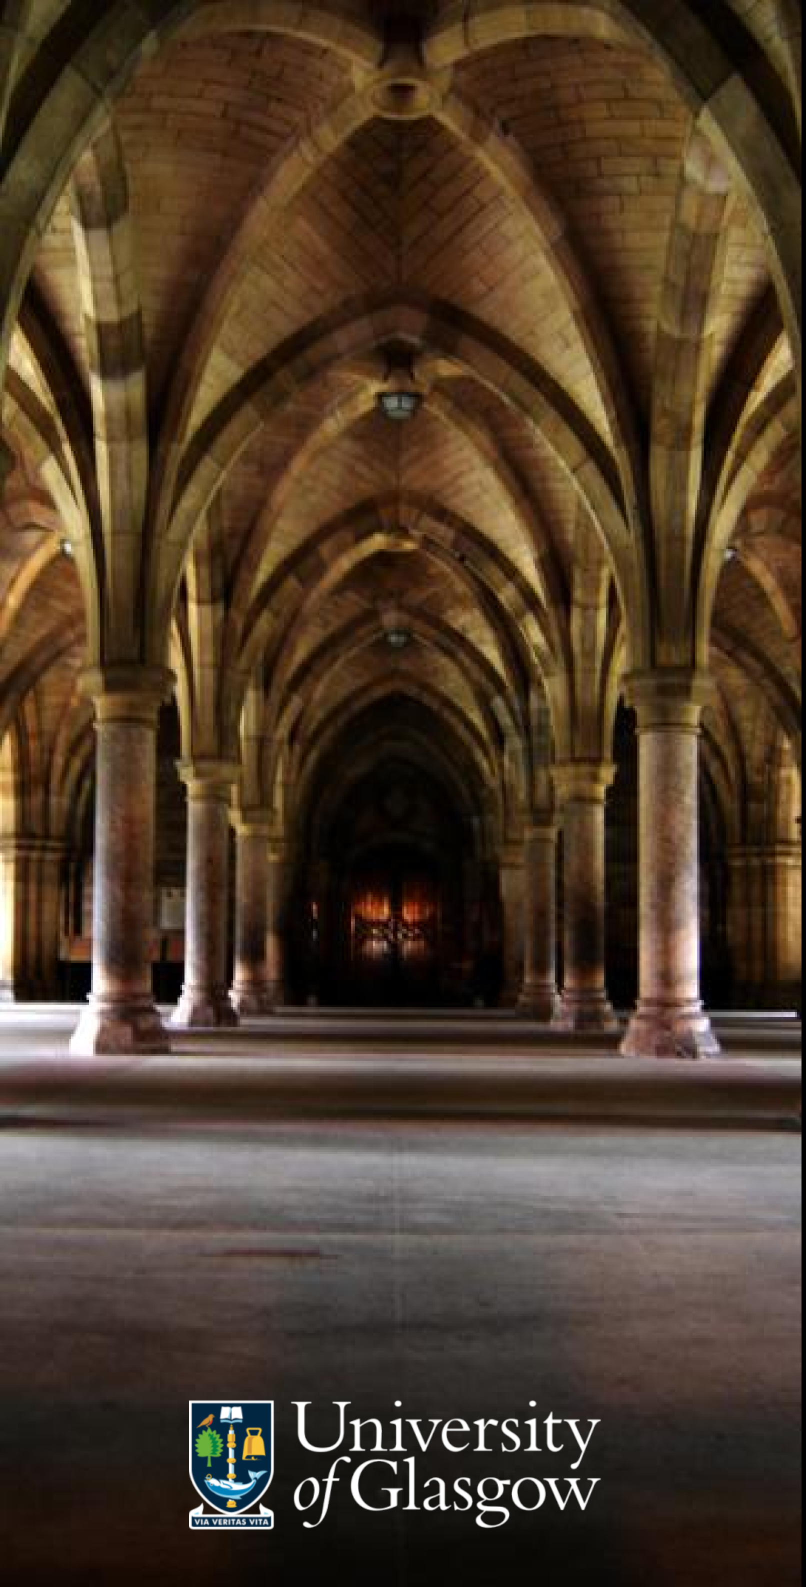

Questions?

Stephen Strowes
sds@dcs.gla.ac.uk