

Tiered Service in Packet-Switched Networks

George N. Rouskas

Laura Jackson (PhD, 2003), Nikhil Baradwaj (MS, 2005),
Ajay Amudalabhasker (MS, 2006), Zyad Dwekat, Srikrishna Khare

Department of Computer Science
North Carolina State University

Outline

- Motivation
- Service tier selection as a directional p -median problem:
 - deterministic demands
 - stochastic demands
 - multiple QoS parameters
 - TDM emulation
- Application: efficient WFQ implementations

Challenges at Data and Control Planes

- **Objective:** support per-flow QoS

Challenges at Data and Control Planes

- **Objective:** support per-flow QoS
- **Control plane:** amount of state information increases with
 - # of users (flows)
 - QoS requirements

Challenges at Data and Control Planes

- **Objective:** support per-flow QoS
- **Control plane:** amount of state information increases with
 - # of users (flows)
 - QoS requirements
- **Data plane:**
 - mismatch between optical capacity and electronic capabilities
 - complexity of QoS functions increases with # of users/constraints

Challenges at Data and Control Planes

- **Objective:** support per-flow QoS
 - **Control plane:** amount of state information increases with
 - # of users (flows)
 - QoS requirements
 - **Data plane:**
 - mismatch between optical capacity and electronic capabilities
 - complexity of QoS functions increases with # of users/constraints
- Maintaining **per-flow** information poses severe **scalability challenges**

Bandwidth-Continuous Networks

Bandwidth-Continuous Networks

- In theory, provider may offer rates of
 - 1.001 Mbps
 - 0.999 Mbps
- network functions designed to accommodate arbitrary rates

Bandwidth-Continuous Networks

- **In theory**, provider may offer rates of
 - 1.001 Mbps
 - 0.999 Mbps
 - network functions designed to accommodate **arbitrary** rates
- **In practice**:
 - How to **distinguish** between two rates?
 - How to **enforce** the two rates?
 - What amount of **state info** is needed?

Bandwidth-Continuous Networks

- **In theory**, provider may offer rates of
 - 1.001 Mbps
 - 0.999 Mbps
 - network functions designed to accommodate **arbitrary** rates
- **In practice**:
 - How to **distinguish** between two rates?
 - How to **enforce** the two rates?
 - What amount of **state info** is needed?

As network capacity ↑, overhead also ↑

Example: Cable TV

- Cable systems support $k > 100$ TV channels
→ $O(2^k)$ combinations

Example: Cable TV

- Cable systems support $k > 100$ TV channels
→ $O(2^k)$ combinations
- **But:** only a few subscription levels offered
 - “basic”
 - “standard”
 - “premium”
 - plus a few “a la carte” channels

Bandwidth-Discrete Networks

- Offer a **small number** of discrete bandwidth tiers

Bandwidth-Discrete Networks

- Offer a **small number** of discrete bandwidth tiers
- Demands mapped to next higher tier → performance **penalty**

Bandwidth-Discrete Networks

- Offer a **small number** of discrete bandwidth tiers
- Demands mapped to next higher tier → performance **penalty**
- **Benefits:** control and data plane operation **simplified**
 - traffic engineering
 - traffic grooming
 - billing, policing, etc.
 - TDM emulation
 - scheduling and QoS support


Outline

- Motivation
- Service tier selection as a directional p -median problem:
 - deterministic demands
 - stochastic demands
 - multiple QoS parameters
 - TDM emulation
- Application: efficient WFQ implementations


Tier Selection: Deterministic Demands

- **Input:** n bandwidth demands, $x_1 \leq x_2 \leq \dots \leq x_n$
- **Output:** $p \ll n$ service tiers, $z_1 < z_2 < \dots < z_p$
- **Objective:**
 - $x_i \rightarrow z_j$ iff $z_{j-1} < x_i \leq z_j$
 - to minimize the total **bandwidth penalty** $\sum_{i=1}^n (z_j - x_i)$

Mapping of Bandwidth Demands to Service Tiers




Mapping of Bandwidth Demands to Service Tiers



- Similar to *p*-median problem on the real line
- **But:** demands mapped to next higher tier
- Dynamic programming algorithm with complexity $O(pn)$

Results: Bandwidth Penalty




Outline

- Motivation
- Service tier selection as a directional p -median problem:
 - deterministic demands
 - stochastic demands
 - multiple QoS parameters
 - TDM emulation
- Application: efficient WFQ implementations


Tier Selection: Stochastic Demands

- **Input:** PDF $f(x)$ of bandwidth demands
- **Output:** p service levels, $z_1 < z_2 < \dots < z_p$
- **Objective:**
 - $x \rightarrow z_j$ iff $z_{j-1} < x \leq z_j$
 - to minimize the **expected bandwidth penalty**
$$\sum_{i=1}^L \int_{z_{i-1}}^{z_i} (z_i - x) f(x) dx$$

Mapping of Bandwidth Demands to Service Tiers




Mapping of Bandwidth Demands to Service Tiers




1. Optimal solution through **nonlinear programming**
2. **Discretize the PDF** \rightarrow reuse previous algorithm

Discretizing the PDF



Results: Bandwidth Penalty



Outline

- Motivation
- Service tier selection as a directional p -median problem:
 - deterministic demands
 - stochastic demands
 - multiple QoS parameters
 - TDM emulation
- Application: efficient WFQ implementations

Tier Selection with Multiple Parameters

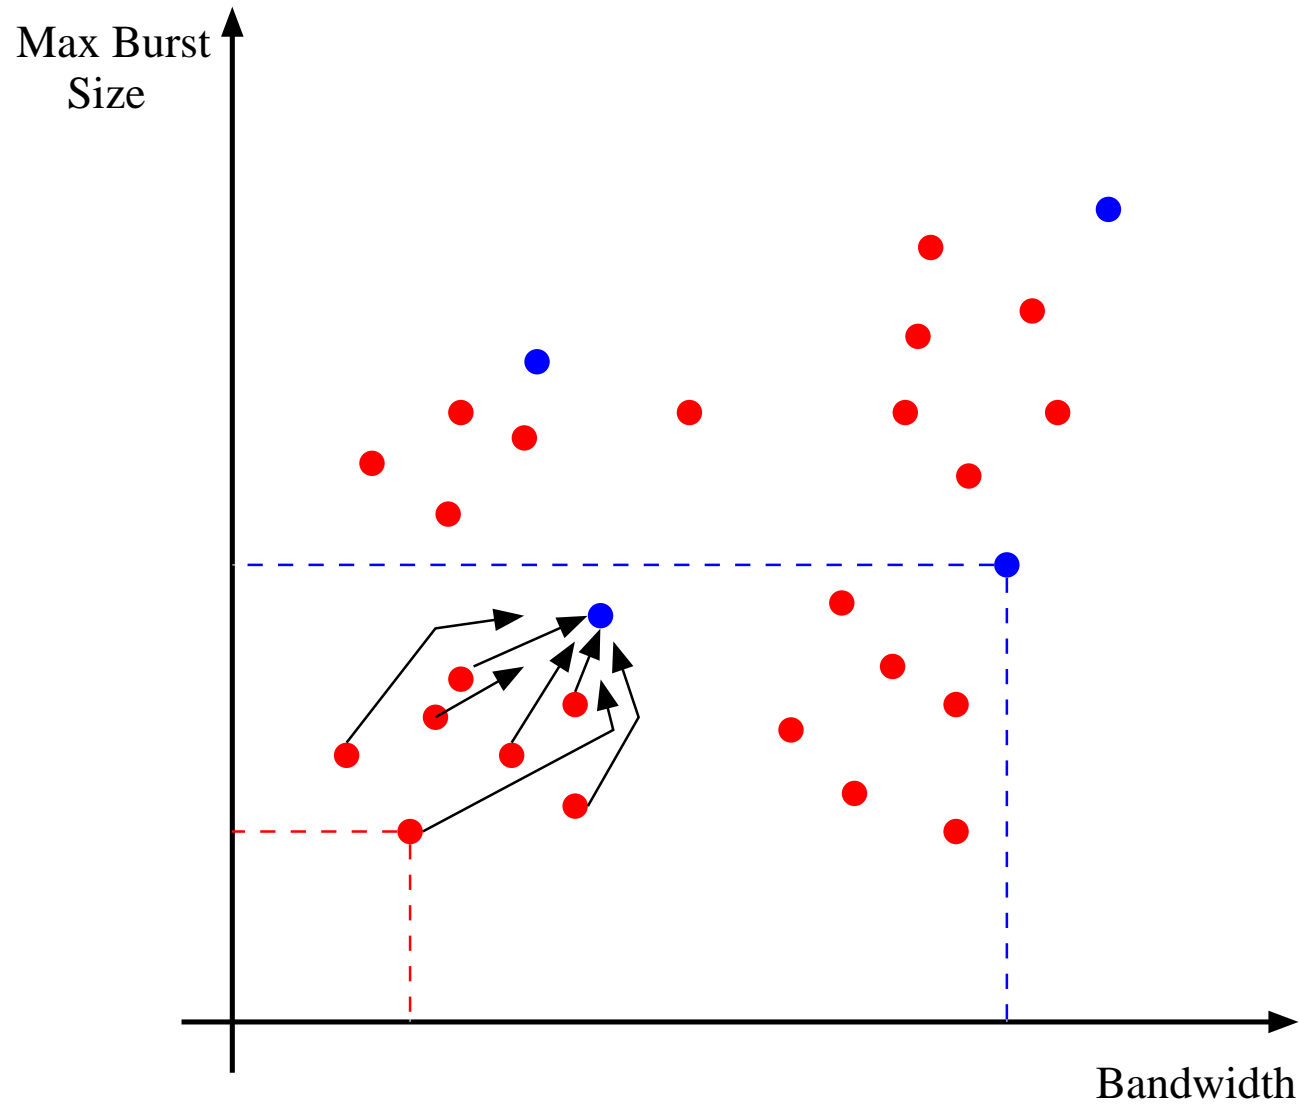
- Traffic demands may have **multiple** attributes or QoS requirements:
 - bandwidth
 - maximum burst size
 - delay bound
 - etc...

Tier Selection with Multiple Parameters

- Traffic demands may have **multiple** attributes or QoS requirements:
 - bandwidth
 - maximum burst size
 - delay bound
 - etc...
- Service tiers in multiple dimensions
 - jointly optimal for a vector of parameters

Tier Selection with Multiple Parameters

- Traffic demands may have **multiple** attributes or QoS requirements:
 - bandwidth
 - maximum burst size
 - delay bound
 - etc...
- Service tiers in multiple dimensions
 - jointly optimal for a vector of parameters
- Problem is NP-Complete for ≥ 2 dimensions
 - reduction from planar 3-SAT, non-polar version

The Directional p -Median Problem in the Plane


Decomposition Heuristic

- Our heuristic:
 - run 1-dimensional algorithm twice
 - once on x -values $\rightarrow p$ best x 's
 - once on y -values $\rightarrow p$ best y 's
 - cross two sets $\rightarrow p^2$ candidate points
 - run the Teitz & Bart (TB) exchange heuristic
 - popular, well-studied, extremely robust

Decomposition Heuristic

- Our heuristic:
 - run 1-dimensional algorithm twice
 - once on x -values $\rightarrow p$ best x 's
 - once on y -values $\rightarrow p$ best y 's
 - cross two sets $\rightarrow p^2$ candidate points
 - run the Teitz & Bart (TB) exchange heuristic
 - popular, well-studied, extremely robust
- Complexity: $O(p^3n)$

Results: Normalized Penalty




Outline

- Motivation
- Service tier selection as a directional p -median problem:
 - deterministic demands
 - stochastic demands
 - multiple QoS parameters
 - TDM emulation
- Application: efficient WFQ implementations

Tier Selection: TDM Emulation

- **Input:** n bandwidth demands, $x_1 \leq x_2 \leq \dots \leq x_n$
- **Output:**
 - $p \ll n$ service tiers, $z_1 < z_2 < \dots < z_p$
 - $z_i = k_i r$ ← additional constraint
- **Objective:**
 - $x_i \rightarrow z_j$ iff $z_{j-1} < x_i \leq z_j$
 - to minimize the total bandwidth penalty $\sum_{i=1}^n (z_j - x_i)$


Mapping of Bandwidth Demands to Service Tiers




Why TDM Emulation?

- Reuse robust control/management functions
- Configurable bandwidth unit (slot) → may be optimized
- Data plane operation **not** affected
 - no “wasted” slots
 - reuse of excess bandwidth

Behavior of the Objective Function




Behavior of the Objective Function



We have developed efficient heuristics

Results: Bandwidth Penalty



Outline

- Motivation
- Service tier selection as a directional p -median problem:
 - deterministic demands
 - stochastic demands
 - multiple QoS parameters
 - TDM emulation
- Application: efficient WFQ implementations

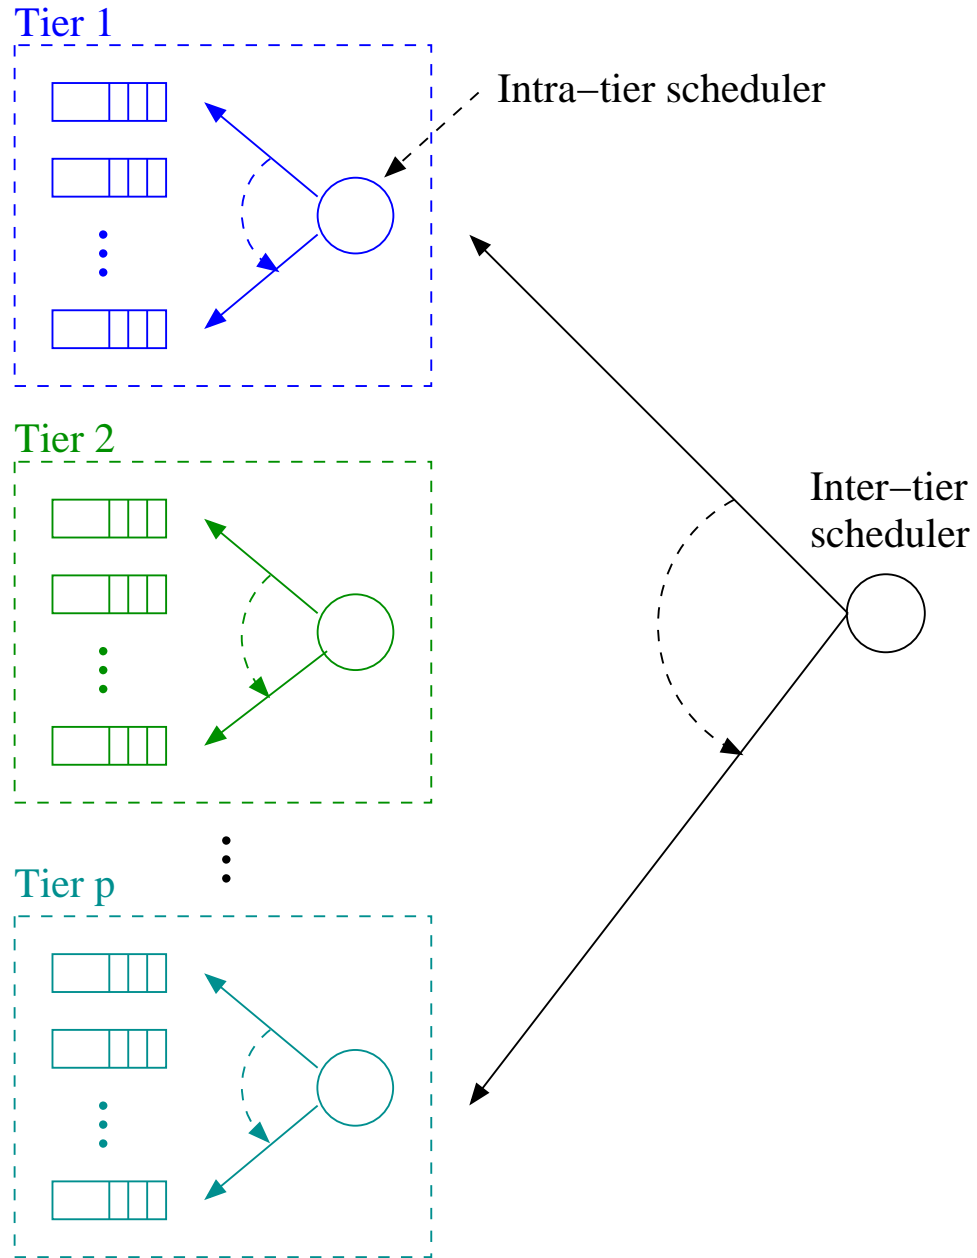
WFQ Packet Scheduling

- Packet departure time based on fluid-flow fair queueing algorithm

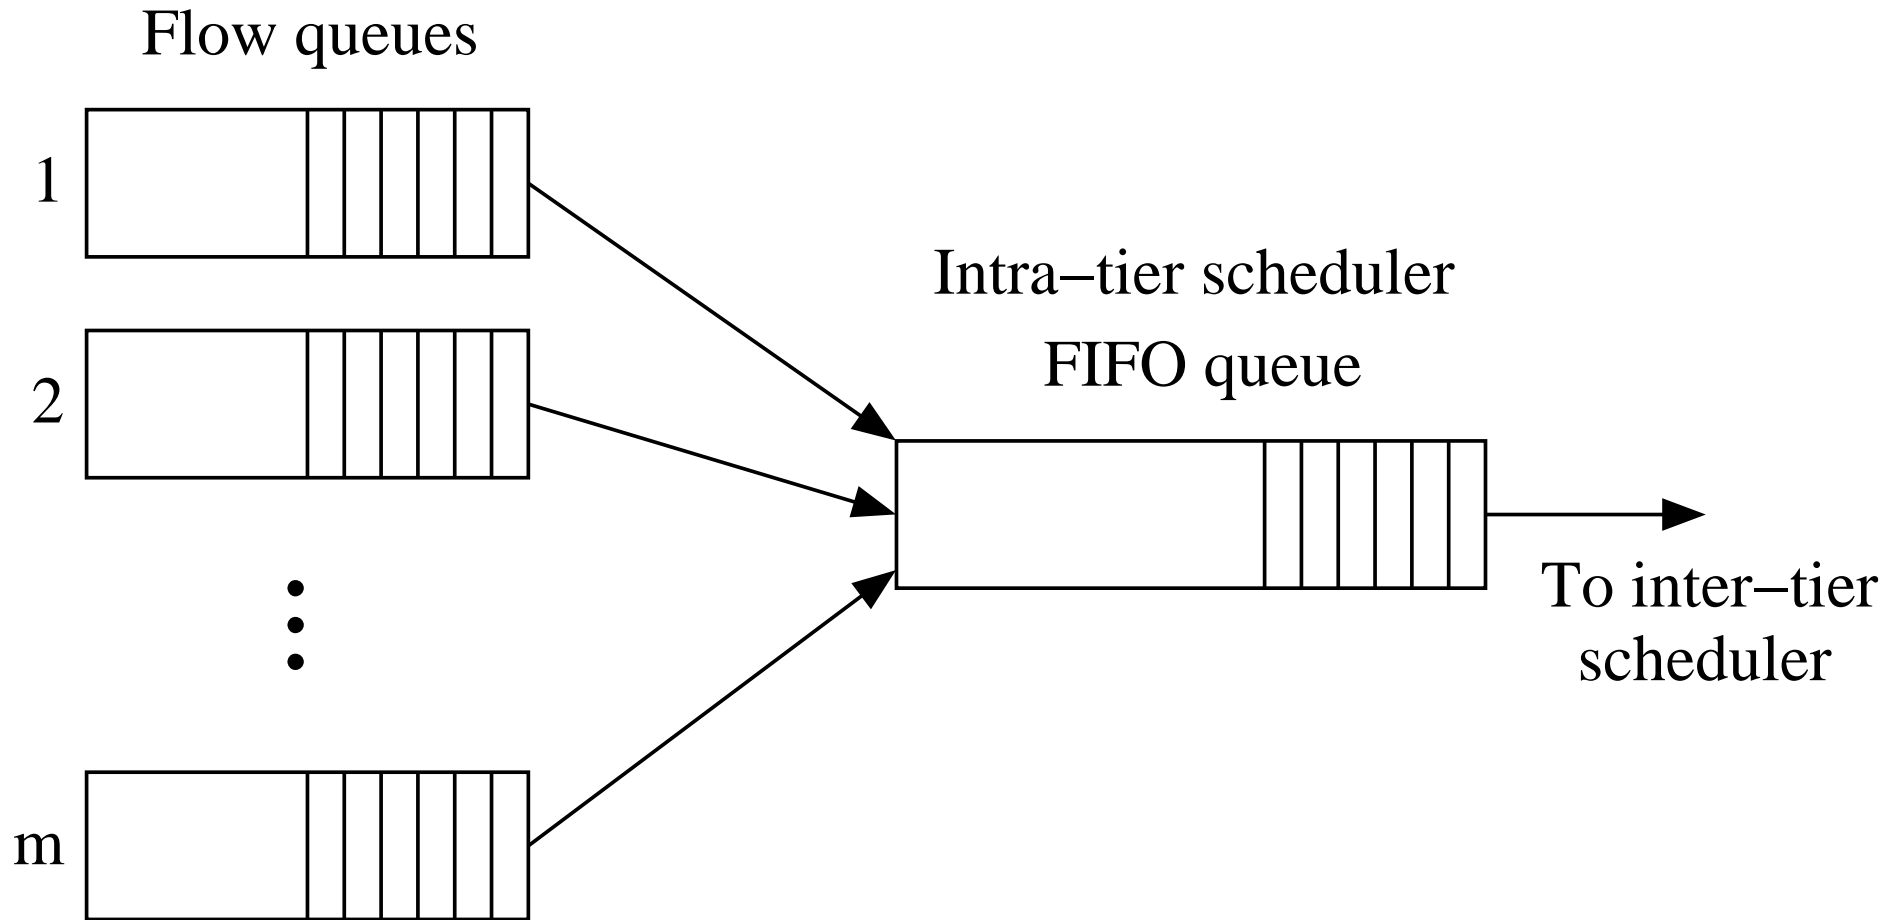
$$F_i^k = \max\{F_i^{k-1}, V(a_i^k)\} + \frac{L_i^k}{\phi_i}$$

- Packet inserted in logical queue
- Queue is **sorted** with respect to departure times
- Complexity:
 - $O(\log n)$ for selecting next packet
 - $O(n)$ for computation of virtual function

TSFQ Inter-Tier Scheduling



TSFQ Intra-Tier Scheduling: Fixed Packets



TSFQ Intra-Tier Scheduling: Variable Packets

