QoS routing and admission control of high-priority connection-oriented flows while protecting TCP traffic

Stephen L. Spitler, USC
Daniel C. Lee, SFU

Highlights: Motivation

- Some ISPs may support both:
  - Traditional destination-based, hop-by-hop forwarding of low-priority, TCP traffic (which we will also refer to as ‘best effort’ traffic)
  - Connection-oriented, end-to-end routing of high priority traffic with QoS requirements
    - E.g., via LSPs in an MPLS network
- Potential problem:
  - At heavily loaded links carrying QoS and TCP traffic, the low-priority TCP traffic may experience delay, bandwidth starvation
- Objective
  - BE-friendly routing of a QoS connection in response to a trunk request; maximize revenue from QoS service

Select a path for the incoming QoS connection

- BE friendly
- Maximize revenue

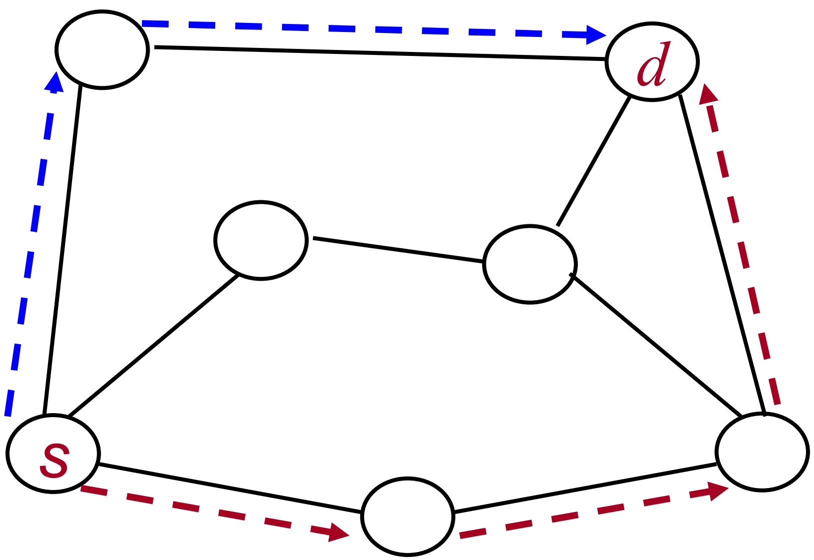

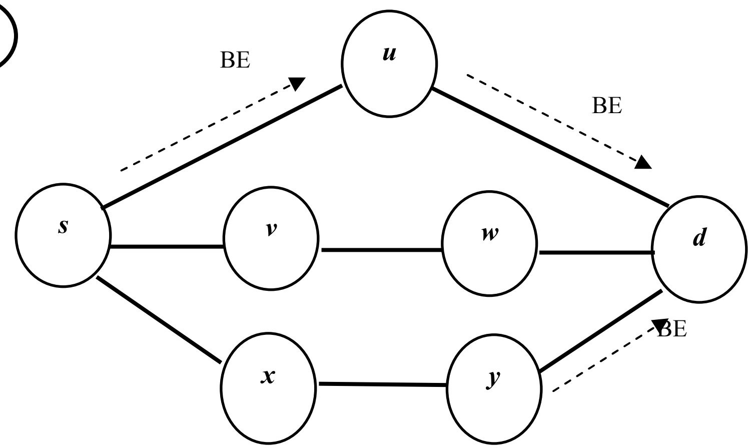

Highlights: Approach

- Link constraint: to be a part of the path
  - A link must have enough unreserved effective bandwidth
  - If a link will have too little bandwidth for BE after accommodating the trunk, that link Is not eligible.
  - Can take advantage of the excessive effective bandwidth

QoS traffic

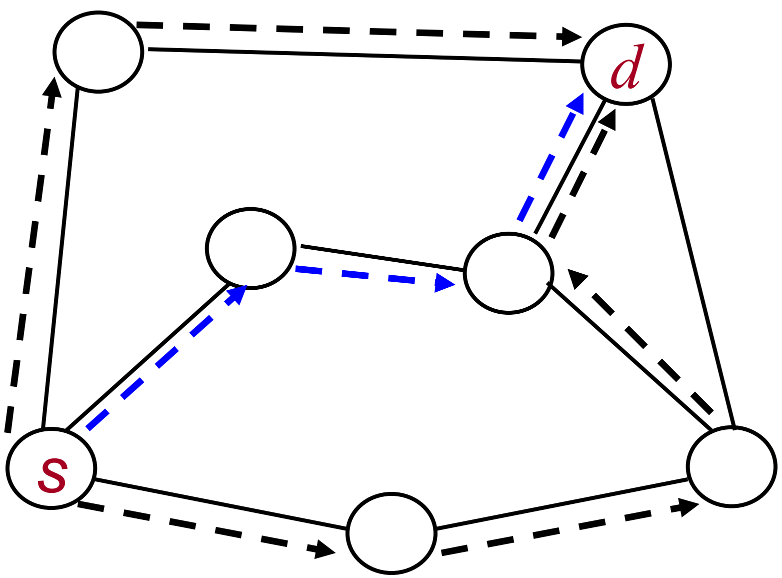

Highlights: Approach

- Two-stage optimization
  - Among the multiple routes with minimum effective bandwidth consumption,
  - Choose the one that hurts the BE the least.
  - We found a method to do this two-stage optimization through a single run of the shortest path finding

Objective

- BE-friendly routing of (trunks of) QoS connections
  - Limit the BE traffic delay that results from routing high-priority QoS trunks
    - Maintain a minimum level of service for BE traffic
  - When multiple paths are equally attractive for routing a QoS trunk, select one that impacts low-priority BE traffic the least

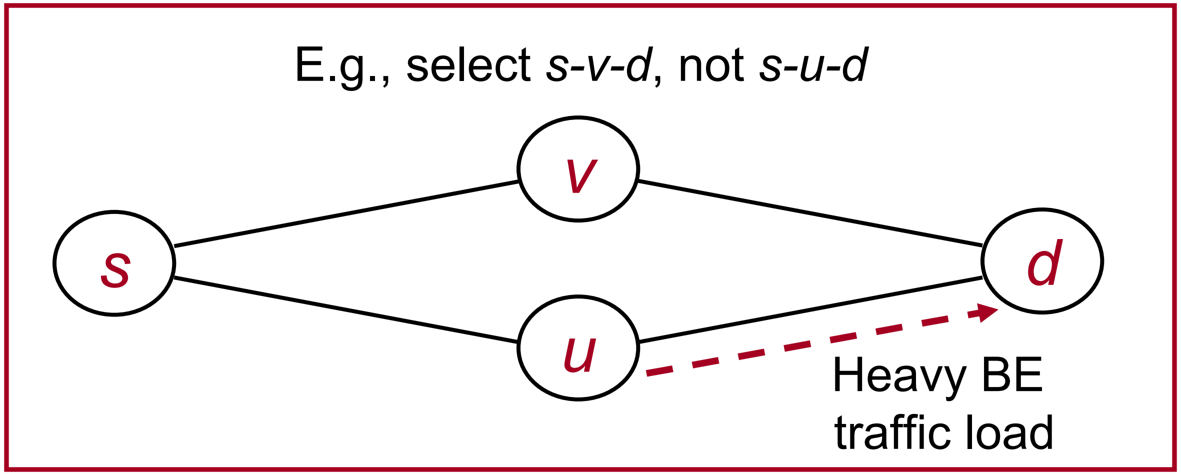

Approach

- BE delay metric
  - An approximate measure of BE traffic delay at each link
- An additional constraint on QoS trunk routing:
  - BE delay metric must not be excessive at any link
- An additional cost of candidate path to support QoS trunk:
  - BE path cost defined in terms of BE delay metric
- A constrained, two-stage optimization to route QoS trunks

BE Delay Metric

- Example BE delay metrics
  - M/M/1-based approximation of time spent by BE traffic at each directed link, *e*

$$\frac{1}{\gamma}\left[\frac{F_e}{\left(C_e^{BE}-F_e\right)}+d_e F_e\right]$$

  - G/G/1-based upper bound on BE queueing delay at each directed link

$$\frac{F_e^2\left(\sigma_a^2-\sigma_b^2\right)}{2\gamma\left(1-F_e / C_e^{BE}\right)}$$

Constraints on QoS Trunk Routing

- A candidate path is feasible if at each link along path:
  - Effective bandwidth of QoS trunk does not exceed residual link capacity

$$\alpha \le C_e - \sum_{\{c | e \in p_c\}} \alpha_{c,e} \equiv R_e^{eff}$$

  - Additional, BE constraint:
    - BE delay metric (that results if path is selected) must not exceed a maximum acceptable limit
    - Implemented as an additional residual link capacity constraint:

(M/M/1 approximation)

$$b \le C_e^{BE} - F_e - \frac{F_e}{(D_{max} - d_e F_e)}$$

Path Costs

- Costs of a feasible candidate path:
  - QoS path cost:
    - Exclusive of effects on BE traffic
    - E.g., added effective bandwidth consumption
  - Additional, BE path cost:
    - Based on BE delay metric
    - Indicative of increase in BE traffic delay if path is selected to support QoS trunk

$$\mathrm{cost}_{BE}(p) \equiv \sum_{e \in p} \frac{F_e b}{\left(C_e^{BE} - F_e - b\right)\left(C_e^{BE} - F_e\right)}$$

(M/M/1 approximation)

QoS Routing Optimization

- A constrained, two-stage optimization for routing QoS trunks
  1. Find feasible candidate paths that minimize QoS cost
  2. Secondary optimization:
     - In case of ties (multiple feasible paths with minimum QoS cost), select a path that minimizes BE cost

Implementation Tricks

- Efficient implementation of QoS routing optimization (Dijkstra's algorithm)
  - Based on observations:
    - Quantization condition may hold
    - Bounded BE cost
- Exploitation of "excess effective bandwidth" to enhance QoS routing performance

Observations

- Quantization condition:
  - The QoS cost of a path may be quantized
    - E.g., if QoS cost is number of hops in path, the quantum, q, is 1
    - E.g., if QoS cost is net bandwidth reservation for path, it may be quantized with a quantum, q, in bits per second
- Bounded BE cost:
  - BE cost of any feasible path is upper-bounded
    - Due to BE constraint which limits BE traffic delay

Efficient Implementation

- Form weighted sum cost of path,

  [*QoS cost*] + w_{BE} [*BE cost*],

  where weighting coefficient w_{BE} is small enough so that

  w_{BE} [*BE cost*] < *q for any path*

- If feasible path, *p*, minimizes weighted sum cost,
  1. Then *p* minimizes QoS cost
  2. Among all feasible minimum QoS cost paths, path *p* is one with minimum BE cost
- Two-stage optimization by Dijkstra with link metric:

$$d(e) \equiv \alpha_e + w_{BE} \frac{F_e b}{\left(C_e^{BE} - F_e - b\right)\left(C_e^{BE} - F_e\right)}$$

(M/M/1 approximation)

Excess Effective Bandwidth

- Bandwidth reserved for a QoS trunk

  = effective bandwidth of QoS trunk

  $\geq$ average bandwidth that trunk actually consumes

- "Excess effective bandwidth"
  - The average amount of bandwidth that is reserved for but unused by a QoS trunk
  - Exploit excess effective bandwidth to support BE traffic
  - Take excess effective bandwidth into account when calculating BE delay metric for each link
    - Eases the BE constraint on QoS routing

Example Simulation and Findings

- BE-friendly LSP routing with path restoration
  - "Path restoration with exact reservations" in 15-node test network [Kodialam, Lakshman '02]
    - Additional BE constraint tends to increase blocking probability of LSP requests to support QoS trunks, but effect is small over wide range of BE traffic loads

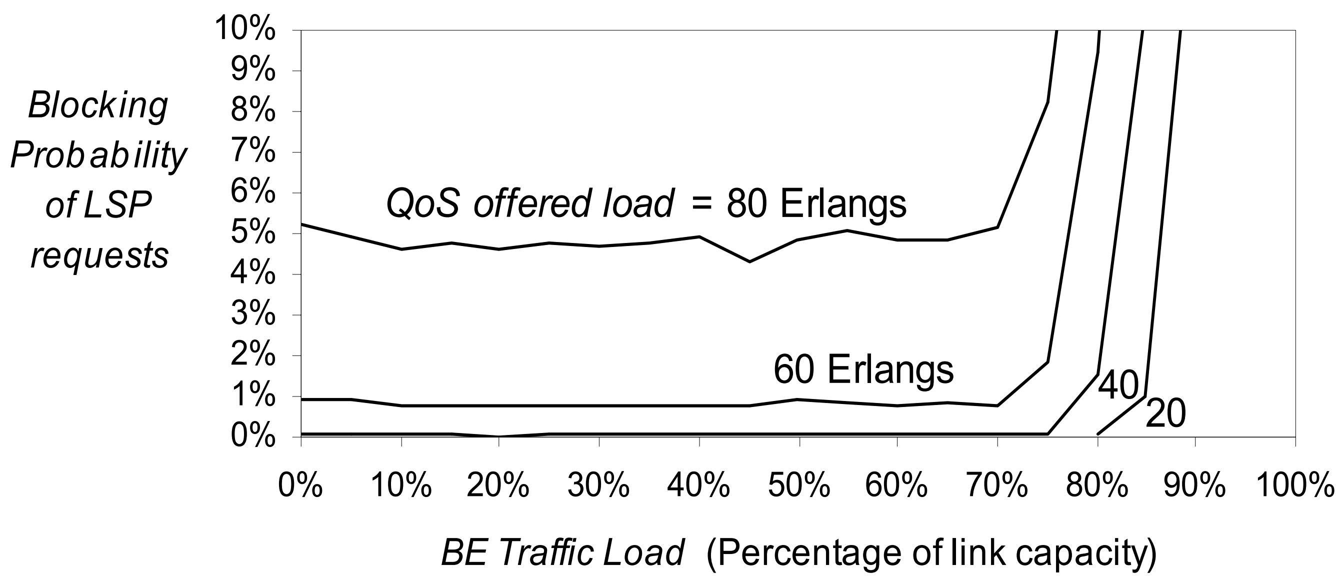

(Case of QoS effective bandwidth equal to twice average bandwidth usage)

Example Simulation and Findings

- Sample US nationwide topology" [Rai,.., Mukherjee, 05] with 24 nodes and 86 directed links.

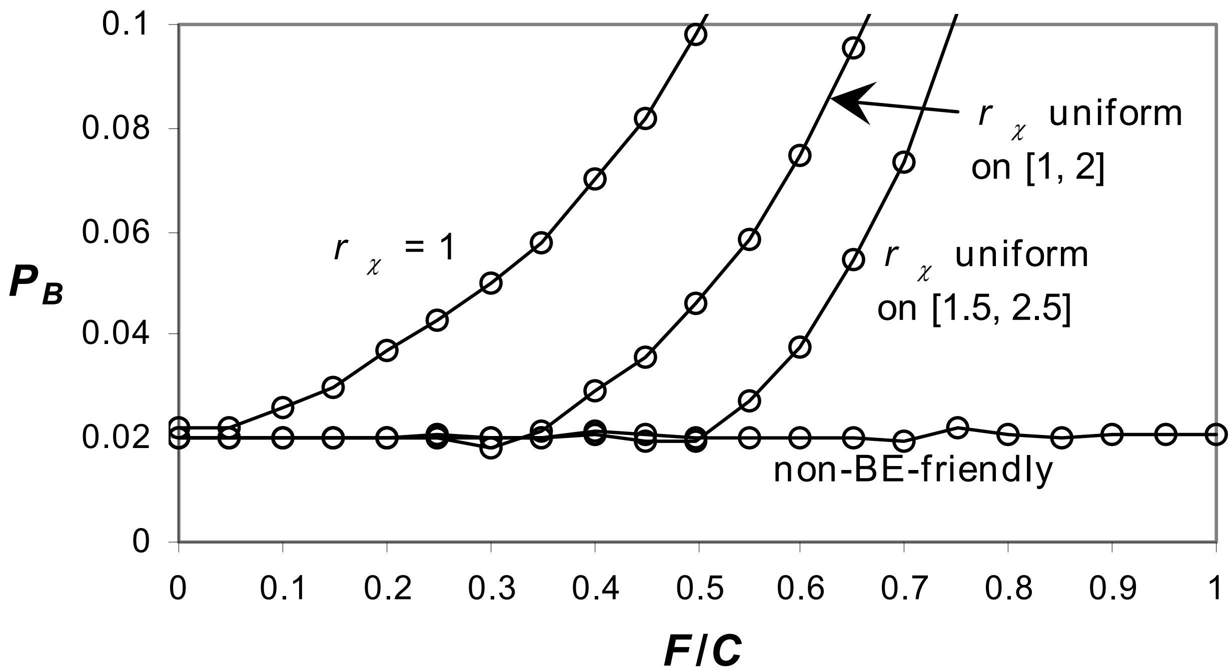

QoS call blocking probability versus F/C for $\rho = 7000$ E .

Example Simulation and Findings

- Sample US nationwide topology" with 24 nodes and 86 directed links.

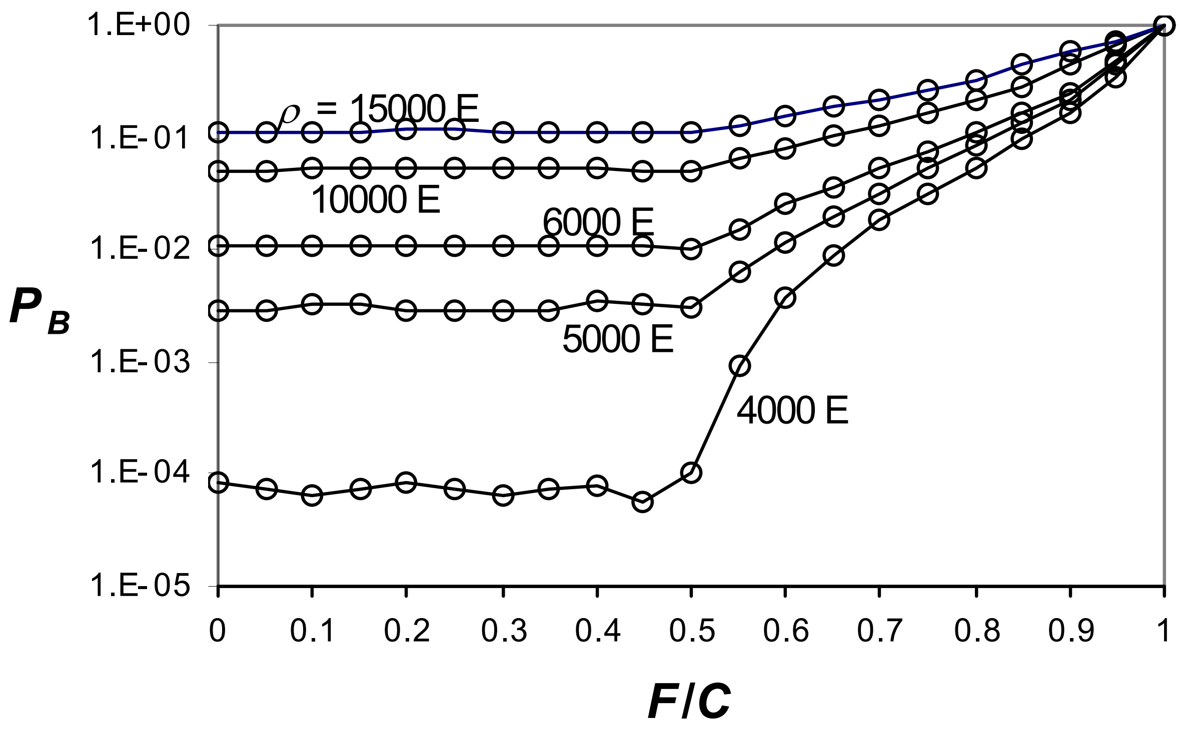

PB versus *F*/*C* with *r* uniform on [1.5, 2.5].

Example Simulation and Findings

- Sample US nationwide topology" with 24 nodes and 86 directed links. Non-uniform s-d pairs (8 nodes are more likely to be an ingress-egress nodes.)

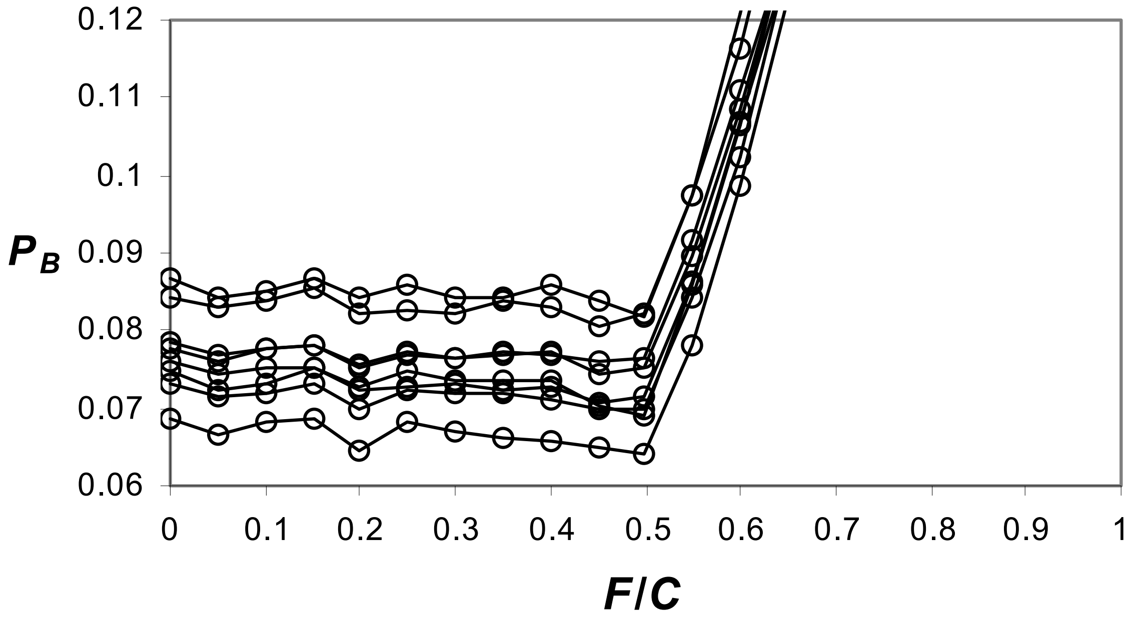

Blocking Prob. versus F/C, $\rho = 10000$ E, with r uniform on [1.5, 2.5].

Discussions

- BE constraint on routing of QoS trunks
  - Guarantee minimum level of service for low-priority TCP (BE) traffic
  - Exploitation of excess effective bandwidth eases BE constraint on QoS routing
  - Simulation results suggest that increase in QoS blocking probability due to BE constraint need not be prohibitive
- Ties between candidate QoS paths, decided by BE cost
- Quantization condition permits efficient QoS routing implementation
  - Dijkstra's algorithm simultaneously implements two stages of routing optimization

Phase transition?

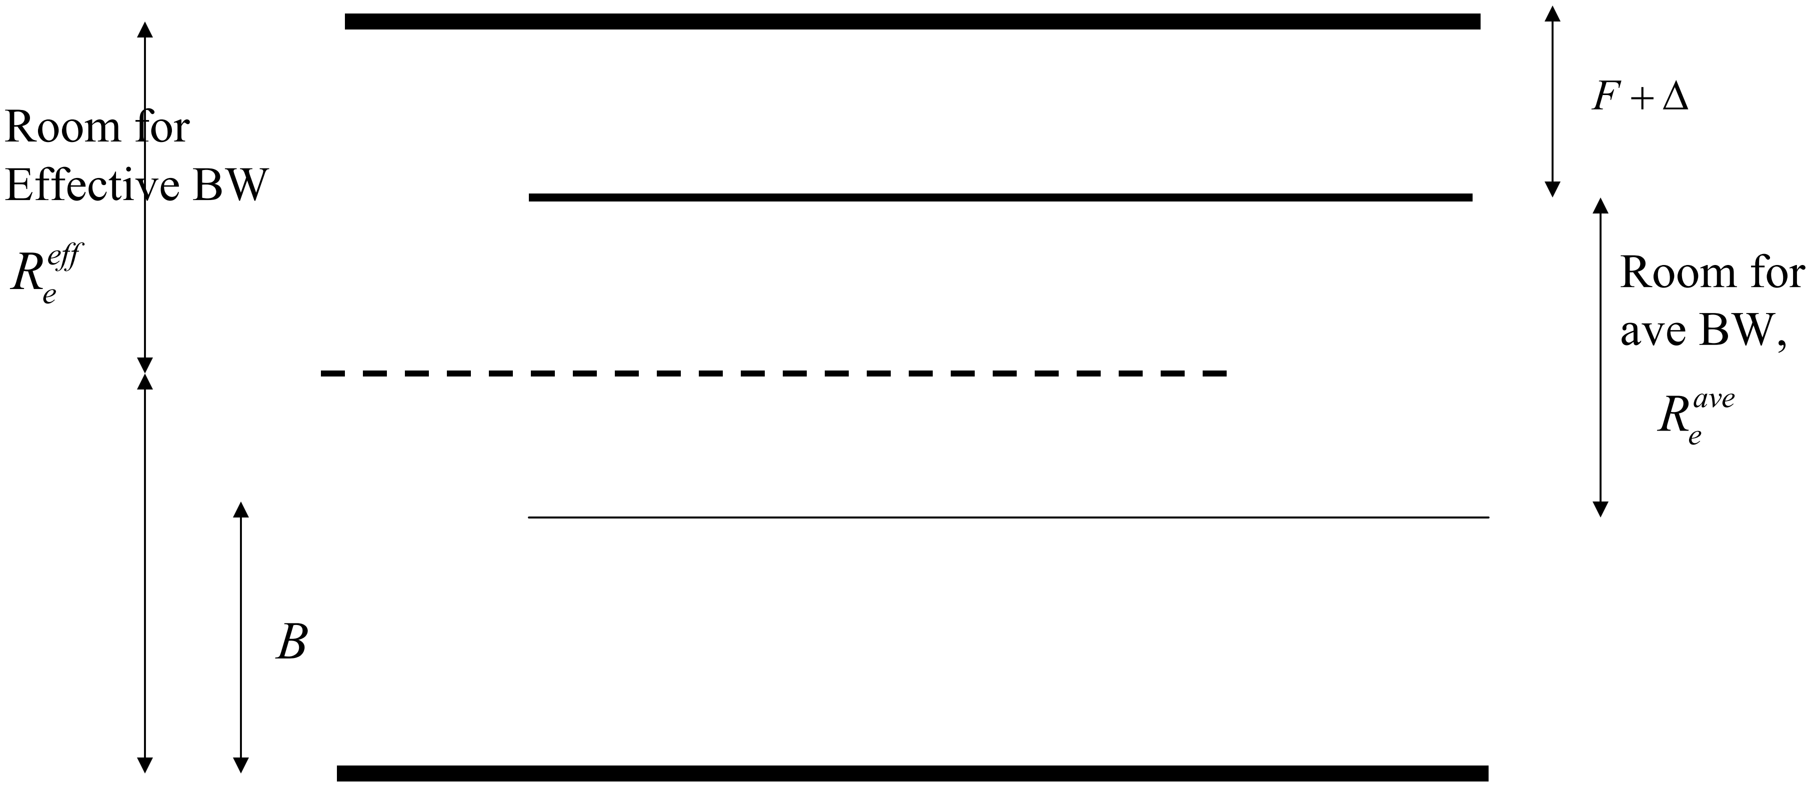

$$\alpha_{\chi,e} \le R_e^{eff}$$

$b_\chi \le R_e^{ave} \Leftrightarrow \alpha_{\chi,e} \le rR_e^{ave}$ (Constraint protecting TCP traffic)

i.e., $\alpha_{\chi,e} \le \min\left\{R_e^{eff}, rR_e^{ave}\right\} = \min\left\{C - rB, r\left(C - F - \Delta - B\right)\right\}$

Constraint protecting TCP traffic is active only when $R_e^{eff} > rR_e^{ave}$;

That is, only when

$C - rB < r\left(C - F - \Delta - B\right)$ or equivalently

$$F > \frac{C(r-1)}{r} - \Delta$$