It's damn hard to decentralise the web

Gareth Tyson

With Aravindh Raman, Sagar Joglekar, Emiliano De Cristofaro, Nishanth Sastry

Queen Mary University of London

It’s damn hard to decentralise the web

“A push towards decentralization…in the end, [became] an extension of centralization”

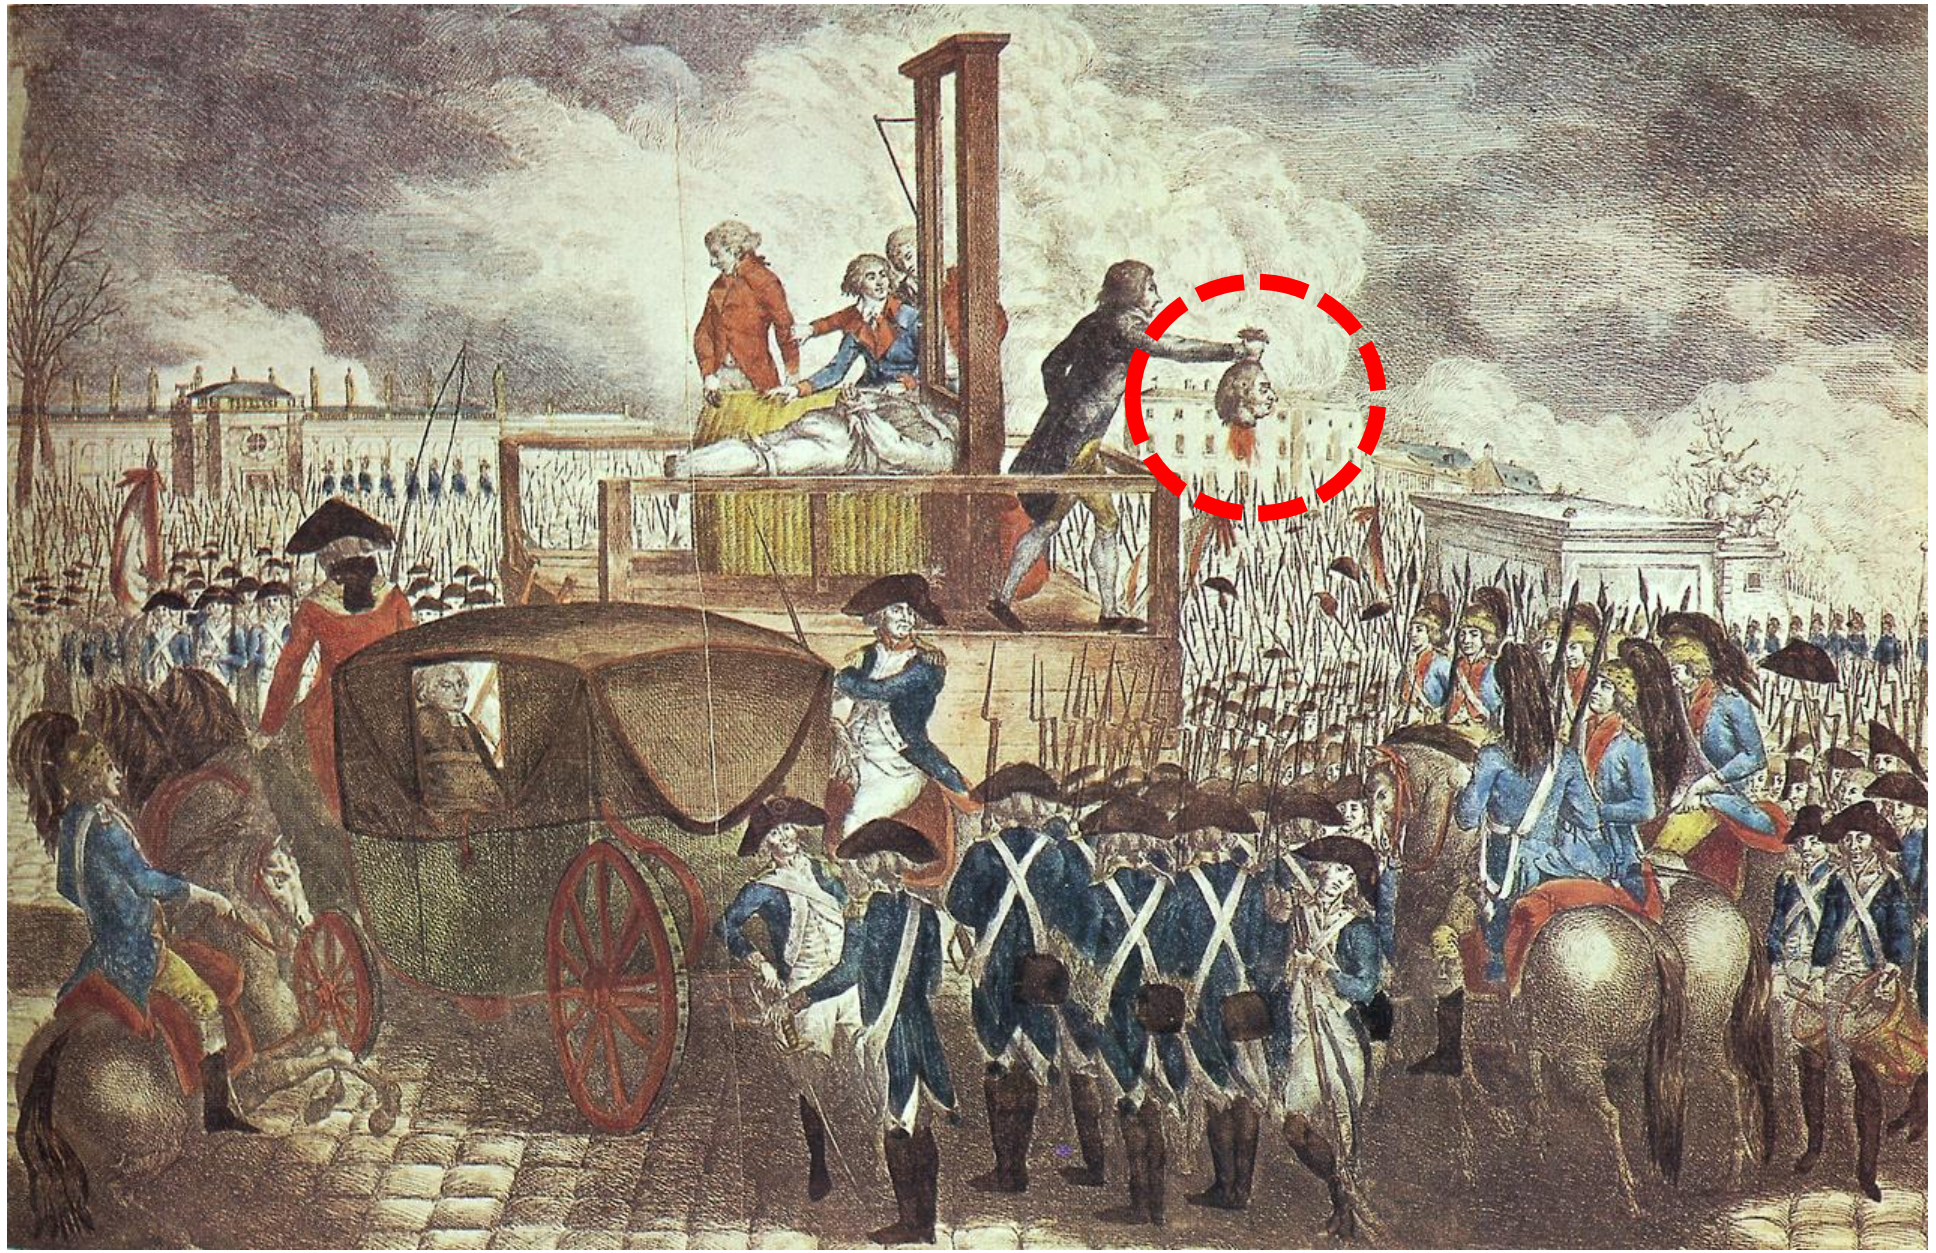

How *did* the web work?

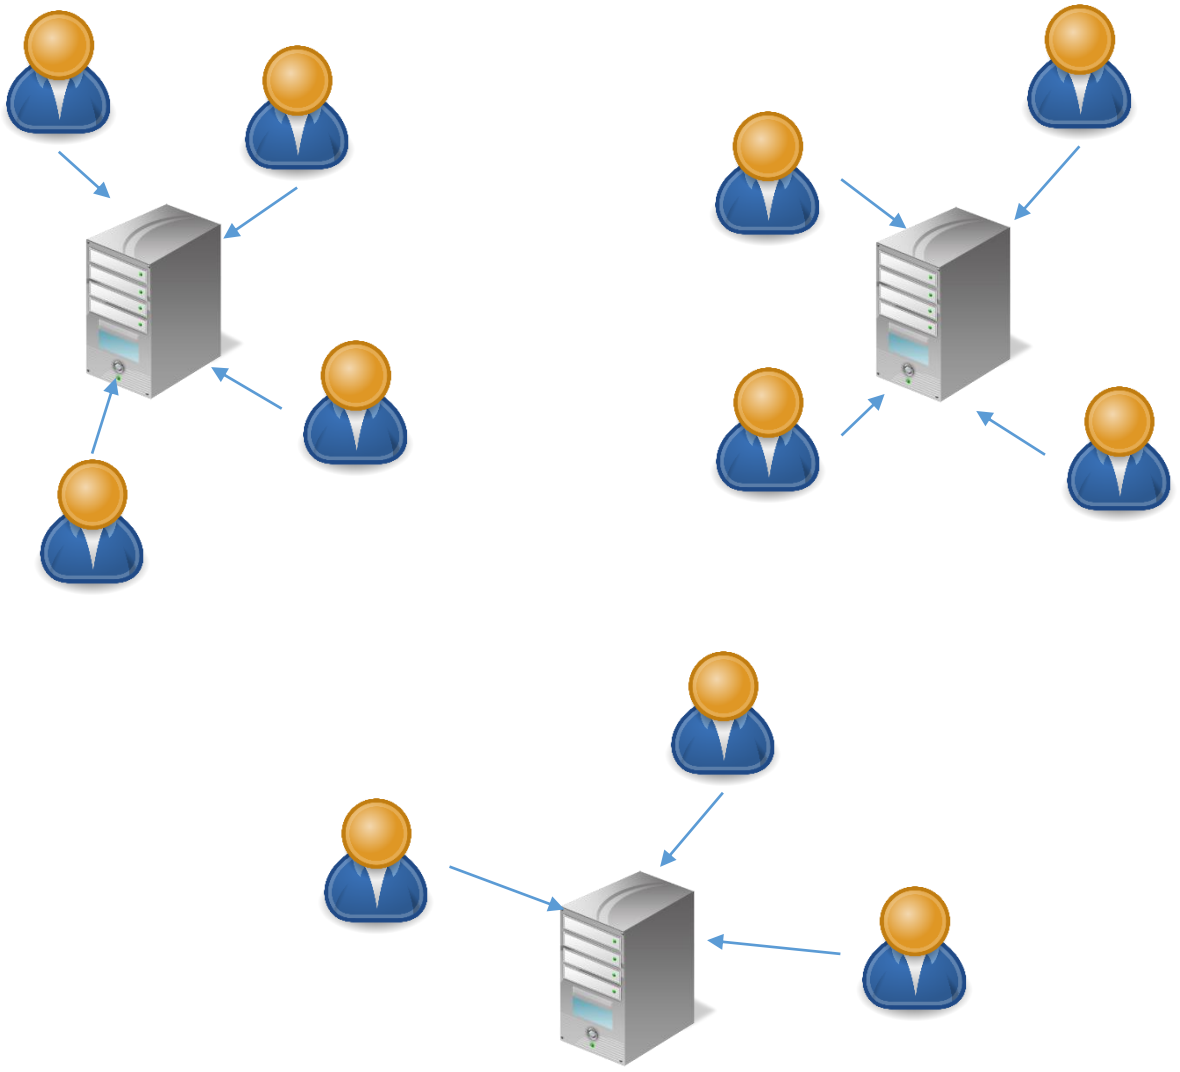

- More symmetrical power structures?
- More resilient?
- No central points of failure?
- Superior privacy?

How *does* the web work?

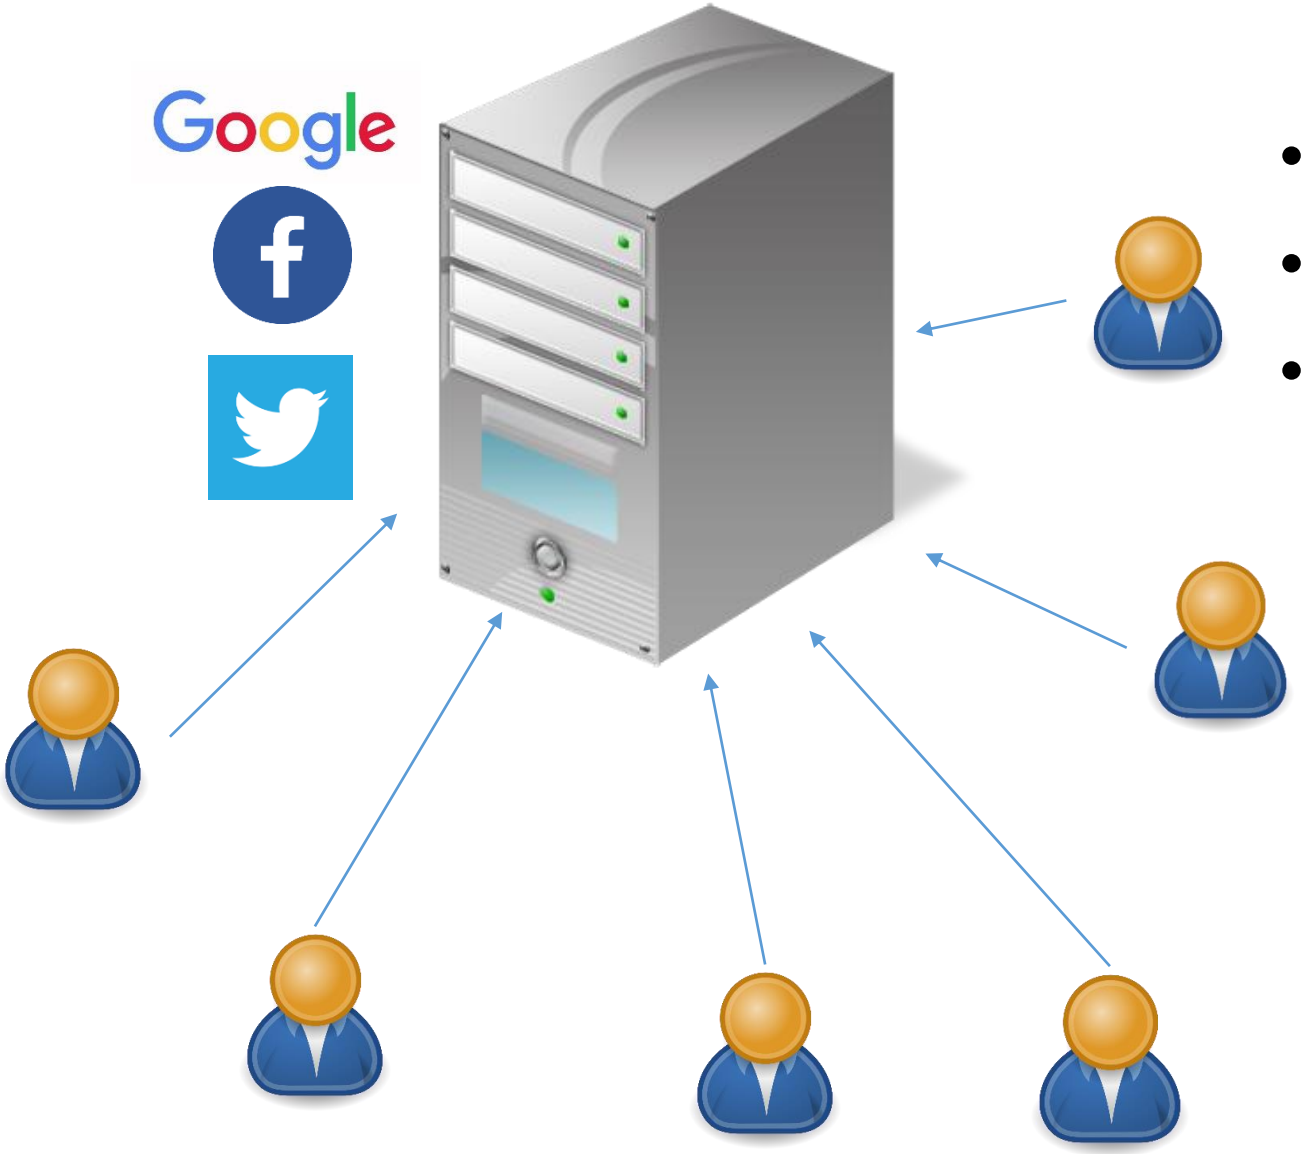

- Critical mass of user?
- Orchestrated deployments?
- Economies of scale?

Can we get the best of both worlds?!

The “**Decentralised Web”** has many open-source implementations

diaspora*

Federated versions of popular services like Facebook, Twitter and YouTube

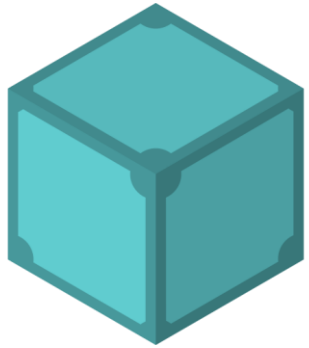

IPFS

How does Mastodon work?

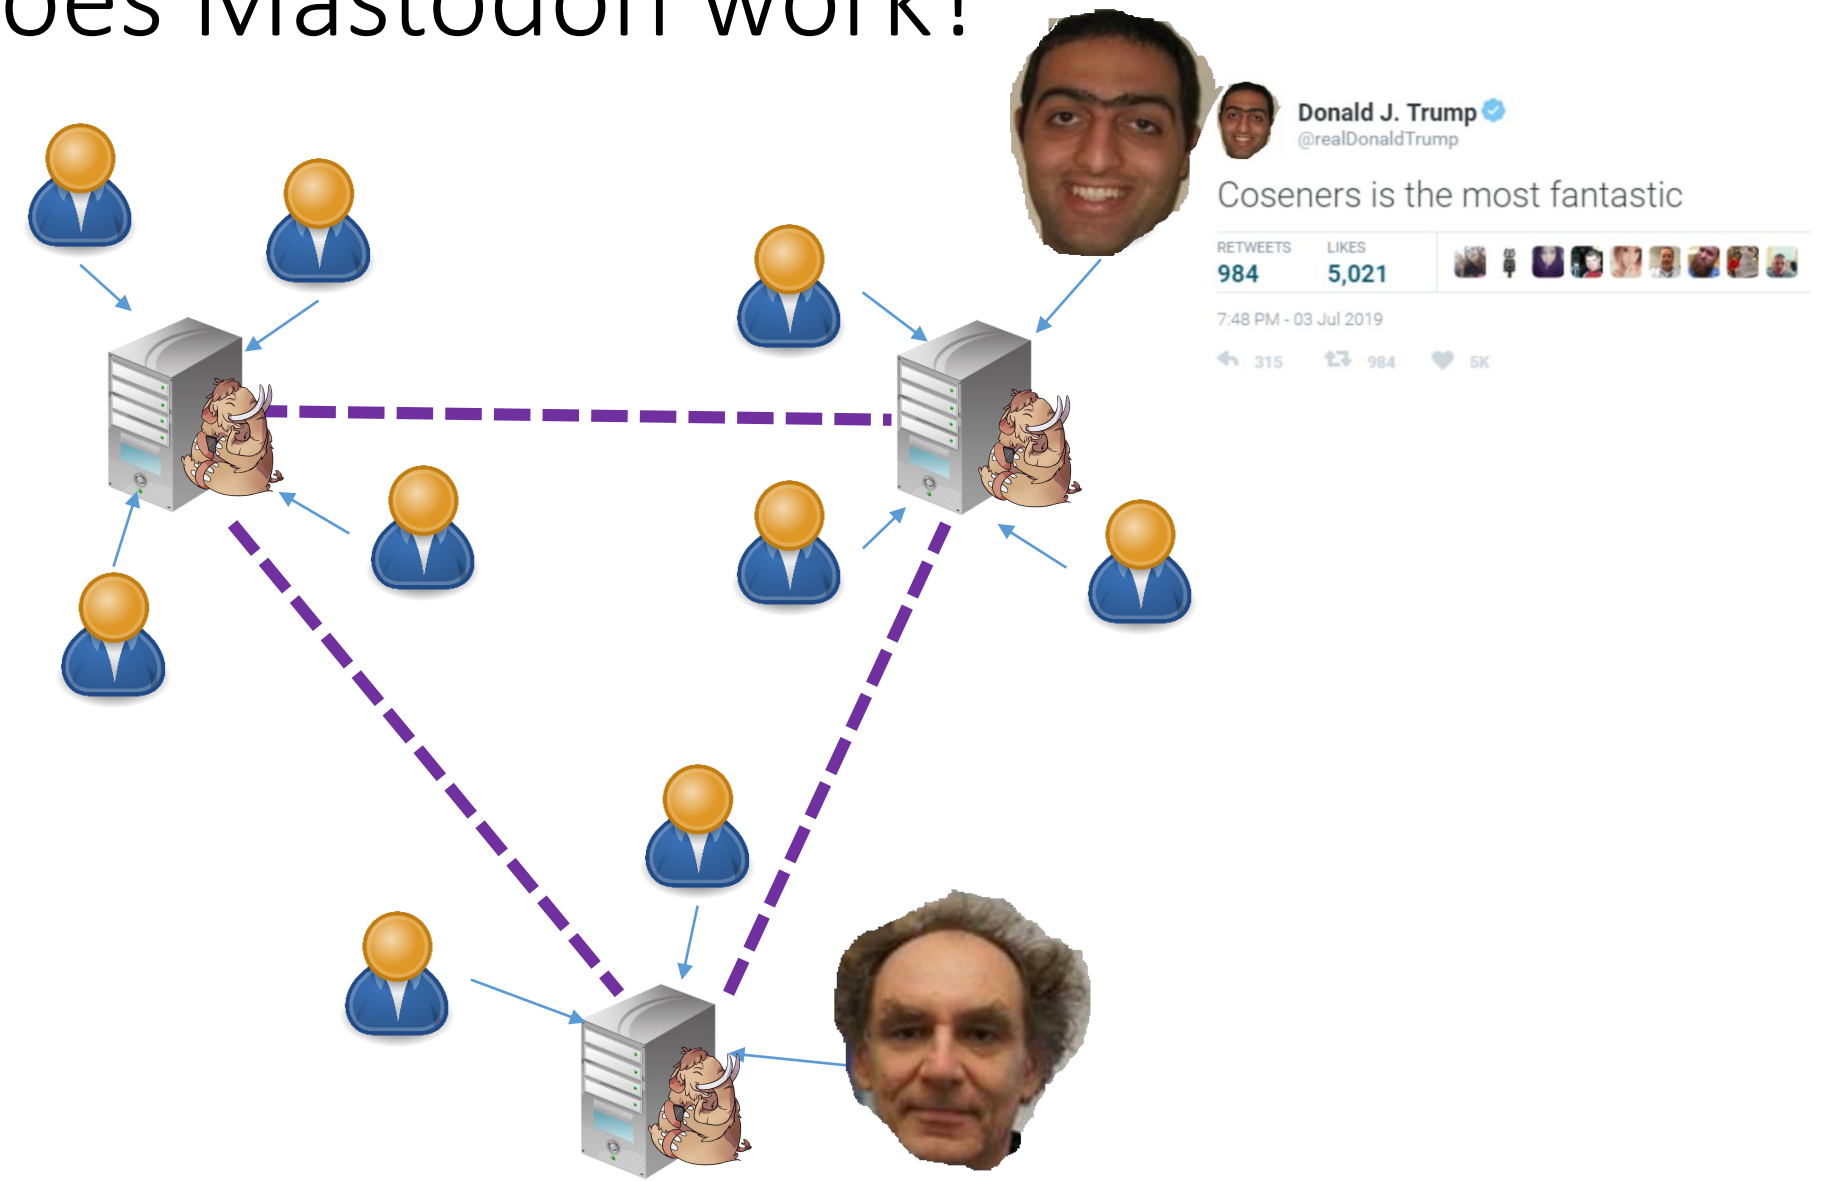

Great!!!

We have decentralised!

No you've not!

Decentralised Louis XIV

Let’s measure Mastodon!

- “Decentralised Twitter”
  - Around 4K servers (aka “instances”)
  - 67M Mastodon “tweets”

- What pressures towards centralisation exist?

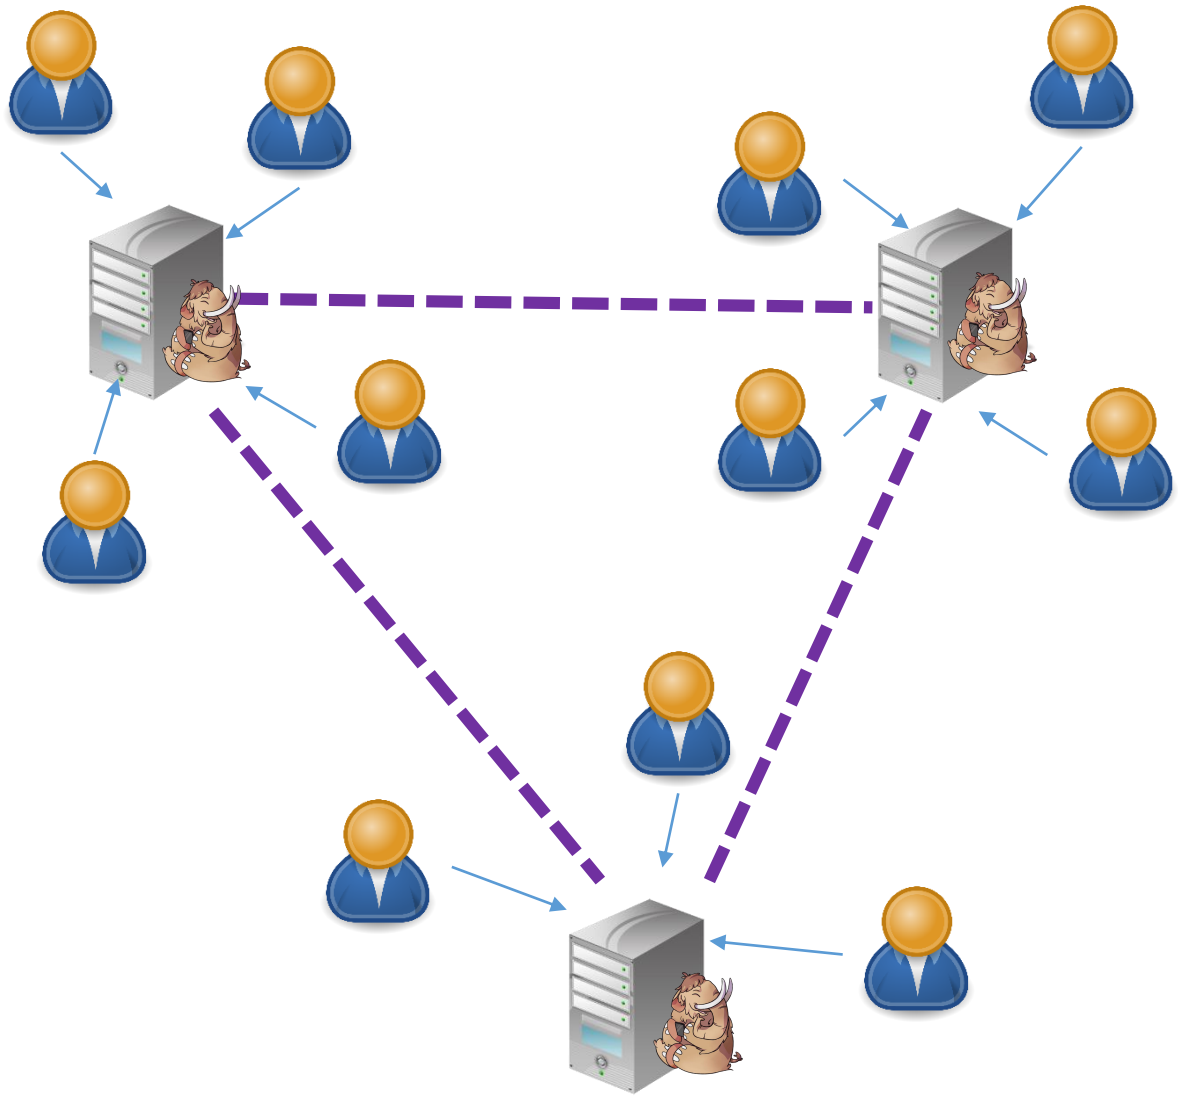

Let’s go through a thought experiment…

Where would you host your server?!

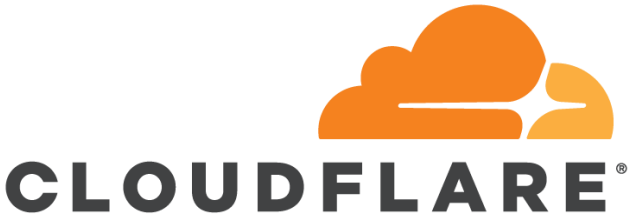

Half of users are reliant on just 3 networks!

…and 2/3 of content

Which instance would you join?!

90% of users join top 5% of the instances

Linux

Top 5% of instances have 95% of content!

But at least decentralized == reliable?!

11% instances inaccessible for half of time

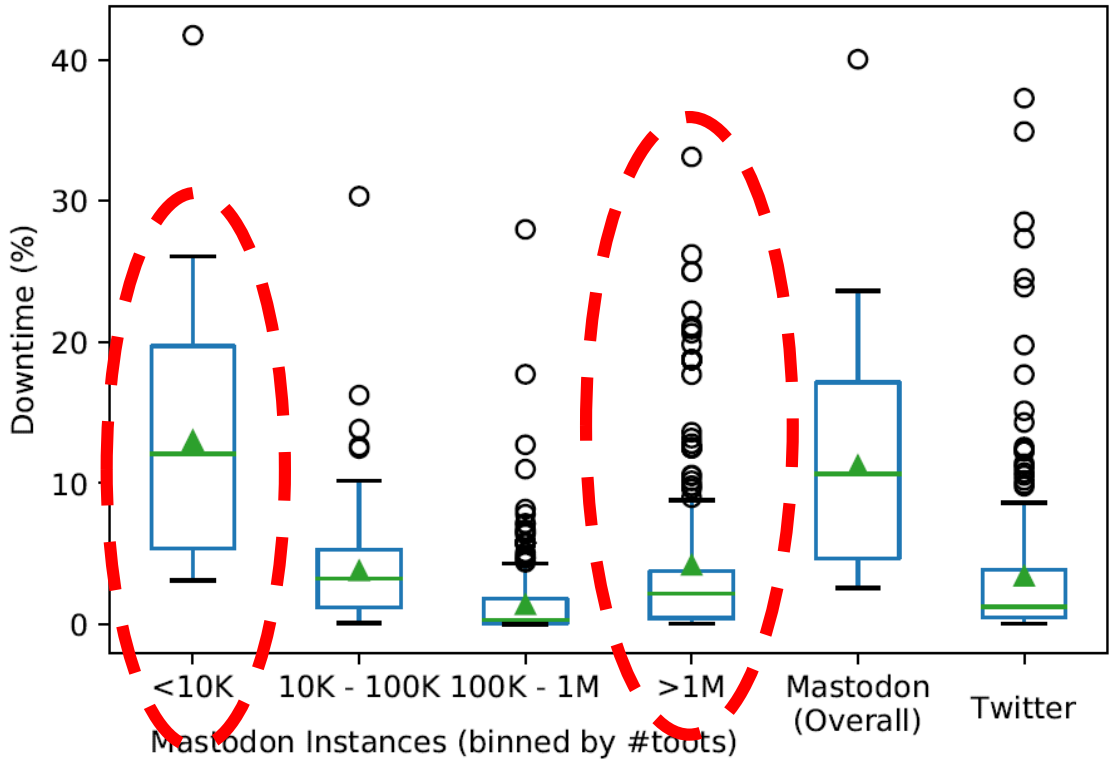

But at least decentralized == reliable?!

11% instances inaccessible for half of time

…still has central dependencies!

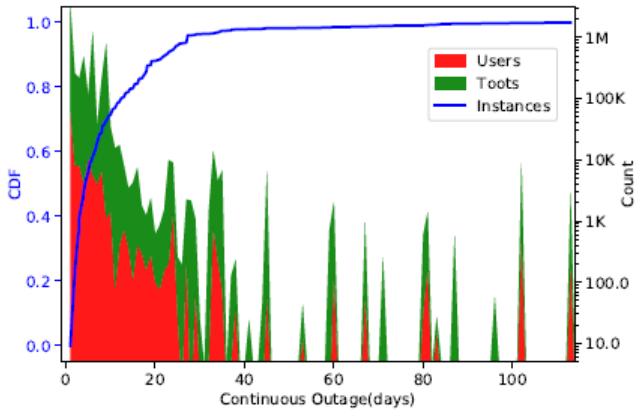

6 entire network outages

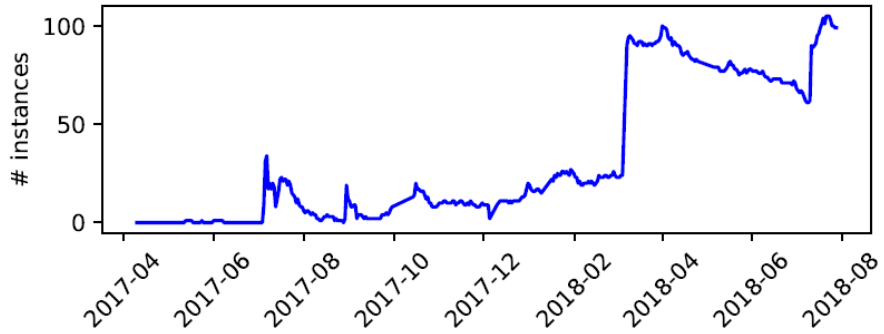

85% of instances use one Certificate Authority

Who you gonna follow?

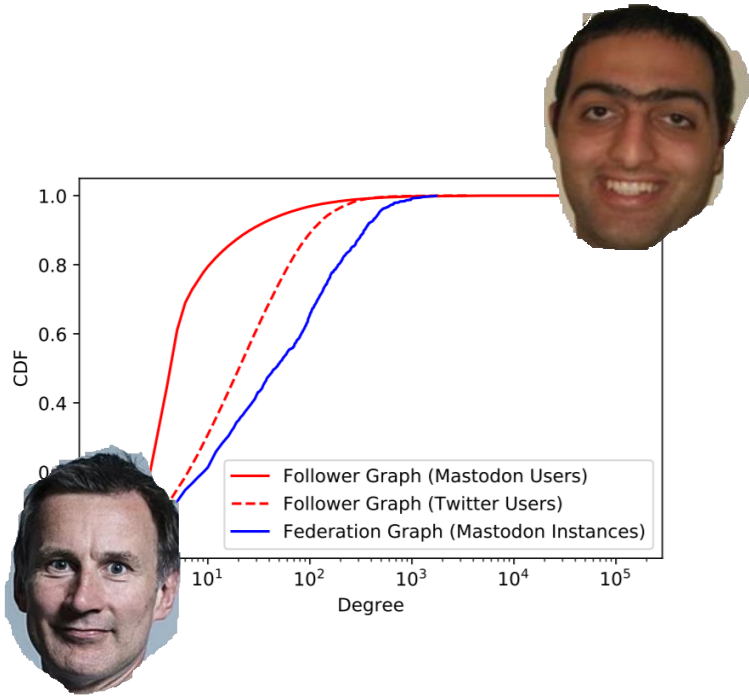

The top 5% accumulate the majority of followers

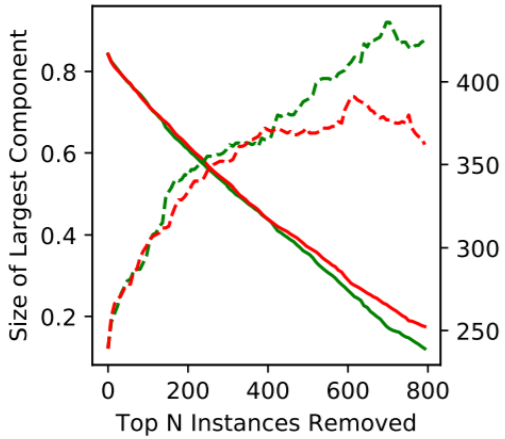

What if Hamed's server goes offline?!

It’s damn hard to decentralise the web

Future Work: Stealing your ideas for the journal version

Gareth Tyson

gareth.tyson@qmul.ac.uk