

# Mid-latency

- ◆ How much latency do you need to add to start seeing end-to-end defense?

# Does it mix?

- ◆ Does low-latency traffic provide cover (“mix”) with mid/high-latency traffic?

# Website fingerprinting

- ◆ Do these attacks work against Tor?
- ◆ Does cell size change things?
- ◆ Does Tor's variable delay change things?
- ◆ What about a little bit of padding, e.g. long-range dummies?

# Location diversity

- ◆ When many nodes are at a single ISP, and many paths are observable by a single ISP, what local algorithms can Alice use to improve (maximize?) her safety?

# Helper nodes and reality

- ◆ What's the right way to do helper nodes in the presence of churn?

# Fragmenting streams

- ◆ Should we fragment streams across multiple paths?

# Congestion attacks

- ◆ Can you “measure” Alice by ICMP pings even if she doesn't relay traffic for you?

# Incentives to relay

- ◆ Is it always unsafe to use your server for your anonymous traffic?



# Pseudonyms/profiles

- ◆ Logging into your gmail account and then posting to Indymedia is bad.
- ◆ But a new circuit for every request is also bad.
- ◆ What's the right compromise/strategy?