Facade: High-Throughput, Deniable Censorship Circumvention Using Web Search

Ben Jones, Sam Burnett, Nick Feamster,
Sean Donovan, Sarthak Grover, Sathya Gunasekaran,
and Karim Habak

Georgia Institute of Technology

FOCI 2014

Censorship is a common problem

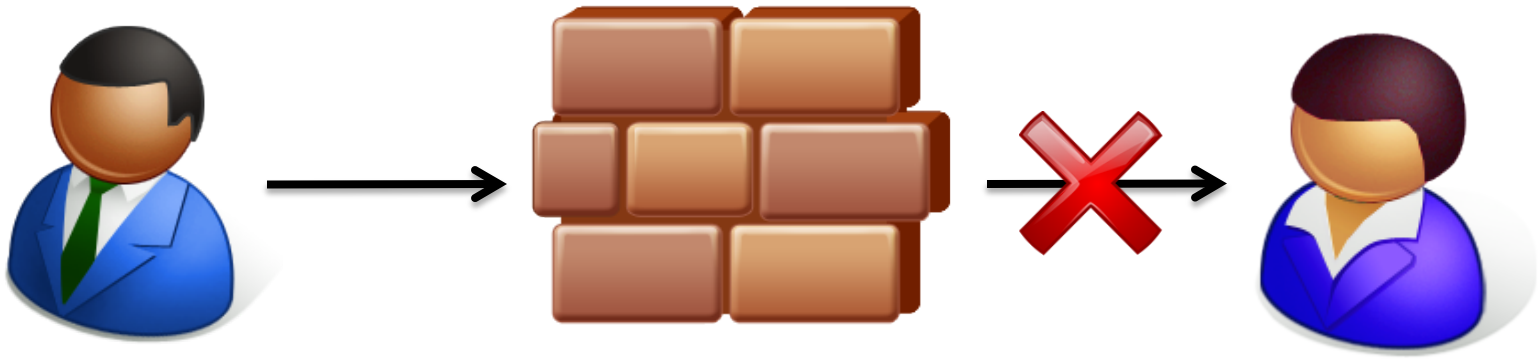

Difficult to hide that you are using Tor

Bob

Alice

HTTP circumvention tools are necessary or they soon will be

Iran partial
traffic

View

Posted February 10th, 201
filtering, positive futures,

Bob

We need to target new points on performance/deniability curve

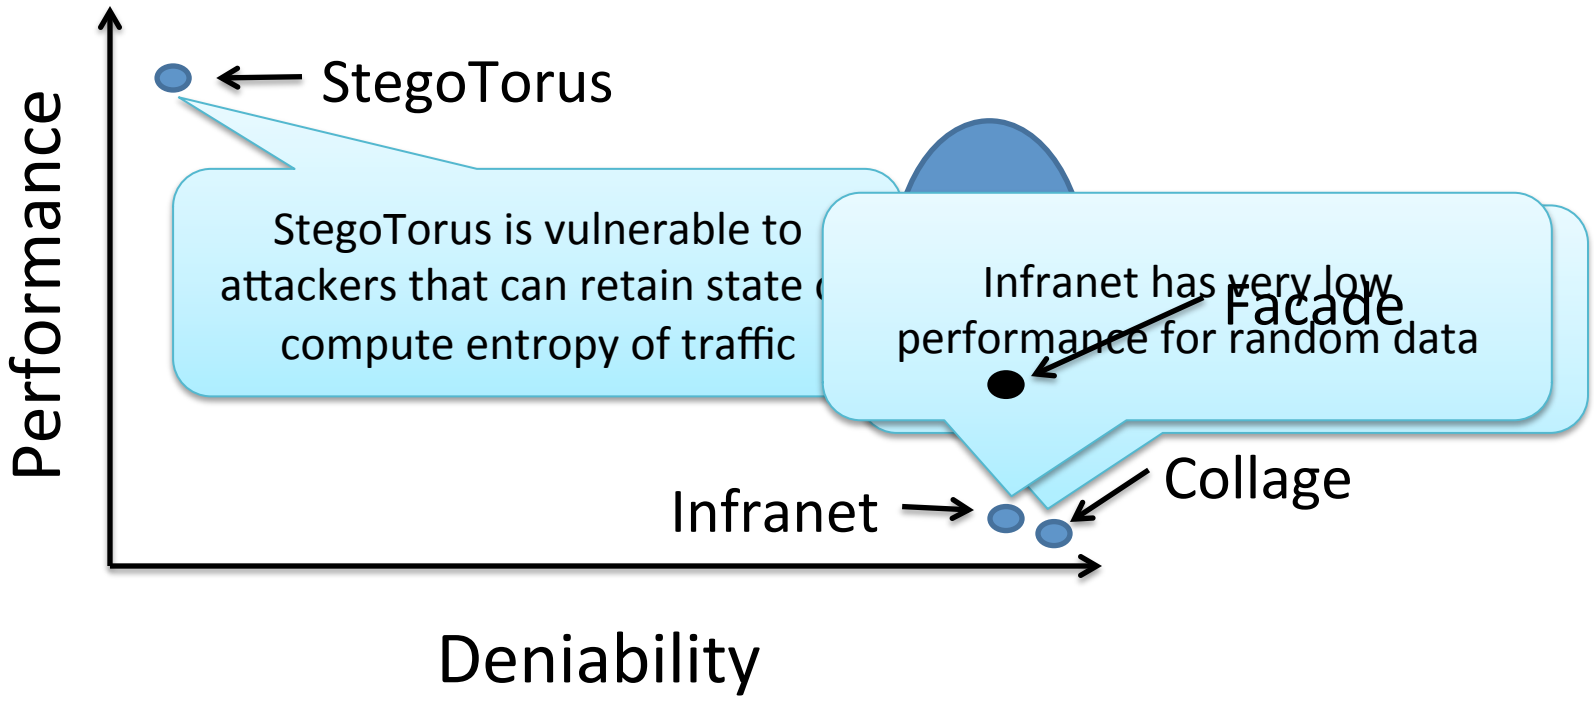

Note: graph drawn for emphasis, not to scale

Research Problems

- How can we create deniable, HTTP covert channels?

- Can we get the deniability of Infranet with better performance for encrypted data uploads?

Our Solution

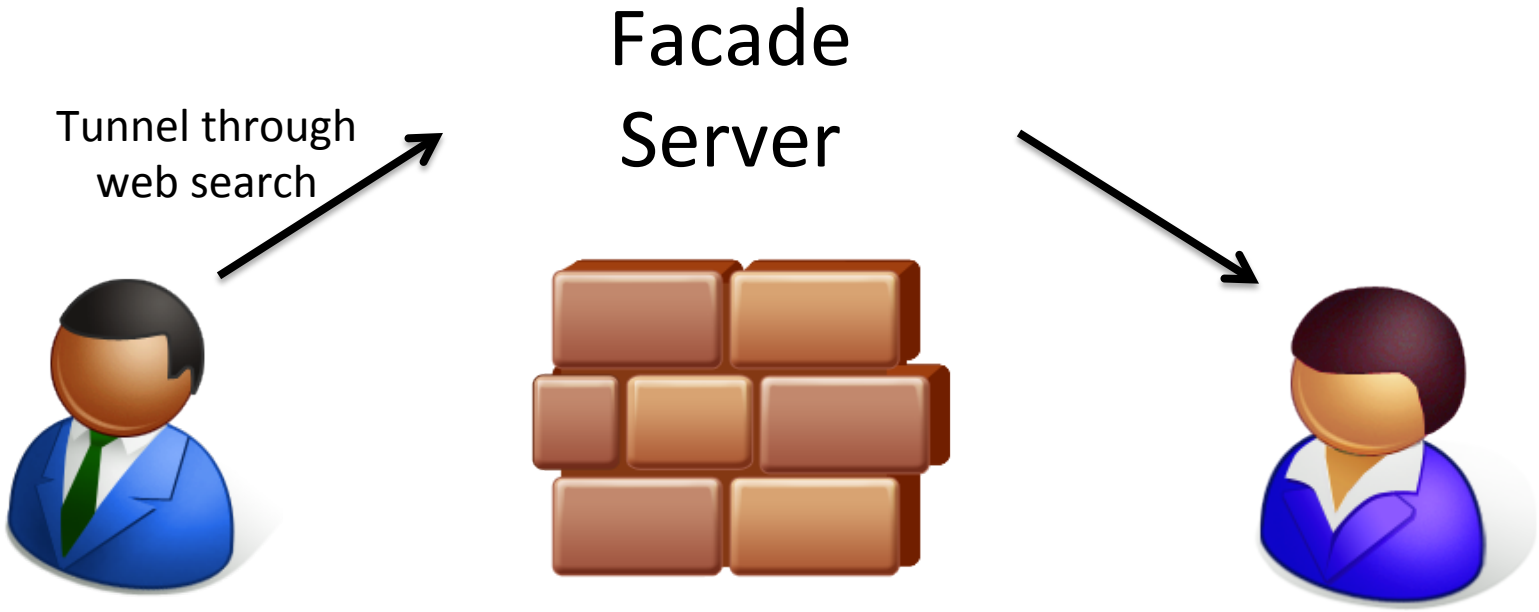

- Everyone searches the web and search has dozens of bits of entropy
- Let's use this entropy to hide information

Outline

- Motivation
- Facade protocol
- Evaluation

Threat Model

- Our target censor can:
  - Detect and block all protocols other than HTTP
  - Store some state (several HTTP request/ response pairs)
    - Ex: detect that information in cookies has not been set by the server
  - Censor can operate in-path
    - Ex: create error conditions to fingerprint client or server

Facade Overview

- Facade encodes information in web search
- Real users browse and search at the same time so Facade encodes information in browsing and search
- Note: Facade server must have sufficient cover search traffic to maintain deniability

Encoding data in search

- Encode information in the path string with a dictionary encoding

- The dictionary is a mapping from data to English

- Example: http://www.example.com/?q=banana+law encodes the string “hello”

Making search deniable with OpenSearch

- What is it?
  - Specification for sending search requests
- How does it work?
  - Encodes query into a URL
- Why are we using it?
  - Widely deployed: Chrome, Baidu, Yandex, etc.

OpenSearch Example

Query: where to buy peanuts

URL: https://duckduckgo.com/?q=where+to+buy+peanuts

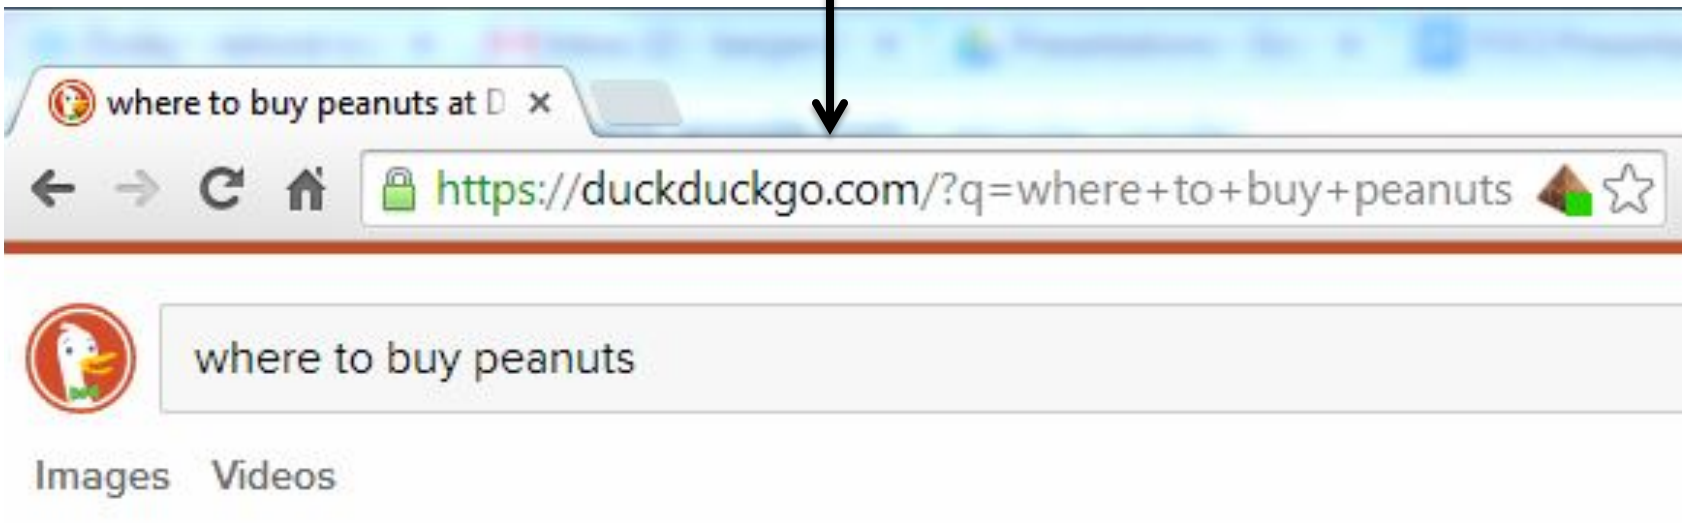

System Overview

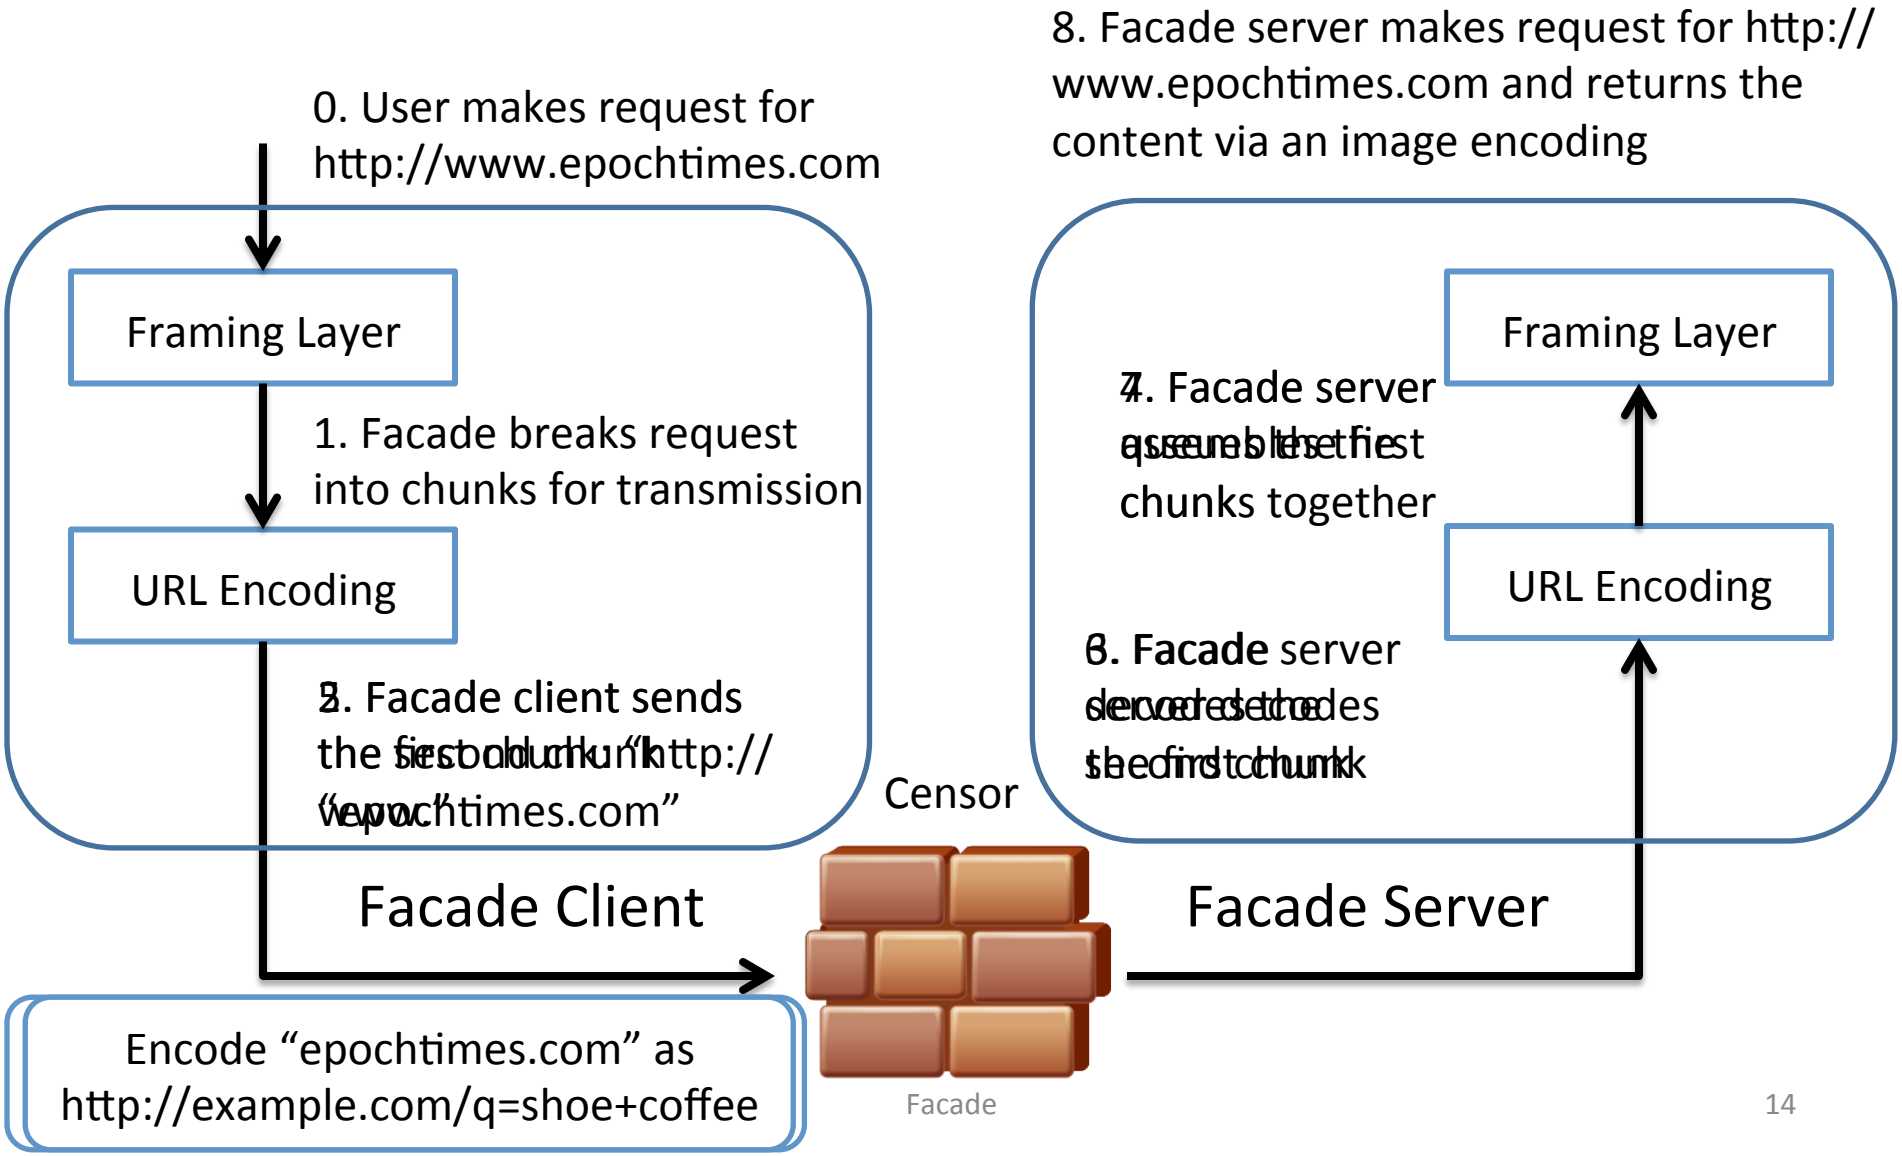

Outline

- Motivation
- Facade protocol
- Evaluation

Performance Evaluation: Methods

- Evaluated entropy of a search request using AOL search corpus
- 20 million queries from 650k users
- Get the average information content (entropy) per query

Performance Evaluation: Results

Tool	Entropy Per Request (bits)	Request Rate to Equal 256kbps	Deniability Set
Facade	[1]~78	3,300	HTTP+Search
Infranet	[2]3	85,300	HTTP+Browse
StegoTorus	[3]12000	21	HTTP

[1] The paper contains an error wherein the entropy is reported with log base e instead of 2

[2] Infranet entropy computed with parameters from paper, i.e. 8 links per page, so $\log_2(8)=3$ bits

[3] StegoTorus entropy calculated based upon Base64 encoding 2000 characters per URL

Future Work: Tradeoffs

- Tune performance/deniability with dictionary choices
  - Per user/site dictionaries
  - Dictionaries with joint PDFs

Conclusion

- Facade: an HTTP covert channel that balances performance and deniability by improving upload performance

- Get Facade (in development) from https://github.com/ben-jones/facade

- Contact me: Ben Jones bjones99@gatech.edu