

Linda Lee*, David Fifield, Nathan Malkin, Ganesh Iyer, Serge Egelman, and David Wagner

A Usability Evaluation of Tor Launcher

Abstract: Although Tor has state-of-the-art anti-censorship measures, users in heavily censored environments will not be able to connect to Tor if they cannot configure their connections. We perform the first usability evaluation of Tor Launcher, the graphical user interface (GUI) that Tor Browser uses to configure connections to Tor. Our study shows that 79% (363 of 458) of user attempts to connect to Tor in simulated censored environments failed. We found that users were often frustrated during the process and tried options at random. In this paper, we measure potential usability issues, discuss design constraints unique to Tor, and provide recommendations based on what we learned to help more users connect to Tor while reducing the time they take to do so. Tor Browser incorporated the changes proposed by this study.

Keywords: Usable Security, User Studies, Tor, Security, Censorship, Anonymity

DOI 10.1515/popets-2017-0030

Received 2016-11-30; revised 2017-03-15; accepted 2017-03-16.

1 Introduction

Tor [13] is an anonymity network that routes Internet traffic through a series of relays to make it difficult to observe sources and destinations. Tor Browser [35] is a modified Firefox browser that routes traffic through the Tor network. Although Tor and Tor Browser were originally designed for anonymity, many now use them to circumvent Internet censorship. Because of this, many

countries block Tor relays specifically to prevent their citizens from circumventing censorship [44].

Tor Browser is bundled with a graphical configuration interface, Tor Launcher, which allows users to bypass censorship by configuring various bridges, proxies, and transport protocols prior to connecting to Tor. Of course, these options are only effective if users understand how to use them. To understand to what extent this is the case, we researched the effectiveness of Tor Launcher as well as how it could be improved.

Improving Tor Launcher has the obvious benefit of helping censored users connect to Tor. It also benefits other Tor users since an increase in the overall number of users increases the anonymity set [12]. Tor does not collect user interaction data from its real users, so we believe that usability testing research is critical to help users in heavily censored regimes and increase user adoption. In this paper, we:

- Perform a qualitative and quantitative usability evaluation of Tor Launcher 5.0.3 under simulated censorship environments
- Contribute design changes that improve users' likelihood of connecting to Tor and also reduce the amount of time users take to connect
- Provide recommendations beyond design changes to improve Tor Launcher, such as leveraging user input (requires research) and building an automated connection scheme (requires Tor infrastructure work).

Specifically, we describe the initial evaluation of the interface (Section 4) and our observations of participants' interactions with it (Section 5). Based on the insights from our observations, we proposed several changes to the interface (Section 6) and measured how successfully and quickly users connect to Tor before and after changes (Section 7). Finally, we provide recommendations based on what we learned (Section 8), we address limitations of our evaluation (Section 9), frame our study in the context of existing work (Section 10), and conclude (Section 11). We hope that our usability evaluation serves as an informative case study and provides guidance on further improving Tor Launcher, as well as other privacy-enhancing technologies.

***Corresponding Author: Linda Lee:** University of California, Berkeley, E-mail: lnl@cs.berkeley.edu

David Fifield: University of California, Berkeley, E-mail: fifield@cs.berkeley.edu

Nathan Malkin: University of California, Berkeley, E-mail: nmalkin@cs.berkeley.edu

Ganesh Iyer: University of California, Berkeley, E-mail: ganesh.v@berkeley.edu

Serge Egelman: University of California, Berkeley and International Computer Science Institute, E-mail: egelman@cs.berkeley.edu

David Wagner: University of California, Berkeley, E-mail: daw@cs.berkeley.edu

2 Technical Background

This section discusses the network components involved in connecting to Tor. An understanding of the various connection options are not required to understand our work, but we refer to these terms and methods throughout the paper.

2.1 Proxies, Relays, Bridges, and Pluggable Transports

A connection to Tor involves many interacting components, illustrated in Figure 1. In the presence of a local firewall and censor, it may involve a proxy and a bridge.

The word “proxy” has a specific meaning in the context of a Tor connection. Generally, a proxy is any computer or service that allows users to make indirect connections to other services. In this sense, any Tor relay and the Tor circuit can also be considered a proxy. However, in this paper, the word “proxy” refers to a SOCKS or HTTP proxy that bypasses local firewalls.

Tor relays are the routers in the Tor network. The Tor network is a group of routers that perform onion routing, a technique for anonymous communication over a computer network. Guard relays receive packets from the user and forward traffic to middle relays. Middle relays forward traffic from an entry relay to an exit relay. Exit relays receive traffic from a middle relay and direct the traffic to its intended destination (e.g., a website). A user uses Tor Browser to connect to the Tor network, which connects to websites on behalf of the user.

Bridges are unlisted alternative entry nodes used to connect to Tor when Tor relays are censored. Bridges also receive packets from the user and forward traffic to middle relays, but may also run a pluggable transport (referred to as “transport” for shorthand). Pluggable transports specify alternative protocols by which Tor Browser may connect to the Tor network, thereby disguising Tor traffic to make it harder to detect and block bridges. When a bridge is used, it takes the place of the guard relay. For convenience, bridges are grouped in the interface by which transport they run.

Tor Launcher is the graphical user interface for configuring an optional proxy, bridge, and transport prior to initiating a connection to the Tor network. Figure 2 shows the various screens of Tor Launcher at the time we began our research.

Most bridges can be blocked by blocking their static IP addresses. These include bridges that run the `fte`

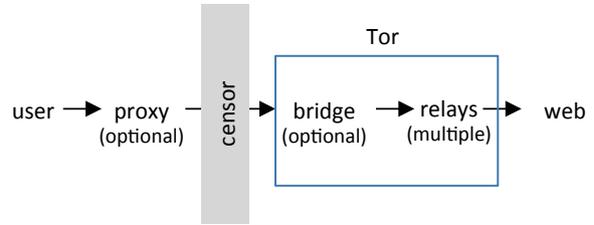


Fig. 1. The chain of components involved in connecting to a website over Tor. Most users do not need a proxy; only users who face a censor need a bridge. In the diagram, “Tor” represents a circuit of hops through the Tor network. We have shown the bridge as a separate component because of the special role it plays. When a bridge is used, it becomes the first Tor relay.

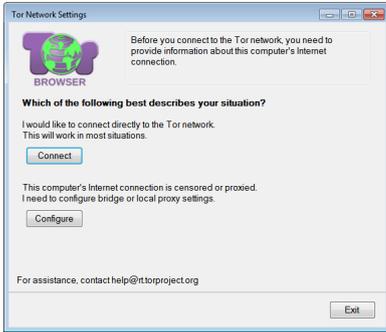
and `fte-ipv6` [14] pluggable transports, which disguise the Tor protocol as another protocol (such as HTTP); these also include the `obfs3` [22], `obfs4` [5], and `scramblesuit` [45] transports, which encrypt or alter the Tor protocol to appear as random noise.

Other bridges are less susceptible to blocking because of how they work. Bridges that run the `flashproxy` [18] and `meeq` [19] transports route traffic through third party web browsers and content delivery networks. The bridges and transports that are accessible in each country vary. Figure 2d shows the bridge and transport options at the time of our study.

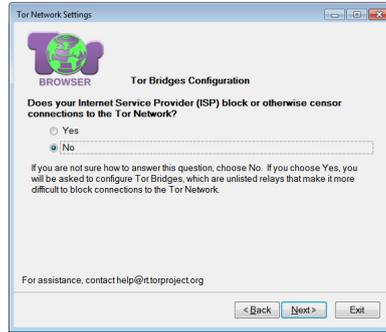
2.2 Connecting to Tor

It is important to note that there are many valid configuration settings that can be used to connect to the Tor network. A user may or may not need a proxy and/or a bridge. A user who does not need a bridge or proxy can still connect with a bridge, proxy, or both, provided that these are configured correctly (though adding unnecessary network hops or obfuscation protocols may slow the connection down).

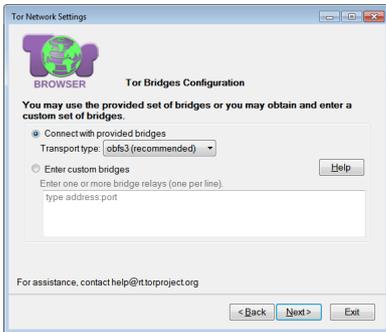
Users can configure a bridge by choosing a transport, which selects the group of bridges that use that transport. For instance, a user choosing the `obfs3` option in Figure 2d will be assigned one of these bridges as their guard relay. If these built-in options do not work, users can obtain bridge configurations through out-of-band channels [8]. Configuring a proxy requires entering in the proxy protocol, IP address, port, and additional optional fields (Figure 2f). The user must obtain this information themselves.



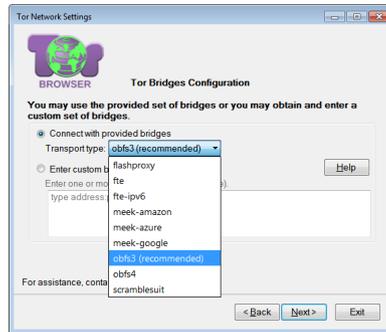
(a) The first screen (F). “Connect” starts a connection. “Configure” leads to the first bridge screen (Figure 2b).



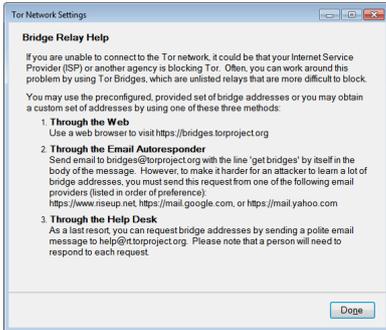
(b) The first bridge screen (B1). “Yes” leads to a bridge configuration screen (Figure 2c) and “no” skips to P1.



(c) The second bridge screen (B2). Users select a bridge here. “Help” directs to bridge help screen (Figure 2e).



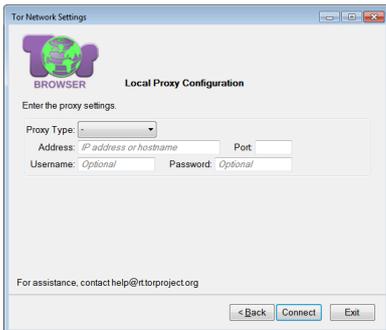
(d) Some hardcoded bridges are listed alphabetically by their transport. The recommended one is obfs3.



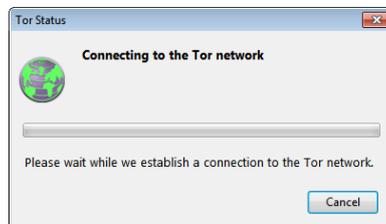
(e) The bridge help screen (BH) describes options for obtaining a custom bridge. “Done” leads back to the previous screen.



(f) The proxy question screen (P1). “Yes” goes to the proxy settings screen (Figure 2g). and “no” starts a connection.



(g) The proxy settings screen (P2). “Connect” starts a connection to Tor with the interface settings.



(h) The progress bar (Pr). The bar fills up as log messages update the connection status.

Fig. 2. The Tor Browser 5.0.3 Tor Launcher GUI. Screens are in the order that they appear.

3 Methodology

This section discusses our IRB-approved usability testing process, including the simulated censorship environments we used to require users to configure different network components to connect to Tor and the evaluation metrics we collected.

3.1 Testing Pipeline

We began by performing a cognitive walkthrough [40, 42] to inspect the interface to anticipate potential problems and formally assess its pros and cons. A cognitive walkthrough is a formal method of inspection performed by researchers without any user involvement. We then conducted qualitative and quantitative usability tests.

Qualitative testing is used to gain an understanding of underlying reasons, opinions, and motivations. We use qualitative testing to develop hypotheses about why people struggle to use the interface and develop ideas for potential design modifications to the interface.

Quantitative testing is used to quantify the problem by way of generating numerical data or data that can be transformed into usable statistics. We use quantitative testing to quantify how people interact with the interface to uncover user patterns.

We believe both methods were necessary for a complete usability evaluation. Qualitative testing illuminates *why or how* the software is or is not usable, while quantitative testing measures *to what extent*. We used our qualitative study observations to design interface changes that would address encountered issues. We used quantitative testing to measure the impact of our proposed changes, as compared to the original interface.

We performed our qualitative experiments one participant at a time, so that we could observe their behaviors in real time and interview them about their experience after use. We recruited five participants per experimental condition to gather qualitative data; this has been suggested as the optimal number of measurements to maximize the cost-benefit ratio [31]. We performed our quantitative experiments at a behavioral laboratory many participants at a time, using instrumentation to capture evaluation metrics. We recruited twenty users per experimental condition to gather quantitative data, based on estimates of the number needed for statistical significance [31].

	E1	E2	E3
websites blocked	X	X	X
public relays blocked		X	X
default bridges blocked			X

Table 1. Summary of our simulated censorship environments. E1 only requires users to connect directly; E2 requires users to connect with a built-in bridge; and E3 requires users to connect with a specific type of built-in bridge or a custom bridge.

3.2 Experimental Setup

We chose to simulate environments for the stability and reproducibility of the experiment, as real censored networks are volatile and complex. Re-routing participant traffic through other countries’ networks also may introduce unforeseen risks to the participants out of our control and performing hundreds of connections in short intervals during experiments may draw unnecessary attention to Tor use in that country. The three simulated censorship environments, which we refer to as E1, E2, and E3 throughout the paper, are summarized in Table 1. E2’s censorship is a superset of E1’s; similarly, E3’s censorship is a superset of E2’s.

In the first environment, E1, particular websites were blocked, which meant that they could be accessed using a default Tor configuration (i.e., bridges or proxies were not needed). In the second environment, E2, we additionally blocked the IP addresses of public Tor relays, which meant that Tor could still be accessed by enabling any of the included bridges. Finally, the third environment, E3, additionally blocked several of the included bridges, which required experimenting with several different bridges—beyond the recommended option—in order to successfully connect.

The simulated environments are not intended to imitate any particular country’s network—though they do use the same techniques that are currently used by various censorship regimes. Rather, we designed simulated environments to have different requirements to connect to Tor. We believe this to be sufficient for the purpose of testing the configuration interface, since this requires the user to take the different interface paths we wanted to test (and matches the actions real users would need to take to circumvent various types of censorship).

We used Tor Browser 5.0.3, the most recent stable release at the time [23]. There were new releases during the experiments, but we used the same version throughout to not introduce confounding factors. We conducted our experiments on Windows machines. To create the environments, we used Windows Firewall rules to block

the IP addresses of public Tor relays and default bridges, and the Windows hosts file to block websites by mapping domain names to localhost.

3.3 Evaluation Metrics

We use the following three industry-standard [3] usability metrics to quantify how easily users completed their task and if they were productive with their time:

- **Completion rate:** the percentage of users that want to connect to Tor and successfully do so. The higher the completion rate, the better. The ideal case would be that everyone is able to connect to Tor if they want to. Even if users do not successfully connect to Tor at first but eventually do so during the experiment, they are considered to have successfully connected to Tor for this metric.
- **Connection time:** the amount of time users take to connect to Tor. The less time it takes for users to connect, the better. The ideal case would be that the connection time is very close to the time it takes for the application to communicate with the servers and establish a circuit of relays for the connection, which can take up to several minutes. Connection time includes this bootstrapping time and any additional time that people spend in the interface.
- **Configuration time:** the amount of time users spend in the interface. The less time it takes for users to configure their preferences, the better. Configuration time is included in connection time, so it will always be less than connection time. We defined configuration time as any time spent on the bridge and proxy screens, since users were most likely configuring bridges and proxies. We excluded any time spent on the progress screen, because users were most likely waiting for connections to establish or time out and not necessarily interacting with the interface. This should be as short as possible.

Completion rate is the most important usability metric, measuring if users can accomplish the task. We defined success as a binary metric that is true if a user connected to Tor and false if a user did not. Task time is a supplemental usability metric that measures user productivity during the task. We use both connection and configuration time to quantify productivity since the bootstrapping time, and hence, time spent waiting for the connection to succeed or fail, can dominate the measurement.

4 Cognitive Walkthrough

We inspected the interface to anticipate potential problems before testing it on study participants.

4.1 Procedure

Two researchers performed the cognitive walkthrough [40] on the Tor Browser 5.0.3 Tor Launcher GUI, which was the most recent version deployed at the time (Figure 2). We first examined each screen in the interface for tasks required, inputs taken, and consequences of possible actions. This was done to map all potential paths through the interface (Figure 3). We then stepped through each user path, systematically tried all possible actions, and recorded our observations.

4.2 Observations

We noticed that the interface leaves a lot of freedom for error while not giving good instructions on how to recover from errors:

- **Technical terms:** The interface uses technical and Tor-specific terms that are likely unfamiliar to an average user (e.g., “Tor network,” “local proxy,” “bridges,” and “transports”).
- **Too many choices:** Users who do not need a bridge or proxy are able to configure them. There are nine choices for bridges and more options to input a custom bridge. Almost any HTTP or SOCKS proxy can be used as the proxy. Custom bridge lines require the bridge information in a specific syntax, but this syntax is not specified in the interface. Proxy fields also do not provide syntactic guidelines. The full list of possible errors we identified can be found in Appendix A.
- **Poor feedback:** Even if there was a syntactic error in the bridge or proxy fields, the user is not notified until after a connection is attempted and fails. If a censored bridge was chosen, the connection timeout takes several minutes. Error messages were often technical and did not suggest how to fix the problem.

These are some of the factors that we believed contribute to users struggling with the interface and therefore may prevent them from connecting to Tor quickly.

5.3 Results

Here we report the main reasons why participants struggled to use Tor Launcher, summarizing general behavioral trends and attitudes, and providing representative quotes from individual subjects.

Participants spent a lot of time going back and forth on screens, looking for proxy information, and waiting for the connection to work or time out. We believe that the reasons for these behaviors were:

- **Unfamiliar concepts:** Participants found technical terms (e.g., relay, bridge, proxy, censor, ISP) confusing, and were often unable to answer the questions that were designed to guide them (Figures 2b and 2f).

P9: *“You can only answer this question correctly if you know what a local proxy is. Otherwise you have no chance...”*

P7: *“[I] was just fiddling, to be honest with you... it was kind of confusing.”*

- **Missing information:** They did not know how to choose a bridge/transport or which bridges worked in their environment (Figures 2d).

P10: *“I have no clue what’s the difference between flash-proxy, fte, etc. And why do I need a custom bridge if there are options built in?”*

- **Bad recovery:** Participants had to wait up to several minutes to know if their connection failed, and were sometimes upset that they had to wait that long to find out that their connection failed.

P13: *“Put in a 30 second timeout! I don’t know, maybe it takes longer than that...”*

P15: *“there doesn’t seem to be any timeout on any of this stuff. Am I waiting long enough? Should I wait 5 minutes?”*

Participants who could connect directly (E1) or use the recommended bridge had an easy time connecting (E2), but the participants who could not use many of the hardcoded bridges (E3) did not know what to do. For a summary of how each participant connected to Tor and what behavioral trends dominated in each environment, see Appendix G.

6 Interface Redesign

We redesigned the interface to address pain points found in our qualitative study. Figure 4 shows our final design.

6.1 Design Considerations

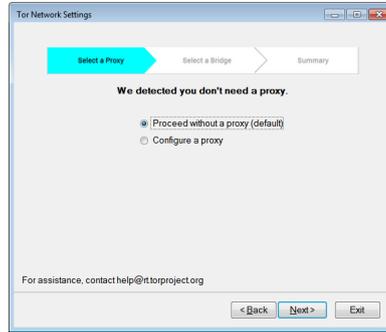
Designing an interface for Tor has unique challenges:

- **User consent:** Some oppressive regimes may detect and prosecute Tor usage, and therefore simply using Tor may carry risks. While we could help users by automating configuration settings, this would involve probing network configurations and sending network traffic that could allow authorities to detect that the user is using Tor.
- **Passive adversaries:** A sufficiently powerful network adversary can detect that a user is connecting to Tor, unless the user chooses some very specific network configurations. We should enable informed users who wish to avoid this risk to do so.
- **Maintenance constraints:** Network environments change all the time. Keeping up with how each country censors and pushing changes before the affected users connect to Tor is difficult.
- **Financial constraints:** The meek bridges are especially effective at avoiding censorship: they are harder to block without imposing significant collateral damage (e.g., to Amazon Web Services) and are harder to detect through passive monitoring [19]. Leveraging these harder-to-block bridges can help users in heavily censored environments, but these relays cost the Tor Project magnitudes more than other relays and using them more would cost even more money.

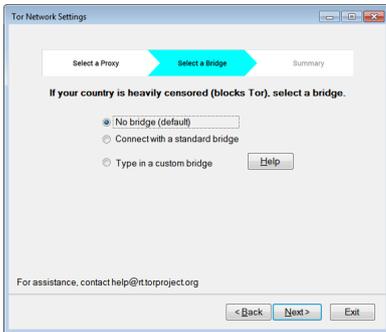
We kept the user involved in the configuration so that they explicitly start a connection to Tor (currently, the user needs to click a “connect” button before Tor Launcher makes a connection to Tor). And although some users in our qualitative study tried to configure unnecessary components in confusion, we still give users the freedom to configure optional components. We limited our design changes to ones that do not require infrastructural changes or cause additional financial burden to make it easier for the Tor project to adopt them.



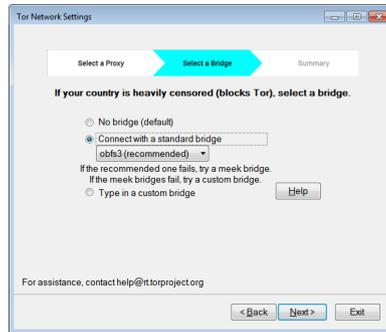
(a) The first screen (F) with reduced, alternative text. We said configuration was required for heavy censorship.



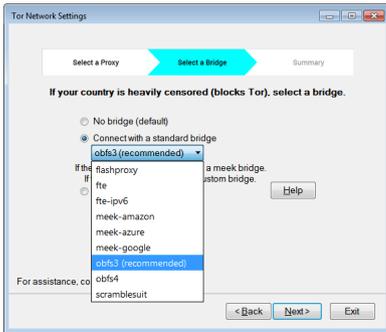
(b) The proxy screen (P). The interface locally detects a proxy and automatically populates the fields as necessary.



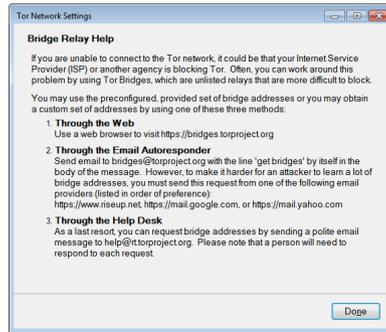
(c) The bridge screen (B). We give users the options to connect with or without a bridge all on one screen, and eliminated the bridge question screen.



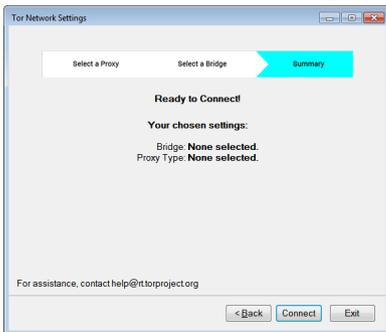
(d) The bridge menu and additional advice shows up when the corresponding option is chosen. The custom bridge input fields are hidden until the custom option is chosen.



(e) This dropdown menu remains the same (see Figure 2d).



(f) The bridge help screen (BH) remains the same as well.



(g) The summary screen (S) summarizes the configuration. This example configuration is connecting directly to Tor.



(h) The progress bar (Pr) shows the status of each component in the connection as the connection is made.

Fig. 4. A redesigned version of the Tor Browser 5.0.3 Tor Launcher GUI. Screens are in the order that they appear.

6.2 Changes Made

In this section, we highlight and explain the specific design changes we made to address the observed pain points while adhering to the design considerations. We:

- reduced the amount of text and technical terms in the text (Figures 4a, 4b, 4c, 4d, 4e, and 4h) to encourage users to read instructions.
- centered the text to make the design compatible for left to right and right to left languages.
- clarified that configuration options were manual and for heavily censored environments (Figure 4a) to discourage users who did not need bridges and proxies from configuring them.
- added a status bar on the proxy and bridge screens (Figure 4b, 4c, 4d, and 4e) to hint at future tasks.
- reordered proxy and bridge screens (Figures 4b and 4c) in a topologically sequential order. This hints that proper proxy configuration is necessary to reach bridges.
- eliminated the bridge and proxy questions (Figure 2b and 2f) to reduce the number of screens.
- added auto-detection for proxies (Figure 4b) in place of the proxy question. In theory, this can be done locally without infrastructural changes or leaking information to the network.
- added the option to not use a bridge on the bridge screen (Figures 4c, 4d and 4e) to replace the function of the bridge question screen (Figure 2b).
- added a recommendation to use one of the three meek-transport bridges (Figure 4d) if the recommended obfs3-transport bridges were unreachable, which happens in heavily censored real environments and our simulated environment E3.
- added a summary screen (Figure 4g) that summarizes the settings before making a connection to Tor, since some participants from the previous study didn't know how they connected to Tor.
- added indicators on the progress screen (Figure 4h) that visualized what components are trying to be reached, and in the case of failure, showed which component was unreachable. This was to help users understand what was going on during connection and to help troubleshoot their configuration.

Our interface was intended as a proof of concept, so we could evaluate these changes in a laboratory testing. We acknowledge that further refinement (aligning text properly, using higher quality visual assets, etc.) would be needed before deployment.

7 Quantitative Testing

We conducted our quantitative experiments at a behavioral laboratory with many participants in parallel, using instrumentation to capture evaluation metrics. These experiments confirmed previously-seen pain points and validated our proposed design improvements.

7.1 Recruitment

We recruited 59 participants from the Xlab mailing list, which is a university-wide pool of participants who have opted in to participating in social science experiments (see Appendix H for recruitment materials). We recruited an additional 65 participants from Craigslist to ensure diversity, since Xlab participants are primarily students and staff from UC Berkeley.

Of our 124 participants, we filtered out 10 participants who did not complete the consent form correctly or who downloaded another version of Tor Browser to use during the experiment. One participant did not report their demographics, but we still used their data. Of the remaining 113, 44.2% identified as male, 54.0% as female, and 1.8% as neither. They were between 18 and 68 years old ($\mu = 28.2$, $\sigma = 12.3$) and highly educated (84.1% had at least a college education).

7.2 Procedure

We ran our experiment at Xlab [46], an experimental social science laboratory. Xlab had 36 workstations with identical Windows 7 laptops, with workstation partitions to discourage participants from looking at others' screens. Before each session, we used a script to revert the computers to a clean image; download Firefox, Chrome, Internet Explorer, VLC media player, and Tor Browser 5.0.3 (instrumented versions that logged user interactions, with or without our proposed UI changes); set up one of the three simulated environments (E1, E2, or E3); and start recording the computer screen. We assigned each participant to one of six conditions (Table 2), which determines their user interface (OLD or NEW) and censorship environment (E1, E2, or E3). We use OLD to refer to the original interface (Figure 2) and NEW to refer the version with our changes (Figure 4).

We received informed consent after describing the study (Appendix I). We instructed participants to use Tor to complete their worksheet (Appendix J) and gave them 40 minutes to complete the task. Researchers

	completion rate (at 40min 8s)		connection time (med)	configuration time (med)
E1-NEW	19/19	100%	0:20	0:06
E1-OLD	19/19	100%	1:01	0:24
E2-NEW	18/18	100%	3:22	0:40
E2-OLD	16/19	84%	5:00	2:04
E3-NEW	13/19	68%	20:25	1:56
E3-OLD	10/20	50%	40:08	9:09

Table 2. An overview of the experiment. Those who did not connect were assigned a time of 40:08, (the maximum time).

maintained minimal interaction with the participants. After the experiment and exit survey (Appendix K), we gave each participant \$30 for their time.

7.3 Results

As seen in Table 2, our design changes helped more participants connect to Tor in less time. However, many participants, including participants who used the new interface, still could not connect to Tor: 63% (72 of 114) of first attempts failed and 79% (363 of 458) of total attempts failed. Those who managed to connect often took a long time to do so (>3 minutes in E2 and >20 minutes in E3). Figure 5 summarizes each participant, showing whether they successfully connected to Tor, how long they took to do so, and where they spent their time. Based on these results, we recommend that our changes be adapted in the meantime, but encourage a drastic redesign of the interface.

7.3.1 Completion Rate

Our changes significantly increased success rates (one-tailed Pearson’s chi-square test; $\chi^2 = 2.808$, $df = 1$, $p < 0.047$, albeit with small effect size $\phi = 0.157$).

Table 3 shows participants’ connections to Tor. As reported in the first two columns, most participants who did not need to configure anything (E1) didn’t, although some optionally used a bridge. Participants who could have used any bridge (E2) all used the recommended bridge. Participants who needed a meek bridge or custom bridge (E3) chose a meek bridge at random; the ones who tried to use a custom bridge did not format their inputs correctly and failed to connect to Tor.

Seventeen percent (19 of 114) were not able to connect to Tor. The most common reasons for failing to connect were not knowing whether to connect directly or configure a bridge or proxy (P73, P75, P89, P91,

	E1-NEW	E1-OLD	E2-NEW	E2-OLD	E3-NEW	E3-OLD
no bridge, no proxy	17	13				
obfs3, no proxy	2	6	18	16		
meek-amazon, no proxy					7	4
meek-google, no proxy					5	4
meek-azure, no proxy					1	1
no bridge, 3rd-party proxy						1
DNF (did not finish)				3	6	10

Table 3. Participants’ connections to Tor. Note that none used the flashproxy, fte, fte-ipv6, obfs4, or scramblesuit bridges to connect. One participant configured a proxy to bypass our simulated censorship environment.

P106, P110), not knowing what they needed and trying to configure proxies (P74, P92, P105, P107, P113), and not knowing what to do when a recommended bridge did not work (P90, P93, P108, P111, P114).

7.3.2 Connection time

Our changes significantly reduced connection time (one-tailed Mann–Whitney; $Z = -1.84$, $p < 0.0328$, $r = 0.172$). We used a one-tailed Mann–Whitney test because the distributions were non-normal, heavily right-tailed, and right-censored.

The simulated censorship environment (and therefore, the difficulty of the configuration) was the strongest factor in determining connection time (Kruskal–Wallis $\chi^2 = 80.5$, $df = 2$, $p < 10^{-15}$). Figure 6 shows that participants in E2 took longer to connect than participants in E1, and participants in E3 took the longest to connect. Ideally, we would prefer an interface that allows users to connect in about the same amount of time, regardless of their configuration requirements; neither the OLD or NEW interface achieves this.

Figure 7 shows the cumulative distribution of connection times. Everyone in E1 was able to connect right away; participants with the NEW interface were faster than participants with the OLD interface. Most participants in E2 were able to connect right away, with participants with the NEW interface and OLD interface taking a similar amount of time. Participants in E3 took various amounts of time to succeed, with participants with the NEW interface connecting significantly faster than those with the OLD interface.



Fig. 5. A visualization of our participants' paths through the interface. The length of the bars show the total time taken to complete the task, except for those we cut off after approximately 40 minutes. Different colors indicate which screen they were on during the experiment. If participants were doing other things (i.e., searching for help in another browser) while off-focus, this time still counts as spending time in that off-focus window. "Not running" are times when Tor Launcher was closed (i.e., restarting the application).

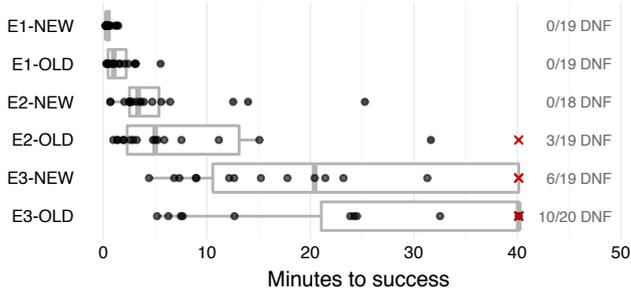


Fig. 6. Connection times. Participants in E2 and E3 took a long time to connect. The “DNF” (did not finish) shows how many were never able to successfully connect.

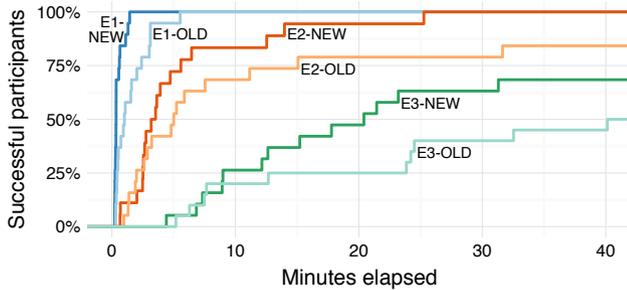


Fig. 7. Cumulative success rates over time. Those who were never able to successfully connect were assigned an arbitrarily large finishing time greater than 40 minutes for the purpose of graphing this cumulative distribution.

7.3.3 Configuration Time

Our changes significantly reduced configuration time (one-tailed Mann–Whitney; $Z = -3.28$, $p < 0.0005$, $r = 0.307$). Again, we used a one-tailed Mann–Whitney test because the distributions were non-normal, heavily right-tailed, and right-censored.

Figure 8 shows that participants’ active configuration time was generally a small fraction of their connection time. The configuration time of participants in E1 was dominated by the time waiting for a connection to be established to Tor because they generally chose the correct configuration right away. The configuration time of participants in E3 was also dominated by the time waiting for a connection to be established to Tor, because they usually needed to try multiple configuration options before connecting, but had to wait for the connection to time out or manually cancel their connection between attempts. The configuration time of participants in E2 was a modest fraction of their connection time, due to a combination of reasons that affected participants in E1 and E3.

Due to the reasons above, participants spent more time on the progress screen than on any other screen

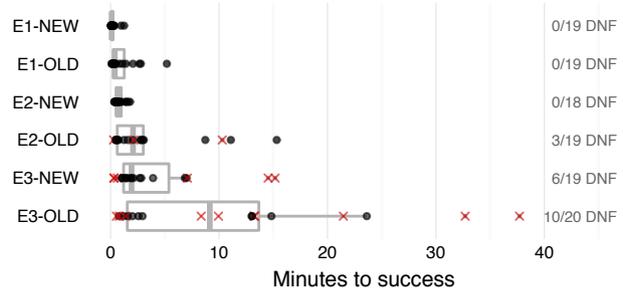


Fig. 8. Configuration times (connection time with time on progress screen subtracted). The differences from Figure 6 (connection times) show that most of the time was spent waiting for the connection to establish or time out.

	First	Proxy	Bridge	Progress
E1-NEW	28%	0%	0%	60%
E1-OLD	30%	0%	0%	29%
E2-NEW	6%	5%	6%	78%
E2-OLD	7%	18%	8%	45%
E3-NEW	3%	5%	5%	77%
E3-OLD	2%	12%	6%	64%

Table 4. The median percent of time spent on each screen, aggregated over participants in each experimental condition.

(Table 4). We defined configuration time as a participant’s connection time spent on all screens but the progress screen, because participants mostly waited on the progress screen. However, we noticed that some participants utilized this time to search for help in another browser while waiting for their connection to time out, so configuration time is not a perfect metric to measure participant engagement.

Although none of the participants needed to configure a proxy to bypass censorship, participants still spent time trying to configure them. Participants the OLD interface spent more time configuring proxies than bridges and participants with the NEW interface spent about an equal amount of time configuring bridges and proxies even when the interface said that they didn’t need one.

8 Recommendations

Our design resulted in shorter connection times and greater success rates, but it could still be improved. We do not recommend our exact design as the optimal design, although we believe it is an improvement. Instead, we offer some suggestions based on what we learned:

- **Assume users do not know technical concepts.** Users should not need to know what a relay, bridge, proxy, or transport is to connect to Tor.
- **Explain why users need to make decisions.** Letting users know why they are required to perform a task may make them more sympathetic to performing that task. For instance, users who need a bridge must choose one manually because different ones work in different countries and Tor does not track users' locations.
- **Leverage user knowledge.** Users know what country they are connecting from, their level of acceptable risk, and what they are using Tor for. Guiding users to make decisions based on this information will work better than asking about their ISP or explaining what obfuscation protocols are.

We did not try any designs that asked for user inputs, due to design constraints. But we see further room for improvement in success rate, connection time, and configuration time that could be made by leveraging user input and automation. We see several promising approaches to improving the Tor launcher:

- **Design changes:** As demonstrated by our design iteration, incremental improvements to the interface can result in significant usability improvements.
- **Naive automation:** The launcher could automatically choose a bridge, wait for a timeout, and try another bridge on failure, so users do not need to do this themselves. We call this “naive” automation, because this scheme does not leverage information about the user or network conditions and essentially just automates the trial and error process that users would otherwise have to do manually.
- **Smart automation:** The launcher could initially connect to Tor via a hard-to-censor bridge, then contact a central server (over this Tor connection) to identify a bridge that will work for the user based on availability and bandwidth constraints, automatically choose that bridge, and re-connect to Tor using that bridge. We call this “smart” automation because this scheme involves communication with Tor servers, which choose a bridge that will work for the user on the first try.

We recommend a gradual approach that starts with deploying design changes, measuring the impact of those changes, and cautiously incorporating automation if the impact of those changes is still unsatisfactory.

9 Limitations

We did not study international users, or the interface in languages other than English. All of our participants were from the United States and were instructed to interact with the English version of Tor Browser.

Participants in a laboratory setting may alter their behavior due to their awareness of being observed [26]. Ours were likely motivated differently compared to real users. We believe that they were likely over-motivated because they were being observed and receiving a monetary reward. If true, this makes our results a conservative estimate of the usability problems.

Our experiment did not directly test configuring a proxy or a custom bridge. From observing participants, we suspect they also struggle with these tasks.

We only tested the interface on the Windows operating system. The interface leverages the native operating system's styling and elements, making the configuration interface on Windows look slightly different from the OS X or Linux equivalents. We acknowledge that participants' unfamiliarity with Windows may have affected our experiment, but we believe that this affected all experimental conditions equally.

10 Related Work

We do not know of any published usability studies of Tor Browser since the introduction of Tor Launcher, which happened in the 3.5 release in December 2013 [34]. Lee and Fifield ran an informal “UX sprint” for Tor Browser in 2015 [25], bringing together developers and users to discover common usability problems. In what was a precursor to the present work, they conducted detailed walkthroughs with five participants to identify common obstacles to using Tor Browser such as hard-to-find downloads, bad interactions with other end-user security software, and misleading error messages.

There are studies that evaluated earlier forms of web browsing over Tor, without a particular focus on configuration. In 2012, Norcie et al. [32] investigated Tor Browser's “stop-points” during installation and use, which they define as places where users' uncertainty about what to do put them at risk of not completing a task. Of their 25 participants, 64% encountered at least one stop-point, such as a long startup time, misidentified icons, or confusion with other web browsers.

In 2010, Fabian et al. [17] estimated how much Tor's added latency harms usability (considered separately

from user interface concerns). They measured the latency of requests to common websites with and without Tor, and mapped the latencies to compute an expected “cancellation rate,” which they define as the fraction of users who would have given up waiting for a request to complete. They estimated a median cancellation rate of 88% for Tor, compared to 14% for direct requests.

In 2007, Clark et al. [9] used cognitive walkthroughs to evaluate four methods of web browsing over Tor (all of which have since been superseded by Tor Browser and Tor Launcher). Evaluating each method against a set of guidelines, they found that none were fully satisfactory, though the ones that used automatic configuration did better. In 2006, Dingedine and Mathewson [12] emphasized that anonymizing networks, by their nature, require lots of users, so usability is an even greater concern than with other security systems.

Usable security aims to improve the security of a system by helping the user make good decisions [33]. This involves informing users about the system, communicating risks in a way that is understandable, clarifying what the security requirements are, and providing guidance with technical tasks [1]. Usable security research explores a variety of security concepts, such as email encryption [20, 43], authentication [10, 28, 38], security indicators [2, 11], and warnings [15, 37]. Our work contributes to this body of knowledge by exploring tools for censorship circumvention.

We used a cognitive walkthrough [40, 42] to inspect our interface [29]. Other usability inspection methods include heuristic evaluation, pluralistic walkthrough and feature inspection [30]. We used interviews and task benchmarking to get qualitative and quantitative data [24]. Qualitative data can also be collected through focus groups, camera studies, and product surveys. Quantitative data can also be collected through eye tracking, clickstream analysis and A/B testing [36].

For our work, we draw on the large body of previous work in secure interaction design, computer-human interaction, and user studies. Work we found useful to our experiment includes research on principles for secure interaction design [47], common problems in human-computer dialogue design [27], user tolerance of security delays [16], treating human attention as a scarce resource when to delegate tasks to users [6], stopping users from installing potentially harmful programs [21], and getting proper user consent [7].

11 Conclusion

Our usability evaluation of Tor Launcher measured why and how users have a hard time connecting to Tor. We found that users struggle because the interface requires them to know technical terms, provides room for error, and does not give proper feedback. Users in heavily censored environments will likely not succeed in connecting to Tor without outside help. We encourage more usability studies on security tools in general as well as other Tor applications, as they have the potential to create user growth, help user retention, and prevent user frustration. Tor Browser version 5.4.5 incorporated changes based on our work. We also discussed the possibility of automating connections to Tor with its developers, and they are now considering this as future work.

12 Acknowledgments

We would like to thank the peer reviewers, especially our shepherd Jens Grossklags, for giving us valuable feedback on this paper. Tor Project employees Georg Koppen, Mike Perry, and Isabela Bagueros provided valuable insights about the history of Tor Launcher and internal constraints. We gratefully acknowledge assistance from the Xlab (Experimental Social Science Laboratory) [46] and financial support from Intel through the ISTC for Secure Computing, the National Science Foundation through award CNS-1318680, the Open Technology Fund through the Freedom2Connect Foundation, the US Department of State, Bureau of Democracy, Human Rights and Labor, the AFOSR under MURI award FA9550-12-1-0040, and the Hewlett Foundation through the Center for Long-Term Cybersecurity.

References

- [1] A. Adams and M. A. Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):40–46, 1999.
- [2] D. Akhawe and A. P. Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*, pages 257–272, 2013.
- [3] W. Albert and T. Tullis. *Measuring the user experience: collecting, analyzing, and presenting usability metrics*. Newnes, 2013.
- [4] Alexa: The top 500 sites on the web. <http://www.alexa.com/topsites>.

- [5] Y. Angel and P. Winter. obfs4 (the obfoursicator), May 2014. <https://gitweb.torproject.org/pluggable-transport/obfs4.git/tree/doc/obfs4-spec.txt>.
- [6] R. Böhme and J. Grossklags. The security cost of cheap user interaction. In *Proceedings of the 2011 workshop on New security paradigms workshop*, pages 67–82. ACM, 2011.
- [7] R. Böhme and S. Köpsell. Trained to accept? a field experiment on consent dialogs. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2403–2406. ACM, 2010.
- [8] BridgeDB. <https://bridges.torproject.org/>.
- [9] J. Clark, P. C. Van Oorschot, and C. Adams. Usability of anonymous web browsing: An examination of Tor interfaces and deployability. In *3rd Symposium on Usable Privacy and Security*, pages 41–51. ACM, 2007.
- [10] R. Dhamija and A. Perrig. Déjà Vu: A user study using images for authentication. In *USENIX Security Symposium*, volume 9, pages 4–4, 2000.
- [11] R. Dhamija, J. D. Tygar, and M. Hearst. Why phishing works. In *SIGCHI conference on Human Factors in computing systems*, pages 581–590. ACM, 2006.
- [12] R. Dingledine and N. Mathewson. Anonymity loves company: Usability and the network effect. In R. Anderson, editor, *Fifth Workshop on the Economics of Information Security*, June 2006.
- [13] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *13th USENIX Security Symposium*, August 2004.
- [14] K. P. Dyer, S. E. Coull, T. Ristenpart, and T. Shrimpton. Protocol misidentification made easy with Format-Transforming Encryption. In *Computer and Communications Security*. ACM, 2013.
- [15] S. Egelman, L. F. Cranor, and J. Hong. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *SIGCHI Conference on Human Factors in Computing Systems*, pages 1065–1074. ACM, 2008.
- [16] S. Egelman, D. Molnar, N. Christin, A. Acquisti, C. Herley, and S. Krishnamurthi. Please continue to hold. In *Ninth Workshop on the Economics of Information Security*, 2010.
- [17] B. Fabian, F. Goertz, S. Kunz, S. Müller, and M. Nitzsche. Privately waiting: A usability analysis of the Tor anonymity network. In *Sustainable e-Business Management*, Lecture Notes in Business Information Processing 58, pages 63–75. Springer, 1 edition, 2010.
- [18] D. Fifield, N. Hardison, J. Ellithorpe, E. Stark, D. Boneh, R. Dingledine, and P. Porras. Evading censorship with browser-based proxies. In *Privacy Enhancing Technologies*, pages 239–258, Berlin, Heidelberg, 2012. Springer-Verlag.
- [19] D. Fifield, C. Lan, R. Hynes, P. Wegmann, and V. Paxson. Blocking-resistant communication through domain fronting. *Privacy Enhancing Technologies*, 1(2):1–19, 2015.
- [20] S. L. Garfinkel and R. C. Miller. Johnny 2: a user test of key continuity management with S/MIME and Outlook Express. In *2005 symposium on Usable privacy and security*, pages 13–24. ACM, 2005.
- [21] N. S. Good, J. Grossklags, D. K. Mulligan, and J. A. Konstan. Noticing notice: a large-scale experiment on the timing of software license agreements. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 607–616. ACM, 2007.
- [22] G. Kadianakis and N. Mathewson. obfs3 (the threebficator), Jan. 2013. <https://gitweb.torproject.org/pluggable-transport/obfsproxy.git/tree/doc/obfs3/obfs3-protocol-spec.txt>.
- [23] G. Koppen. Tor Browser 5.0.3 is released, Sept. 2015. <https://blog.torproject.org/blog/tor-browser-503-released>.
- [24] K. Krol, J. M. Spring, S. Parkin, and M. A. Sasse. Towards robust experimental design for user studies in security and privacy. In *Learning from Authoritative Security Experiment Results (LASER)*, pages 21–31, San Jose, CA, 2016. USENIX.
- [25] L. Lee and D. Fifield. UX Sprint 2015 wrapup. <https://blog.torproject.org/blog/ux-sprint-2015-wrapup>, Feb. 2015. Accessed: 2015-10-5.
- [26] R. McCarney, J. Warner, S. Illiffe, R. Van Haselen, M. Griffin, and P. Fisher. The Hawthorne Effect: a randomised, controlled trial. *BMC medical research methodology*, 7(1):1, 2007.
- [27] R. Molich and J. Nielsen. Improving a human-computer dialogue. *Communications of the ACM*, 33(3):338–348, 1990.
- [28] R. Morris and K. Thompson. Password security: A case history. *Communications of the ACM*, 22(11):594–597, 1979.
- [29] J. Nielsen. Usability inspection methods. In *Conference companion on Human factors in computing systems*, pages 413–414. ACM, 1994.
- [30] J. Nielsen. Summary of usability inspection methods, 2016. <https://www.nngroup.com/articles/summary-of-usability-inspection-methods/>.
- [31] Nielsen Norman Group. Why you only need to test with 5 users. <http://www.nngroup.com/articles/how-many-test-users/>.
- [32] G. Norcie, K. Caine, and L. J. Camp. Eliminating stop-points in the installation and use of anonymity systems: A usability evaluation of the Tor Browser Bundle. In *5th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETS)*, 2012.
- [33] B. D. Payne and W. K. Edwards. A brief introduction to usable security. *IEEE Internet Computing*, 12(3):13–21, 2008.
- [34] M. Perry. Tor Browser Bundle 3.5 is released, Dec. 2013. <https://blog.torproject.org/blog/tor-browser-bundle-35-released>.
- [35] M. Perry, E. Clark, and S. Murdoch. The design and implementation of the Tor Browser. Technical report, Tor Project, Mar. 2013. <https://www.torproject.org/projects/torbrowser/design/>.
- [36] C. Rohrer. When to use which user-experience research methods, 2016. <https://www.nngroup.com/articles/which-ux-research-methods/>.
- [37] S. E. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor's new security indicators. In *IEEE Symposium on Security and Privacy (S&P)*, pages 51–65, 2007.
- [38] X. Suo, Y. Zhu, and G. S. Owen. Graphical passwords: A survey. In *21st Annual Computer Security Applications Conference (ACSAC'05)*, pages 10–pp. IEEE, 2005.
- [39] D. Travis. Writing the perfect participant screener. <http://www.userfocus.co.uk/articles/screeners.html>. Accessed:

2016-04-06.

- [40] User Experience Professionals Association. Usability body of knowledge: Cognitive walkthrough, 2016. <http://www.usabilitybok.org/cognitive-walkthrough>.
- [41] VLC media player. <https://www.videolan.org/vlc/>.
- [42] C. Wharton, J. Rieman, C. Lewis, and P. Polson. The cognitive walkthrough method: A practitioner's guide. In *Usability inspection methods*, pages 105–140. John Wiley & Sons, Inc., 1994.
- [43] A. Whitten and J. D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Usenix Security*, 1999.
- [44] P. Winter and S. Lindskog. How the Great Firewall of China is blocking Tor. *Free and Open Communications on the Internet*, 2012.
- [45] P. Winter, T. Pulls, and J. Fuss. ScrambleSuit: A polymorphic network protocol to circumvent censorship. In *Workshop on Privacy in the Electronic Society*. ACM, 2013.
- [46] Xlab: Experimental Social Science Laboratory. <https://xlab.berkeley.edu/>.
- [47] K.-P. Yee. User interaction design for secure systems. In *International Conference on Information and Communications Security*, pages 278–290. Springer, 2002.

A Success and Error States in Tor Launcher 5.0.3

Here we define the end state codes that appear in Figure 3. `e#` and `e#'` are equivalent, except that `e#` uses a bridge while `e#'` does not.

success states:

- s0: no bridge, no proxy
- s1: valid bridge, no proxy
- s2: valid bridge and valid proxy (blank port)
- s3: valid bridge and valid proxy (specified port)

error states:

- e0: need a bridge or proxy
- e1: hardcoded bridge blocked
- e2: custom bridge: field left blank
- e3: custom bridge: syntax error
- e4: custom bridge: invalid address
- e5: custom bridge blocked
- e6: proxy type is not selected
- e7: proxy: IP address left blank
- e8: proxy: syntax error
- e9: proxy: invalid IP address, blank port
- e10: proxy: invalid IP address, good port
- e11: proxy: invalid IP address, bad port
- e12: proxy: valid IP address, bad port
- e13: proxy: valid IP address, blank port

The same inputs can lead to different end states, depending on the network environment. For example, Tor Launcher will connect to a proxy successfully with the port left blank if the proxy uses port 80 (s2) but will not if the proxy uses another port (e12).

B Qualitative User Study Recruitment Posting

We are recruiting participants for an in-person research study at the University of California, Berkeley. You will need to come in to our lab and perform tasks on a computer for an hour or less. You will be compensated \$30 for participating. No special knowledge and no technical experience is required. If you are interested, fill out the survey at [<survey link>](#).

C Qualitative User Study Prescreening Survey

We are recruiting participants for an in-person research study at the University of California, Berkeley. You will need to come in to our lab and perform tasks on a computer for an hour or less. You will be compensated \$30 for participating. No special knowledge and no technical experience is required.

1. Please select when you are available. We will assign you an hour experiment time slot during one of those times.
2. I am able to provide my own transportation to the University of California, Berkeley campus.
3. Thank you for your interest! Please provide an email address where we can contact you to share more logistical details.
4. we are looking for a very small number of participants, so unfortunately, we may not be able to accommodate everyone who applies. Would you like us to let you know about future opportunities?
5. What is your gender?
6. What is your age?
7. Please select your highest completed (or current) level of education.
8. What is your occupation?
9. Do you speak any languages other than English fluently?

10. If you have a personal computer, what kind do you use?
11. Which of the following terms have you heard of? *<Answer choices: a checklist of the following terms: malware; proxy services; phishing; SSL; X.511 certificates; Tor.>*
12. How often do you use the following software or features? *<Answer choices: a grid of radio buttons. Software/features (rows): HTTPS on web pages; proxies or other censorship circumvention tools; virtual private networks (VPN); file or whole-disk encryption; anonymity systems (e.g., Tor); email encryption (e.g., PGP); chat or instant messaging encryption; voice communication encryption. Frequency (columns): never; less than once a month; a few times a month; several times a week; daily.>*

Thank you for filling out this form. You are now done!

D Qualitative User Study Introduction Script

Imagine you live in an oppressive country that censors part of the Internet. We have simulated this in the laboratory by blocking certain websites and services. The purpose of this experiment is to evaluate the use of Tor browser, which is a browser that can circumvent censorship and let you visit blocked websites. Currently, torproject.org is blocked (you can check this by going to torproject.org on a standard browser, like Firefox, Chrome, or Internet Explorer).

To circumvent censorship successfully, you will need to set up Tor browser correctly and use it to get to Wikipedia. If you are able to reach the website, then you know that you have successfully circumvented censorship. Fill out the question on the worksheet. This isn't intended to be hard, just write what you see. We want to just check you saw the website.

Before you start, do you have any questions about what you are asked to do?

E Qualitative Study Worksheet Text

Imagine you live in an oppressive country that censors part of the Internet. We have simulated this in the laboratory by blocking certain websites and services. The

purpose of this experiment is to evaluate the use of Tor browser, which is a browser that can circumvent censorship and let you visit blocked websites. For instance, www.torproject.org is blocked. Check this by going to the site on a standard browser, like Firefox, Chrome, or Internet Explorer. It will fail to load, when you can visit other sites.

To complete this worksheet, you will need to set up Tor browser (on your desktop) correctly and use it to get to blocked site. If you can visit wikipedia, then you know that you have successfully circumvented censorship.

F Post-Experiment Standard Interview Questions

We asked our participants these questions after they were given time to configure Tor Browser.

1. Can you talk us through what you did along with what you were thinking at the time?
2. What was most challenging part of connecting?
3. Were there any unfamiliar terms?
4. How did you decide which options to choose?
5. What did you think about using Tor?
6. What is one change you would recommend?
7. Did you need any additional information?

In addition to these questions, we asked our participants about specific questions based on their observation, usually regarding a specific choice in action, a particular screen they seemed stuck on, and any errors they encountered during the configuration process.

G Qualitative Testing Observations

We summarize each participant's session.

G.1 E1

In E1, any connection to Tor worked. Configuring bridges and proxies was optional. Observations in E1:

P1 (new user, direct, 0:39): *They connected directly.*

P2 (new user, direct, 1:20): *They spent time reading the text on the start screen, and connected directly because they were intimidated by the other option.*

P3 (new user, obfs3, 1:39): *They connected with the recommended bridge. They were also going to use a proxy, but decided not to when they saw the proxy input fields.*

P4 (new user, obfs4, 8:56): *They chose the recommended bridge and then tried to configure a proxy. After not being able to fill out the proxy input fields, they started over and connected with an obfs4 bridge. Note that they did not try connecting with the obfs3 bridge, which would have worked.*

P5 (previous Tor user, obfs3, 1:02): *They connected with the recommended bridge.*

Participants in E1 were able to connect to Tor, but most were unsure if their actions were correct. Three of the five participants unnecessarily configured bridges.

G.2 E2

In E2, a bridge was required to connect. Any of the hardcoded bridges from the dropdown menu, as well as any custom bridges acquired out of band would work. Configuring a proxy was optional. Observations in E2:

P6 (new user, obfs3, 4:04): *They tried connecting directly. After watching the progress bar not make any progress for a couple minutes, and gave up on the direct connection. Then they connected with the recommended bridge.*

P7 (new user, obfs3, 9:22): *They tried connecting directly. They restarted Tor Launcher and tried to connect directly again—and repeated this two more times. On the fourth restart, they tried to configure a proxy, but gave up. On the fifth restart, they connected using the recommended bridge.*

P8 (new user, direct, 10:40): *They tried connecting directly. After waiting a couple minutes, they figured that did not work and looked at proxy settings. But after looking at the settings and being intimidated, they tried connecting directly again. Then, they tried to configure a proxy, and gave up. They tried connecting directly again, which worked unexpectedly. They found a bug in our setup that we fixed for the later experiment. E2 required a bridge to connect, but new public relays came online after we ran our firewall rules and they were able to connect using those new relays.*

P9 (previous Tor user, obfs3, 4:08): *They tried connecting directly. And they tried connecting directly through a different user path by answering ‘no’ to the bridge and proxy questions in the interface, connecting from the proxy screen. Then, they connected using the recommended bridge.*

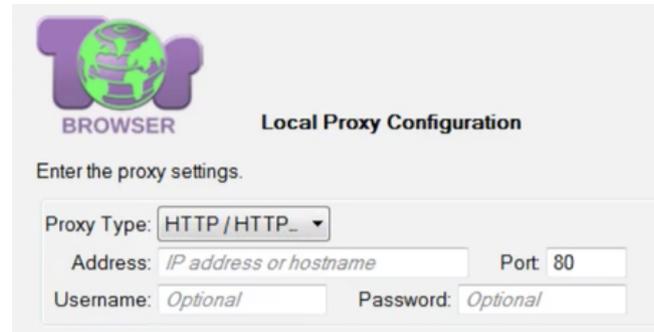


Fig. 9. Proxies were not necessary to connect to Tor, but many spent time trying to configure one when they couldn't connect. Here, we show P10's unsuccessful attempt.

P10 (previous Tor user, failed to connect): *They tried a direct connection. Then, they tried a connection with the recommended bridge, which should have worked. However, our setup experienced a clock drift. Tor does not allow connections with bad clocks. They spent the rest of their time trying different bridges and proxies in vain (Figure 9). This participant made the most connection attempts, using a variety of bridges and proxies.*

Participants in E2 tried a direct connection to Tor before *eventually* connecting with the recommended bridge. They spent a fair amount of time waiting before determining that connections failed. Some stated that they connected directly because the interface says that it would work in “most situations” (Figure 2a). Our participants did not know if they needed bridges or proxies, so they only tried to configure one when the direct connection failed. Since they did not know whether they needed a bridge or proxy, they often tried to configure proxies when they did not need one and should have configured bridges instead.

G.3 E3

In E3, a bridge was required to connect. Most of the hardcoded bridges in the dropdown menu would not work, with the exception of meek-amazon, meek-azure, and meek-google. Custom bridges can also work. Configuring a proxy was optional. Observations in E3:

P11 (new user, failed to connect): *They tried connecting directly, and then with the recommended bridge. Then, they retried connecting directly and retried connecting using the recommended bridge. After that, the participant spent the rest of their time trying to configure a proxy and retrying connections with the recommended bridge.*

P12 (new user, failed to connect): *They tried connecting with the recommended bridge. They decided that they need a proxy along with the recommended bridge, and spent the rest of their time trying to configure a proxy.*

P13 (new user, failed to connect): *They tried connecting directly, and then with the recommended bridge. They retried connecting with the recommended bridge. When that didn't work, they tried to configure a proxy. Then, they gave up on the task entirely.*

P14 (new user, meek-google, 26:48): *They tried a direct connection by answering 'no' to the bridge and proxy questions. After waiting a while, they canceled the connection and tried connecting with the recommended bridge. They retried connecting with the recommended bridge two more times, then tried to connect with a flashproxy bridge, an obsf4 bridge, a scramblesuit bridge. After retrying the recommended bridge again, they made a connection using a meek-google bridge.*

P15 (new user, meek-azure, 7:31): *They tried connecting directly, and then with the recommended bridge. After that, they tried a connection using a meek-azure bridge, but gave up after before it could connect. They went back to the first screen (Figure 2a) and clicked connect, which made a connection using a meek-azure bridge. The connect button makes a connection to Tor with the toggled settings in the interface. When the interface is initialized to not use a bridge or proxy, clicking connect on the first page starts a direct connection. But since the participant had chosen a meek-azure bridge, the connect button made a connection to a meek-azure bridge, even though the text in the first screen says that they will "connect directly." Because of this, the participant thought they made a direct connection. Although they succeeded in connecting to Tor, they did not realize what they had done or why connecting to a meek-azure bridge worked. They chose that bridge at random.*

P16 (previous Tor user, custom bridge, 22:06): *They tried and retried connecting directly, then they tried connecting with the recommended bridge. After briefly looking at error messages in the system log, they retried connecting directly two more times. After trying to configure a proxy for a while, they instead decided not to use one and tried connecting with a meek-google bridge, which would have worked. But due to a bug in the progress bar that prevents it from updating correctly on subsequent attempts, they gave up before it could connect. They made a connection with the custom bridge by following the instructions from the bridge help button (they created a throwaway email account, emailed the bridge responder with poor syntax and did not get a response, emailed the bridge responder again with correct syntax and got a response, and then typed in the bridge information into the custom bridge field).*

Participants in E3 also generally tried a direct connection to Tor. They did not know if they needed a bridge

or proxy, and like the participants in E2, only configured them when the direct connection failed. Of the participants that failed, they all tried the recommended bridge but not any others. The participants that did manage to connect to Tor did not know which bridges worked and which ones did not, or which ones to try next in case the recommended one failed. P14 brute-forced options until something worked, P15 did not think that they were using a bridge, and P16 gave up on a hard-coded bridge that would have worked and used a custom bridge instead.

H Quantitative User Study Recruitment Posting

We are recruiting up to 40 participants for a user study at UC Berkeley. The experiment will involve basic Internet browsing tasks. You are not eligible if you have participated in our previous sessions.

Payment: \$30 Amazon gift card

Duration: 1 hour

Where: Xlab at Hearst Memorial Gymnasium

<list of sessions>

To be eligible, you must be an adult (18 or older). This is to comply with university policies on research.

If you are interested: 1. Email lnl@berkeley.edu with the sessions you are able to attend. We will confirm your participation and assign you a session. 2. Come to Xlab at the appointed time for the experiment.

I Quantitative User Study Introduction Script

Imagine you live in an oppressive country that censors part of the Internet. We have simulated this in the laboratory by blocking certain websites and services. The purpose of this experiment is to evaluate the use of Tor browser, which is a browser that can circumvent censorship and let you visit blocked websites. Currently, torproject is blocked (you can check this by going to torproject.org on a standard browser, like Firefox, Chrome, or Internet Explorer).

To circumvent censorship successfully, you will need to set up Tor browser correctly and use it to get to

Wikipedia. Tor is already installed for you. On the desktop, you should see a globe icon that says “Start Tor Browser.” If you are able to reach the website, then you know that you have successfully circumvented censorship. Fill out the question on the worksheet. This isn’t intended to be hard, just write what you see. We want to just check you saw the website.

Afterward, we ask you to take a short survey to collect some information about you. The link is also on your worksheet. We will give you time to complete this task. If you finish early, we ask that you sit at your desk until the remainder of the hour. Since we are recording your screen, we ask that you don’t do anything personal afterward, like checking your email.

Before you start, do you have any questions about what you are asked to do?

J Quantitative User Study Worksheet Text

Imagine you live in an oppressive country that censors part of the Internet. We have simulated this in the laboratory by blocking certain websites and services.

The purpose of this experiment is to evaluate the use of Tor browser, which is a browser that can circumvent censorship and let you visit blocked websites. For instance, www.torproject.org is blocked. Check this by going to the site on a standard browser, like Firefox, Chrome, or Internet Explorer. It will fail to load, when you can visit other sites.

To complete this worksheet, you will need to set up Tor browser (on your desktop) correctly and use it to get to blocked site. If you can visit wikipedia, then you know that you have successfully circumvented censorship.

1. Visit Wikipedia’s main page (https://en.wikipedia.org/wiki/Main_Page). What is the topic of today’s selected article?
2. After you have completed the tasks, fill out this survey: <http://bit.ly/tor-survey>

K Quantitative User Study Exit Survey

We’d like to know more about you. All of your answers will be stored separately from any identifying information in order to protect your confidentiality.

This survey is part of a research project being conducted by the University of California, Berkeley. If you have any questions about your rights or treatment as a research participant in this study, please contact the University of California at Berkeley’s Committee for Protection of Human Subjects at 510-642-7461, or email subjects@berkeley.edu. If you agree to participate, please click Next below.

1. What is your participant ID? (This can be found on the sticker on the left hand corner of the desk you are currently sitting at.)
2. What is your gender?
3. What is your age?
4. Please select your highest completed (or current) level of education.
5. What is your current occupation?

Thank you for participating in our experiment. You are now done! Please sit at your desk for the remainder of the experiment. Our researchers will formally announce the end of the experiment.