Compact Routing the AS Graph

James Madeley
Email j.i.madeley@lboro.ac.uk
Supervisors Prof. Iain Phillips & Dr. Posco Tso

Figure 1: IPv4 BGP Routing Table

900,000 active BGP entries

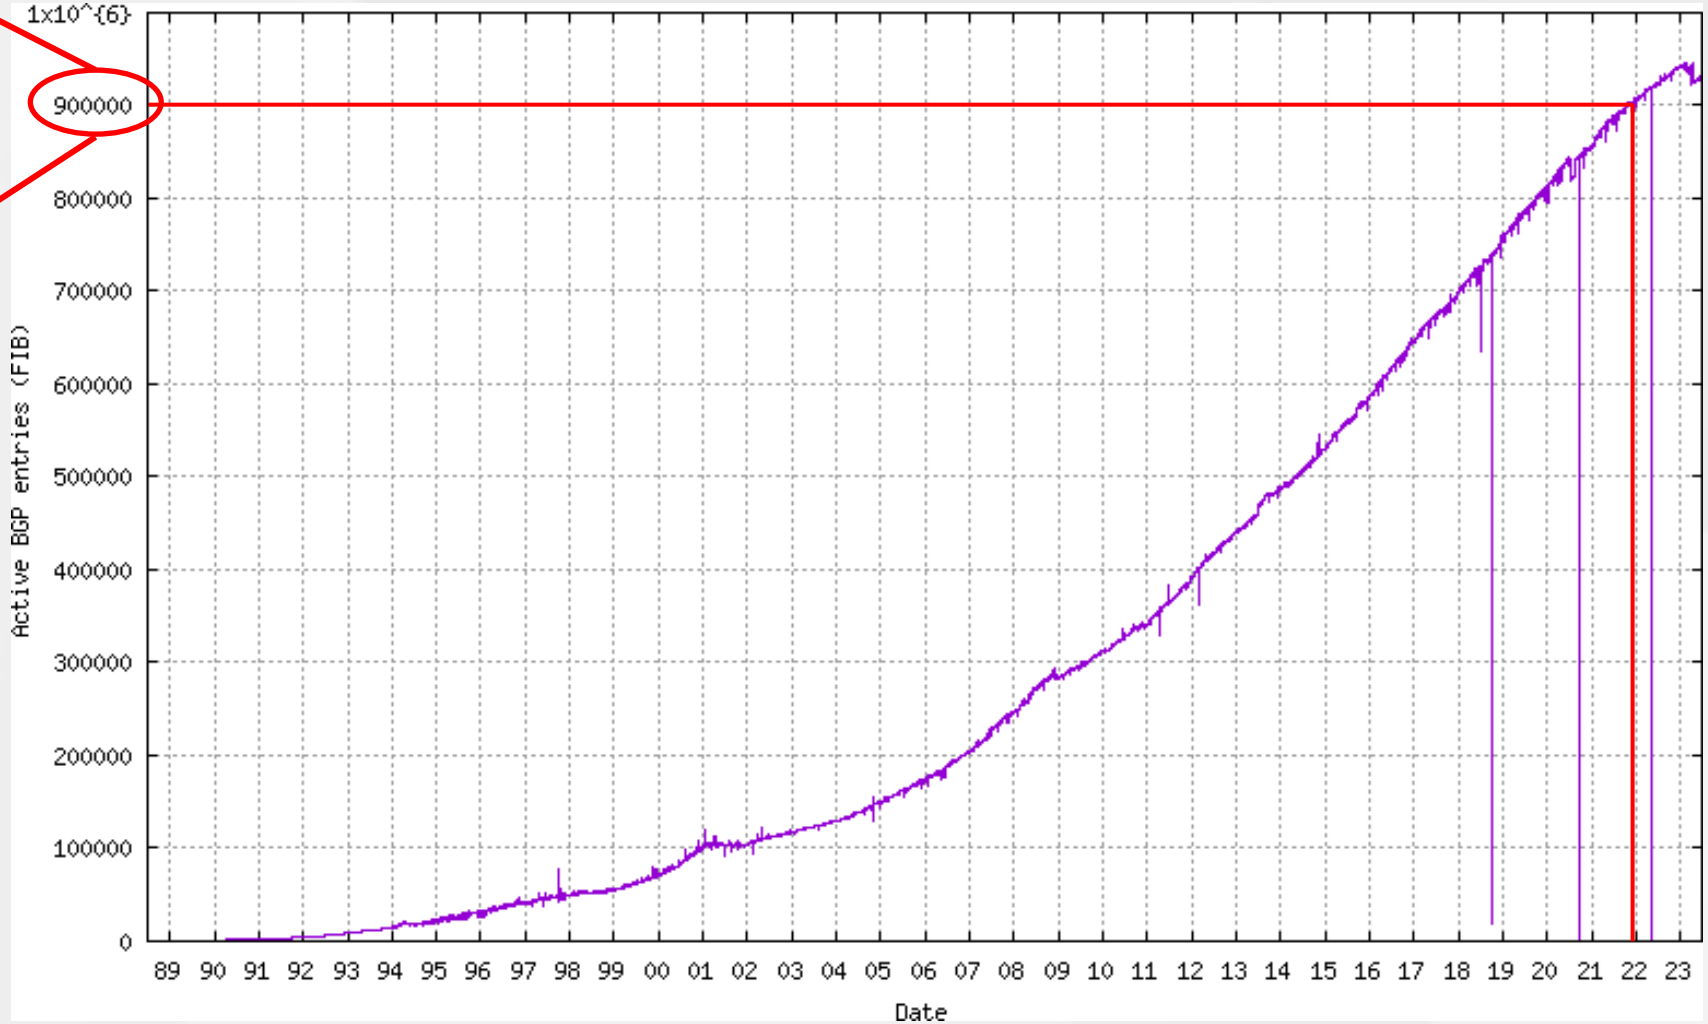

By Geoff Huston from CIDR-Report at https://bgp.potaroo.net/as2.0/bgp-active.html

What does a larger routing table mean?

Increase in memory usage

Increase in search time

What is Compact Routing?

Graph Theory → Reduces routing table size

Compact Routing Steps

1. Establish hierarchy across nodes
2. Define landmarks (optional – set up local clusters)
3. Nodes associate with a landmark
4. Route to a node using its landmark (optional – use local cluster shortcuts)

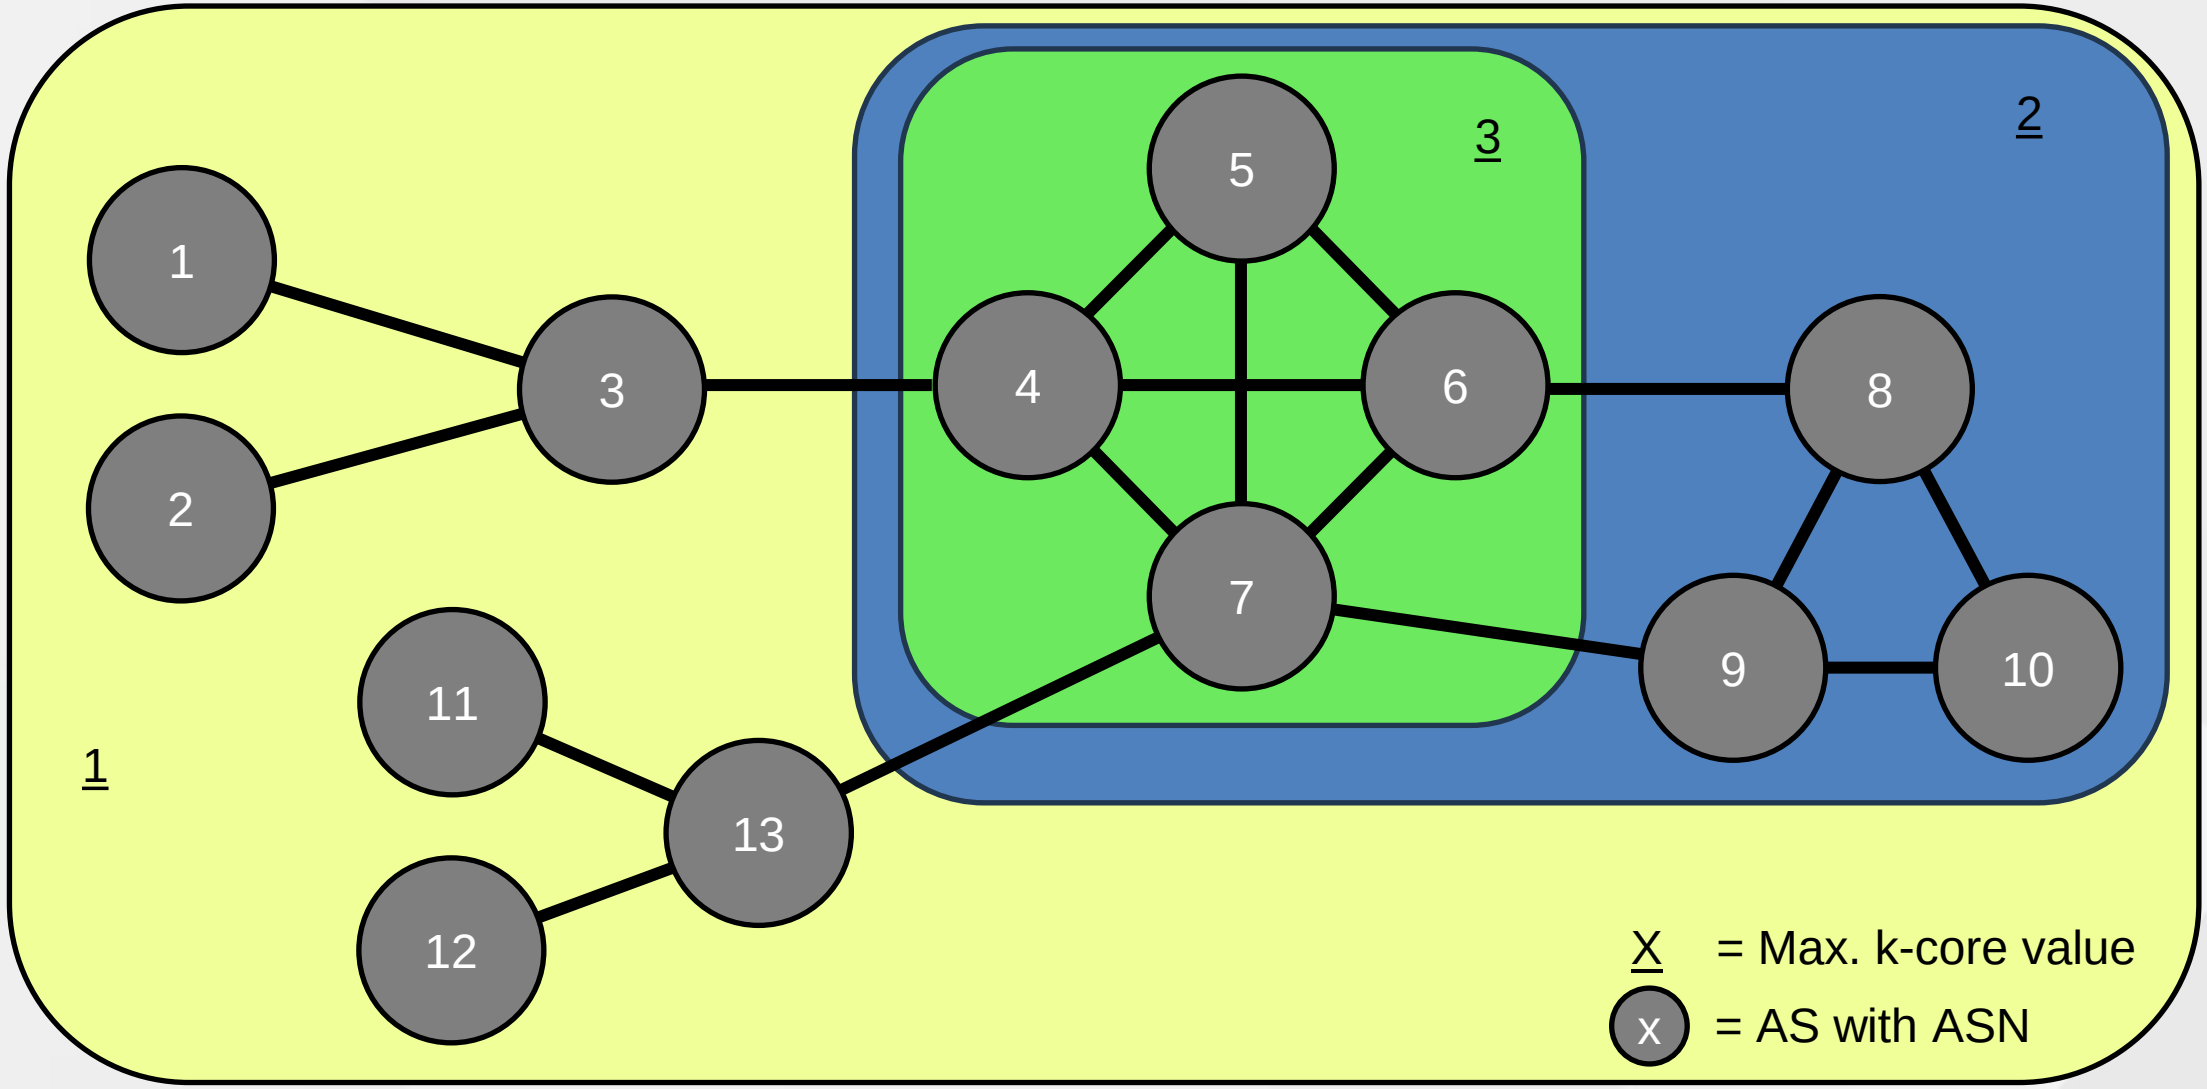

Example based on Jakma's protocol [2]

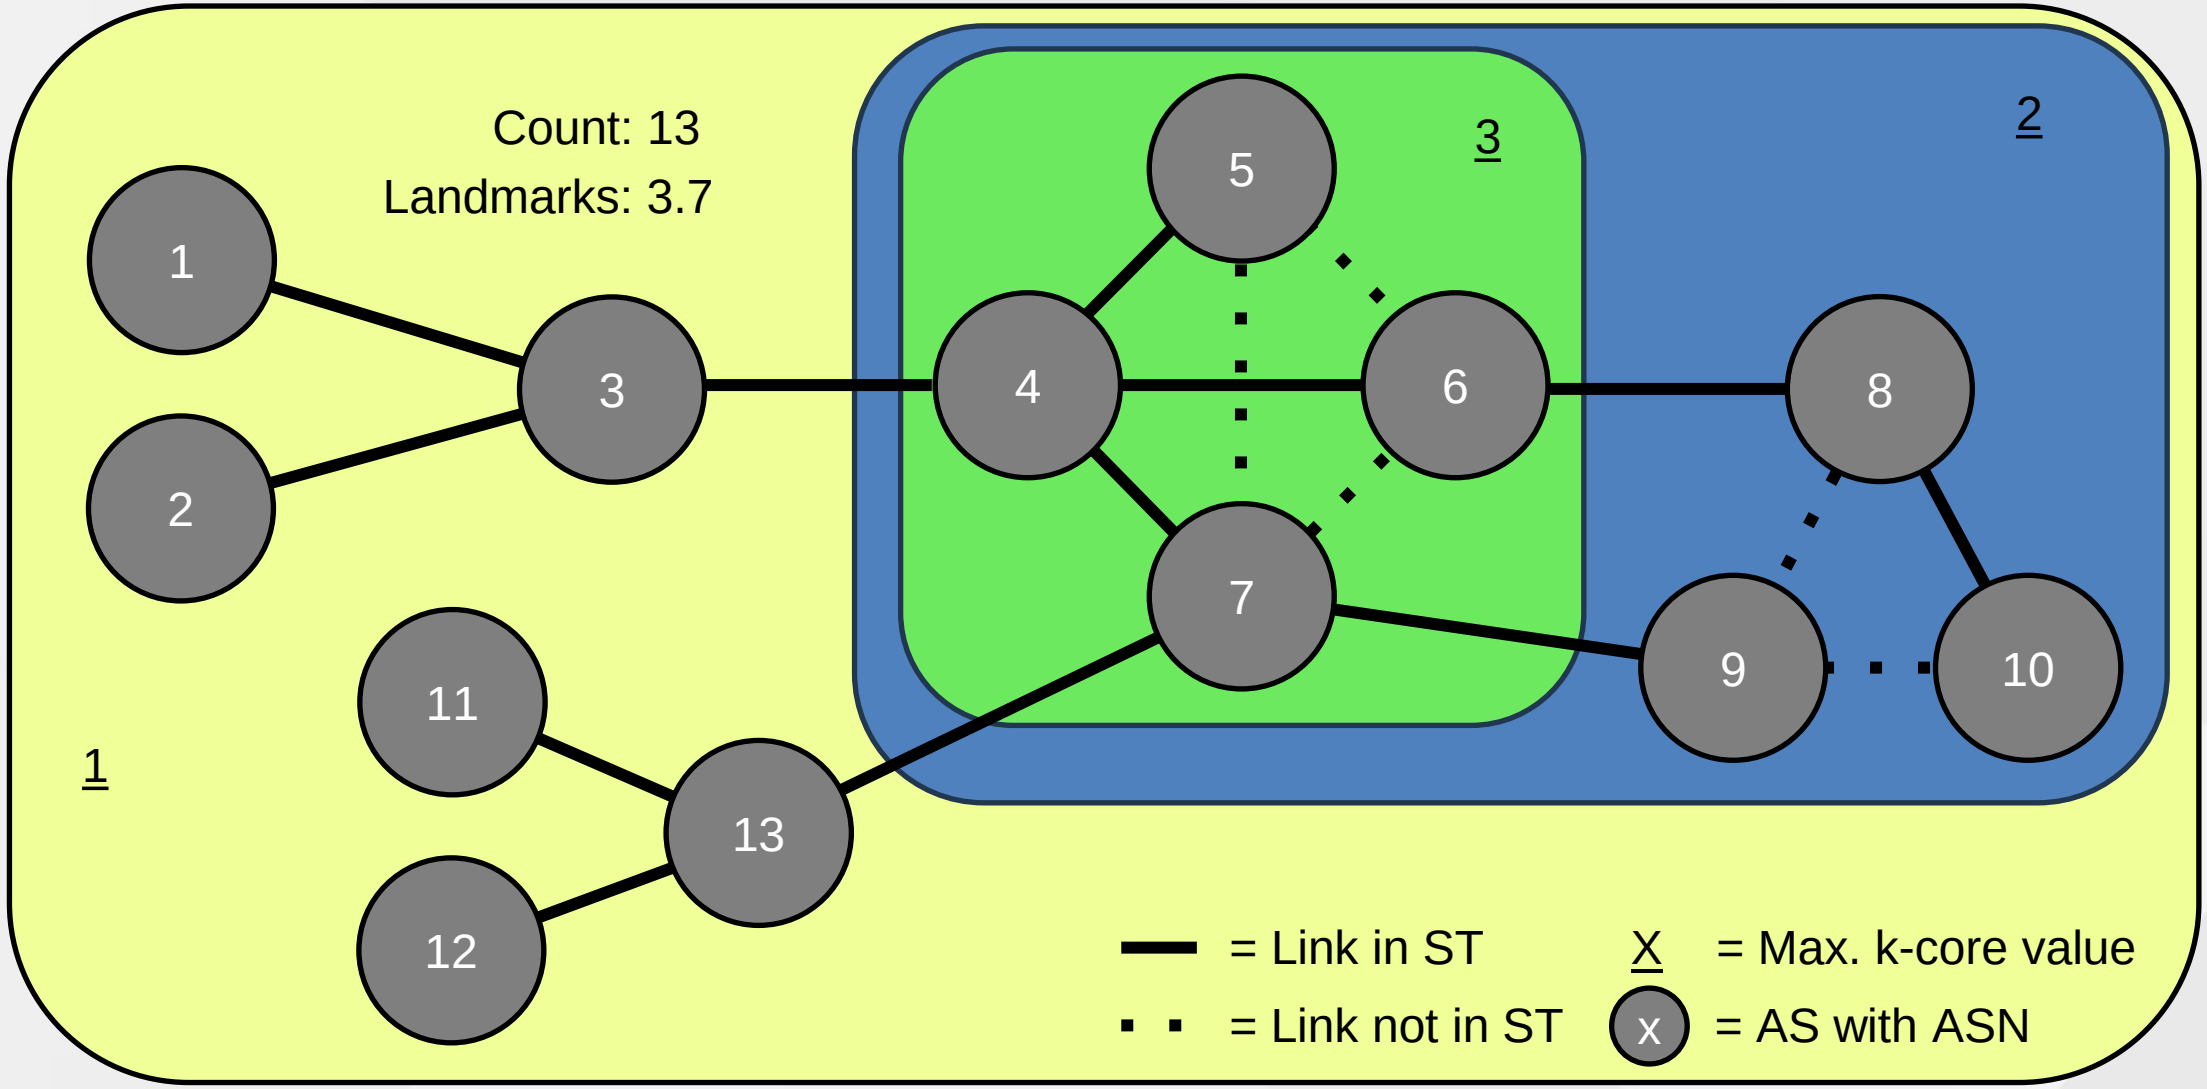

Example based on Jakma's protocol [2]

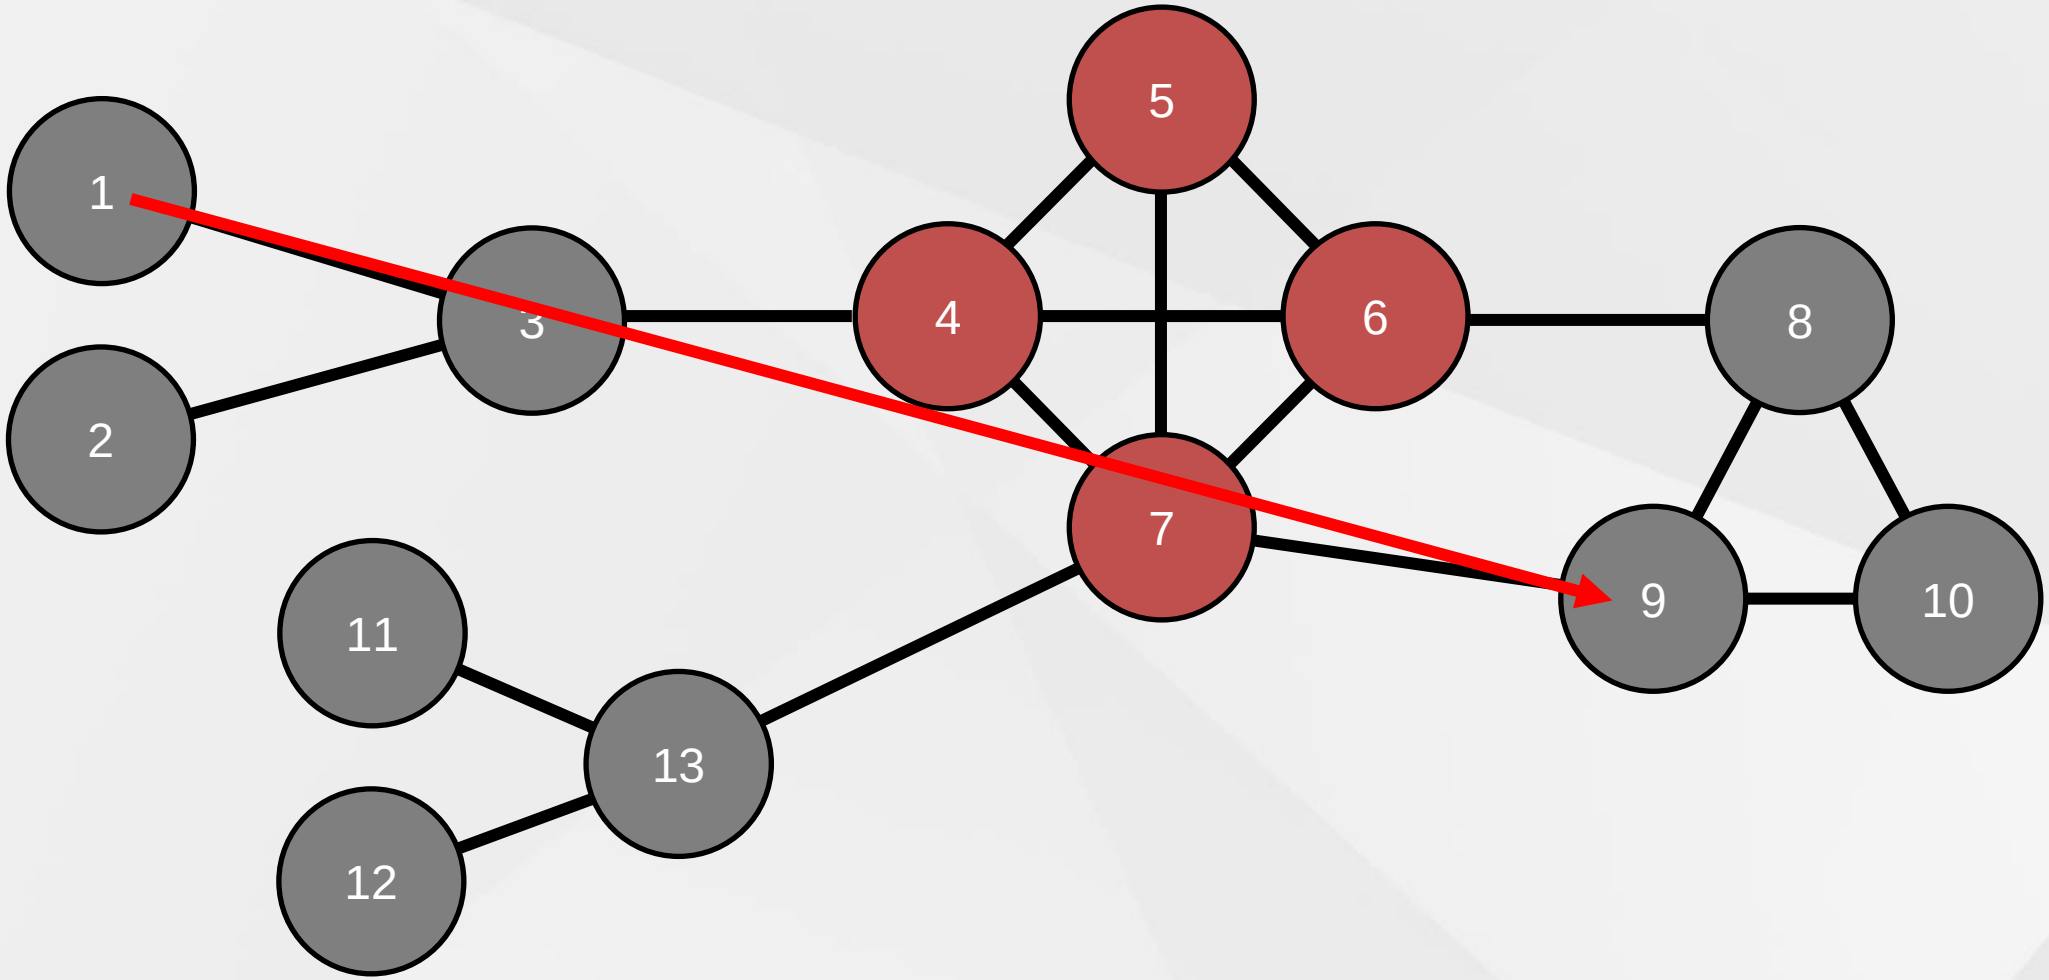

Example based on Jakma's protocol [2]

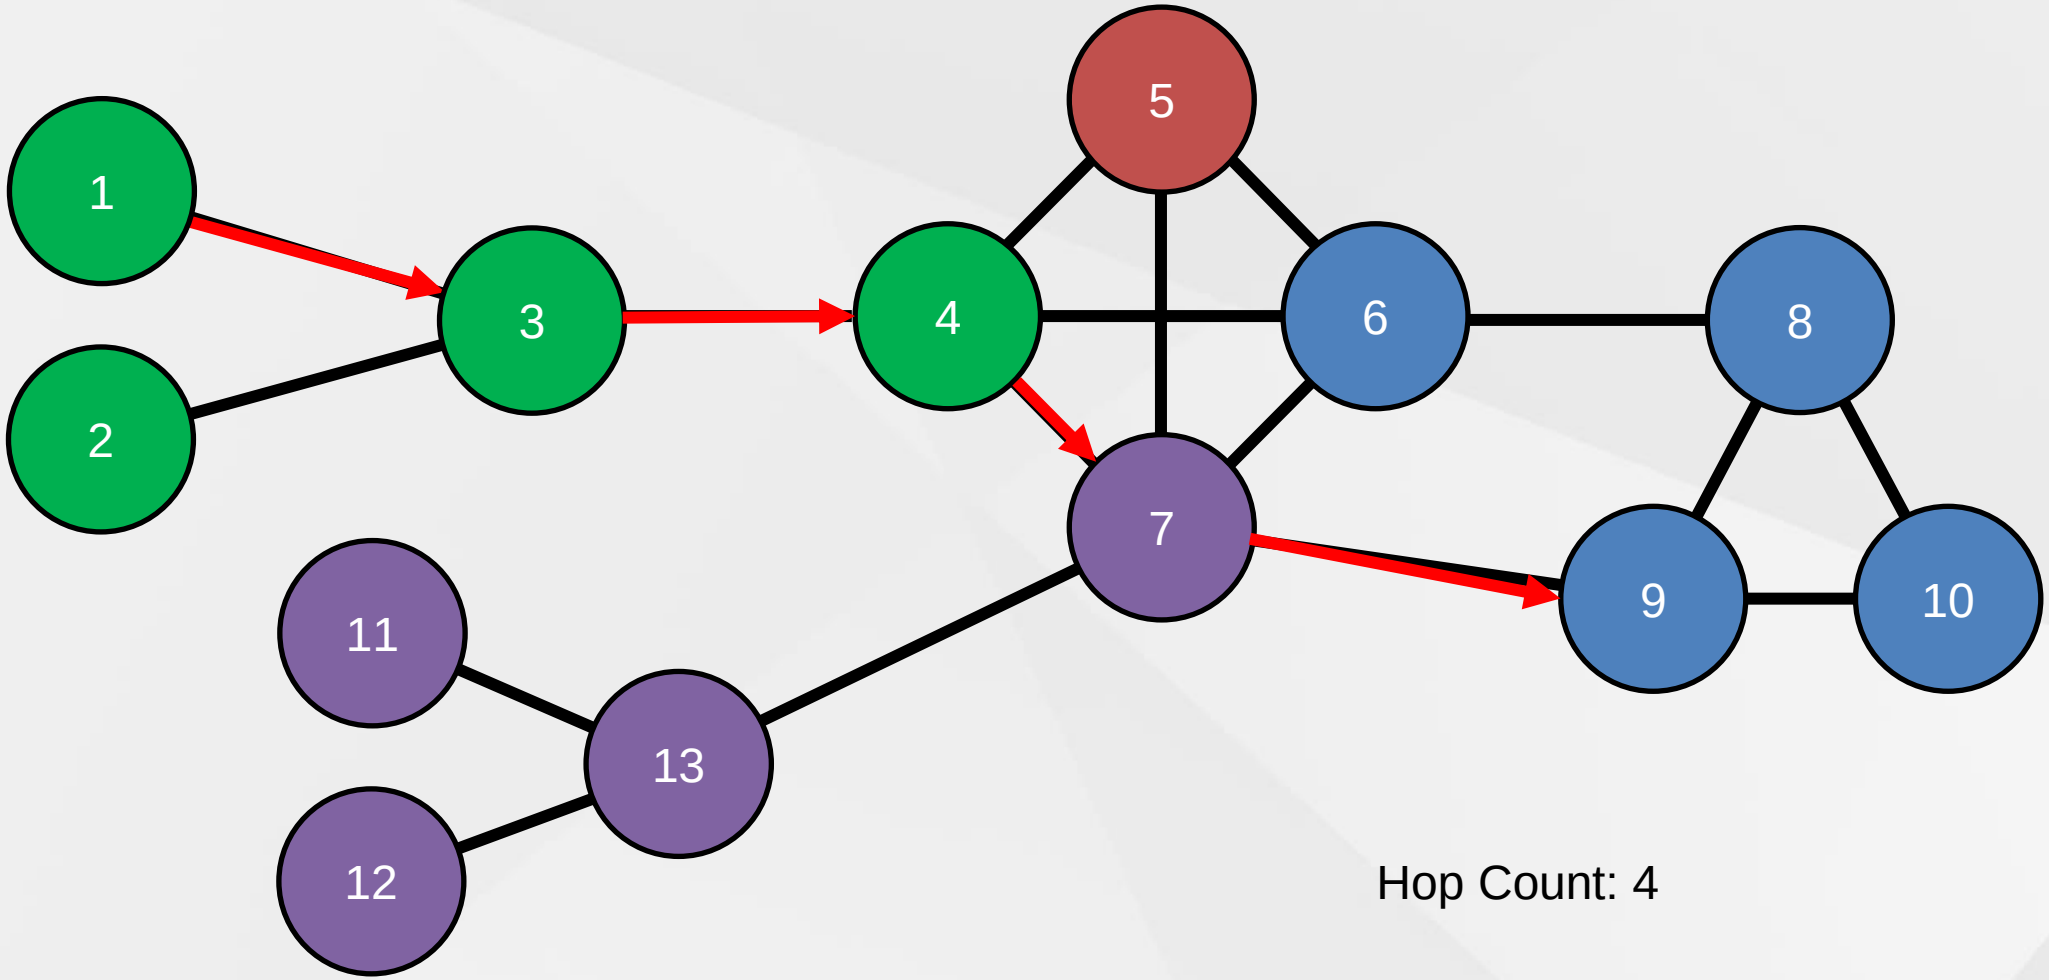

Hop Count: 4

Example based on Jakma's protocol [2]

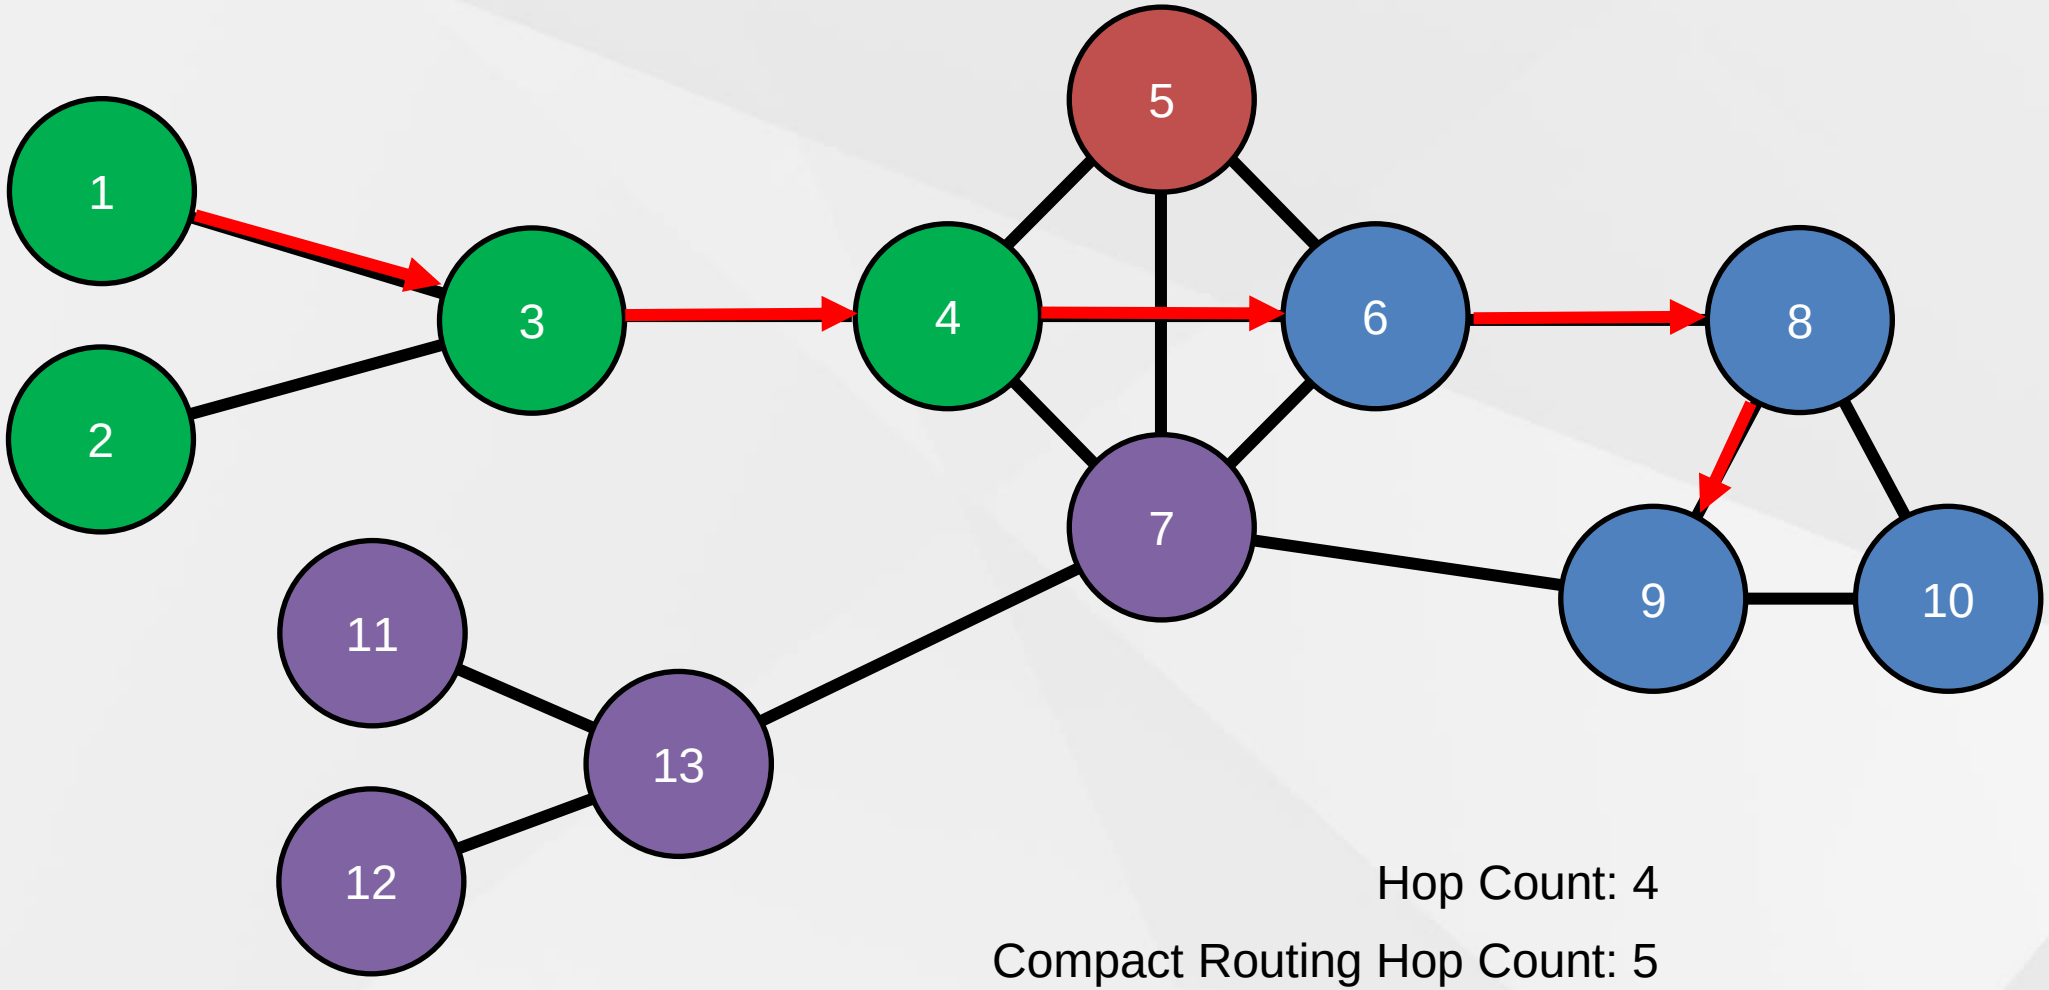

Hop Count: 4

Compact Routing Hop Count: 5

Example based on Jakma's protocol [2]

Table Size: 12

Compact Routing Table Size: ~~4~~ + 2 = 6

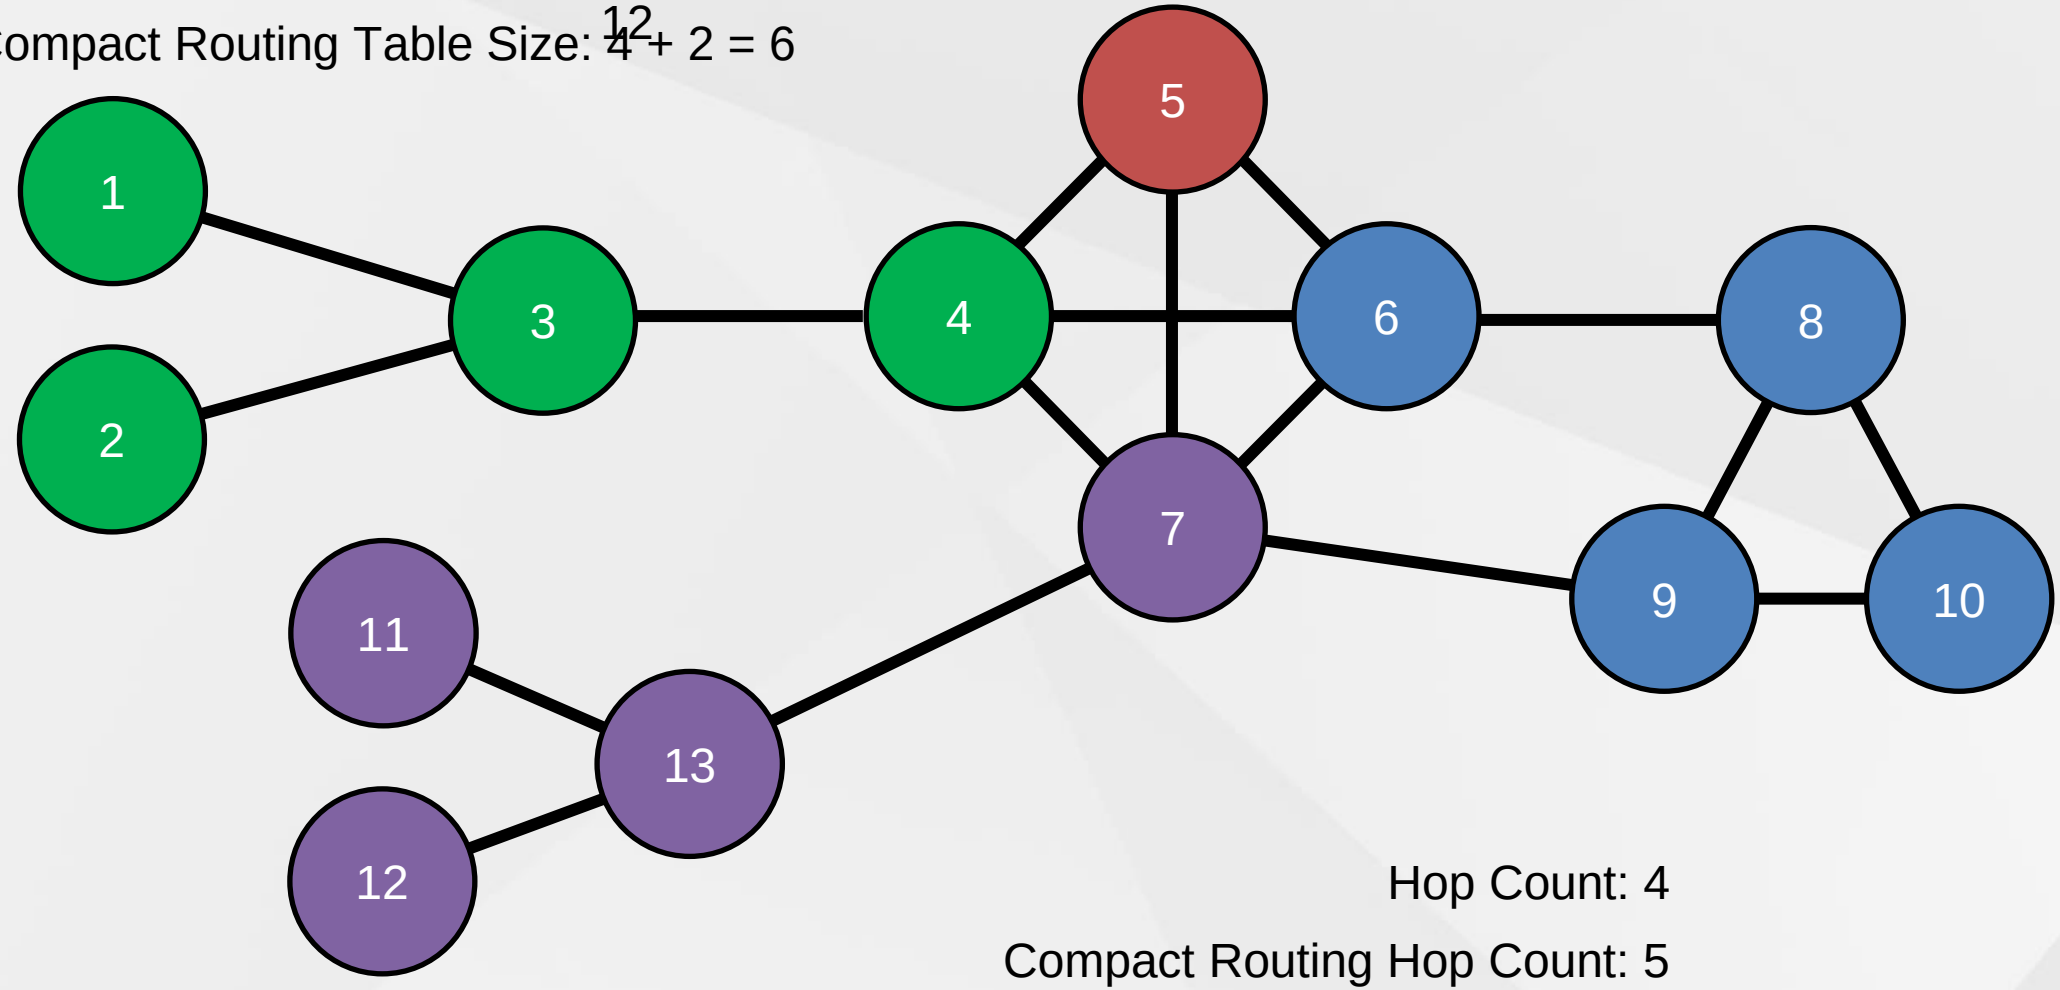

Hop Count: 4

Compact Routing Hop Count: 5

Example based on Jakma's protocol [2]

Our Work

Last year:

- Generating configuration files for large-scale emulations

This year:

- Golang simulator of Jakma’s protocol
- Developing a compact routing daemon

Our Work cont.

Upcoming:

- Run the protocol on large-scale network emulations
- Start gathering results

Measurements:

- Table size (certain algorithms claim >99% reduction [5])
- Path length
- Convergence time
- Resilience
- Number of messages sent

Implementation Challenges

AS Connectivity Fluctuation

The Internet is a dynamic network

Strowes showed k-core stability [3,4]

Upcoming work:

- Further test the Jakma protocol [2]

AS Peering

AS relationships are crucial to Internet operations

Proposed solution:

- ‘controllable clusters’

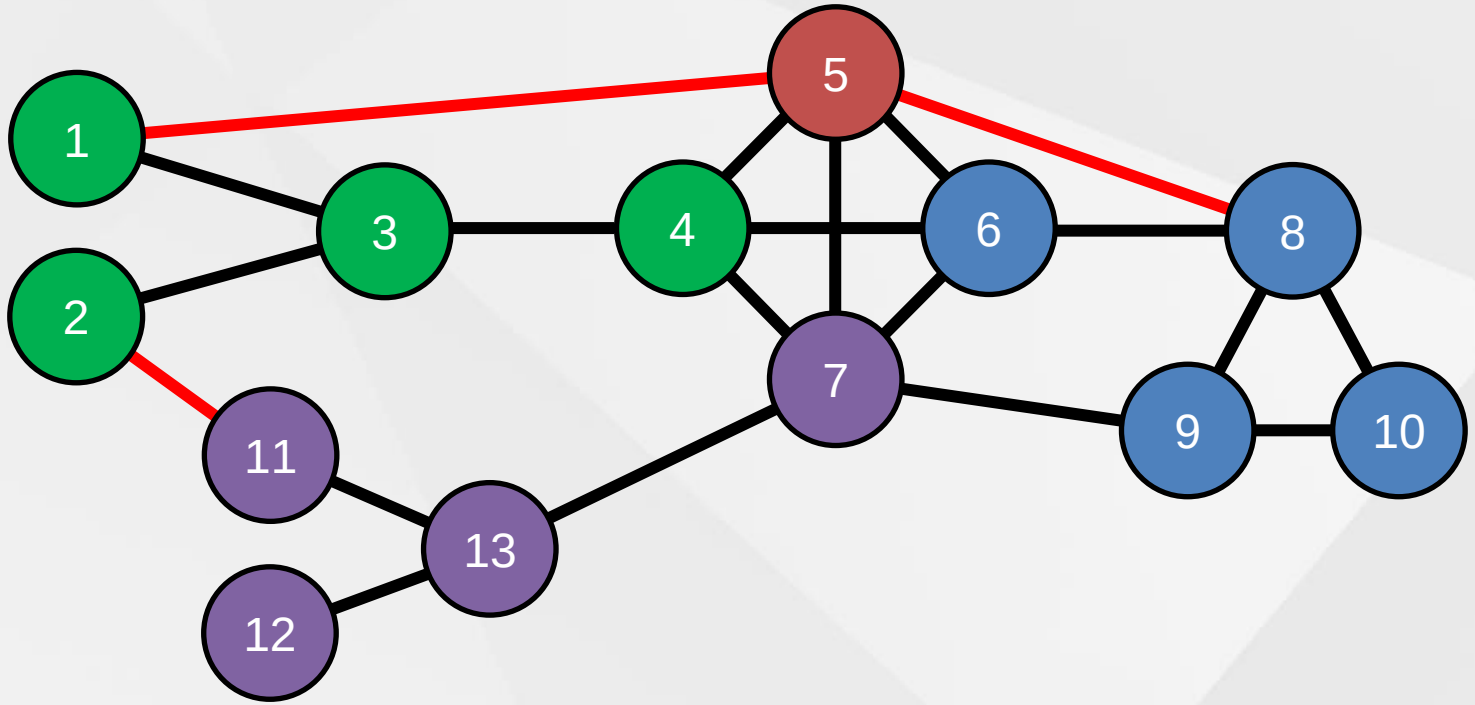

Landmark Selection

There is a cost for being a landmark
ASes may not want to act as a landmark

Proposed solutions:

- Human-chosen landmarks for flexibility
- Allow ASes to volunteer as landmarks

What if no ASes volunteer?
Simply, compact routing can’t be done

Forwarding

Each AS needs identifier containing landmark address

These identifiers must be shared prior to routing

Our solution:

- Share information via DNS

Forwarding cont.

Use IPv6 extension headers

Ver	Traffic Class	Flow Label	
Payload Length		Next Header = EH1	Hop Limit
Source Address			
Destination Address			
Next Header = UL	Landmark Address		
Upper Layer (UL) Header			
Payload			

Testing at scale

The structure of the Internet is constantly changing (‘hypergiants’, CDNs, etc.)

Data sets are out-of-date

Proposed solutions:

- CAIDA internet topology, PeeringDB, etc.
- Create a topology generator

Summary

- Compact routing trades shortest path for smaller tables
- We have a simulator becoming a daemon
- We will test it on large emulated networks
- There are many implementation considerations
- We need representative synthetic Internet topologies

Any questions?

References

[1] Lenore J. Cowen. 1999. Compact Routing with Minimum Stretch. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '99).* Society for Industrial and Applied Mathematics, USA, 255–260.

[2] Paul Jakma. 2016. *A distributed, compact routing protocol for the Internet.* Ph.D. Dissertation. University of Glasgow.

[3] Stephen B. Seidman. 1983. *Network structure and minimum degree.* Social Networks 5, 3 (1983), 269–287, https://doi.org/10.1016/0378-8733(83)90028-X

[4] Stephen D. Strowes. 2012. *Compact routing for the future internet.* Ph.D. Dissertation. University of Glasgow.

[5] Mikkel Thorup and Uri Zwick. 2001. Compact Routing Schemes. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA '01).* Association for Computing Machinery, New York, NY, USA, 1–10. https://dl.acm.org/doi/10.1145/378580.378581

Additional slides

A Short History of Compact Landmark Routing

- 1999: First algorithm proposed by Cowen [1]
- 2001: Improved by Thorup & Zwick (TZ) [5]

- 2012: Strowes found that, for landmark selection, the k-core decomposition works well [3,4]

- 2017: Jakma presented a fully distributed landmark routing protocol [2]