
ATM Forum Document Number: ATM_Forum/98-0885R1

TITLE: Proposal to move the text on buffer management scheme for GFR from living list to baseline

SOURCE: Rohit Goyal, Raj Jain, Sonia Fahmy, Bobby Vandalore, Mukul Goyal
The Ohio State University,
Department of Computer and Information Science,
2015 Neil Ave, DL 395, Columbus, OH 43210
Phone: 614-688-4482
{goyal,jain}@cis.ohio-state.edu

This work was sponsored in part by NASA Lewis Research Center

DISTRIBUTION: ATM Forum Technical Committee
Traffic Management Working Group

DATE: February 1999, Atlanta

ABSTRACT: This contribution proposes to move the GFR text labeled Section VII.2.3 from the living list to the baseline text.

NOTICE: This document has been prepared to assist the ATM Forum. It is offered as a basis for discussion and is not binding on the contributing organization, or on any other member organizations. The material in this document is subject to change in form and content after further study. The contributing organization reserves the right to add, amend or withdraw material contained herein.

1 Motion

The text in the living list GFR section labeled VII.2.3, be moved to the baseline text.

The modified baseline text sections will look like the following:

VII.2.1 GFR Implementation using Weighted Fair Queuing and per-VC accounting

(Unchanged)

VII.2.2 GFR Implementation Using Tagging and FIFO Queue

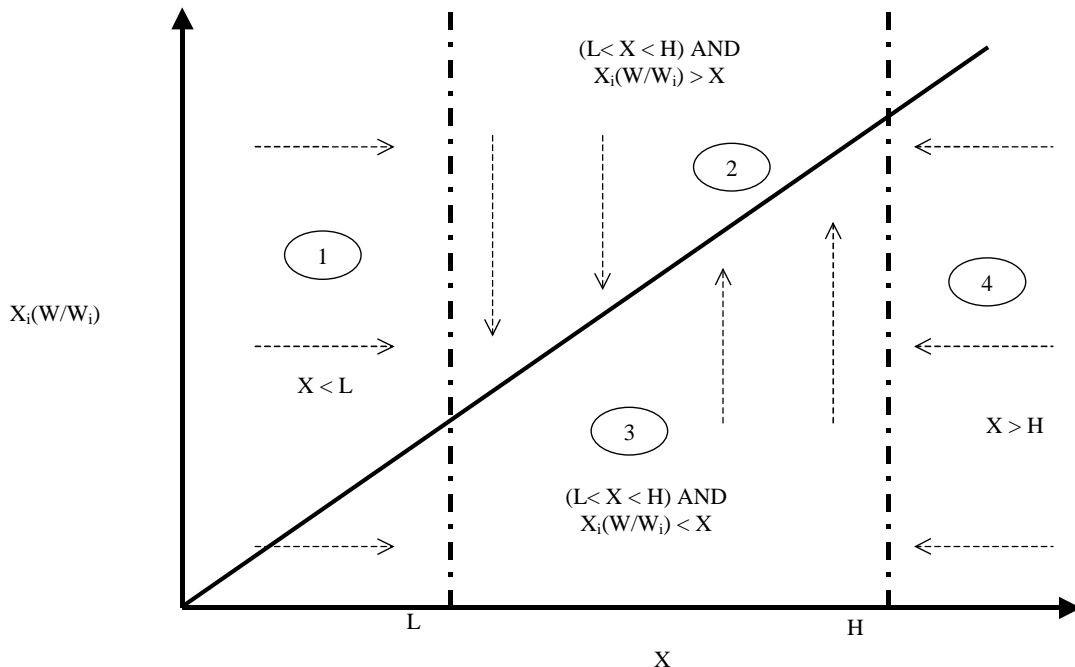
(Unchanged)

VII.2.3 GFR Implementation Using Differential Fair Buffer Allocation

Differential Fair Buffer Allocation (DFBA) uses the current queue length as an indicator of network load. The scheme tries to maintain an optimal load so that the network is efficiently utilized, yet not congested. In addition to efficient network utilization, DFBA is designed to allocate buffer capacity fairly amongst competing VCs. This allocation is proportional to the MCRs of the respective VCs. The following variables are used by DFBA to fairly allocate buffer space:

- X = Total buffer occupancy at any time
- L = Low buffer threshold
- H = High buffer threshold
- MCR_i = MCR guaranteed to VC_i
- W_i = Weight of $VC_i = MCR_i / (\text{GFR capacity})$
- $W = \sum W_i$
- X_i = Per-VC buffer occupancy ($X = \sum X_i$)
- Z_i = Parameter ($0 \leq Z_i \leq 1$)

DFBA tries to keep the total buffer occupancy (X) between L and H . When X falls below L , the scheme attempts to bring the system to efficient utilization by accepting all



incoming packets. When X rises above H , the scheme tries to control congestion by performing EPD. When X is between L and H , DFBA attempts to allocate buffer space in proportional to the MCRs, as determined by the W_i for each VC. When X is between L and H , the scheme also drops low priority (CLP=1) packets so as to ensure proportional buffer occupancy for CLP=0 packets.

The figure above illustrates the four operating regions of DFBA. The graph shows a plot of the current buffer occupancy X versus the normalized fair buffer occupancy for VC_i . If VC_i has a weight W_i , then its target buffer occupancy should be $X \cdot W_i / W$. Thus, the normalized buffer occupancy of VC_i is $X_i \cdot W / W_i$. The goal is to keep this normalized occupancy as close to X as possible, as indicated by the solid line in the graph. Region 1 is the underload region, in which the current buffer occupancy is less than the low threshold L . In this case, the scheme tries to improve efficiency. Region 2 is the region with mild congestion because X is above L . As a result, any incoming packets with CLP=1 are dropped. Region 2 also indicates that VC_i has a larger buffer occupancy than its fair share (since $X_i > X \cdot W_i / W$). As a result, in this region, the scheme drops some incoming CLP=0 packets of VC_i , as an indication to the VC that it is using more than its fair share. In region 3, there is mild congestion, but VC_i 's buffer occupancy is below its fair share. As a result, only CLP=1 packets of a VC are dropped when the VC is in region 3. Finally, region 4 indicates severe congestion, and EPD is performed here.

In region 2, the packets of VC_i are dropped in a probabilistic manner. This drop behavior is controlled by the parameter Z_i , whose value depends on the connection characteristics. This is further discussed below.

The probability for dropping CLP=0 packets from a VC when it is in region 2 depends on several factors. The drop probability has two main components – the fairness component, and the efficiency component. Thus, $P\{drop\} = fn(\text{Fairness component, Efficiency component})$. The contribution of the fairness component increases as the VC's buffer occupancy X_i increases above its fair share. The drop probability is given by

$$P\{drop\} = Z_i \left(\alpha \frac{X_i - X \times W_i / W}{X(1 - W_i / W)} + (1 - \alpha) \frac{X - L}{H - L} \right)$$

The parameter α is used to assign appropriate weights to the fairness and efficiency components of the drop probability. Z_i allows the scaling of the complete probability function based on per-VC characteristics.

The following DFBA algorithm is executed when the first cell of a frame arrives at the buffer.

BEGIN

IF ($X < L$) THEN

 Accept frame

ELSE IF ($X > H$) THEN

 Drop frame

ELSE IF (L < X < H) AND (X_i ≤ X*W_i/W) THEN

Drop CLP1 frame

ELSE IF (L < X < H) AND (X_i > X*W_i/W) THEN

Drop CLP1 frame

Drop CLP0 frame with

$$P\{drop\} = Z_i \left(\alpha \frac{X_i - X \times W_i / W}{X(1 - W_i / W)} + (1 - \alpha) \frac{X - L}{H - L} \right)$$

ENDIF

END

VII.2.4 Evaluation Criteria

(From VII.2.3 in the baseline text document.)