Counting daily bridge users

Karsten Loesing
karsten@torproject.org

Tor Tech Report 2012-10-001
October 24, 2012

Abstract

As part of the Tor Metrics Project, we want to learn how many people use the Tor network on a daily basis. Counting users in an anonymity network is, obviously, a difficult task for which we cannot collect too sensitive usage data. We came up with a privacy-preserving approach for estimating directly connecting user numbers by counting requests to the directory mirrors and deriving approximate user numbers from there. In this report we describe a modified approach for estimating the number of users connecting via bridges by evaluating directory requests made to bridges. We compare this new approach to our current approach that estimates bridge user numbers from total unique IP addresses seen at bridges. We think that results from the new approach are closer to reality, even though that means there are significantly fewer daily bridge users than originally expected.

1 Introduction to our new approach to count bridge users

In this report we describe a new approach for estimating the number of daily users connecting to the Tor network via a bridge. This new approach uses counts of directory requests made to bridges as its main data sources. This is similar to how we estimate daily directly connecting users that connect to the Tor network via a non-bridge relay. Our current approach for estimating daily bridge users is to count unique IP addresses of connecting clients at bridges. We refer to our earlier report [1] for an overview of estimating user numbers in the Tor network.

We estimate daily bridge users by first summing up directory requests per day reported by bridges (Section 2). We extrapolate these reported requests to the expected total number of directory requests in the network (Section 3). We then assume that there is an average number of 10 directory requests that every client makes per day and derive daily user numbers by dividing by that average number (Section 4). We further derive users per country by including country information of connecting IP addresses (Section 5). There are at least two ways to remove unwanted artifacts from results: we may have to ignore reports from bridges that have been running as non-bridge relays and that might still report directly connecting users (Section 6); and we may need to ignore days when there were problems with the consensus process, leading to an increase in directory requests which is likely not caused by an actual increase in users (Section 7).

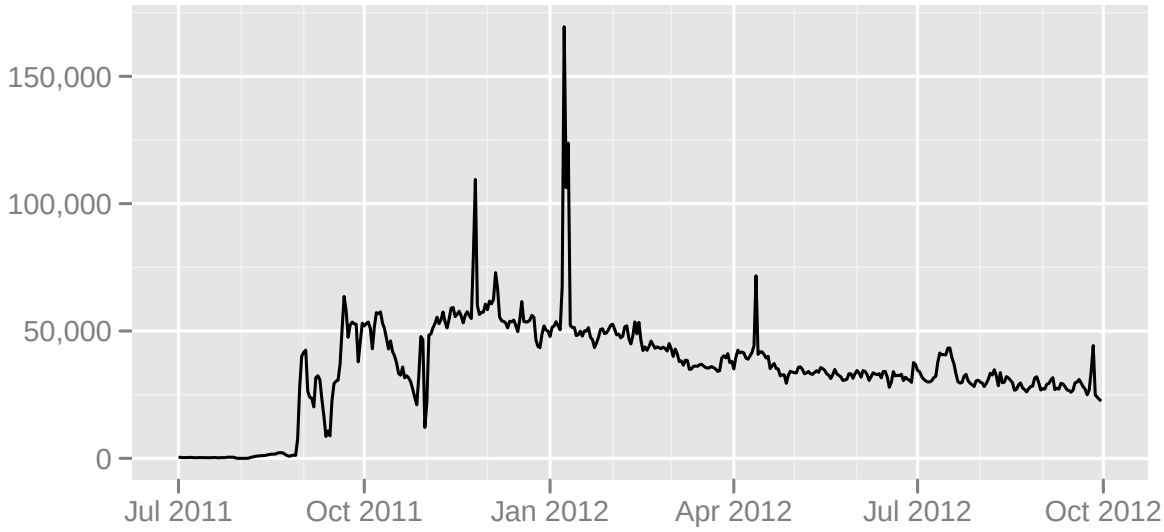

Figure 1: Daily sums of directory requests reported by all bridges

2 Counting reported directory requests to bridges per day

All relays running a recent enough Tor version contain code to collect and report statistics on processed directory requests over 24-hour periods. We refer to our previous work [3] for more details on aggregating usage statistics in Tor.

Bridges were originally not supposed to report directory request statistics, because bridge users were estimated based on unique IP address counts, and we wanted to avoid collecting any more data than needed. But due to a bug, only part of directory request statistics were suppressed when running in bridge relay mode. As a result, bridges report a trimmed version of directory request statistics, which are however still enough to estimate daily users. As an example, the following directory request statistics were reported by a bridge in September 2012.

```
extra-info goinpostal 7363FF835F5D79EA1F0CC2EB757B03866D4515F7
dirreq-stats-end 2012-09-18 15:26:38 (86400 s)
dirreq-v3-resp ok=5040,not-enough-sigs=0,unavailable=0,not-found=0,not-modified=0,busy=0
```

From this example we learn that this bridge successfully processed 5,040 version 3 directory requests in the 24 hours preceding September 18, 2012, 15:26:38 UTC. We have no information how many of these requests happened in the 8.5 hours of September 17 or in the 15.5 hours of September 18. We assume a uniform distribution of requests over the 24-hour interval and count 1,797 requests for September 17 and 3,243 requests for September 18. We extract these data points for all bridges publishing their descriptors to the bridge authority and sum up responses per day. Figure 1 contains the number of reported requests per day. A few observations:

1. There were hardly any reported requests until September 2011. The likely explanation is that, before Tor version 0.2.3.1-alpha, collecting and reporting directory request statistics

was disabled by default. This default was changed in Tor version 0.2.3.1-alpha which was released on May 5, 2011. With more and more bridges upgrading to the 0.2.3 series, the number of reported directory requests increased, too.

2. The reported request numbers in September and October 2011 have quite high volatility, making it difficult to use these request numbers for actual user number estimates.

3. There are at least three unusual spikes in request numbers in November 2011, January 2012, and April 2012, which were unlikely caused by sudden increases and decreases in user numbers.

4. From November 2011 on, there is a general downward trend from around 60,000 requests per day to just 30,000 in September 2012.

3 Extrapolating to total directory requests in the network

So far, we only know how many directory requests were processed by the bridges reporting them. We need to take into account that not all bridges report these statistics for various reasons: bridges may not be configured to report directory request statistics, which in particular applies to bridges running an earlier version than 0.2.3.1-alpha; bridges may run for less than 24 hours, thus not finishing a 24 hour statistics interval and discarding requests processed up to that time; bridges may have finished a 24-hour statistics interval, but went offline before publishing statistics to the bridge authority. (We analyzed in more detail what fraction of our bridges are not reporting usage statistics in [2].) As a result, we need to extrapolate reported requests to what we expect as the total number of requests in the network.

A straight-forward way to extrapolate to the total number of directory requests in the network would be to make the following assumption: every bridge that does not report directory request statistics, on average, processes as many directory requests as a bridge that reports them. Under this assumption, we could count the number of running bridges per day and the number of bridges reporting directory request statistics, compute the fraction of reporting bridges, and divide reported requests by that fraction. This assumption works fine as long as the variance between request numbers processed by bridges is small, or as long as the fraction of reporting bridges is high. However, the former is not the case, because there are some hard-coded bridge addresses in bundles distributed on the Tor website, leading to these bridges processing and maybe reporting far more directory requests than others. The latter is not always the case either, in particular before September 2011, as we could see in Figure 1.

A better way to extrapolate to total requests in the network is to consider a second statistic published by a subset of bridges: the number of bytes written to respond to directory requests. We assume that this number is proportional to the number of processed directory requests, even though we do not assume an exact linear relation. The subset of bridges reporting byte statistics is not necessarily the same subset thas is reporting directory request statistics. By taking into account byte histories, we can better estimate how many directory requests have not been reported by bridges that at least have reported byte histories. An example for reported written directory bytes, coming from the same extra-info descriptor as the example above, is as follows:

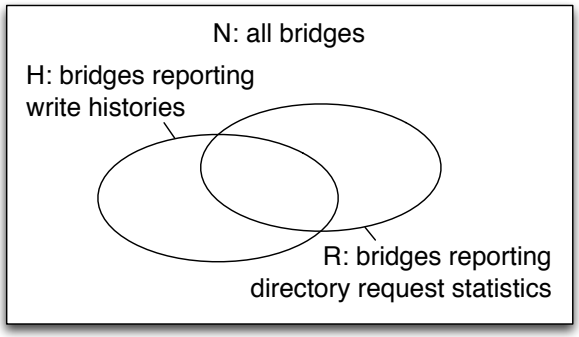

Figure 2: Subsets of bridges reporting write histories and directory request statistics

```
extra-info goinpostal 7363FF835F5D79EA1F0CC2EB757B03866D4515F7
dirreq-write-history 2012-09-19 05:17:32 (900 s) 13370368,10539008,56751104,27235328,
14555136,10524672,63341568,37339136,24343552,22490112,29155328,17792000,3502080,[...]
```

From these write histories we can extract how many bytes a bridge has spent on answering directory requests on a given day. By looking at the reported directory request history values, we can learn how many bytes were written while the bridge was also collecting and later reporting directory request statistics, and we can learn how many bytes were written outside of those statistics intervals. Similarly, we can learn how many directory requests were processed at times when the bridge did not report byte histories.

On a side note, as one can see from the example, byte history intervals are only 15 minutes long and thereby much shorter than directory request statistics intervals. It might be that this level of detail has privacy implications, in particular on bridges with only very few users. We probably don't need byte histories on this level of detail. We leave the analysis what level of detail is required as future work. Results could lead to increasing the history interval length on bridges to 1 hour or more.

In the following, we define R as the subset of bridges reporting directory request statistics, H as the subset reporting byte histories, and N as the entire set of bridges in the network. Figure 2 illustrates these subsets and the variable names. Also, we define $r()$ as the number of directory requests reported by a given set of bridges, $h()$ as the number of bytes reported by the bridges in a given set, and $n()$ as the absolute number of bridges in a set.

Knowing the number of reported directory requests, $r(R)$, we can extrapolate to the expected total number of directory requests, $r(N)$, by multiplying with the reciprocal of the fraction of written directory requests that got reported to us, $\frac{h(N)}{h(R)}$:

$$r(N) = r(R) \times \frac{h(N)}{h(R)} \tag{1}$$

Estimating total written directory request bytes in the network, $h(N)$, is easy. We assume here that the bridges that didn't report directory request bytes wrote the same number of bytes per bridge on average as reporting bridges.

$$h(N) = h(H) \times \frac{n(N)}{n(H)} \tag{2}$$

Estimating written directory request bytes by the bridges that reported directory request statistics, $h(R)$, is somewhat harder. We first split the set into the set of bridges reporting both statistics and the set of bridges reporting only statistics and no write histories.

$$h(R) = h(R \cap H) + h(R \setminus H) \tag{3}$$

The first number is something we can read from the descriptors. The second number requires us to apply the same assumption from above, namely that bridges that didn't report byte histories wrote the same number of bytes, on average, as reporting bridges. (Note how this equation is very similar to equation 2.)

$$h(R \setminus H) = h(H) \times \frac{n(R \setminus H)}{n(H)} \tag{4}$$

Putting everything together, we come up with a way to compute estimated directory requests in the network:

$$r(N) = r(R) \times \frac{h(H) \times n(N)}{h(R \cap H) \times n(H) + h(H) \times n(R \setminus H)} \tag{5}$$

Figure 3 shows reported directory requests, estimated fraction of directory requests that got reported by bridges, and estimated total directory requests in the network. A few observations:

1. The top-most graph is the same as in Figure 1 which we already discussed on page 2.

2. The middle graph shows an upwards trend of the fraction of bridges reporting directory request statistics. Fractions of under 25% as seen until end of October 2011 make it difficult to extrapolate to the total number of requests in the network. These fractions also explain the observed volatility of reported requests until end of October 2011. Beginning with November 2011, fractions are at 50% or higher, exceeding 75% in most of 2012.

3. The bottom-most graph is the result of dividing request numbers in the top-most graph by fractions in the middle graph. The low fractions in late August and early September 2011 lead to very high estimated total requests in the network. We'll want to treat these surprisingly high numbers with care, but so far, there is no reason to believe they're totally wrong. From November 2011 to September 2012, the continuously decreasing reported request numbers combined with the increasing fraction of reported requests lead to a continuous decrease in estimated directory requests in the network. From this graph it seems that bridge usage has steadily decreased in the past 11 months.

4 Dividing by 10 for estimating number of users

With the estimated number of daily directory requests in the network we can now estimate the number of daily users. We make the assumption that there is an average number of directory requests per day that every client makes to keep their network information up-to-date. As of writing this report, network status consensuses are fresh for three hours, requiring clients to download a new document every 2 to 3 hours. Hence, a client that is online all day would

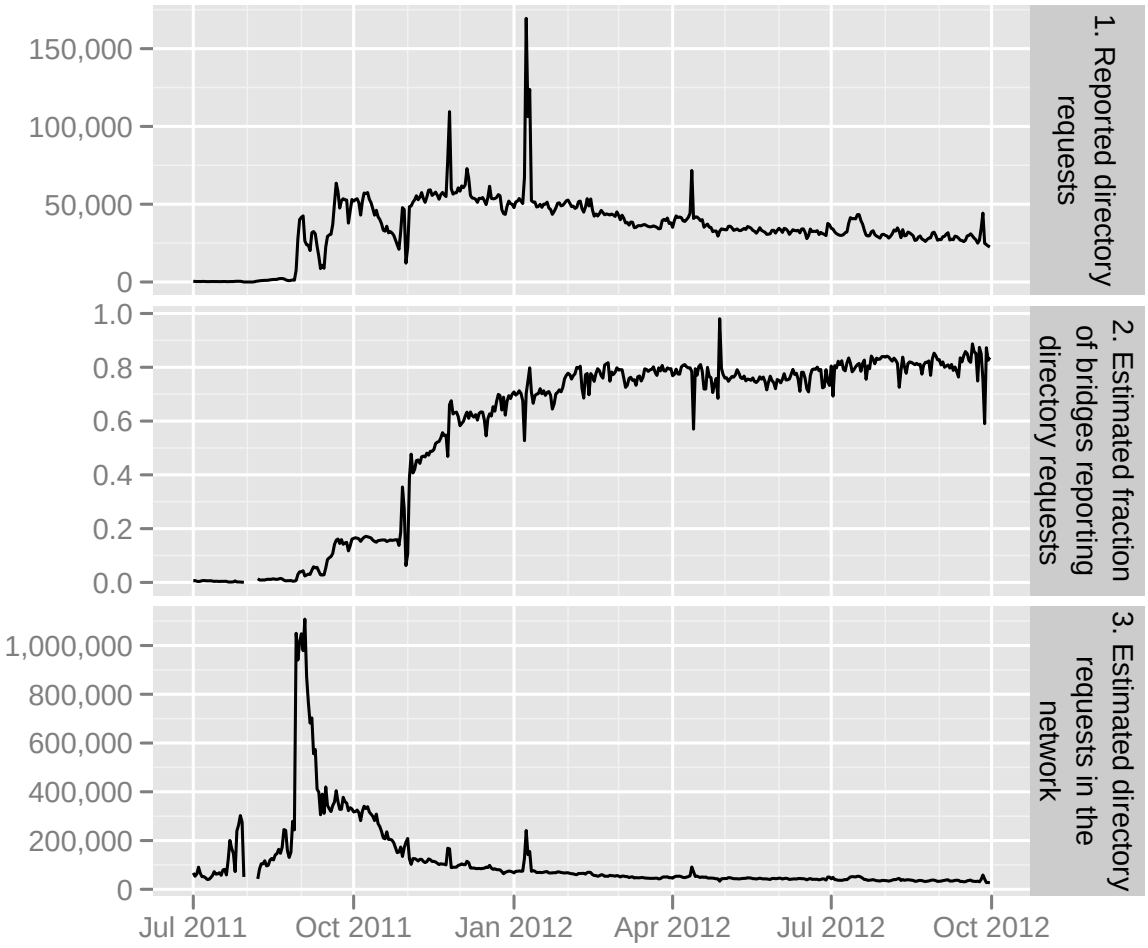

Figure 3: Reported directory requests, estimated fraction of bridges reporting directory requests, and estimated directory requests in the network

make 8 to 12 directory requests on that day. Not all clients are online all day, thus reducing the average number of directory requests made by clients. We, somewhat arbitrarily, chose 10 as the number of directory requests that the average client makes every day. 10 is also the number that we use in directly connecting user statistics, so that both estimates for non-censored and censored users will be easy to compare. We could evaluate whether 10 is a good number by asking volunteers to have their Tor clients record directory request numbers made on a given day, and use these actual numbers to come up with a better number. But given that we apply the same number of requests per client to all days, the actual value does not influence development over time, allowing us to observe trends over time and still have a rough idea of absolute numbers.

Now we can derive user numbers from total directory requests in the network. Figures 4 and 5 show the total number of users connecting via bridges over time, for the entire period for which we have data and for the third quarter of 2012. These graphs contain the same data as

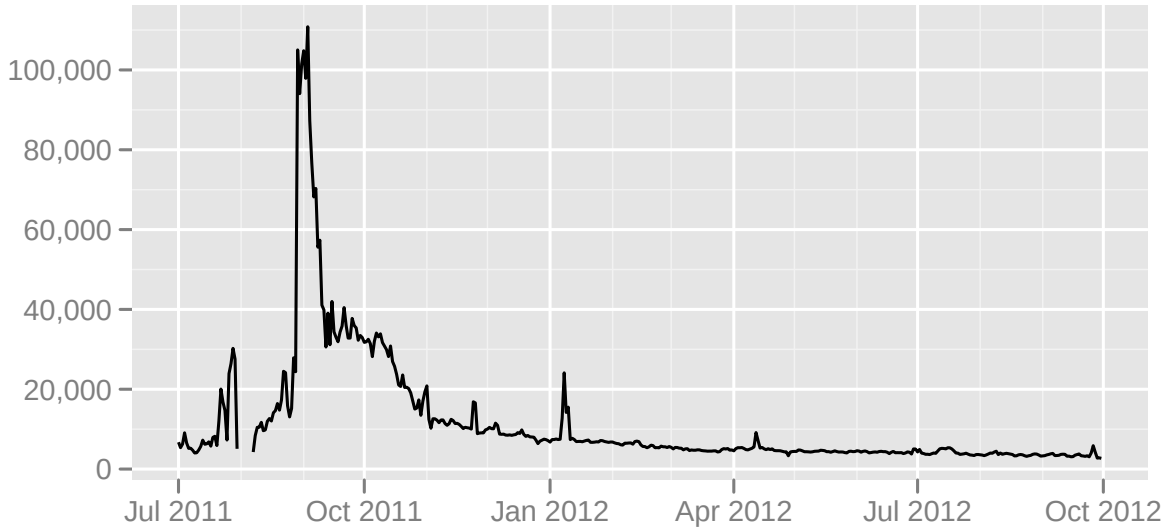

Figure 4: Estimated daily bridge users from all countries from July 2011 to September 2012

the bottom-most graph in Figure 3, but divided by 10. The most surprising result is that there are only 3,500 daily bridge users these days.

5 Breaking down to user numbers by country

So far, we only have an estimate of daily bridge users in the network, but no user numbers per country. In contrast to directory request statistics reported by relays, bridges do not report request numbers by country code. But we still have statistics on originating countries in the unique IP address statistics reported by bridges. These are the statistics we have been using for estimating daily bridge users so far. We assume that the country distribution of connecting bridge clients is similar to bridge clients downloading directory requests. As an example, these are the statistics on connecting clients reported by the same bridge from earlier examples:

```
extra-info goinpostal 7363FF835F5D79EA1F0CC2EB757B03866D4515F7
bridge-stats-end 2012-09-18 15:27:00 (86400 s)
bridge-ips ir=32,??=16,at=8,eg=8,jp=8,lv=8,pk=8,us=8
```

We sum up unique IP addresses and calculate a fraction of IP addresses for every country and day. (We could also have weighted country information of reporting bridges with the bridge's fraction of written directory request bytes, but this seemed like overkill for this analysis, so we left it for future work.) We multiply the estimated number of total users in the network with the fraction of unique IP addresses coming from a country and come up with the estimated number of users in that country. Figure 6 shows the result for estimated daily bridge users coming from Syria. We chose Syria for this example, because that's one of the countries with most bridge users these days. The approach would work for all other countries, too.

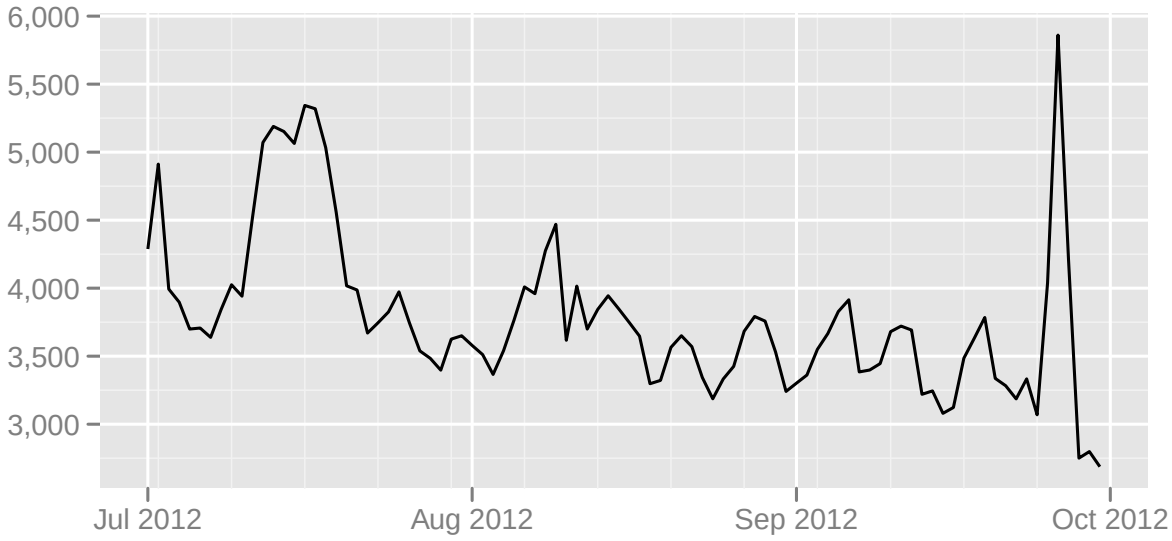

Figure 5: Estimated daily bridge users from all countries in the third quarter of 2012

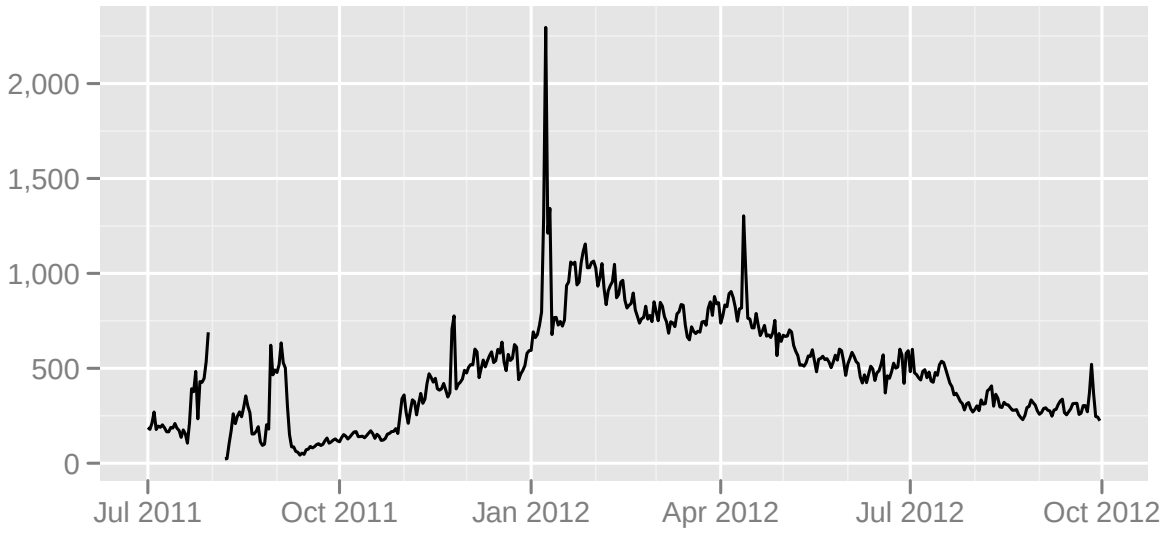

Figure 6: Estimated daily bridge users from Syria

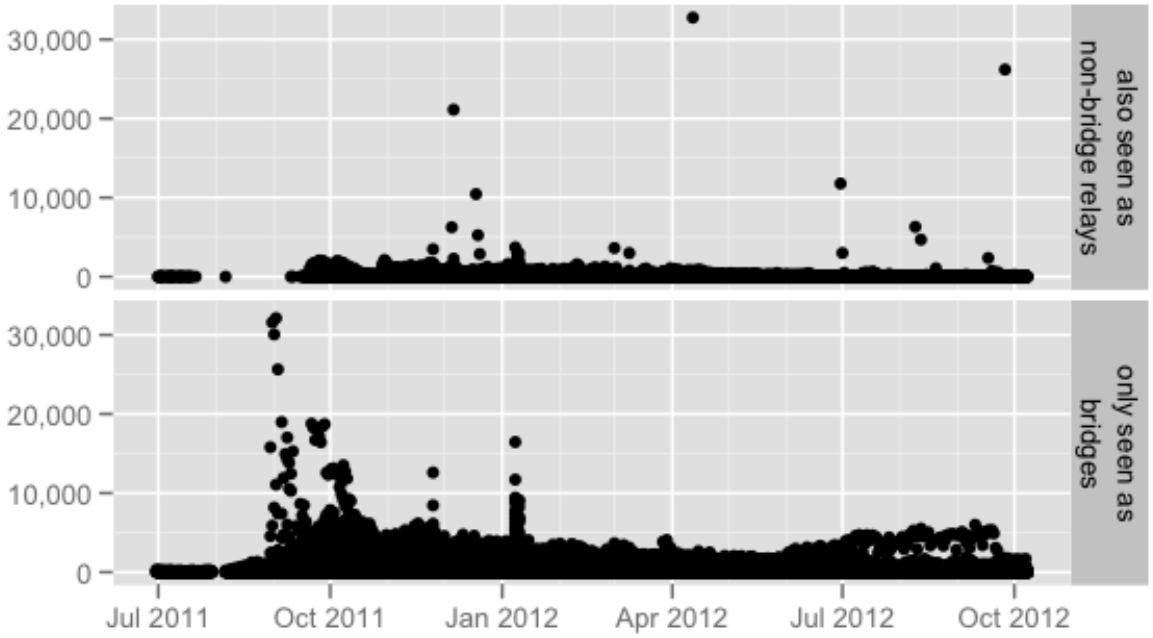

Figure 7: Reported directory requests by bridges that have or have not been seen as non-bridge relays

6 Ignoring bridges that were running as non-bridge relays

One aspect that we have been ignoring so far is that there are bridges that were running as non-bridge relays before. It is likely that these bridges report directory requests coming from directly connecting clients of which there are far more than bridge clients.

Figure 7 shows reported directory requests of bridges that have or have not been seen as non-bridge relays. The focus here is on data points which are seemingly outliers. Two noteworthy examples are the 30,000+ requests in April 2012 and the 25,000+ requests in September 2012.

In our current approach to count daily bridge users, we ignore any such bridge, because they could skew results. We'd like to exclude data points coming from bridges that report unrealistic statistics. However, it seems that ignoring all bridges that have been seen as non-bridge relays would mean removing too many data points. More research is needed to define criteria when a data point probably contains directory requests by non-bridge clients and should be ignored.

7 Ignoring days with too few network status consensuses

A closer look at spikes in total estimated directory requests and at archives of network status consensuses reveals an interesting correlation: in 4 out of 5 cases when less than 20 consensuses were published on a given day, the number of directory requests went up a lot. Figure 8 shows the number of published consensuses per day.

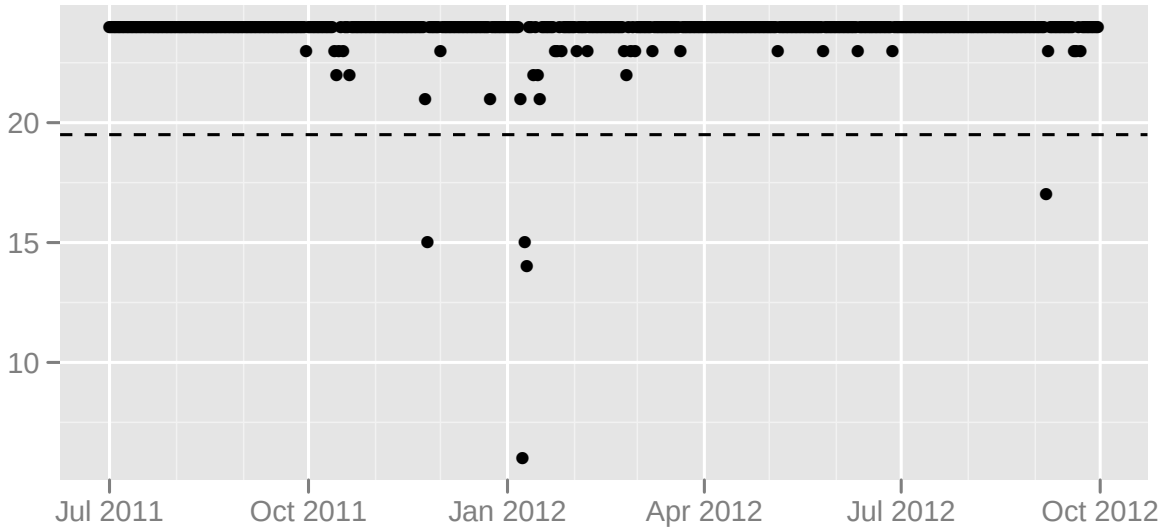

Figure 8: Published consensuses per day

By default, the directory authorities publish a new consensus every hour. Missing consensuses indicate a problem with the consensus process, meaning that the available consensuses become outdated and that clients send out more requests to get a recent enough consensus, thus raising the number of directory requests in the network.

We could decline providing bridge usage statistics for days when the archives had less than 20 consensuses. This would fix the spike in late November 2011 and especially the huge one in early January 2012. However, we decided not to remove these days yet and left it as future work to analyze how we can detect problems with the consensus process leading to higher directory request numbers.

8 Comparing old and new approaches to count bridge users

We briefly compare results from our new approach based on directory requests to results from our existing approach based on unique IP addresses. Figures 9 and 10 show the estimates of daily bridge users in the new and in the old approach, for the entire observed period and for the third quarter of 2012. The general trend is about the same, though the new approach only outputs about one tenth as many daily bridge users as the old approach did. We think that results from the new approach are closer to reality, because the reasoning behind it is much more plausible than the design of the old approach. We refer to our earlier report [1] for details of the old approach to count daily bridge users including a discussion of its weaknesses.

9 Suggesting next steps

We identified a few starting points for further improving the described approach in this report:

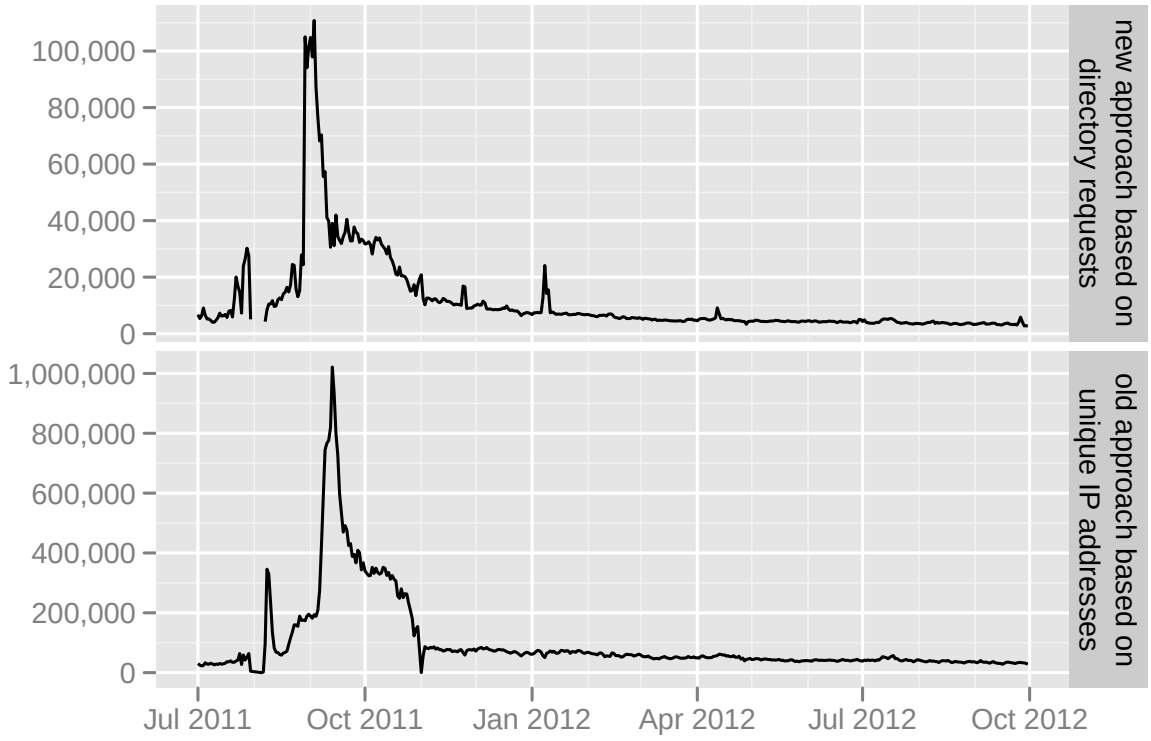

Figure 9: Estimated daily bridge users in the new and in the old approach from July 2011 to September 2012

- We should re-run the analysis under the assumption that bridges use a larger time period for byte histories than 15 minutes. There are potential privacy problems with seldomly used bridges reporting usage statistics on such a high detail level. The described approach to estimate daily users should work as well with byte histories on a detail of a few hours. Once we know what detail is required, we should change the default in the Tor sources, and we could update the bridge descriptor sanitizing code to increase the byte history interval in sanitized bridge descriptors.
- We should evaluate whether 10 is a reasonable average number of directory requests made by a client per day. One way to do this evaluation is to ask volunteers to have their Tor clients record directory request numbers made on a given day.
- We should look closer at weighting country information reported by bridges. Maybe we'll have to weight unique IP address fractions by country with the reporting bridge's written directory request byte fraction to get more accurate user numbers by country.
- We should further investigate how bridges that have been seen as non-bridge relays affect the results. If we need to ignore reported statistics by these bridges, we'll want to make sure to only exclude as few reports as necessary.

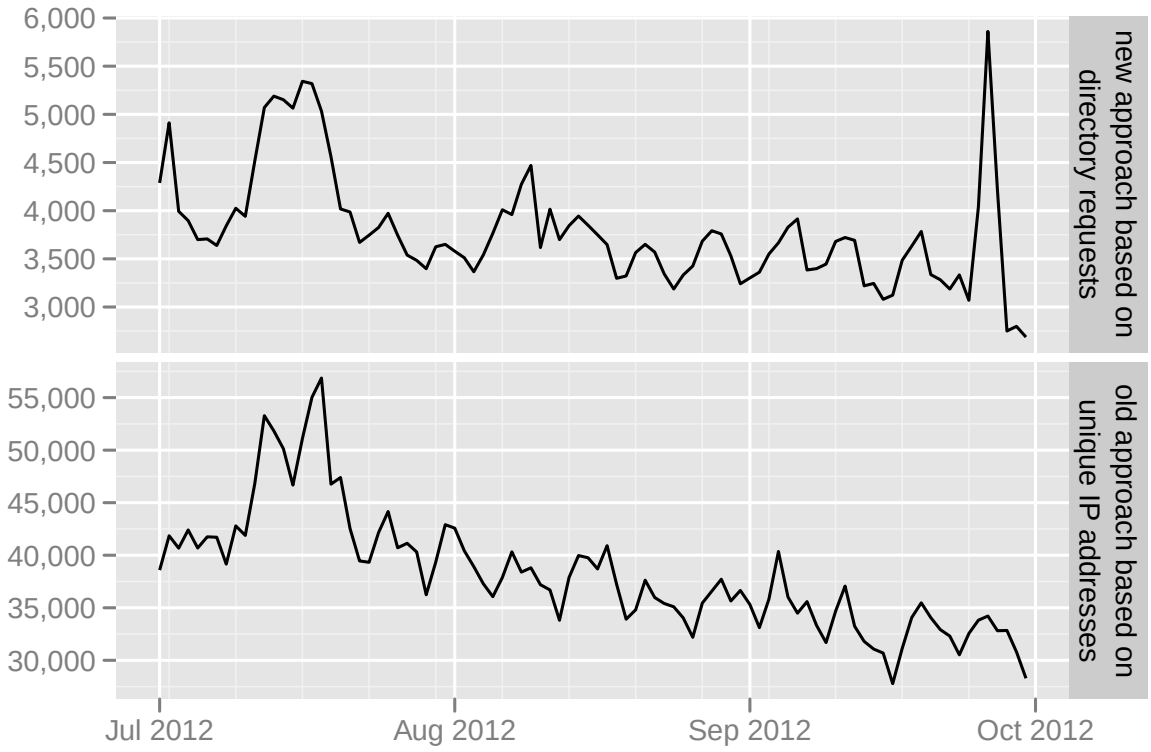

Figure 10: Estimated daily bridge users in the new and in the old approach in the third quarter of 2012

- We should further analyze problems with the consensus process and how they affect directory request numbers. Both statistics of daily directly connecting and of daily bridge users could benefit from new insights here.

Acknowledgements

Thanks to George Kadianakis for his valuable input on a draft of this report, especially by reviewing the math part of extrapolating reported requests to total requests in the network, and for pointing out the potential privacy problem of too short byte history intervals.

References

[1] Sebastian Hahn and Karsten Loesing. Privacy-preserving ways to estimate the number of Tor users. Technical Report 2010-11-001, The Tor Project, November 2010. https://research.torproject.org/techreports/countingusers-2010-11-30.pdf.

[2] Karsten Loesing. What fraction of our bridges are not reporting usage statistics? Technical Report 2012-04-001, The Tor Project, April 2012. https://research.torproject.org/techreports/bridge-report-usage-stats-2012-04-30.pdf.

[3] Karsten Loesing, Steven J. Murdoch, and Roger Dingledine. A case study on measuring statistical data in the Tor anonymity network. In *Proc. Workshop on Ethics in Computer Security Research*, Tenerife, Canary Islands, Spain, January 2010. https://metrics.torproject.org/papers/wecsr10.pdf.