Measuring the Tor Network

Alternative Requirements for Relay Flags

Sebastian Hahn, Karsten Loesing, Steven J. Murdoch

April 11, 2009

Abstract

This document describes the simulation results of alternative requirements for relays to obtain the `Fast`, `Stable`, and `Guard` flags. The simulation is based on the directory archives of descriptors between January 2008 and February 2009. All scripts and a howto for performing the evaluation can be found under: `git://git.torproject.org/~karsten/git/metrics/`

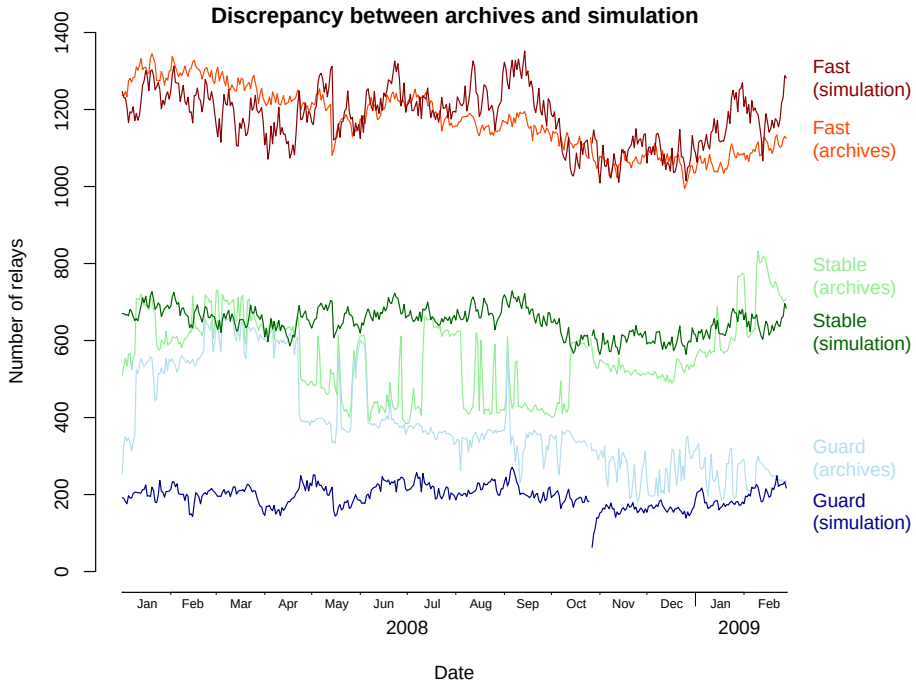

Figure 1: This first graph shows the inevitable discrepancy between simulation results and reality. There are at least three sources of error: First, only a sample of 1/8 of all relays was considered in the simulation for performance reasons. Second, the simulation is limited to relay uptimes as referenced from hourly snapshots which are more coarse-grained than continuous connectivity information. Third, the number of relays with `Stable` and `Guard` flags in the archive data varies by up to 200 relays due to a bug in the consensus process which is not contained in the simulation.

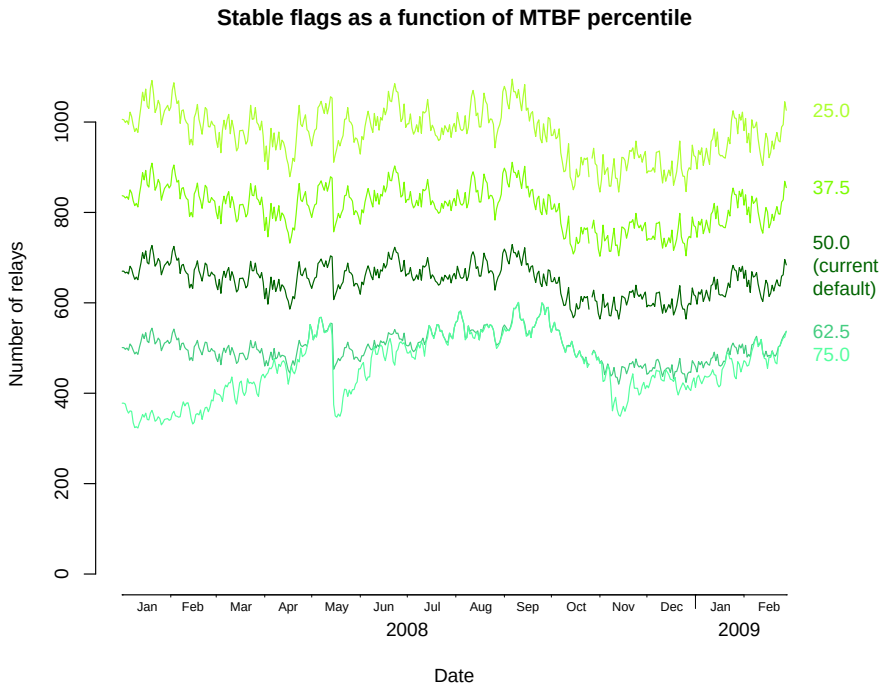

Figure 2: One requirement for being assigning the `Stable` flag is a relay's weighted mean time between failures (MTBF). A relay gets the `Stable` flag if it has a higher MTBF than the median of all active relays. The graph shows alternative percentiles between 25 and 75%. Obviously, if the requirement is relaxed to having a MTBF higher than only 25% of all relays, the number of `Stable` flags increases; if the requirement is raised to 75%, the number of `Stable` flags decreases. The lines for 62.5 and 75% overlap in parts, because of a fixed MTBF limit of 5 days which some relays reach even though not being in the top 25% of all relays.

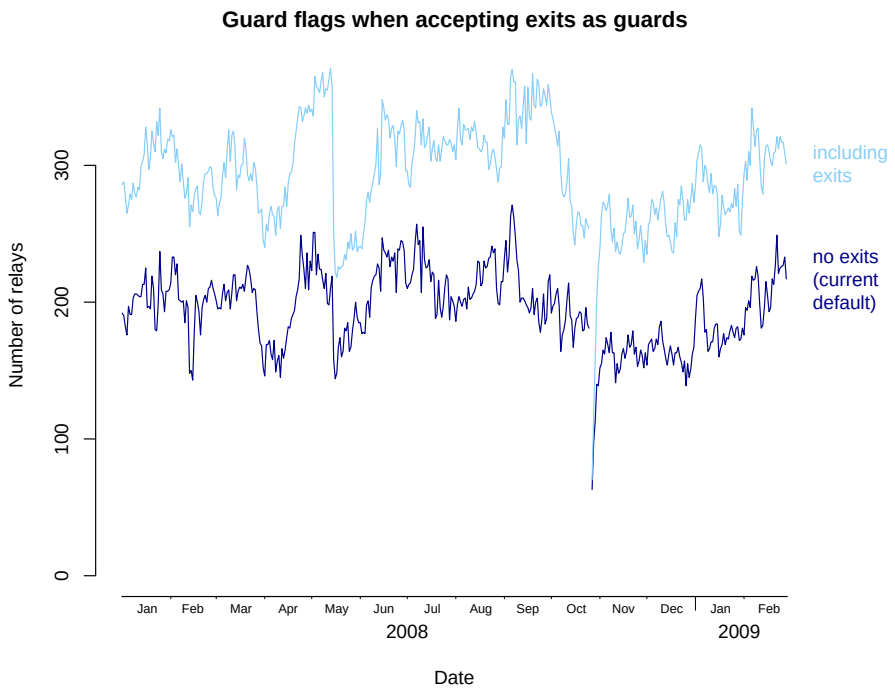

Figure 3: Currently, a relay cannot be assigned both the `Guard` and the `Exit` flag. The reason for the two sets being distinct is the attempt not to overload the rare exit nodes with guard traffic. Unsurprisingly, if this requirement was dropped, the number of relays with the `Guard` flag would increase significantly.

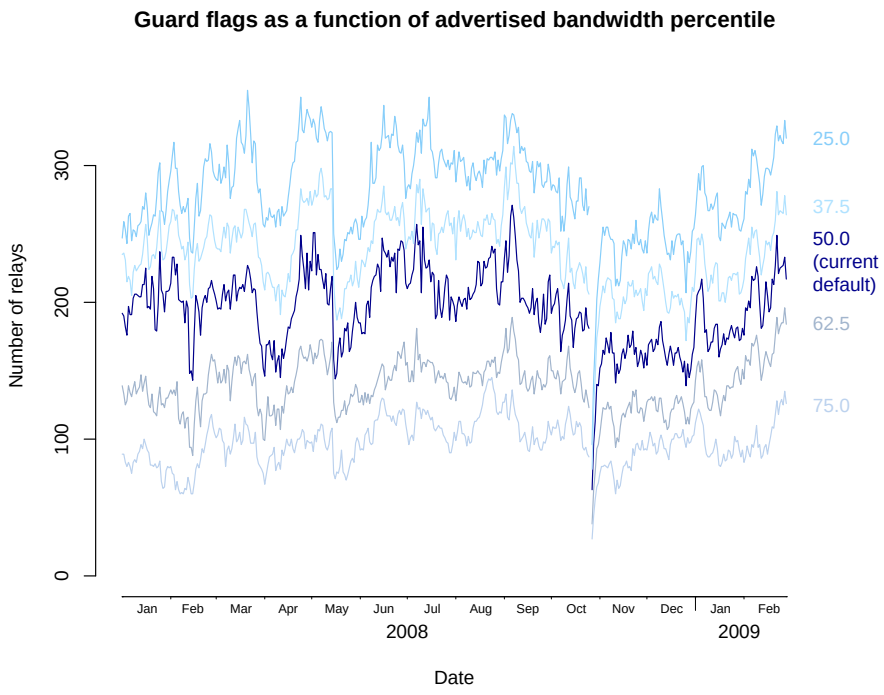

Figure 4: One requirement for obtaining the `Guard` flag is having an advertised bandwidth that is at least the median of all "familiar" relays (familiar means being in the set of 7/8 of relays with highest weighted time known). If the requirement of having an advertised bandwidth of the median of familiar relays is relaxed to the advertised bandwidth only 25% of familiar relays, the number of `Guard` flags increases. Likewise, requiring 75% leads to decrease of `Guard` flags.

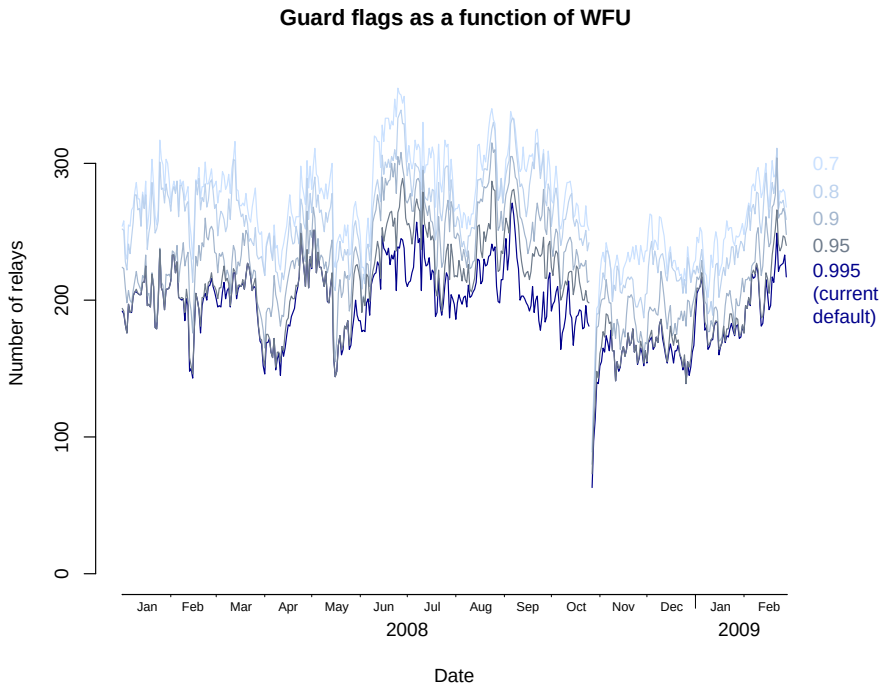

Figure 5: Another requirement for the `Guard` flag is a weighted fractional uptime of at least 0.995, i.e., that a relay was available for at least 99.5% of the time it is known to a directory. If this requirement is relaxed to 0.95, 0.9, 0.8, or even 0.7, the number of `Guard` flags increases as shown in this graph.