Data Packet Loss in a Queue with Limited Buffer Space

Jose Gil, Motorola Research

Background: A system that occurs in several contexts in computer and data networks has the generic form shown in Figure 1. There are N sources, and each source n transmits bursts of data at speed R_n packets
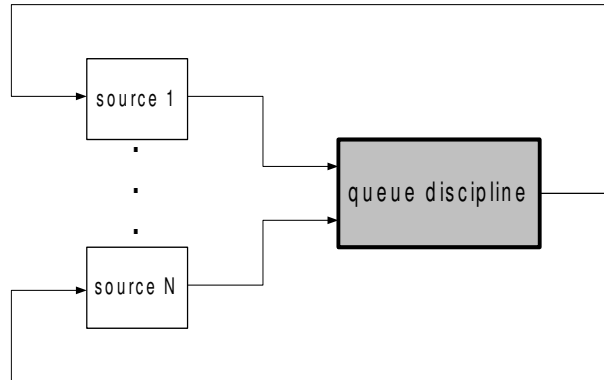

Figure 1: Diagram of system

a second, and they share a medium to which access is controlled by a queueing discipline. The queue has a limited buffer space B, and so in some conditions packets are dropped. The system feeds back to source n the proportion p_n of its packets that were dropped, and the source then adjusts its sending rate according to the formula

$$R_n(p_n) = \frac{K}{\mathrm{RTT}_n \sqrt{C p_n + D}}. \tag{1}$$

Here K, C and D are to be treated as known dimensionless constants, and RTT_n is the round-trip-time of the packets, from the time source n sends them to the time it receives the feedback that a proportion p_n of them were dropped.

The problem is to find the equilibrium point $(R_n(p_n), p_n, \mathrm{RTT}_n)$, and also information about the statistics of the distribution, *e.g.* the variance of the packet loss distribution.

Detailed models:

Sources: There are 2 cases to consider:

1. The N sources *always* have data to send.

2. The N sources each switch between ON/OFF according to independent Markov chains, and when source n is ON, it sends data in busts at rate R_n.

Packet discard policies: There are 2 packet discard policies that might be in operation:

1. Drop-tail: this simply means that if a packet arrives when the buffer is full then it, and any subsequent packets of the same burst, are dropped.

2. Random Early Discard (RED): If the buffer occupancy exceeds some threshold B_t, $(B_t < B)$, then arriving packets are dropped with a probability varying linearly between 0 at B_t and 1 at B, different packets being independent.

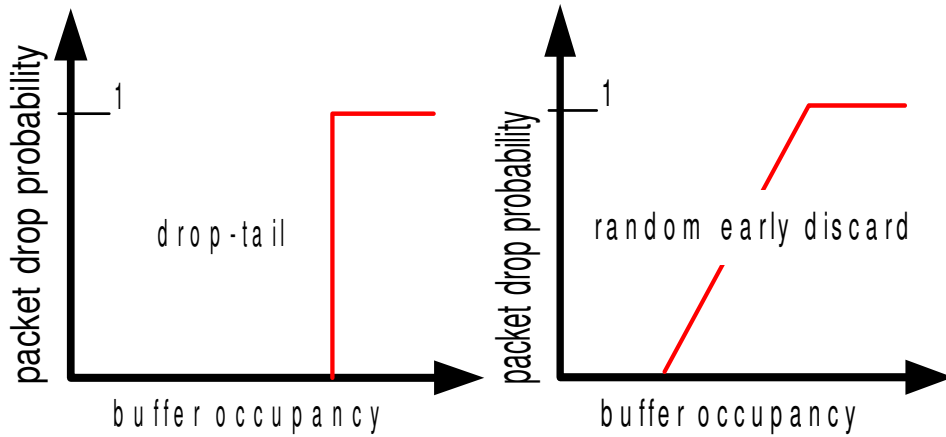The graphs of packet drop probability against buffer occupancy in these 2 cases are illustrated in Figure 2.

Figure 2: Packet loss probability for the two packet discard policies

Queueing disciplines: There are 3 queueing disciplines of interest:

1. FIFO: First in—first out. This is the simplest case in which all packets form a single queue and are dealt with in order of arrival.

2. Round Robin: Here, there are N queues, one for the packets from each source. Packets are serviced one from each queue in turn, in a fixed cyclic order, skipping any queues that are empty.

3. Giving priority to flows with the least service time: In this case, the queue manager system has information about the average transmission rates for the packets from each source, and gives some weighted priority to those packets that can be transmitted fastest.

The queueing patterns in the first 2 cases are illustrated in Figure 3.
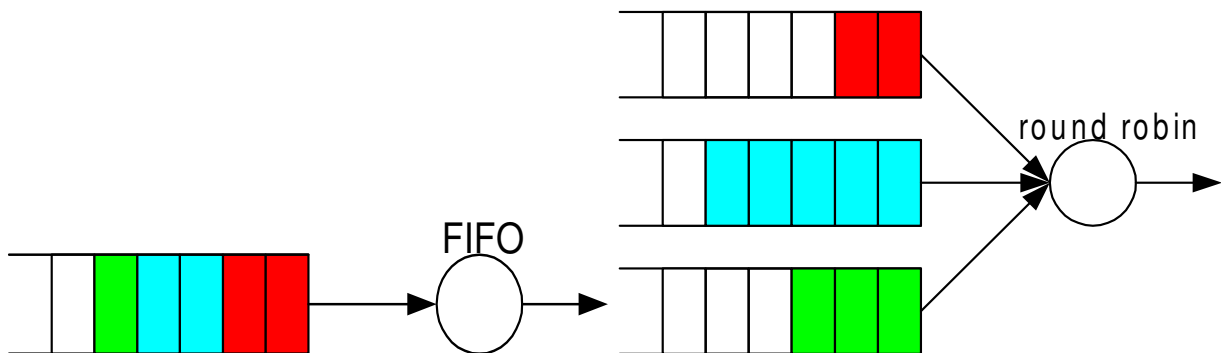

Figure 3: FIFO and round robin queueing systems.

Service times: The service time per data packet is random, and for the purposes of this study it can be treated as exponentially distributed with a known mean service rate μ_n, different packets being independent.

Priorities for the Study Group: The combinations of cases, in order of most interest, are

1. Drop-tail and FIFO

2. Drop-tail and Round Robin

3. Random Early Discard and FIFO

4. Random Early Discard and Round Robin