CBFC vs Reactive Congestion Control:
A Taylor Expansion Analogy

Ankush Jain and A Consortium of Math-Literate LLMs

2025-10-06

1 The Core Insight

Credit-Based Flow Control (CBFC) directly observes and enforces buffer constraints locally at each hop. Reactive end-to-end congestion control must infer and predict these same constraints from delayed, aggregated feedback. The relationship is analogous to knowing a function versus approximating it with a Taylor series: you need progressively more terms (coefficients) to match the original.

2 What Each Approach Knows

2.1 CBFC: Direct Observation

At hop i, the switch tracks:

- Queue occupancy $q_i(t)$
- Buffer limit B_i
- Free space $B_i - q_i(t)$

Decision rule: The transmitter may send only while its link-local credits $C_i(t)$ are $\geq$ the packet/flit size. Credits reflect downstream free buffer and are returned as buffers drain (per packet/flit).

This is enforced locally with a single-hop control loop (wire + SerDes + receiver update + credit return). Safety comes from proactive credit accounting, including pipeline/headroom margins, so the sender cannot inject without credits—rather than reacting after a threshold breach.

2.2 Reactive: Delayed Inference

The sender, after RTT delay d, receives feedback signal $s(t-d)$ that depends on network state. The sender must reconstruct:

$$q_i(t) \approx f(s(t-d), \text{additional info})$$

and predict:

$$q_i(t+d) \approx g(q_i(t), \dot{q}_i(t), \ddot{q}_i(t), \ldots)$$

to decide the safe sending rate right now.

3 The Taylor Expansion Analogy

Think of the "true" control signal as a function $Q(t)$ representing network state (queue occupancies across all hops). CBFC samples $Q(t)$ with zero delay at each hop.

Reactive control at time t only has access to $Q(t-d)$ and must approximate $Q(t)$ and predict $Q(t+d)$.

3.1 Order 0: Threshold

Use only current (delayed) value:

$$Q(t) \approx Q(t-d)$$

This is a constant approximation. To stay safe, you need huge margin: set threshold $\tau \ll B - C \cdot d$ where C is link capacity. Result: severe underutilization.

3.2 Order 1: Rate of Change

Add derivative:

$$Q(t) \approx Q(t-d) + \dot{Q}(t-d) \cdot d$$

Now you need two coefficients:

- $Q(t-d)$: current congestion level
- $\dot{Q}(t-d)$: rate of change

Better prediction, but still aggregated across all hops.

3.3 Order 2: Acceleration

Add second derivative:

$$Q(t) \approx Q(t-d) + \dot{Q}(t-d) \cdot d + \frac{1}{2}\ddot{Q}(t-d) \cdot d^2$$

Three coefficients needed. Helps with changing traffic patterns.

3.4 Multi-Hop Decomposition

For path with n hops, the true state is a vector:

$$\mathbf{Q}(t) = [q_1(t), q_2(t), \ldots, q_n(t)]$$

Each hop has independent constraint $q_i(t) \leq B_i$. But the sender receives one aggregate signal. To properly approximate:

$$\mathbf{Q}(t) \approx \mathbf{Q}(t-d) + \dot{\mathbf{Q}}(t-d) \cdot d + \ldots$$

you need:

- n coefficients for current state: $q_1(t-d), \ldots, q_n(t-d)$
- n coefficients for derivatives: $\dot{q}_1(t-d), \ldots, \dot{q}_n(t-d)$
- Cross-traffic terms: what others are sending at each hop
- Higher-order terms if traffic is bursty

The number of coefficients scales with: (system complexity) $\times$ (prediction horizon) $\times$ (accuracy required).

4 Why CBFC Needs Zero Coefficients

CBFC doesn't approximate or predict. At each hop i:

$$q_i(t) \leq B_i$$

is enforced directly by observing $q_i(t)$ right now, at this hop. No delay means no Taylor expansion needed. The function is sampled at the point where it's used.

5 The Equivalence

To make reactive control approximate CBFC performance:

1. Need per-hop state visibility (n coefficients)
2. Need rate of change per hop (n more coefficients)
3. Need cross-traffic awareness
4. Need accurate delay estimation
5. May need higher derivatives if traffic is non-smooth

In the limit of perfect information and infinite terms, reactive control converges to CBFC behavior—at which point you've reconstructed local flow control remotely.

6 Practical Implication

When systems claim to achieve high utilization without loss using "perfect network knowledge" or "careful traffic management," they're implicitly using many coefficients. They've either:

- Gathered rich per-hop state (approaching CBFC information)
- Constrained traffic to safe levels (trading utilization for safety)
- Accepted occasional loss (admitting approximation errors)

There's no magic. Delay creates an information gap that requires either more coefficients to bridge or conservatism to avoid.

7 Summary

CBFC: Direct observation, zero delay, simple mechanism.
Reactive: Delayed observation, requires approximation. Number of coefficients needed grows with:

- Number of hops (system dimension)
- Delay magnitude (prediction horizon)
- Traffic variability (higher derivatives)
- Desired accuracy (approximation order)

The Taylor expansion analogy captures why reactive approaches grow in complexity: they're approximating what CBFC knows directly.