# Privacy and Anonymity Properties of HTTP/2, SPDY, QUIC, and TLS 1.3

Mike Perry

The Tor Project

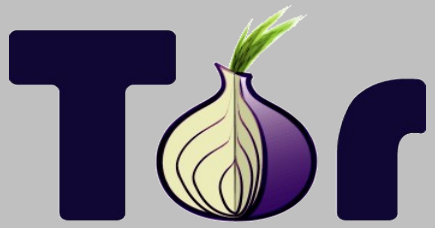

# Tor Basics

- TCP Overlay Network; Stream abstractions
  - TCP SOCKS Proxy
- ~3 million daily users
  - Not the same users every day!
  - ~1 million users update the browser within 1 week
  - ~5 million Android installs
- Tor Browser is a small team
  - 5 engineers total
  - Standards participation is difficult for us

# Abstract Privacy and Anonymity Issues

- Linkability sources
  - State management (supercookies/identifiers)
  - Browser fingerprinting
- Traffic integrity and confidentiality
- Traffic analysis
  - Traffic fingerprinting
  - Correlation
  - Confirmation
  - Route manipulation and analysis

# Terminology Normalization

- “Linkability”
  - The ability to associate one user action with another
  - Types: “PBM”; “CPD”; “Fingerprinting”; “3rd party”
- “Fingerprinting” != “Identifier storage”
  - Identifiers are content-accessible browser state (aka “supercookies”)
  - Fingerprinting is any stateless vector
- “First Party Isolation”
  - Bind all content-accessible browser state to the URL bar domain
  - AKA “Double-Keying”

# First Party Relationships

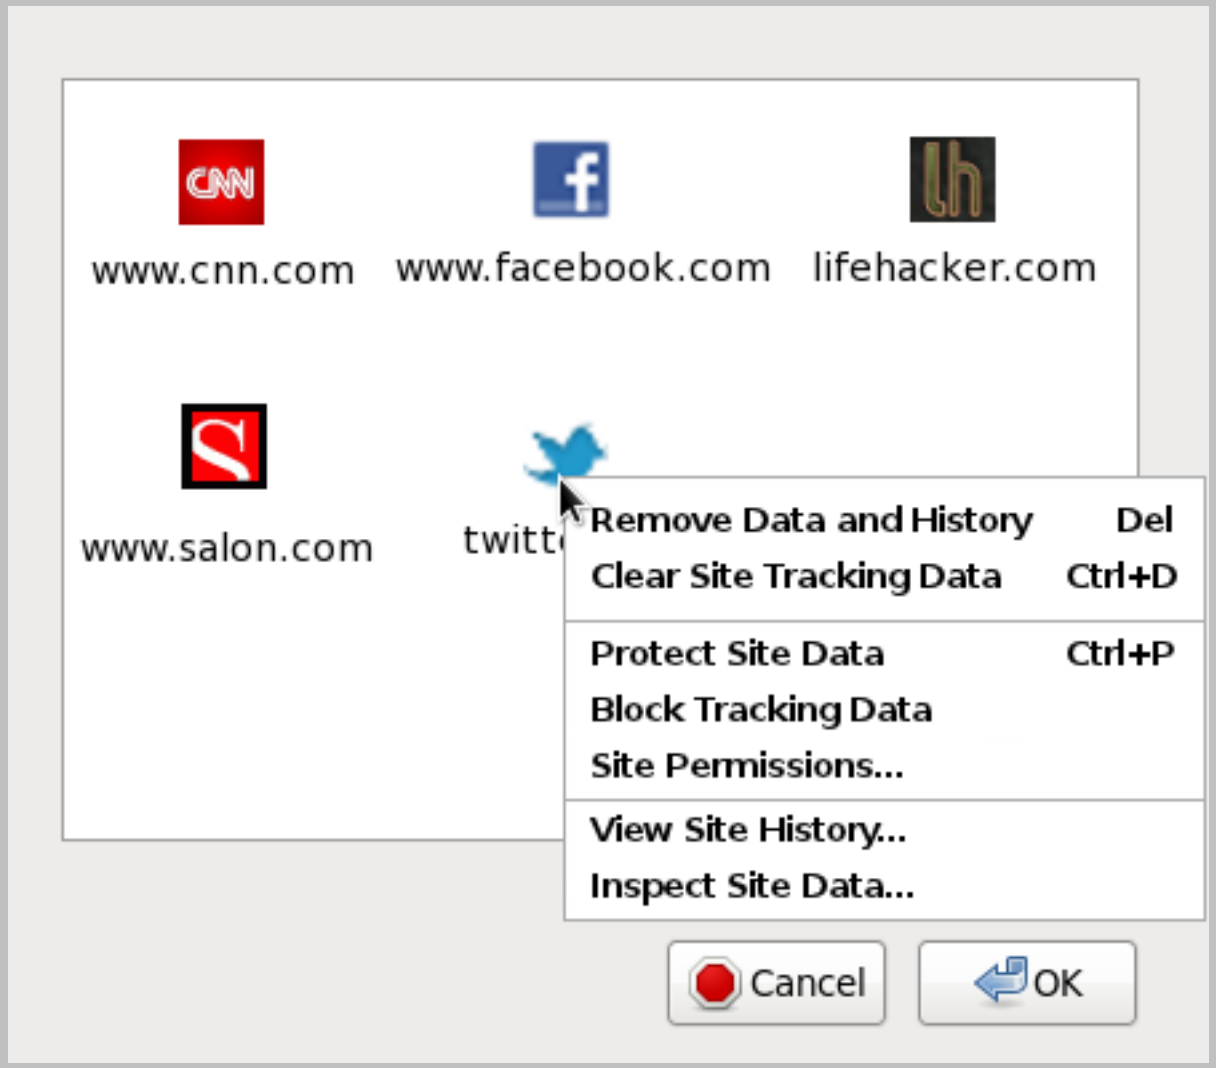

# Identifier Storage in HTTP/2 & SPDY

- Alternative-Services Header caching
- ALPN and NPN successes cached to govern initial connection counts
- Server PUSH response caching
- SETTINGS caching

# Identifier Storage in QUIC

- Superset of HTTP/2, plus:
  - 0-RTT state caching
  - IP spoofing protection opaque cookie
  - Discovery and Alternate-Protocol state
  - 64bit connection-id (for third parties)
  - Congestion window information?

# Traffic Analysis

- Confirmation, Correlation, Route Estimation
  - PING and SETTINGS are a concern
    - We may limit number of responses and introduce delays
  - RTT estimation in QUIC
- Website Traffic Fingerprinting
  - TLS: 'Side-Channel Leaks in Web Applications'
    - Padding ~256bytes mitigates many cases
  - Very sensitive to base rate: More pages → less accuracy and less padding
  - Tor's 512 byte cell size helps, as does multiplexing
  - Interestingly, pipelining may hurt

# Effects of the Base Rate Fallacy

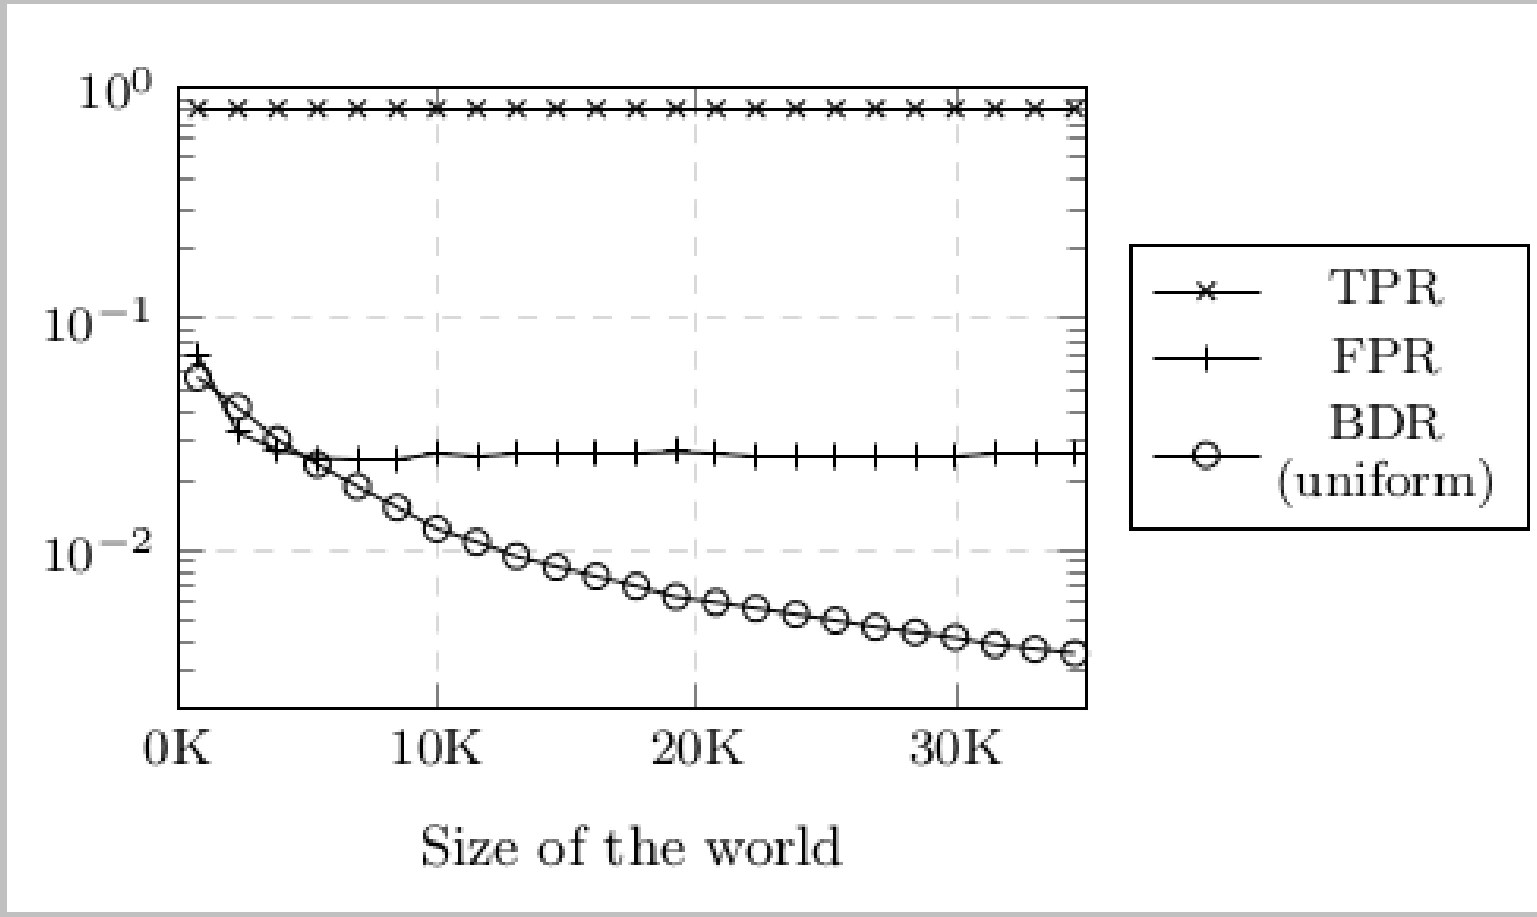

# TLS 1.3 Wishlist

- Optional Encrypted SNI/Extensions?
  - Helps with traffic fingerprinting and censorship
  - Some users may want to burn an RTT for this...
- Update the PSK ticket id on every resume?
  - Eliminate observer linkability capability
- Can 0-RTT also get a PSK ticket id?
  - Enables Perspectives-style multipath verification
  - Help servers guard against key theft/MITM
- Padding?

# Tor's View of Fingerprinting

- “Active” vs “passive” distinction insufficient
  - “Direct” vs “Inferred” might be better?
- Sources of fingerprinting in order of concern:
  - End-user configuration details
  - Device and hardware characteristics
  - Operating System vendor and version differences
  - User behavior
  - Browser vendor and version differences (ignored)
- Fingerprinting is dependent on user base size
  - Let's standardize private network usage!

# Fingerprinting

- QUIC
  - Timestamps in ACK, NONC
  - Local link property inference?
- HTTP/2 and SPDY
  - Couldn't find anything.. Did I miss anything?

# Other Discussion

- Most specifications typically tie state management to clearing all cookies
  - Can we also specify third party state management?
  - If a user disabled third party cookies, they probably don't want other third party state either..
- Third party connection and protocol discovery isolation
- Did I miss any identifier/fingerprinting vectors?

# Thanks

Mike Perry <mikeperry@torproject.org>

C963 C21D 6356 4E2B 10BB  335B 2984 6B3C 6836 86CC

https://www.torproject.org/projects/torbrowser/design/