

Table 7: Effect of pre-acceptance responses on good and junk mail. Data from server1.

type	number	percentage
good delayed by 4 hours	33,021	13% of good
good delayed by 12 hours	34,899	13% of good
junk mails rejected	565,922	95% of junk

5 Testing the responses

This section provides some evidence for how much effect the responses described above would have on reducing the amount of spam processed, and on reducing the effect of the volume of junk mail on the flow of good mail through a mail server.

Firstly, we consider the effect of temporarily failing new servers and servers predicted to send junk. Unfortunately it is very difficult to test this just using log data, as the response is to request the remote sending server to retry later, and it is difficult to predict the behavior of the sending server.

A best-case estimate would be to assume that spammers and virus-infected machines do not retry, but that good senders do. Thus out of 855,228 messages accepted by server 1, if a 4 hour delay was used for new senders and a 12 hour delay for predicted junk mail, the effect would be as in Table 7. Only 26% of good mail would be delayed, and as discussed above, a significant proportion of this mail is likely to be misclassified junk. If this sort of system were widely deployed, it is likely that spammers would implement retries. This would cause the amount of junk mail rejected to decline considerably.

It is much easier to predict the effect of post-acceptance responses. We do this by constructing a model of a mail server that allows us to calculate the time taken to process a mail message, under different amounts of loading. We can then alter that model to incorporate prioritization schemes and predict the final performance of the system.

The initial model of the system is shown in Figure 6 (a). This is a generic model of a mail server that includes a mail scanner or filter. The incoming mail is handled by an SMTP process that writes the mail to a local disk or spool q_{MS} , taking time t_{IN} . The mail is then loaded, scanned and placed in a second spool q_{OUT} by the mail scanner, marking the mail accordingly (normally by writing a header), and taking on average t_{SCAN} . The mail is then taken from this second spool and delivered

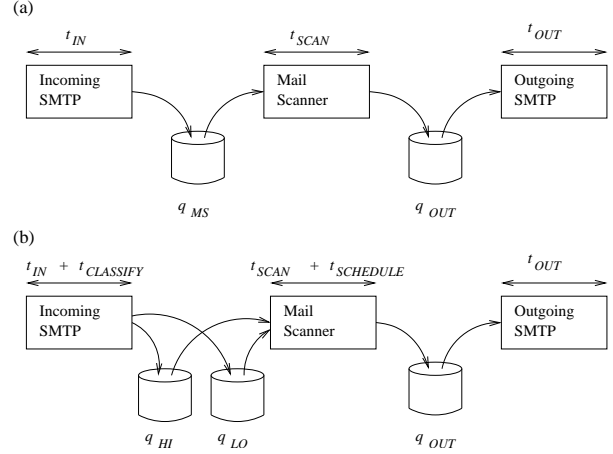


Figure 6: (a) Model of basic mail server. Incoming mail is handled by an SMTP process, before being scanned for viruses and spam by the mail scanner, and being delivered using a second SMTP process. (b) Model of mail server with prioritization. The incoming process places mail in one of two queues depending on the predicted message type. The mail scanner selects messages from the queues using a scheduling algorithm.

by the outgoing SMTP process, taking t_{OUT} . The overall time will be $t_{IN} + t_{SCAN} + t_{OUT}$ plus the time that mails spend on the queues waiting to be processed. As each of the spools is effectively a queue, this system can be analyzed using Queueing Theory [10].

Prioritization can be implemented by having two queues for the mail scanner, one for high priority or good mail (q_{HI}) and the other for low priority or junk mail (q_{LO}). When a mail arrives it is classified into good or junk (this extra processing taking $t_{CLASSIFY}$), and placed into one of these queues. The mail scanner then uses a simple scheduling algorithm to select which message to process next. The simplest algorithm is “absolute priority”, where the scanner always takes from the high priority queue unless it is empty, when it services the low priority queue [19]. The scheduling operation is modelled as taking $t_{SCHEDULE}$. As before, the total time is the sum of the individual times plus the time spent queuing.

We simulated both models using the DEMOS system performance modelling tool [4]. This allows synthetic traffic to be “played” through the model and the system performance evaluated. The parameters of the model (the various service times t_{IN} , t_{SCAN} , t_{OUT}) were estimated from the log data. These parameters are difficult to estimate, because the times in the logs include mes-