Validation of P2P overlay models within OverSim

Jamie Furness, Mario Kolberg, Marwan Fayed

The need for P2P simulation

- The strength of P2P networks is their ability to scale to huge numbers of users.
- Changes need to be well tested before being deployed to an existing network.
- PlanetLab provides a ~1000 node network spread around the world which can be used for testing - but still small scale and tedious to use.

Introducing OverSim

- OverSim is a P2P network simulation engine, providing models of most commonly used P2P overlays.
- Built on-top of OMNeT++, a discrete event simulator designed for network protocol simulation.
- Supports abstraction or full simulation of the underlying network.
- Depending on available hardware, supports simulating networks of over 100,000 nodes.

Blind search over DHTs

- My research has been looking at performing complex queries over varying DHTs.
- Implemented as an application which sits on-top of the DHT models within OverSim.
- This work looks at the validation of these underlying DHT models.

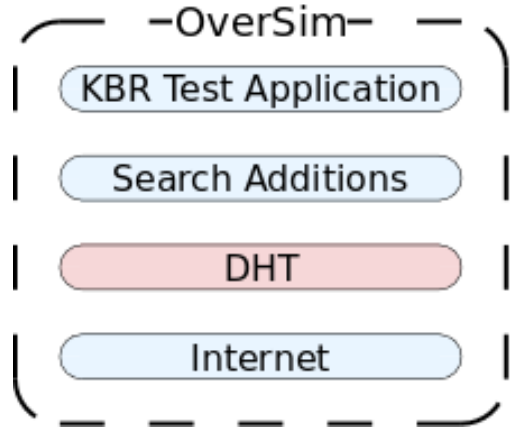

Background - Distributed Hash Tables (DHTs)

- Network of nodes, each assigned a unique identifier.
- Often logically structured in a ring, each node ordered based on their identifier.
- Data to be stored is also assigned a unique identifier, and usually stored at the node closest in the ring.
- Known as a Distributed Hash Table (DHT), effectively provides a Hash Table interface to each node, allowing them to put and retrieve data by sending it to or retrieiving it from the responsible node.

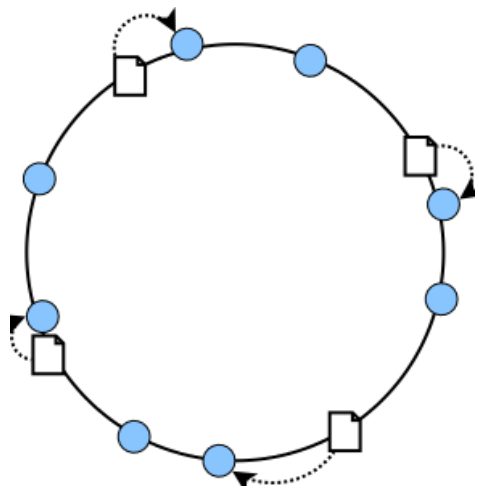

Routing in a DHT

- DHTs can be split into different types: one-hop, variable-hop, and multi-hop.
- In most cases a nodes routing table will only contain addresses for a specific subset of other nodes within the network.
- Clearly the average number of hops required depends on the size and contents of these routing tables.

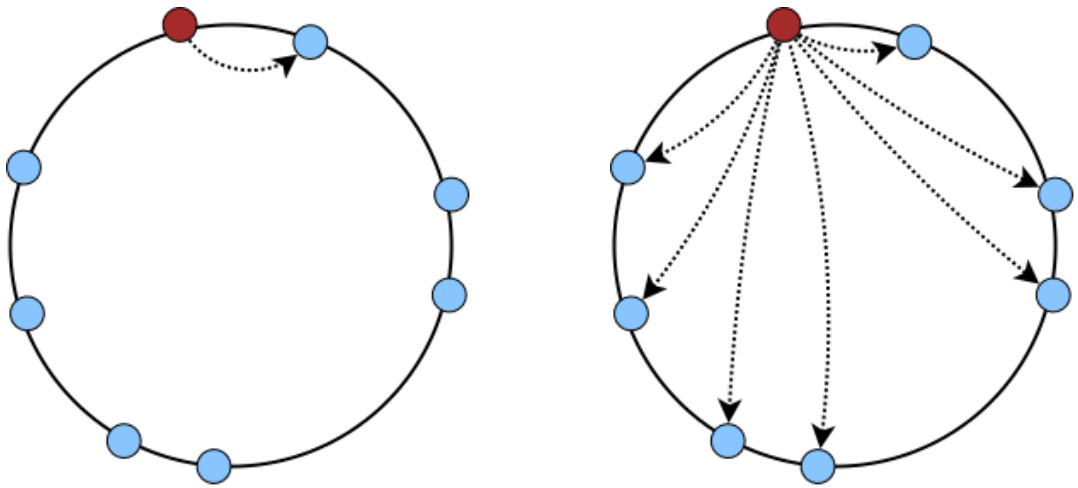

Validating simulation models

- Simulating a P2P network in a realistic environment is a challenging task. It is important to ensure the overlay models used are valid representations of the modelled DHT.
- There are two main aspects we are interested in validating:
  1. The routing algorithm.
  2. The maintenance algorithm.
- This can be done by comparing:
  1. The lookup hop count with the expected hop count.
  2. The node failure rate with the lookup failure rate.

Chord

Chord routing tables include the immediate neighbours, and a selection of nodes at logarithmically increasing distance around the ring.
The maximum expected hop count, assuming an up-to-date routing table, is $O \log_2 N$.

If neighbouring nodes fail, routing tables are updated immediately. Further away nodes are probed periodically and updated if needed.

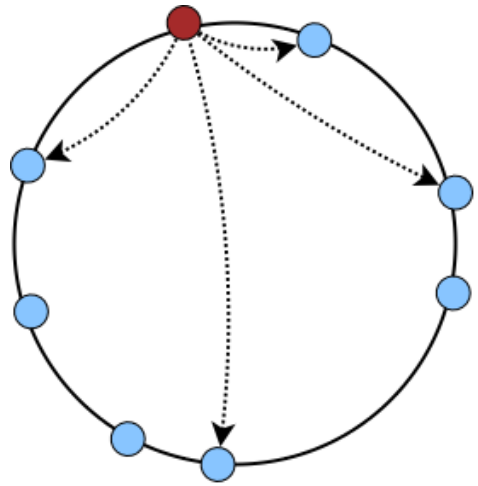

Network size vs average hop count

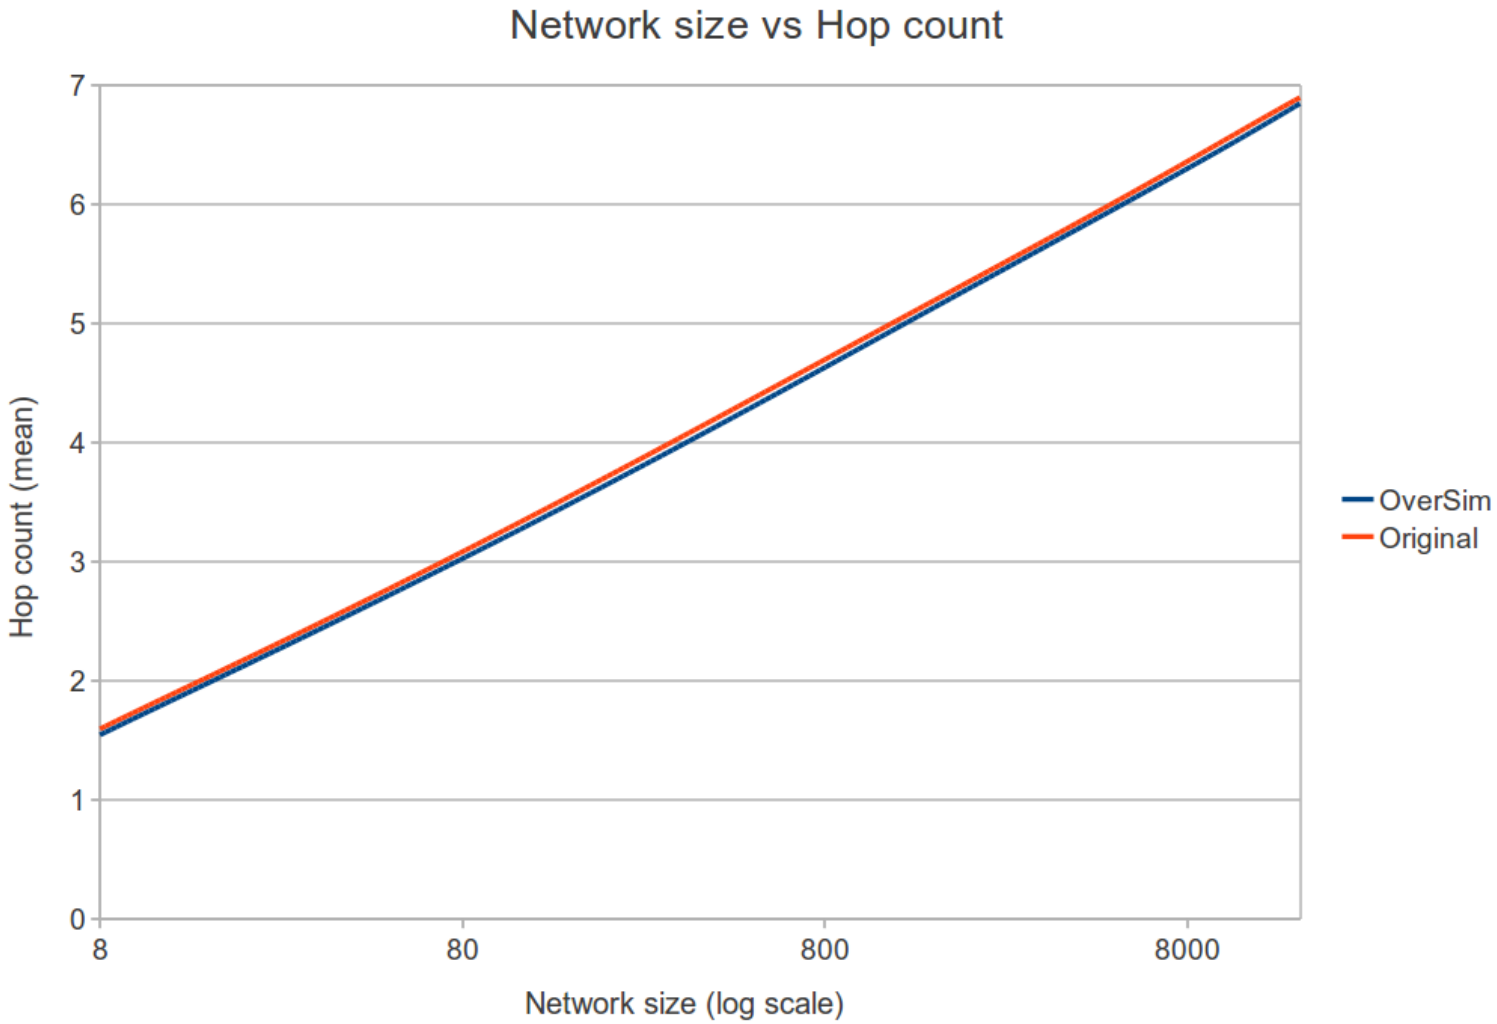

Hop count distribution

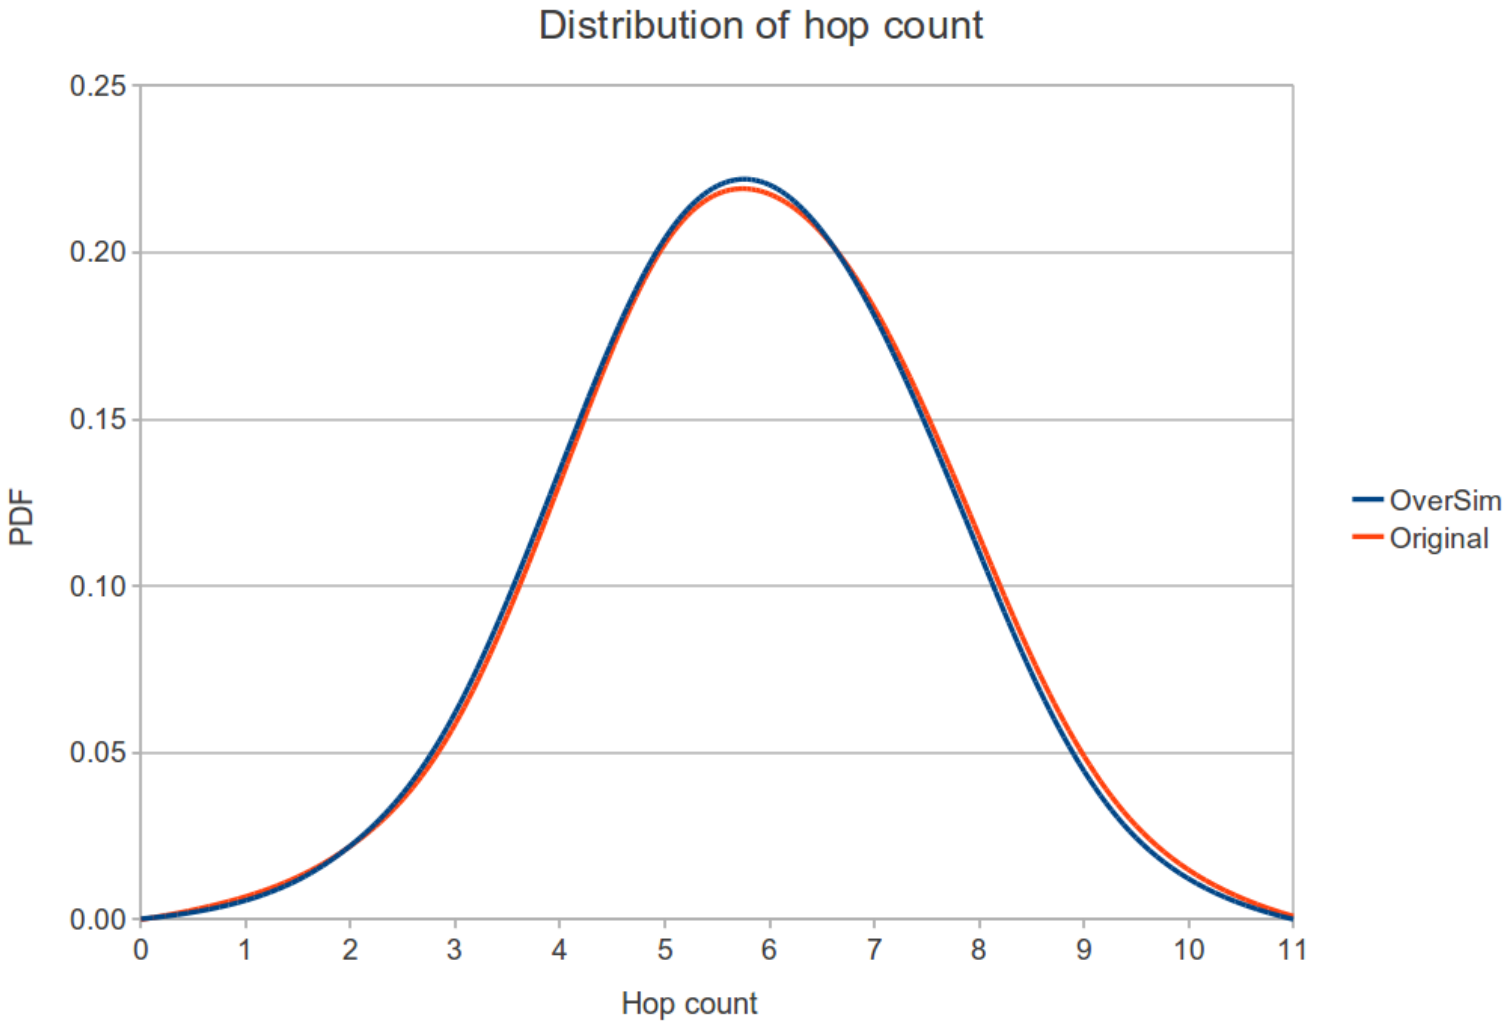

Node failure vs lookup failure

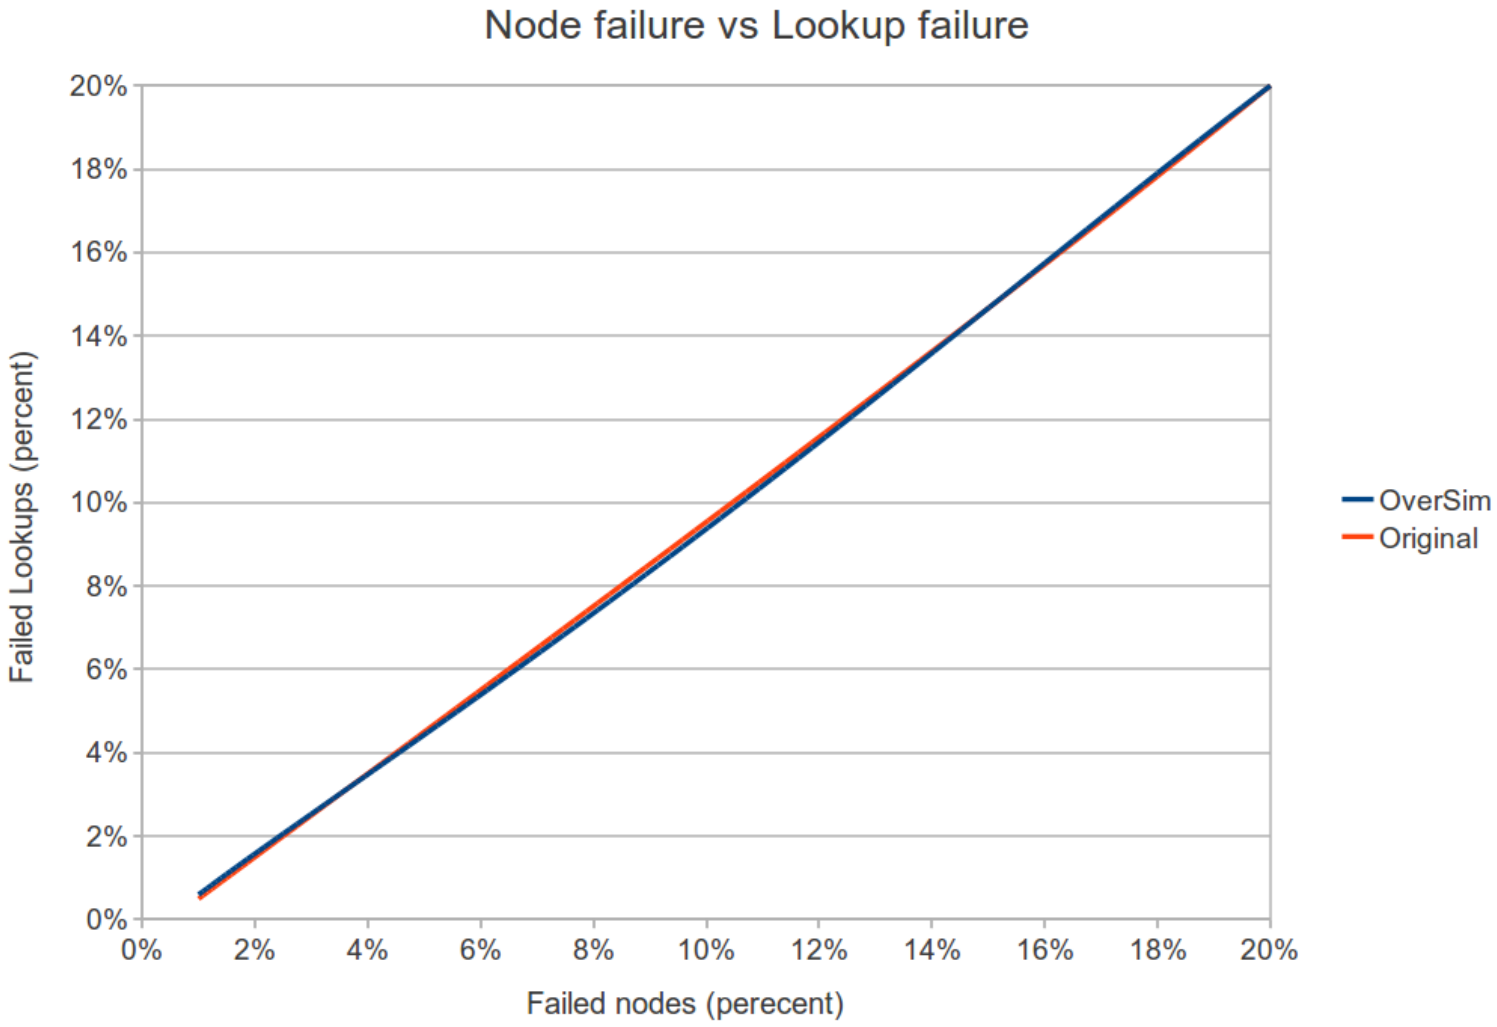

Summary

- OverSim allows for much easier testing of P2P algorithms, on a much larger scale than otherwise possible.
- For existing DHT models we have validated both the routing and maintenance algorithms against their original descriptions.
- This includes Chord, Pastry, Kademlia, and EpiChord.