

Characterizing Global Web Censorship: Why is it so hard?

Phillipa Gill

The Citizen Lab/Stony Brook University

Work done in collaboration with:

**Masashi Crete Nishihata, Jakub Dalek, Sharon Goldberg,
Adam Senft and Greg Wiseman**

Overview

Large-scale politically driven Internet outages are well known...

- ...but what happens **within** countries is less well understood

We leverage data gathered by an interdisciplinary group (Open Net Initiative) to bootstrap analysis

- 77 countries, 286 distinct ISPs, measured from 2007-2012
- **Advantages:** context about **what, when,** and **where** to measure
- **Disadvantages:** dearth of technical data/raw measurements


Our results highlight important challenges for censorship research!

Background

- Where censorship can happen:

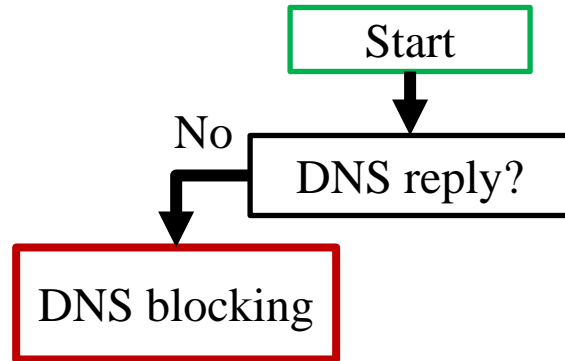
Background

- Where censorship can happen:



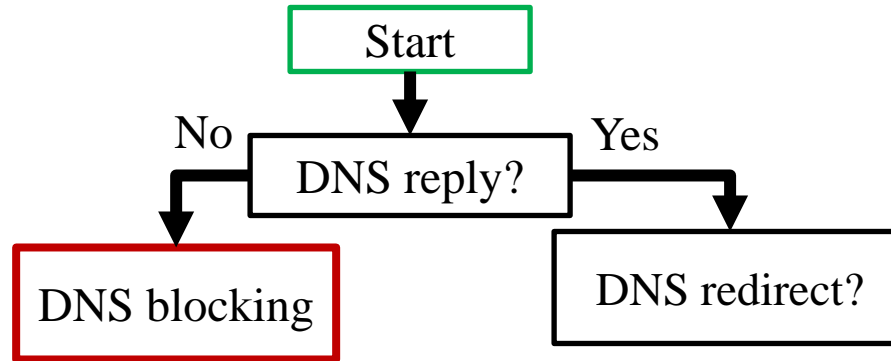
Background

- Where censorship can happen:



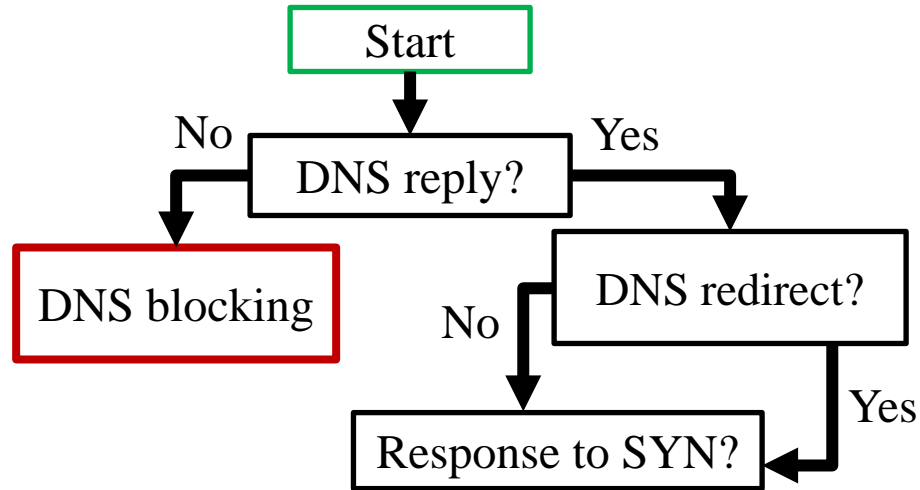
Background

- Where censorship can happen:



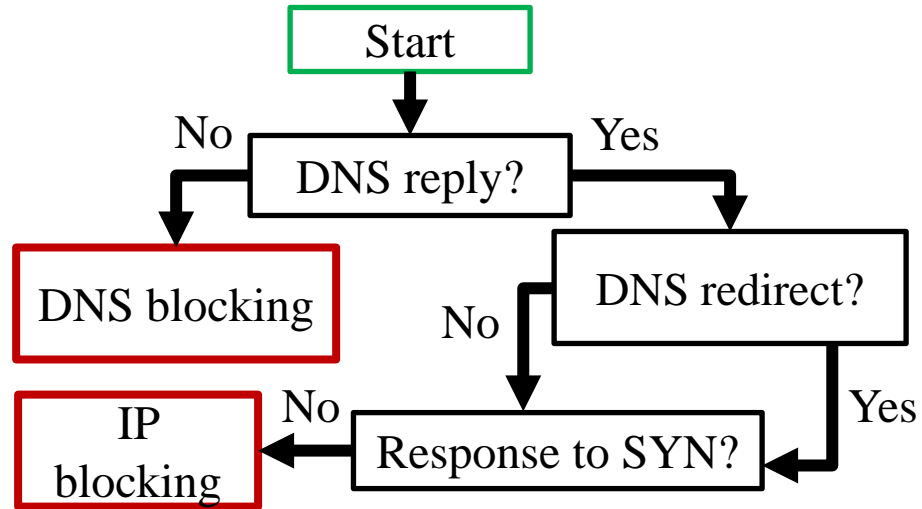
Background

- Where censorship can happen:



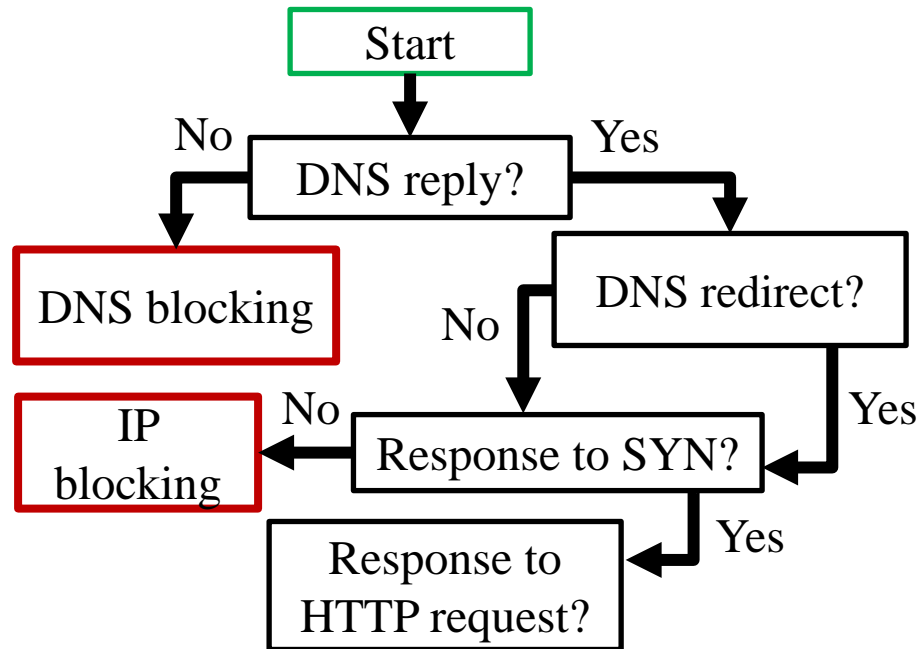
Background

- Where censorship can happen:



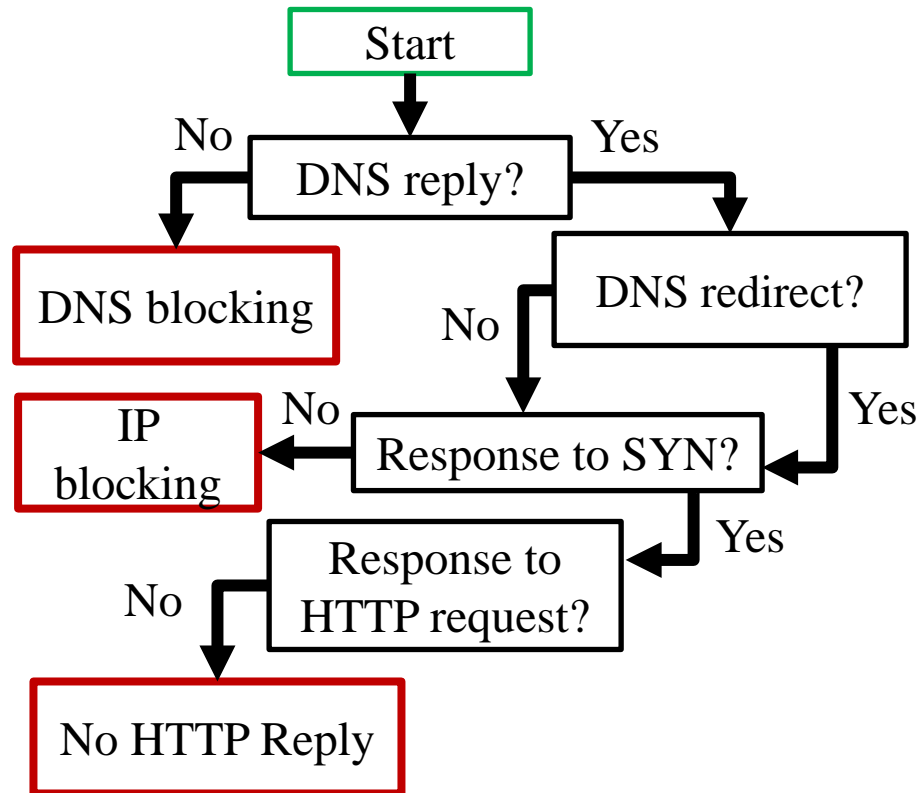
Background

- Where censorship can happen:



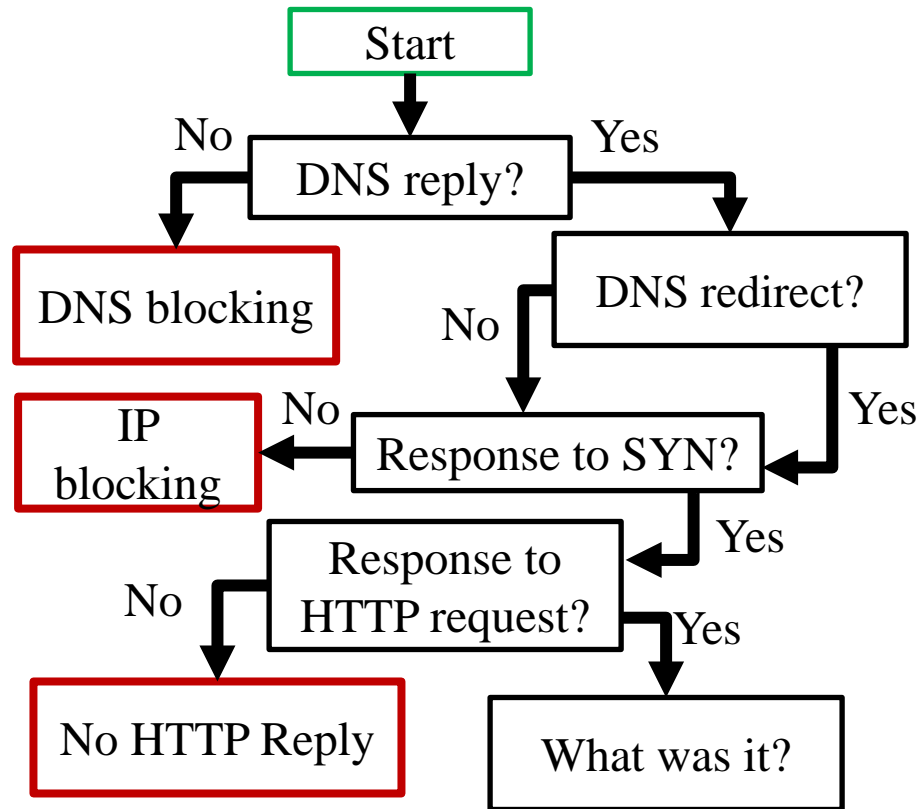
Background

- Where censorship can happen:



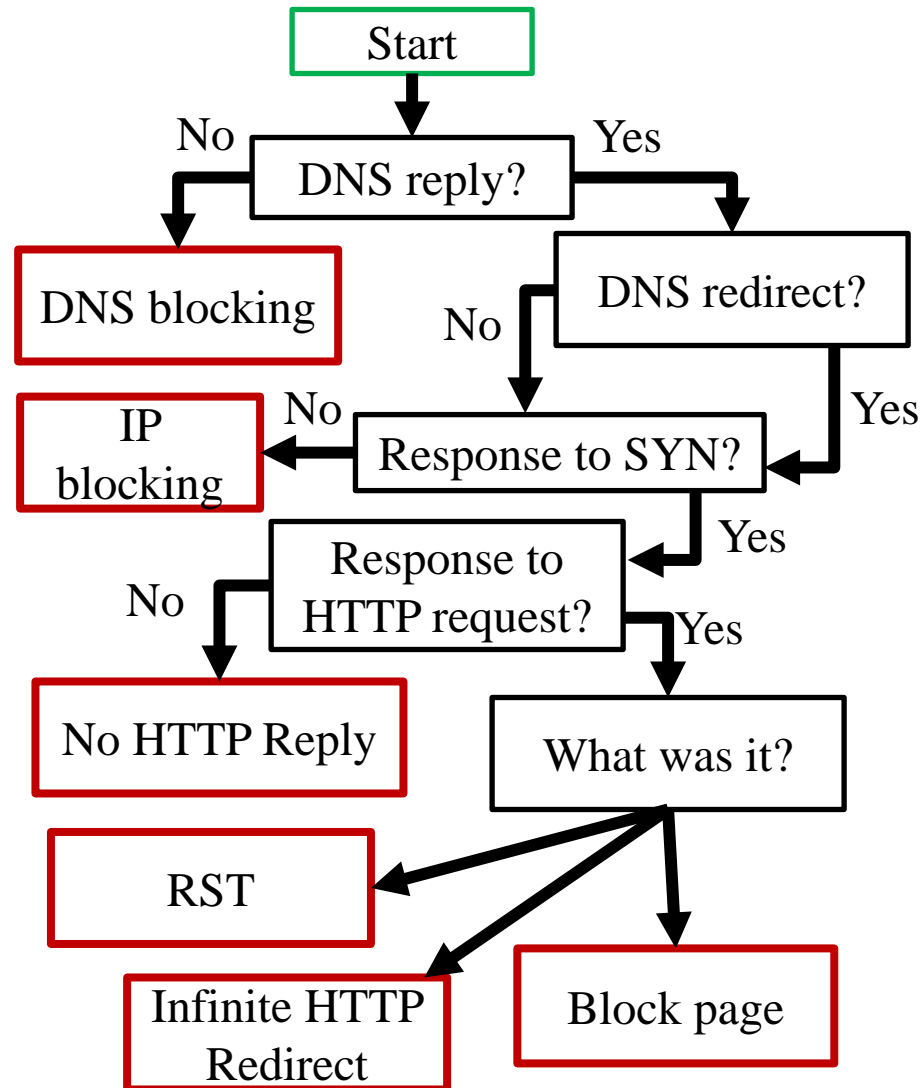
Background

- Where censorship can happen:



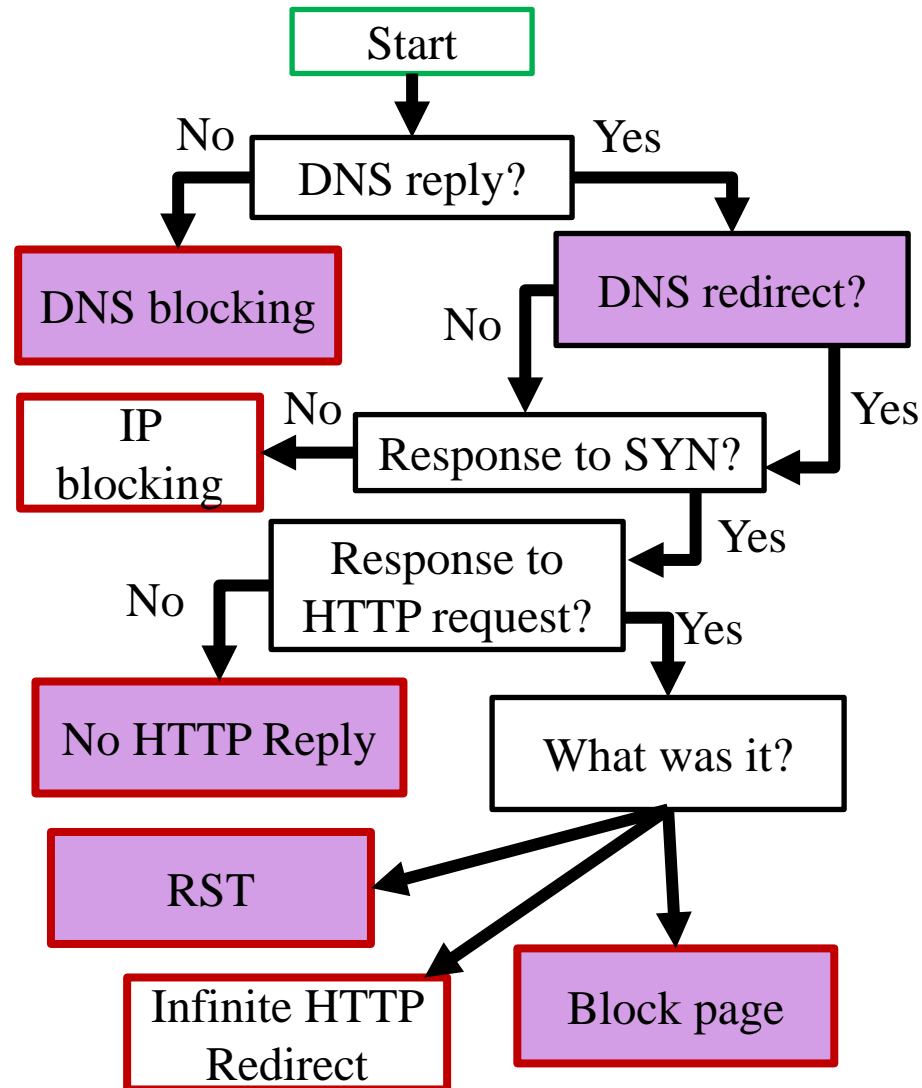
Background

- Where censorship can happen:



Background

- Where censorship can happen:




Methodology

- **Basic idea:** Issue requests for a consistent set of sites in the **field** and a control location (**lab**)
- Software synchronizes the requests between lab and field
- Once both lab and field have completed, results sent back to the lab for more analysis
- **What is tested:**
 - Sites that are likely to trigger censorship
 - Determined in collaboration with regional groups
- **Where are tests run:**
 - Combination of targeted/opportunistic testing
 - Performed by regional collaborators after informed consent meeting


Challenges for censorship research

1. Variation between countries




■ No DNS Reply ■ DNS Redirection ■ No HTTP Reply
■ RST ■ Blockpage

1. Variation between countries




■ No DNS Reply ■ DNS Redirection ■ No HTTP Reply
■ RST ■ Blockpage

1. Variation between countries




1. Variation between countries



There is no such thing as a “representative” country


2. Variation between ISPs

Decentralized blocking in UAE



2. Variation between ISPs


Decentralized blocking in UAE



“Du” ISP does not censor prior to April 2008

2. Variation between ISPs


Decentralized blocking in UAE



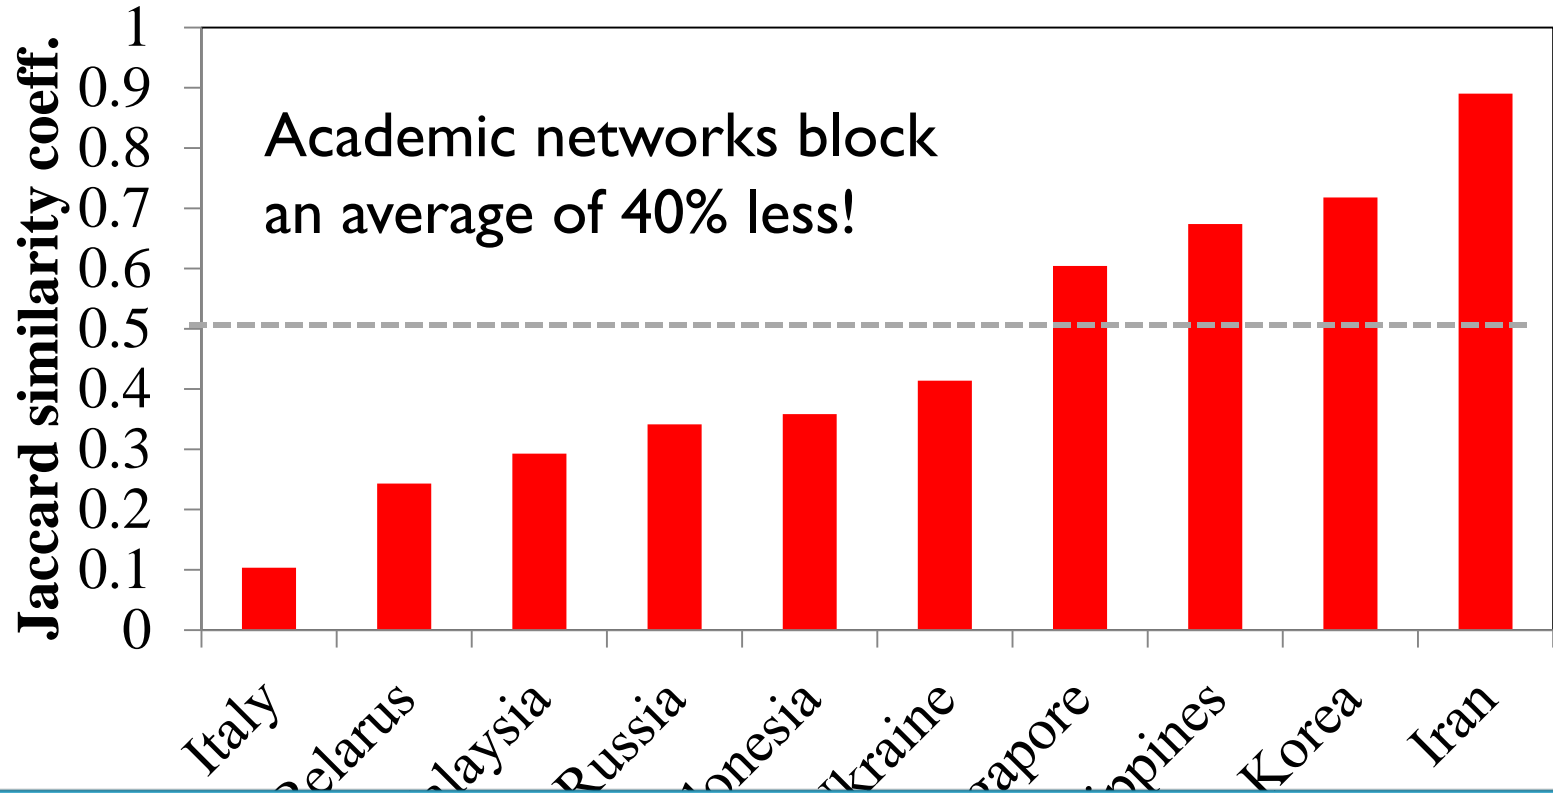
Censorship is a per-ISP property
(when censorship is decentralized)

2. Variation between types of networks

2. Variation between types of networks




2. Variation between types of networks



Academic networks are not representative!

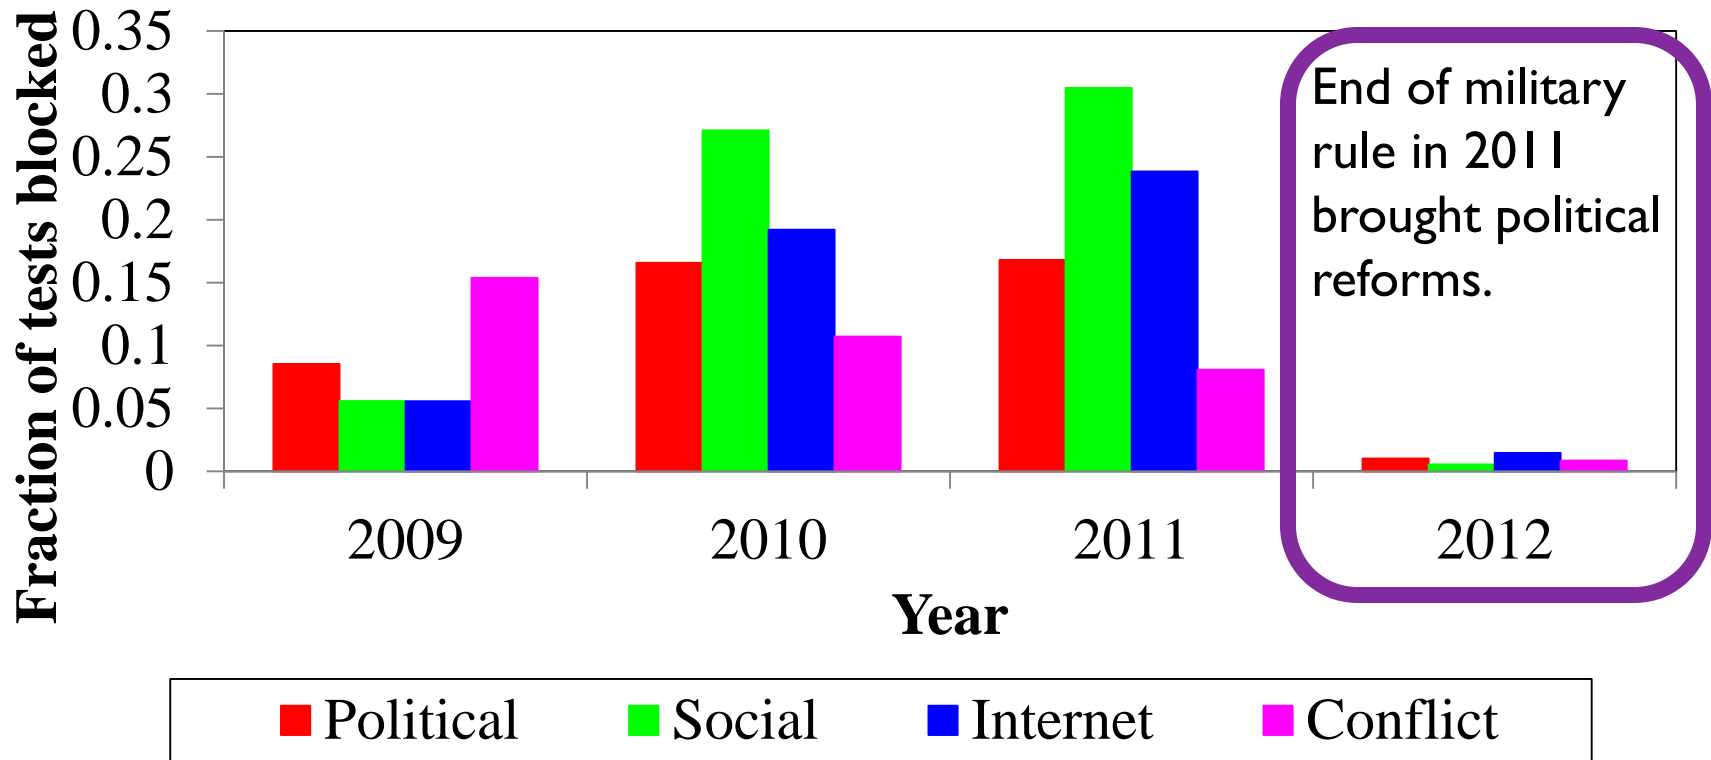
3. Sudden temporal shifts in blocking

Censorship in Burma over time




3. Sudden temporal shifts in blocking

Censorship in Burma over time



3. Sudden temporal shifts in blocking


Censorship in Burma over time



Need to measure over time and correlate with political changes


4. Stealthy blocking of certain content

Censorship of content in Yemen



4. Stealthy blocking of certain content

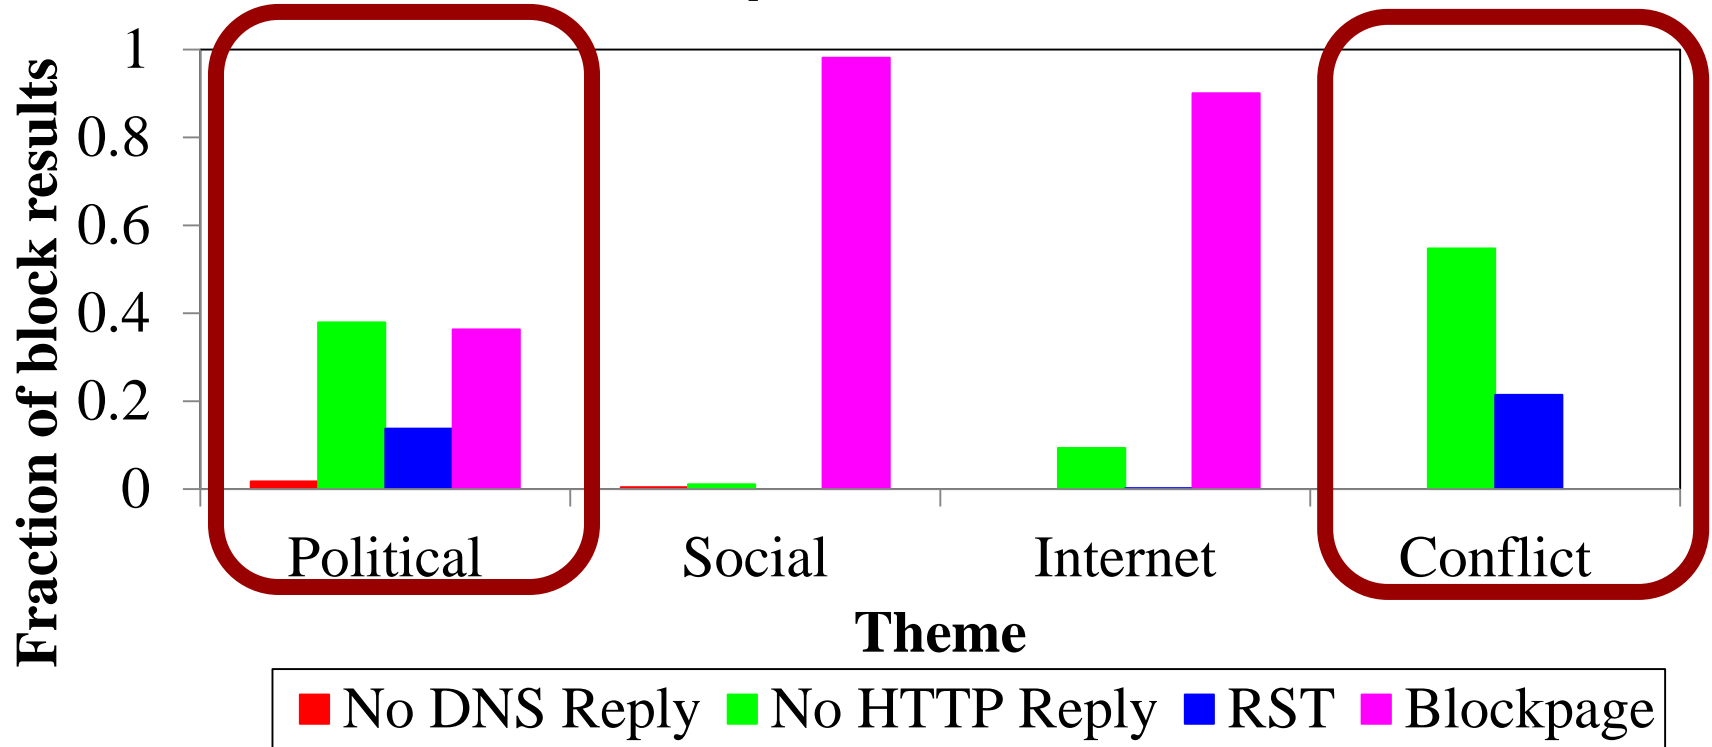
Censorship of content in Yemen



Transparent blocking of
social and Internet content

4. Stealthy blocking of certain content

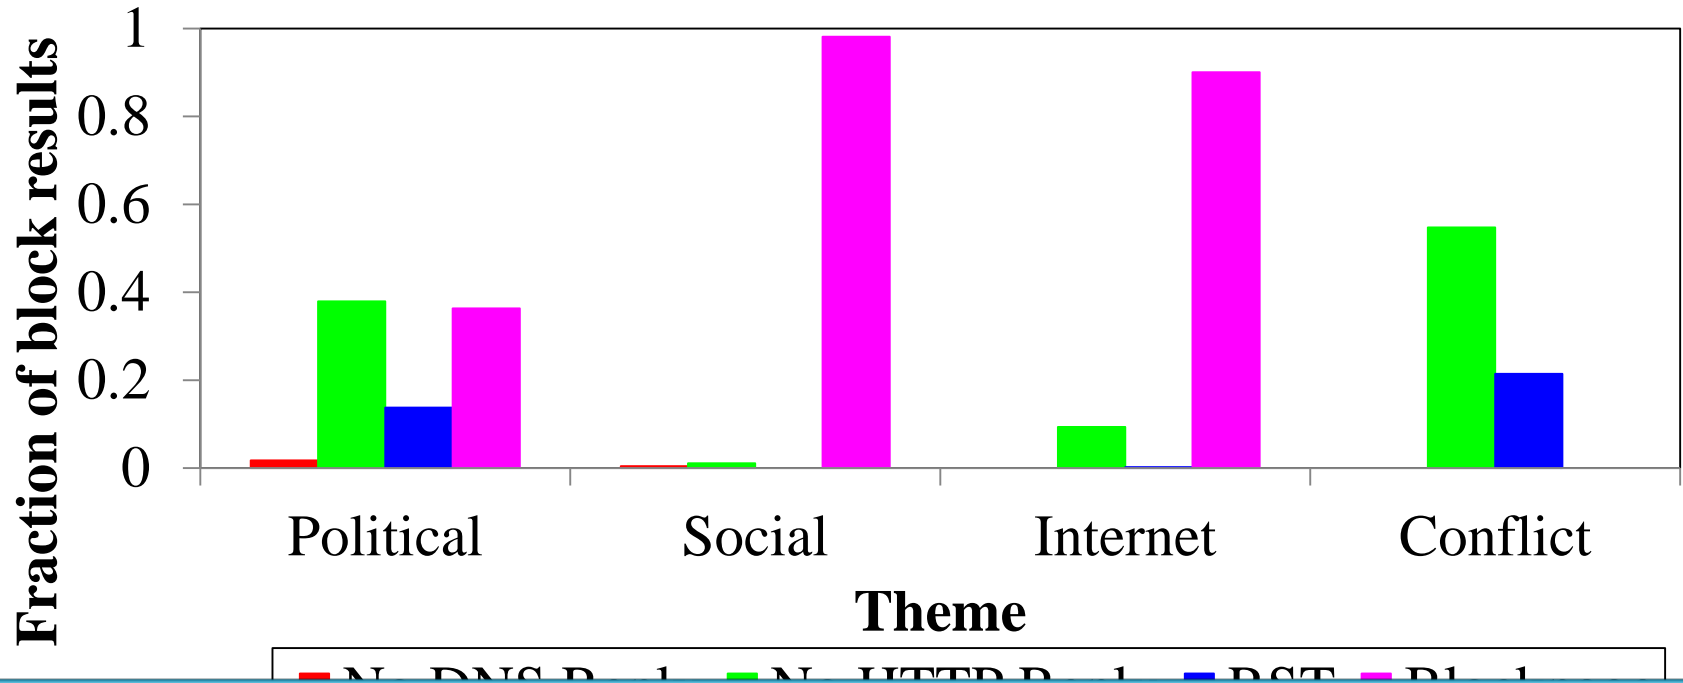
Censorship of content in Yemen



“Stealthy” blocking of political and conflict related content


4. Stealthy blocking of certain content

Censorship of content in Yemen




Measurement needs to be robust to distinguish failure from censorship

5. The type of content tested matters

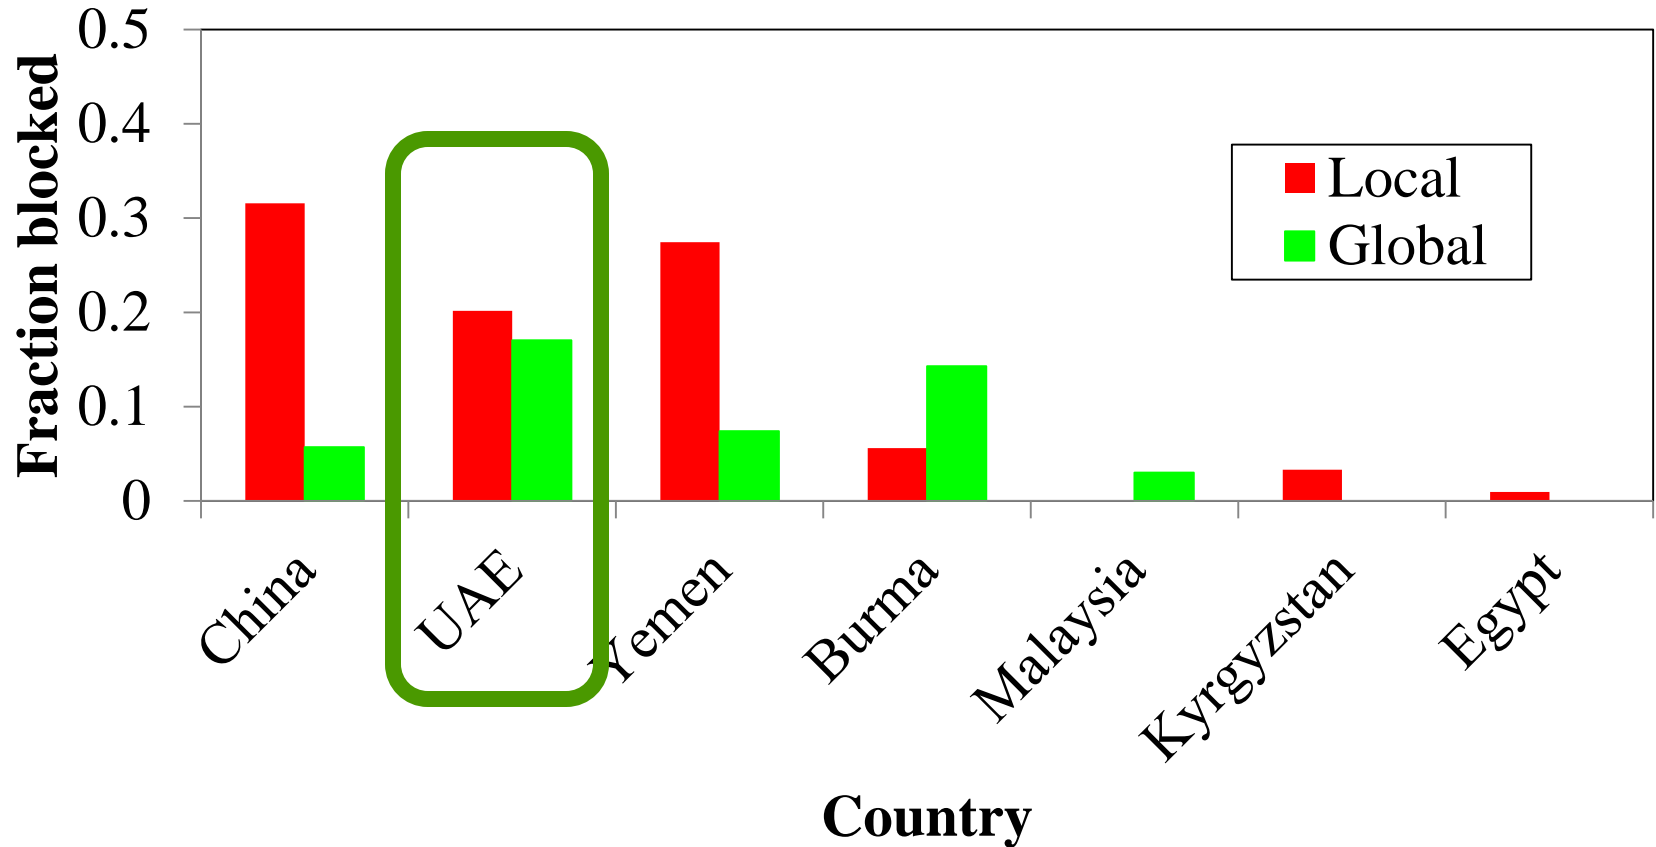


5. The type of content tested matters



3-5X more blocking of local content in China/Yemen
* most blocked content is political

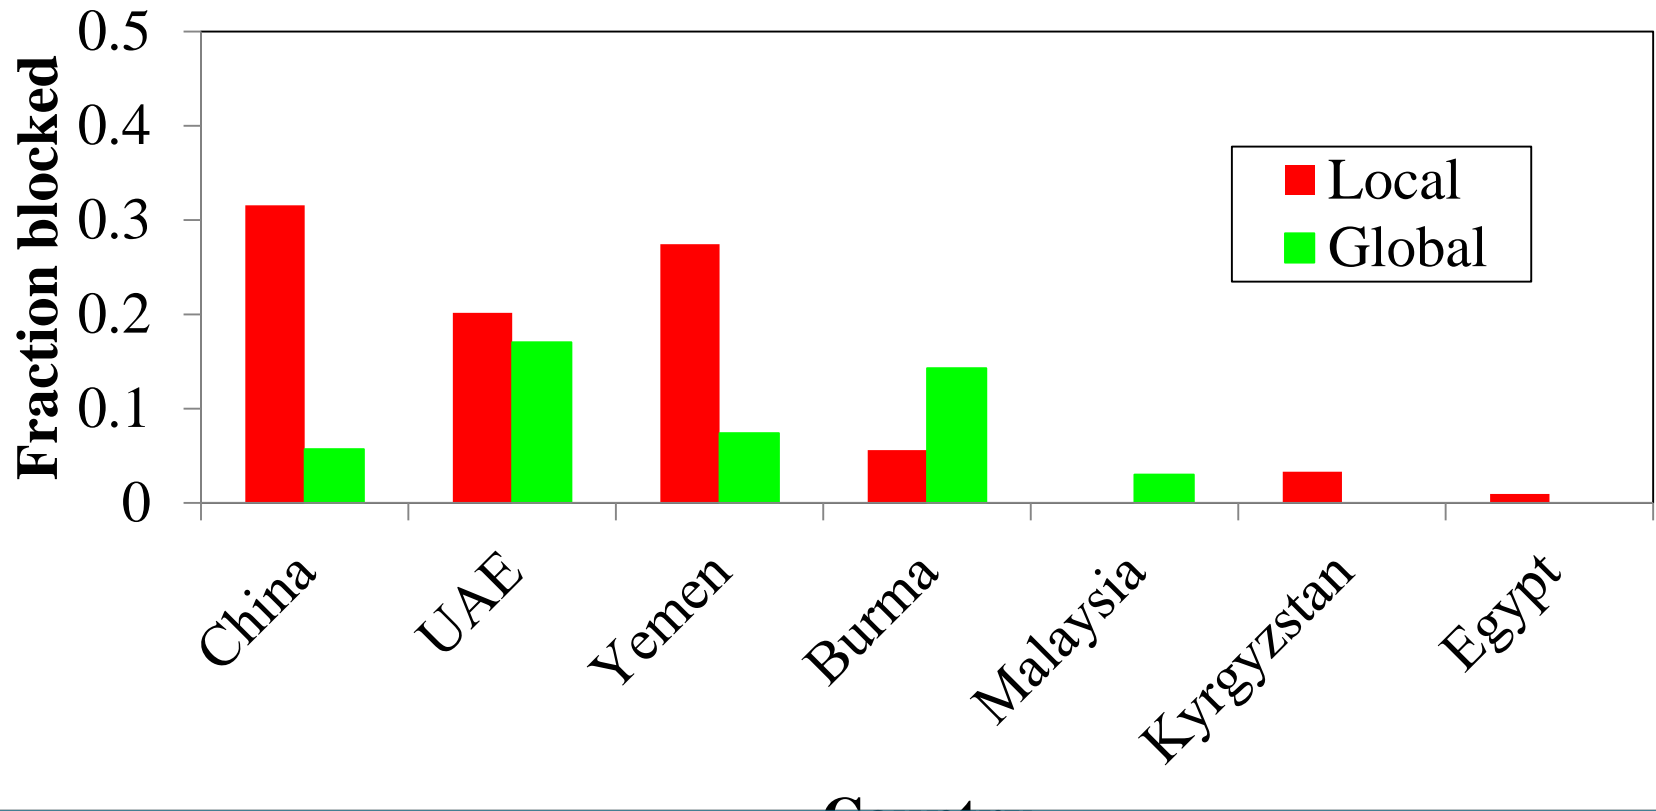
5. The type of content tested matters



Less discrepancy in UAE

* most blocked content is social

5. The type of content tested matters



Need to take an interdisciplinary approach to determine **what** content to test

Challenges for censorship research:

1. Variations between technology used by countries
2. Variations between ISPs and between ISPs and institutions
3. Sudden temporal shifts in blocking
4. Stealthy blocking of certain content
5. Locally relevant content is more likely to be blocked

And more!

... maintaining infrastructure across funding cycles/staff turn over

... informed consent/preserving user privacy when testing can pose a **physical risk!**

What's next?

More measurements, taking an interdisciplinary approach to tackle the problem:

- Rigorous measurements + political context

Data sharing?

- **Short answer:** we're working on it.
- **Longer answer:** this project has laid the foundation in terms of unifying the data and removing PII.
 - Anticipate releasing data in the next ~4 months

What I hope to get out of this workshop

- Discuss how existing platforms may be used for censorship research
 - Particularly interested in:**
 - Platforms with visibility into the network edge
 - DNS/BGP measurements
- Discuss how a large scale, long-term censorship measurement platform may be built
- Discuss how we might distinguish transient failures/TCP bugs from actual censorship