Why Credit Control for ABR Traffic?

H. T. Kung
Harvard University
kung@das.harvard.edu

January 24, 1995

Prepared for:
Debate: Rate-Based vs. Credit-Based Flow Control
ATM Year Three Conference
February 27-March 2, 1995, San Jose, California

Flow Control Problem

- Congestion occurs at a switch when input rates are greater than output rate.
- Sources must be told directly or indirectly to slow down.
- Fundamentally it is a feedback control problem. Feedback delay must be properly handled.
- Technical goals:
  - Zero or very low cell loss rate
  - Instant grabbing of all available bandwidth
  - High utilization
  - Fair share
  - Performance robustness
  - No complex parameters to adjust

Two Traffic Models

1. Smooth traffic
   - Constant, predictable load, or load change in time scales larger than feedback control delay
   - Flow control is easy. No flow control method should ever fail on this traffic.
2. Bursty traffic
   - No predictability: "Fractal" behavior as measured on some data traffic
   - Examples of bursty usages: Interactive Mosaic users, high-level protocol packets, and RPCs
   - No steady state: E.g., most TCP sessions transfer less than a few dozen KBytes and thus require only a few ms transmission time on a OC-3 link.
   - Claim: Credit-based flow control can handle this difficult, real-world traffic well.

Credit Approach

Per VC, Link-by-Link Flow Control

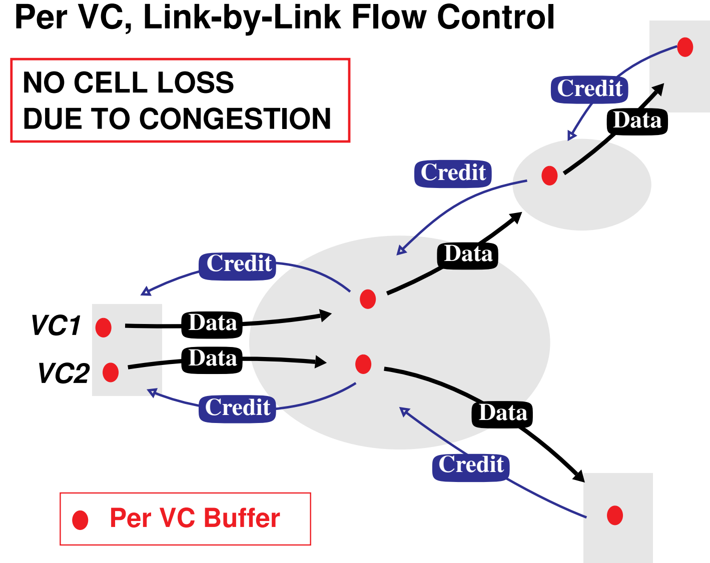

Credit Update Protocol (CUP)

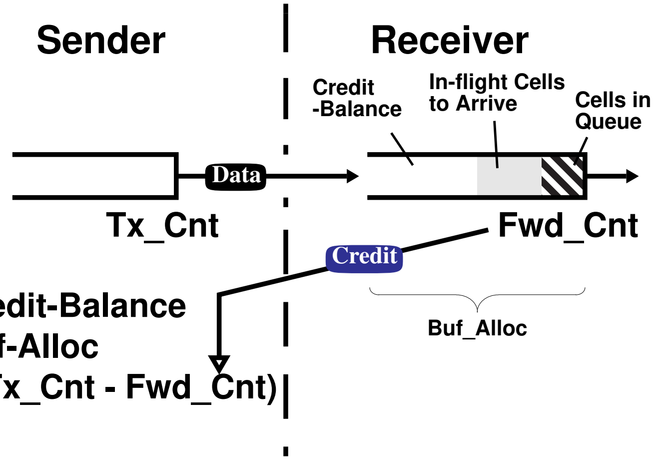

Credit-Balance
= Buf-Alloc
- (Tx_Cnt - Fwd_Cnt)

After having forwarded e.g. 10 data cells of a VC, the receiver transmits a credit record upstream containing the current Fwd_Cnt value of the VC.

Credit Approach Summary

- Link-by-link, per VC, window-based flow control
- CUP is simple, and achieves all the technical goals listed earlier. (It is a lower level and lighter weight sliding-window protocol than X.25 and TCP. E.g., CUP does not retransmit lost data cells and its data cells do not carry sequence numbers.)
- Adaptive buffer allocation can be used to save memory. It is sufficient that Mem = 4*RTT + 2*#VC.
- Validated by extensive simulation tests involving hundreds of VCs over links of 100-fold varying bandwidths and propagation delays

Rate Approach

- Works well as long as rates can be set perfectly or near optimally.
- However, bursty traffic implies that rates can never be set correctly. So adaptive rate setting is necessary.
- Adaptation can not be precise for inherent reasons:
  Incomplete and out-of-date information
  Variable control delay of the adaptation
- Adaptation should not be precise either:
  Rates should be set higher than fair share to allow opportunistic grabbing of available bandwidth (see next slide)
- Since adaptive rate setting *can not* and *should not* be precise, bounding the liability of overrunning switch buffers is the real, first order issue.

Over Allocation Is Necessary for Achieving High Efficiency

Precise Allocation:

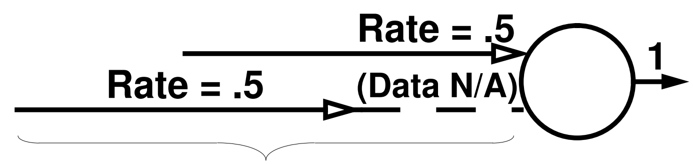

Transmission Time = 3X

Over Allocation:

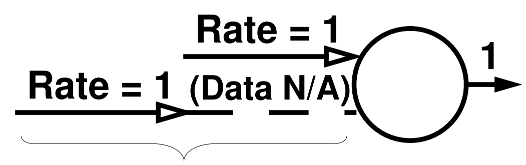

Transmission Time = 2X

Credit Is Inherently Superior in Preventing Buffer Overflow

- Credit flow control is explicit about how much data a sender may transmit without receiving further credit.
- Lost or delayed feedback messages will not hurt; the sender would just use the previous allowance.
- When necessary, the sender's transmission can be even stopped completely, that is, rate = 0.

Thus, credit approach can afford to be aggressive by allocating large buffers, or equivalently high rates, to achieve high efficiency.

Rate's Approximation to Credit Is a Challenge

- Effective approximation of the above credit properties by rate control has been a major challenge and a subject of many Ph.D. theses.
- Just imagine how hard it would be to make TCP work reliably and efficiently by using rates instead of windows control over the Internet.
- Rate control could also use link-by-link feedback and per-VC queueing to provide some relief (with increased cost), but inherently will never be as efficient and as robust in performance as credit control for highly bursty traffic.

What to Do Next?

Working in the rate control framework, we have the following options:

1. Do nothing.
2. Investigate *rate control in credit style* to achieve zero cell loss and bound switch buffer size. See next slide.
3. Include a *sequence number (SN) field* in rate-management messages. SN can capture Fwd_Cnt and thus allow CUP implementation. Independently, SN would be useful for policing anyway. E.g., after a switch receives a SN originally generated by it and later sent back by the source, the switch can start policing. How SN would be interpreted can be negotiated at configuration time.

Rate Control in Credit Style

This method can achieve zero cell loss, but still needs careful parameter tuning (e.g., on rate changing speeds) to achieve fairness and efficiency under bursty traffic.

Notations:

UM	Unoccupied Memory in the buffer at the receiver
CR	Committed Rate for a VC
TCR	Total Committed Rate for all VCs sharing the same buffer
MR	Measured Rate for a VC at the output. (To simplify presentation, measurement code is omitted here.)
TMR	Total Measured Rate for all VCs
RTT	Round-Trip Time in cell cycles

Switch Rate Setting Code (in Sketch):

```
begin
if (UM ≤ 2*RTT*Cur_TCR) or (UM > 8*RTT*Cur_TCR)
   Update rates committed to senders:
      New_TCR = UM/(4*RTT);
      Compute new committed rate for each VC:
      CR = (MR/TMR)*New_TCR;
   Send updated CRs upstream
end
```

Conclusions and Suggestions

1. **Since traffic is expected to get burstier over time as network speed increases, rate control alone will not be sufficient and must incorporate or complement with credit-like buffer-based methods or ideas.**
2. **As our field experience of using ATM networks expands, we will have much to learn, especially on interaction of ATM flow control with higher-lever protocols.**
3. **Recognizing the above points, standards and implementations should provide enough flexibility, e.g., by allowing negotiated use of different flow-control protocols, and additional fields in rate-management messages that could be used to add desirable credit features.**

References on Credit Control /1

[1] H. T. Kung, T. Blackwell and A. Chapman, "Credit-Based Flow Control for ATM Networks: Credit Update Protocol, Adaptive Credit Allocation, and Statistical Multiplexing," *Proc. ACM SIGCOMM '94 Symposium on Communications Architectures, Protocols and Applications*, pp. 101-114.

[2] H. T. Kung and K. Chang, "Receiver-Oriented Adaptive Buffer Allocation in Credit-Based Flow Control for ATM Networks," *Proc. INFOCOM '95*, April 1995.

[3] H. T. Kung and A. Chapman, "The FCVC (Flow-Controlled Virtual Channels) Proposal for ATM

References on Credit Control /2

Networks," Version 2.0, 1993. A summary appears in *Proc. 1993 International Conf. on Network Protocols*, pp. 116-127. (Postscript files of this and other related papers by the authors and their colleagues are available via anonymous FTP from virtual.harvard.edu:/pub/htk/atm.)

[4] H. T. Kung and R. Morris, "Credit-Based Flow Control for ATM Networks," *IEEE Network Magazine*, March 1995.

[5] C. Ozveren, R. Simcoe, and G. Varghese, "Reliable and Efficient Hop-by-Hop Flow Control," *Proc. ACM SIGCOMM '94 Symposium on Communications Architectures, Protocols and Applications*, pp. 89-100.