# Shaping Packet Delay in Backpressure Stacks Using LIFO Service Priority

Scott Moeller and Bhaskar Krishnamachari
ANRG Tech Report
Viterbi School of Engineering
The University of Southern California

February 2, 2010

## 1 Problem Statement

We will consider some arbitrary network of queues $\mathbf{Q}$, whose backlogs update as a result of control decisions (e.g. routing, forwarding between queues). For any particular queue in this network, $Q_q(t) \quad q \in \mathbf{Q}$, if there is some minimum recurrent backlog, then we claim that the waiting time for packets serviced by this queue is shaped through LIFO service priority; the result of which at low loading is unboundedly better delay performance for serviced packets when compared to FIFO discipline.

In order to quantify the potential delay advantages of LIFO over FIFO within backpressure routing, we will consider solutions using the performance optimal Lyapunov networking framework and under general penalty functions that in this framework result in link usage penalties. Let $\overline{\mathcal{R}}_{i\to j}(t)$ be the link rate, $\Delta\mathcal{Q}_{i,j}(t)$ the queue differential between queues $i$ and $j$, $\theta_{i\to j}(t)$ be some link utilization penalty, all varying with t. Let V be a constant trading average queue backlog for link penalty minimization. We then have the following link usage weights which will be maximized across the network by the backpressure stack:

$$w_{i,j} = (\Delta\mathcal{Q}_{i,j}(t) - V \cdot \theta_{i\to j}(t)) \cdot \overline{\mathcal{R}}_{i\to j}(t) \tag{1}$$

This setting is consistent with our backpressure collection protocol (BCP [?]), where $\theta_{i\to j}$ was the link Expected Transmit Count (ETX) which is greater than or equal to 1. It can easily be shown then that for any queue, the minimum recurrent backlog is at least equal to the node's minimum hop count from the sink.

## 2 Delay Analysis

We will now formalize the concept of delay shaping in the context of stochastic network optimization.

**Definition** Queue $U_i(t)$ is defined as having a *stabilized permanent backlog* $b_i^{min} > 0$ if there exists a $t^*$ such that for all $t \geq t^*$, $\mathcal{Q}_i(t) \geq b_i^{min}$ and there exists an infinite sequence of time slots $t^* \leq t_0 \leq t_1 \leq \cdots$ for which $\mathcal{Q}_i(t_j) = b_i^{min}$. Formally:

$$\left[\exists t^* \text{ s.t. } \forall t \geq t^* \quad \mathcal{Q}_i(t) \geq b_i^{min}\right] \cap \left[\exists\, (t^* \leq t_0 \leq t_1 \leq \cdots) \text{ for which } \mathcal{Q}_i(t_j) = b_i^{min}\right] \tag{2}$$

**Definition** The *average delivered packet delay* is defined as the average delay for packets passing through a queue but not trapped indefinitely within.

It is not useful to consider average delivered packet delay for arbitrary queueing systems, as without stability the volume of indefinitely trapped packets may grow to infinity. Within the context of a stabilized permanent backlog queue operating with LIFO service priority, however, this metric is meaningful. Though the *average packet delay* through such a queue is unchanged, we can improve the *average serviced packet delay* by permanently trapping and effectively discarding $b_i^{min}$ packets.

**Theorem 2.1.** *(The LIFO Delay Advantage for Constantly Backlogged Queues) Let $\mathcal{Q}_i(t)$ be a queue with stabilized permanent backlog $b_i^{min}$ and arrival rate $\lambda_i$. Then the time average delivered packet delay relationship under FIFO ($\overline{W}_{FIFO}$) and LIFO ($\overline{W}_{LIFO}$) queueing disciplines is exactly:*

$$\overline{W}_i^{FIFO} = \overline{W}_i^{LIFO} + \frac{b_i^{min}}{\lambda_i} \tag{3}$$

*Proof.* $\mathcal{Q}_i(t)$ has stabilized permanent backlog $b_i^{min} \implies$ (2) holds.

**Case LIFO:**
Under a LIFO discipline, any data arriving to find backlog greater than or equal to $b_i^{min}$ will be emptied infinitely often per (2). The oldest $b_i^{min}$ packets within the LIFO queue at time $t^*$ are trapped indefinitely, and therefore are not considered in calculation of average delivered packet delay. Beyond time $t^*$, the average delivered packet delay of the LIFO queue is therefore equivalent to the average packet delay of a LIFO queue operating with the oldest $b_i^{min}$ packets removed. Let $\overline{N}_i^{LIFO}$ be the time average number of packets in LIFO queue i after removal of the $b_i^{min}$ trapped packets.

**Case FIFO:**
Under a FIFO discipline, the average delivered packet delay is always equal to the average packet delay, as every arriving packet is eventually serviced. Let $\overline{N}_i^{FIFO}$ be the time average number of packets in FIFO queue i.

As a result of the LIFO queue discipline, and the stabilized permanent backlog, we then find:

$$\begin{aligned}\overline{N}_i^{FIFO} = \overline{N}_i^{LIFO} + b_i^{min} &\implies \frac{\overline{N}_i^{FIFO}}{\lambda_i} = \frac{\overline{N}_i^{LIFO}}{\lambda_i} + \frac{b_i^{min}}{\lambda_i} \\ &\implies \overline{W}_i^{FIFO} = \overline{W}_i^{LIFO} + \frac{b_i^{min}}{\lambda_i} \qquad \text{(Little's Applied Twice)}\end{aligned}$$

Where in the final step we use the fact that $\overline{N}_i^{FIFO}$ is serviced with FIFO service priorty and that the modified LIFO queue empties infinitely often, therefore Little's Theorem applies for both queues. □