Freie Universität  Berlin

Last-Minute Paper Submissions, Forgotten Passwords and Greylisting – an Interesting Dilemma, and How to Solve It

Philipp Reinecke and Katinka Wolter
{philipp.reinecke, katinka.wolter}@fu-berlin.de

Scenario

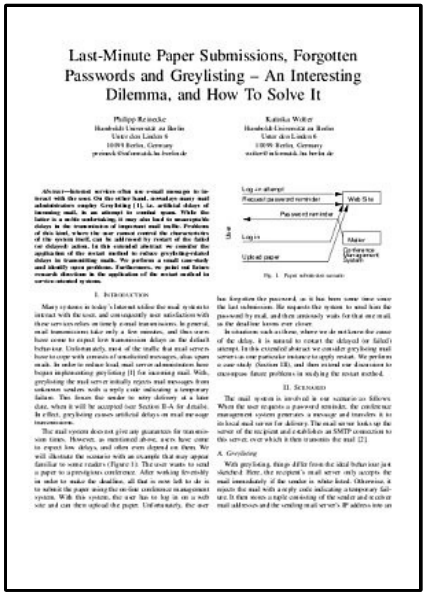

Last-Minute Paper Submissions, Forgotten Passwords and Greylisting – An Interesting Dilemma, and How To Solve It

Philipp Reinecke
Humboldt-Universität zu Berlin
Unter den Linden 6
10099 Berlin, Germany
preinecke@informatik.hu-berlin.de

Katinka Wolter
Humboldt-Universität zu Berlin
Unter den Linden 6
10099 Berlin, Germany
wolter@informatik.hu-berlin.de

Abstract—Internet services often use e-mail messages to interact with the user. On the other hand, nowadays many mail administrators employ Greylisting [1], i.e. artificial delays of incoming mail, in an attempt to combat spam. While the latter is a noble undertaking, it may also lead to unacceptable delays in the transmission of important mail traffic. Problems of this kind, where the user cannot control the characteristics of the system itself, can be addressed by restart of the failed (or delayed) action. In this extended abstract we consider the application of the restart method to reduce greylisting-related delays in transmitting mails. We perform a small case-study and identify open problems. Furthermore, we point out future research directions in the application of the restart method in service-oriented systems.

I. Introduction

Many systems in today's Internet utilize the mail system to interact with the user, and consequently user satisfaction with these services relies on timely e-mail transmission. In general, mail transmissions take only a few minutes, and thus users have come to expect low transmission delays as the default behaviour. Unfortunately, most of the traffic that mail servers have to cope with consists of unsolicited messages, also spam mails. In order to reduce load, mail server administrators have begun implementing greylisting [1] for incoming mail. With, greylisting the mail server initially rejects mail messages from unknown senders with a reply code indicating a temporary failure. They lessen the sender to retry delivery at a later date, when it will be accepted (see Section II-A for details). In effect, greylisting causes artificial delays on mail message transmissions.

The mail system does not give any guarantees for transmission times. However, as mentioned above, users have come to expect low delays, and often even depend on them. We will illustrate the scenario with an example that may appear familiar to some readers (Figure 1). The user wants to send a paper to a prestigious conference. After working feverishly in order to make the deadline, all that is now left to do is to submit the paper using the on-line conference management system. With this system, the user has to log in on a web site and can then upload the paper. Unfortunately, the user has forgotten the password, as it has been some time since the last submission. He requests the system to send him the password by mail, and then anxiously waits for that mail as the deadline looms ever closer.

In situations such as these, where we do not know the cause of the delay, it is natural to restart the delayed (or failed) attempt. In this extended abstract we consider greylisting mail servers as one particular instance to apply restart. We perform a case-study (Section III), and then extend our discussion to encompass future problems in studying the restart method.

II. Scenario

The mail system is involved in our scenario as follows. When the user requests a password reminder, the conference management system generates a message and transfers it to its local mail server for delivery. The mail server looks up the server of the recipient and establishes an SMTP connection to this server, over which it then transmits the mail [2].

A. Greylisting

With greylisting, things differ from the ideal behaviour just sketched. Here, the recipient's mail server only accepts the mail immediately if the sender is white listed. Otherwise, it rejects the mail with a reply code indicating a temporary failure. It then stores a triple consisting of the sender and receiver mail addresses and the sending mail server's IP address into an

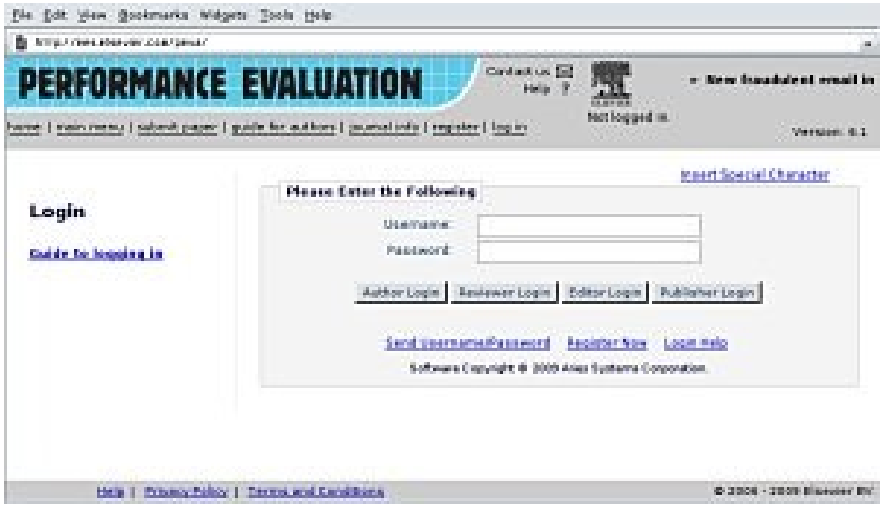

Scenario: How it is supposed to be

Author

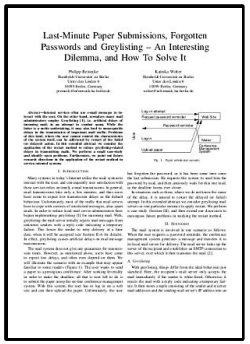
Last-Minute Paper Submissions, Forgotten Passwords and Greylisting – An Interesting Dilemma, and How To Solve It

Request Reminder

Reminder

Login

Submission System

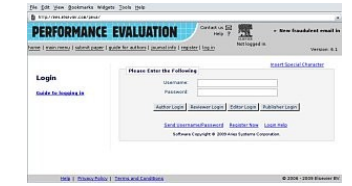

Scenario: How it *is* with Greylisting

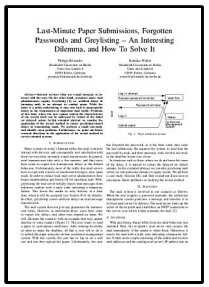
Last-Minute Paper Submissions, Forgotten Passwords and Greylisting – An Interesting Dilemma, and How To Solve It

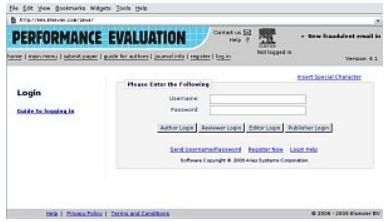

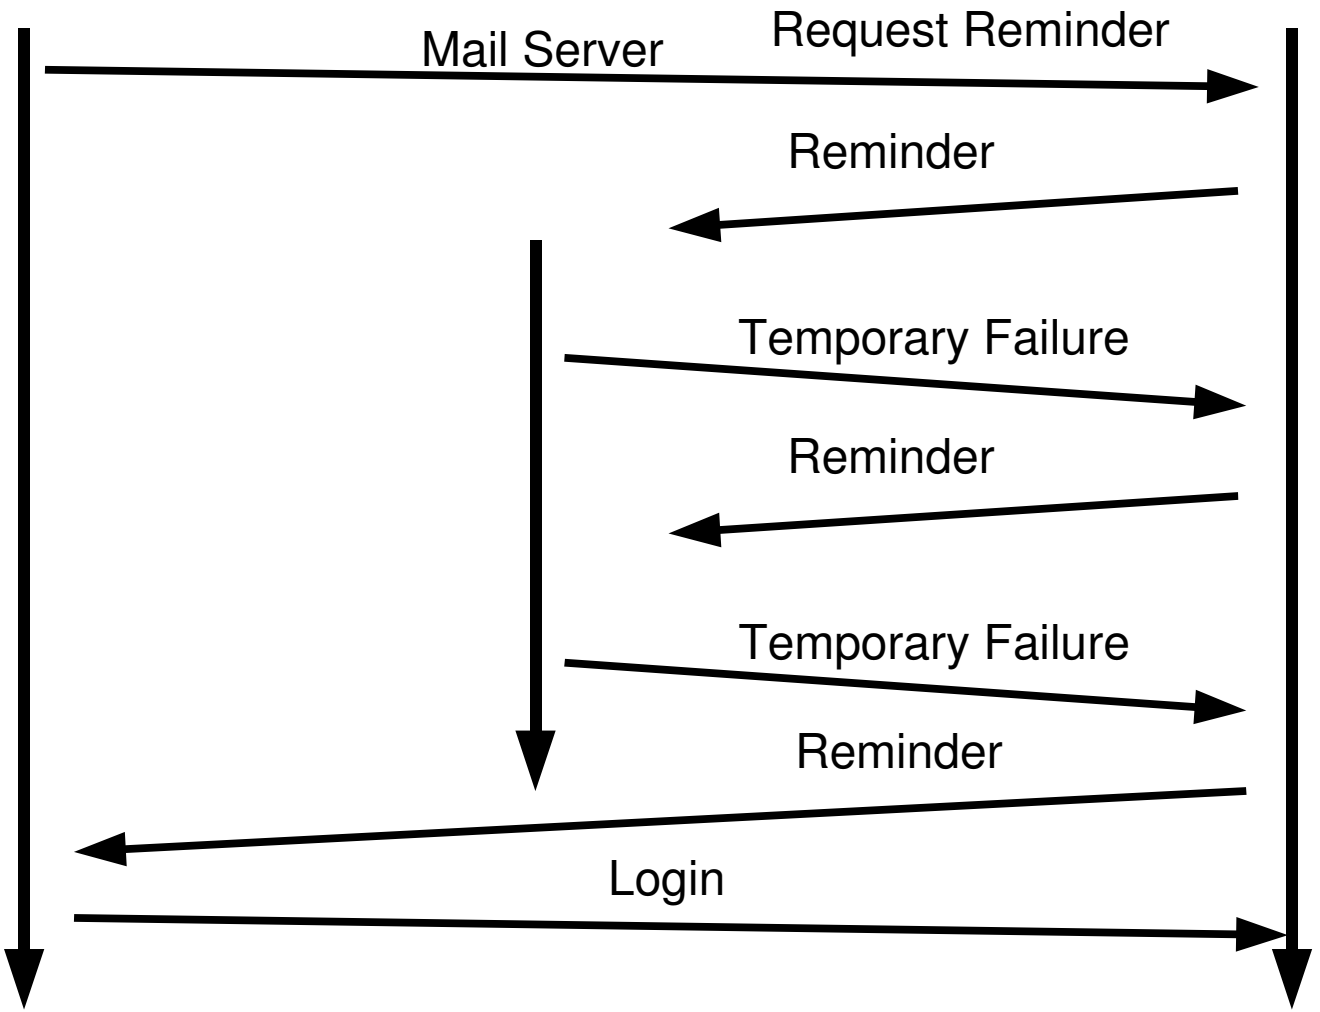

Scenario: What you see

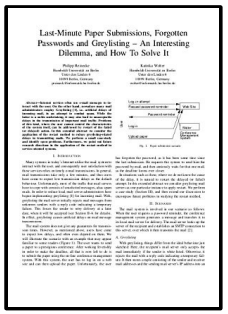
Last-Minute Paper Submissions, Forgotten Passwords and Greylisting – An Interesting Dilemma, and How To Solve It

Author

Request Reminder

Submission System

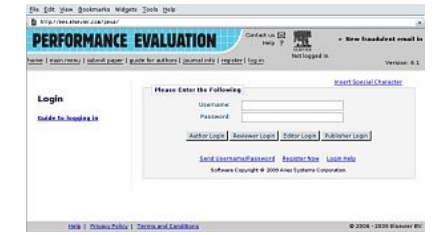

Deadline

Reminder

Login

Problem Statement

- Greylisting introduces delays into reminder mail delivery
- Details are hidden from the user
- Is it sensible to repeat the password reminder request (i.e. should you restart the process)?
- And if so, when?
- Criterion for applicability of restart: High variability (SCV)
- Optimal timeout: Computed using algorithm from [vMW04, vMW06]

Case-Study: Large Educational Institution

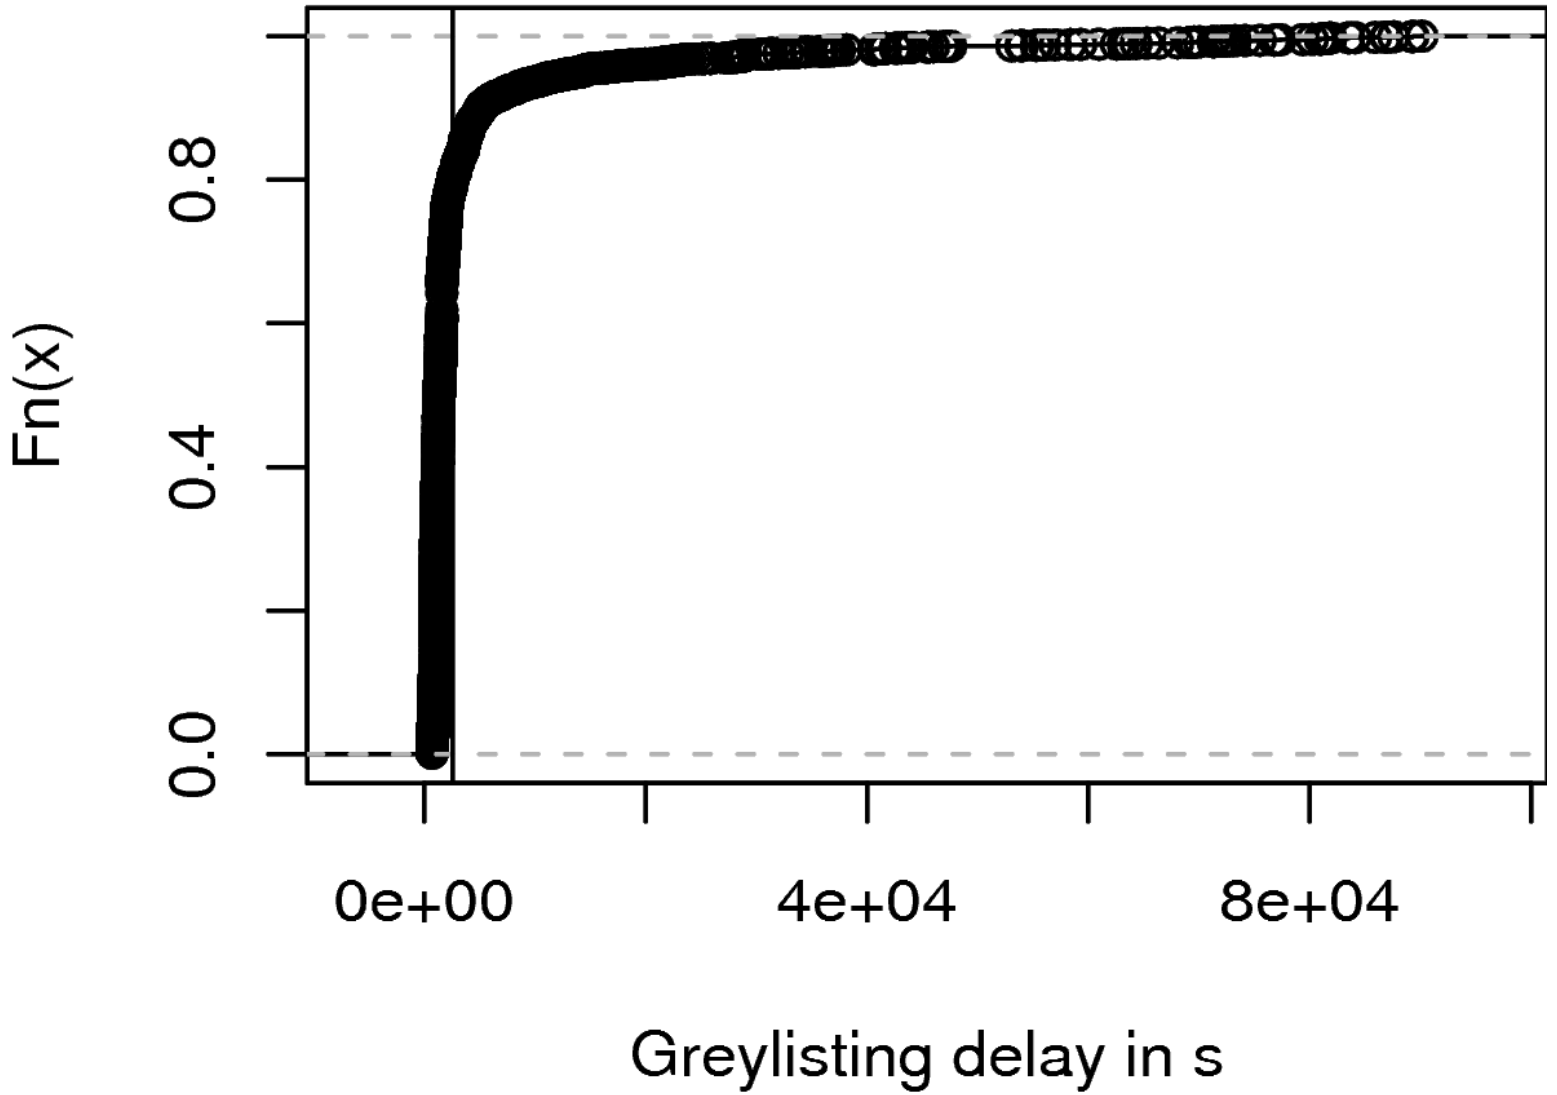

Case-Study: Large Educational Institution

- Parse X-Greylist header of received mails for greylisting delays
- After removal of outliers: 3692 samples (mails from the period 24 October 2007-11 June 2009)
- Minimum 720s (12 min), maximum 90131s (~25 h)
- Mean 3784.3 s (~1 h), median 1353 s (~22 min)
- SCV 6.6
- Optimal timeout: 2592 s (43 min)

Models for Analysis

- Model M1:
  - Single random variable for all observed mail servers
  - Restart may be handled by different server
- Model M2:
  - One random variable X for each mail server
  - Restart is handled by the same server
- Model M3:
  - Cluster servers by behaviour
  - Weighted choice for initial attempt, restart handled by server from same cluster

M2: Split Data by Sender

- 1151 individual senders (by sender address)
- only 4 sent more than 100 messages

Sender 1149

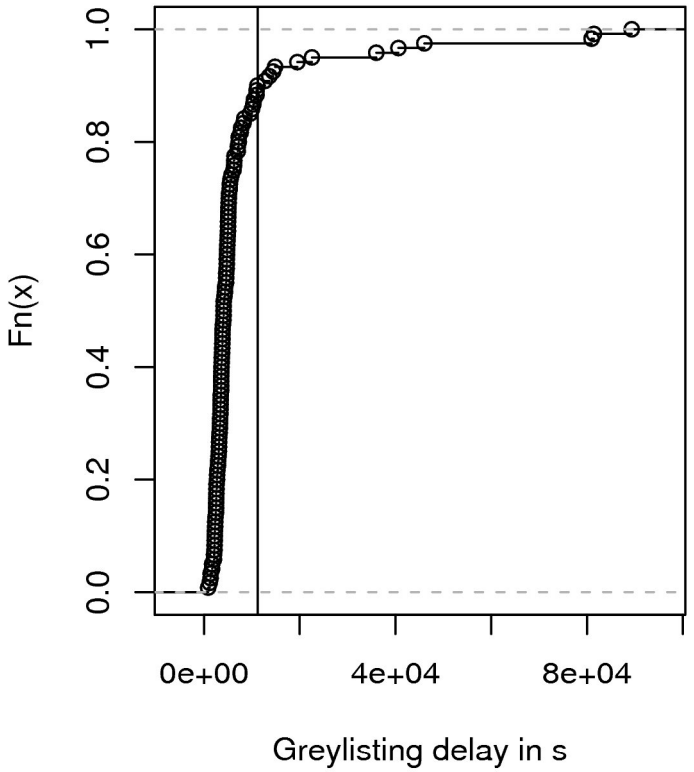

- 120 Samples
- Mean: 2:12h
- SCV: 3.1
- Timeout: 3:07h

Sender 1150

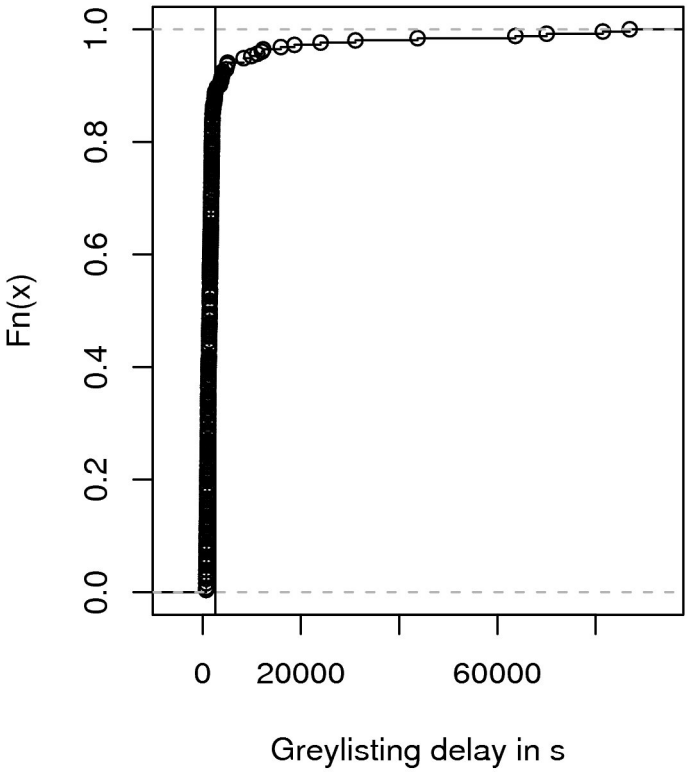

- 253 Samples
- Mean: 56min
- SCV: 8.9
- Timeout: 43min

M3: Split Data by Clustering

- Cluster senders by mean greylisting delay
- Base probability of choosing a cluster on cluster size

Cluster 3

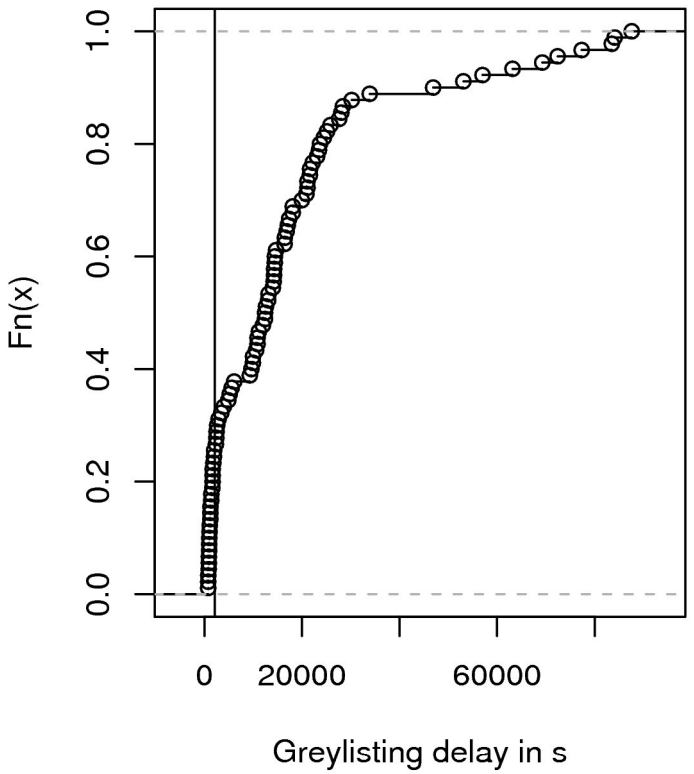

- 90 Samples
- α=0.024
- Mean: 4:54h
- SCV: 1.4
- Timeout: 34 min

Cluster 4

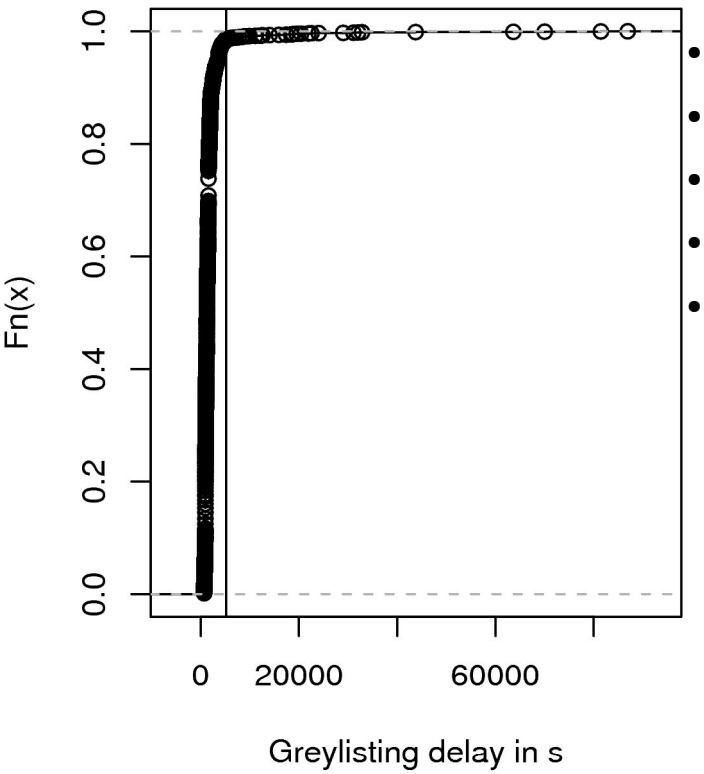

- 2820 Samples
- α=0.896
- Mean: 29min
- SCV: 3.9
- Timeout: 1:26h

Results and Observations

- M1: Restart seems applicable (SCV = 6.6)
- M2 and M3: Restart applicable with some senders
- Optimal timeouts vary between models and depending on senders
- M2 and M3 may be more appropriate for the scenario, but suffer from data scarcity

Future Work

- Study correlation of samples: Difficult, because spammers might be the only ones to send mail in this manner
- Obtain more data for analysis using models M2 and M3

Any volunteers?

- Try to test the approach in practice
- Apply the concepts for restart in SOA systems, where on-line algorithms also suffer from data scarcity

Fin.