

Master's Thesis: Censorship-resistant Collaboration with a Hybrid DTN/P2P Network

Philipp Hagemeister

Institut für Informatik
Heinrich-Heine-Universität Düsseldorf

29.3.2012

Threat Model

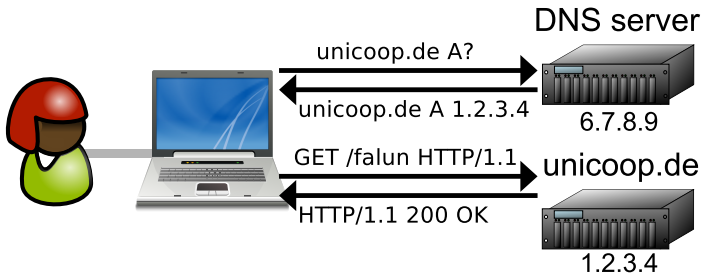
- Collaboration systems proliferate free speech
- Attacker does not want free speech
- ⇒ Attacker goal: Disrupt collaboration systems
- Attacker controls ISP and national infrastructure



Figure: Attacker (representation)


Requests in Current Collaboration Systems

- adocracy, echo, LiquidFeedback, UniCoop are web applications
- Request diagram:



DNS Censorship

- Attacker controls default DNS server
- Contemplated in Germany and US
- Used in Belgium, Denmark, Italy, Turkey, Burma, China, ...
- Easily circumvented (→ Alessandro Lenzen, 2011)
- Long-term solution: client-side DNSSec



IP Censorship

- Attacker can drop packets from or to specific IP addresses
- Used in China, Egypt, Libya, Pakistan, Thailand




Deep Packet Inspection

- Attacker filters packets for search terms
- Used in China, Iran
- Prevented by encryption



Physical Attacks

- Attacker physically seizes or takes over server
- Happened in Germany!
 - In 2011, servers of the Piratenpartei were confiscated
- Defense: Multiple servers



Peer-To-Peer (P2P) Networks

- Multiple servers alone are not sufficient
- Eliminate all single points of failure!
- We need a truly decentralized system
- ... a **Peer-to-peer (P2P)** network



Bootstrapping

How do we get the address of a peer?

Bootstrapping

How do we get the address of a peer?

- Hardcoded
- Human input

Bootstrapping

How do we get the address of a peer?

- Hardcoded
- Human input
- DNS
- HTTP(S)

Bootstrapping

How do we get the address of a peer?

- Hardcoded
- Human input
- DNS
- HTTP(S)
- IP multicast

Bootstrapping

How do we get the address of a peer?

- Hardcoded
- Human input
- DNS
- HTTP(S)
- IP multicast
- Email / SMS

Bootstrapping

How do we get the address of a peer?

- Hardcoded
- Human input
- DNS
- HTTP(S)
- IP multicast
- Email / SMS
- Decoy routing

Bootstrapping

How do we get the address of a peer?

- Hardcoded
- Human input
- DNS
- HTTP(S)
- IP multicast
- Email / SMS
- Decoy routing

Bootstrapping: Solvable

The number of bootstrapping schemes allow us to evade all but the most sophisticated censorship systems.

Other P2P considerations

- Structured vs unstructured
- Sybil and other active attacks
- Broadcasting
- NAT traversal


Other P2P considerations

- Structured vs unstructured
- Sybil and other active attacks
- Broadcasting
- NAT traversal
- Privacy
 - Solved by **anonymization networks**
 - Examples: I2P, Tor, Freenet
 - Need to be integrated
 - → Paul Baade

P2P: Conclusion

A P2P network can provide an adequate defense against censorship.

Back to the Threat Model



Total Internet Shutoff

- Attacker can turn off Internet access
- Happened in 2011 in Egypt and Libya
- Arguably permanently in Cuba and North Korea



- Transfer data with USB thumb drives
- Delay-Tolerant Networks (**DTNs**) do not require continuous connection
- Fields of use:
 - Interplanetary communication
 - Developing nations
 - Military/naval
 - Sneakernet in Cuba

DTNs allow communication even in the case of a Internet shutoff

Revision Control in DTNs

- Challenge in DTNs: Distributed consensus is not possible
- Nevertheless, we want want **revision control**
 - .. primarily for history, accountability, and change management

Revision Control in DTNs

- Challenge in DTNs: Distributed consensus is not possible
- Nevertheless, we want want **revision control**
 - .. primarily for history, accountability, and change management
- Graph-based revision control systems: git, mercurial, bazaar, PlatinVC
 - Need to be adapted for DTNs (→ Janine Haas, 2012)

Revision Control in DTNs

- Challenge in DTNs: Distributed consensus is not possible
- Nevertheless, we want want **revision control**
 - .. primarily for history, accountability, and change management
- Graph-based revision control systems: git, mercurial, bazaar, PlatinVC
 - Need to be adapted for DTNs (→ Janine Haas, 2012)
- Patch-based revision control systems: darcs
 - Complex, not yet practical
 - Could be the silver bullet

Revision Control in DTNs


- Challenge in DTNs: Distributed consensus is not possible
- Nevertheless, we want want **revision control**
 - .. primarily for history, accountability, and change management
- Graph-based revision control systems: git, mercurial, bazaar, PlatinVC
 - Need to be adapted for DTNs (→ Janine Haas, 2012)
- Patch-based revision control systems: darcs
 - Complex, not yet practical
 - Could be the silver bullet
- Document-oriented revision control: CouchDB, MongoDB
 - Simple, but weak guarantees

Revision Control in DTNs

- Challenge in DTNs: Distributed consensus is not possible
- Nevertheless, we want want **revision control**
 - .. primarily for history, accountability, and change management
- Graph-based revision control systems: git, mercurial, bazaar, PlatinVC
 - Need to be adapted for DTNs (→ Janine Haas, 2012)
- Patch-based revision control systems: darcs
 - Complex, not yet practical
 - Could be the silver bullet
- Document-oriented revision control: CouchDB, MongoDB
 - Simple, but weak guarantees
- Common base technology: **Content-Adressable Storage(CAS)**
 - Stores a set of bytes, accessed with hash(bytes).
 - No conflicts, $sync(CAS1, CAS2) = CAS1 \cup CAS2$
 - Can store (almost) all of the revision control system data

Architecture

- *Transports* abstract the specific communication channel
 - P2P over TCP
 - DTN over USB thumb drive
 - P2P over anonymization network
 - DTN over facebook
- Requirement: Application should be available on every device
⇒ web application



Web Application Fallback

- Best experience (DTN) if system is locally installed
- Public web servers for the masses
- If a web server becomes unavailable, switch to another one

Web Application Fallback

- Best experience (DTN) if system is locally installed
- Public web servers for the masses
- If a web server becomes unavailable, switch to another one
- Alternative: Continue working offline (→ Tim van Cleef)
- Future: Whole application in the browser

Conclusion

- Censorship resistance is important for collaboration software
- Censorship-resistant P2P network
- In case of total shutoff: DTN
- Future reasearch and implementation required

Questions?

Questions?

This presentation: <http://phihag.de/2012/mtpres.pdf>

Thesis: <http://phihag.de/2012/mt.pdf>

Source code: <http://phihag.de/2012/d2p/>

Demo



Warning: Experimental Prototype!

Future Work

- General code quality, documentation, and testing
- Automated unit and functional tests
- Simulation framework
- P2P bootstrap implementation and analysis
- NAT traversal for the P2P transport
- Structured P2P implementation with efficient broadcast
- Integration into DTN standards (RFC 4838 ...)
- Research into partial replication
- Robust thumb drive storage formats
- Steganography and cryptography
- Ports to other platforms, in particular android, *BSD, iOS, Mac OS X, WebOS, Windows, Windows Phone

Future Work (continued)

- Project search functionality
- User Management
- Extend functionality of the main policy drafting application
 - A WYSIWYG editor
 - Comments to specific lines or paragraphs (→ Julius Römmler)
 - Better usability
- Demonstrate and develop a client-side application
- Prototype browser-to-browser P2P with WebRTC
- Create a decentralized security framework
- Allow closed groups as well as read-only ones
- Allow voting applications
- Extend revision control
 - Integrate graph- and/or patch-based revision control systems
 - Improve the CAS performance
- Integration with other platforms (such as adhocracy)
- Integration with PKIs such as German ID card

- Problem: Where do we store keys
- Browser integration problematic (→ Evgeni Golov, 2012)
- Option: private key = hash(password)

- Problem: Where do we store keys
- Browser integration problematic (→ Evgeni Golov, 2012)
- Option: private key = hash(password)
- √ project:
 - Project ID = hash(project public key, security specification)
 - Allow private projects by encrypting everything with a symmetric key
 - Symetric key is stored alongside project data, encrypted with users' public keys

- Problem: Where do we store keys
- Browser integration problematic (→ Evgeni Golov, 2012)
- Option: private key = hash(password)
- √ project:
 - Project ID = hash(project public key, security specification)
 - Allow private projects by encrypting everything with a symmetric key
 - Symetric key is stored alongside project data, encrypted with users' public keys
 - Allow read-only projects by requiring changes to be signed by a key ...
 - ... which in turn is signed by the project's key

- Distributed verifiable anonymous voting is not possible!
- Requires trusted intermediaries
- Or trusted voting registrars

Extended Threat Model

- Assumption so far: *User can run arbitrary software on her device.*
- Assumption: User *has access to* a device
- Assumption: User controls (general-purpose) device.
 - May be restricted with UEFI Secure Boot
 - Signed firmware required on Apple iPad, iPhone, iPod
 - Signed firmware required on some android devices
- Attacker may also physically go after users
 - ⇒ Anonymity/Pseudonymity required

Extended Threat Model

- Assumption so far: *User can run arbitrary software on her device.*
- Assumption: User *has access to* a device
- Assumption: User controls (general-purpose) device.
 - May be restricted with UEFI Secure Boot
 - Signed firmware required on Apple iPad, iPhone, iPod
 - Signed firmware required on some android devices
- Attacker may also physically go after users
 - ⇒ Anonymity/Pseudonymity required
- Attacker can use **malware** to gain control of the device
 - Happened in Germany: Staatstrojaner
 - Blackberry malware in UAE

DPI in China

- Chinese network-level DPI searches for keywords like *falun gong*
- Injects an RST packet
- Blocks all packets between the peers for a couple of minutes


192.168.1.13	192.168.1.1	DNS	70 Standard query A pku.edu.cn
192.168.1.13	192.168.1.1	DNS	70 Standard query AAAA pku.edu.cn
192.168.1.1	192.168.1.13	DNS	120 Standard query response
192.168.1.1	192.168.1.13	DNS	102 Standard query response A 162.105.129.21 A 162.10
192.168.1.13	162.105.129.21	TCP	74 56558 > http [SYN] Seq=0 Win=14600 Len=0 MSS=1460
162.105.129.21	192.168.1.13	TCP	58 http > 56558 [SYN, ACK] Seq=0 Ack=1 Win=3840 Len=
192.168.1.13	162.105.129.21	TCP	54 56558 > http [ACK] Seq=1 Ack=1 Win=14600 Len=0
192.168.1.13	162.105.129.21	HTTP	128 HEAD /falux_gXng HTTP/1.1
162.105.129.21	192.168.1.13	TCP	54 http > 56558 [ACK] Seq=1 Ack=75 Win=5840 Len=0
162.105.129.21	192.168.1.13	TCP	259 [TCP segment of a reassembled PDU]
192.168.1.13	162.105.129.21	TCP	54 56558 > http [ACK] Seq=75 Ack=206 Win=15544 Len=0
192.168.1.13	162.105.129.21	HTTP	128 HEAD /falun_gong HTTP/1.1
162.105.129.21	192.168.1.13	TCP	54 http > 56558 [RST, ACK] Seq=206 Ack=149 Win=1923

P2P: Structured vs Unstructured

- Structured networks are stable
- But may be easier to disrupt!

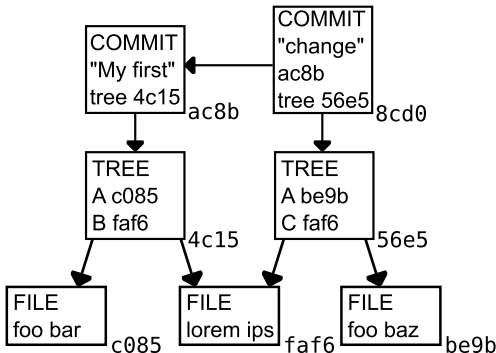
P2P: Structured vs Unstructured

- Structured networks are stable
- But may be easier to disrupt!
- Broadcasting much more efficient in structured networks




Graph-based Revision Control Systems

- Every file, tree, commit is mapped to a block of content
- Block is stored in a CAS
- Accessible only by `hash(block)`




Terms in Revision Control Systems

- Every change is recorded in a *commit*
- Commits form a *DAG*:



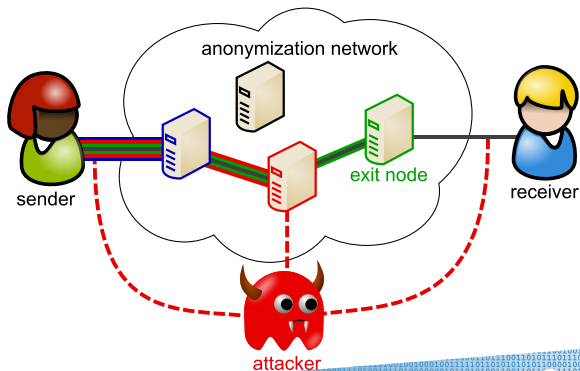
Problems in Graph-based Revision Control Systems

- Assumption: Always one common HEAD
- Problem: **D**elays mean that automatic merging can go on forever



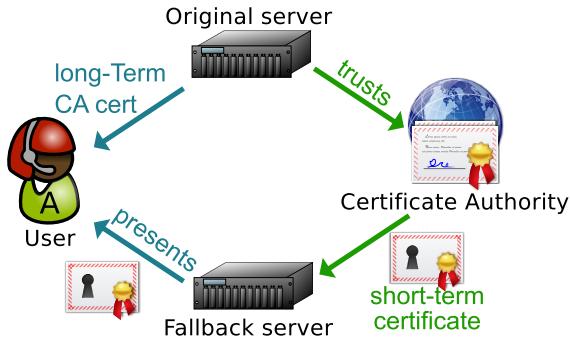
Anonymization Frameworks

- Use a user-chosen combination of *mixes*
- Tor (bidirectional, TCP-like)
- I2P (unidirectional, UDP-like)
- GnuNet (only storage)



Web Fallback Verification

- Problem: What if attacker compromises a server?
- Solution: Short-term certificates
- CA(might be blocked) does never interact with user



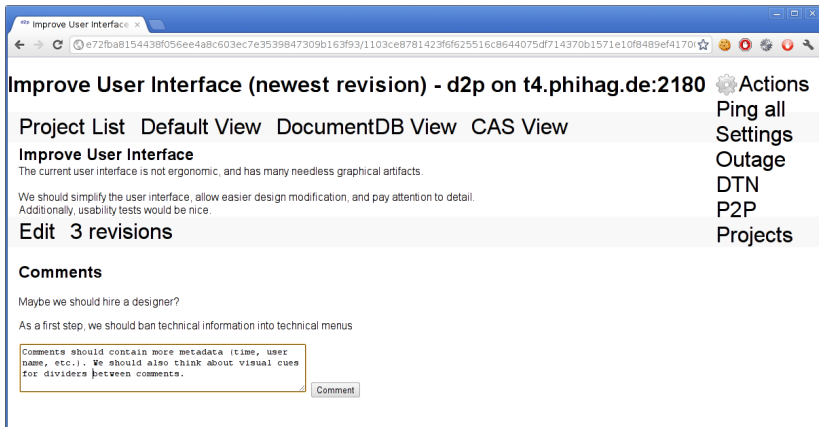
Implementation Considerations

- Code (especially views) must be **portable**
 - Required for offline version (→ Tim van Cleef)
 - We may also want to reimplement/compiler the application for the browser
 - Mustache: Logic-less web templates

Implementation Considerations

- Code (especially views) must be **portable**
 - Required for offline version (→ Tim van Cleef)
 - We may also want to reimplement/compiler the application for the browser
 - Mustache: Logic-less web templates
- Python **3** for clean code (Why not 2? bytes vs `string`)
- **Tornado** as asynchronous framework
- Modern web technologies (WebSocket, WebRTC, HTML5 semantic elements)
- Automated tests, simulation

Screenshots (1)



Improve User Interface (newest revision) - d2p on t4.phihag.de:2180

Project List Default View DocumentDB View CAS View

Improve User Interface

The current user interface is not ergonomic, and has many needless graphical artifacts.

We should simplify the user interface, allow easier design modification, and pay attention to detail. Additionally, usability tests would be nice.

Edit 3 revisions

Actions
Ping all
Settings
Outage
DTN
P2P
Projects

Comments

Maybe we should hire a designer?

As a first step, we should ban technical information into technical menus

Comments should contain more metadata (time, user name, etc.). We should also think about visual cues for dividers between comments.

Comment

Screenshots (2)

Transcend (/dev/sdb1) - DTN endpoint - d2p on t4.phihag.de:2180

Back to Transport Overview

Transcend (/dev/sdb1) Disable
JetFlash_Transcend_8GB_GCRS9U8Z.0.0

Projects

- . d Import
- . a
- . b Import
- . c Import

Project X Add Project

Actions
Ping all
Settings
Outage
DTN
P2P
Projects