Resource Management in QoS-aware Wireless Cellular Networks

Zhi Zhang

Dept. of Electrical and Computer Engineering
Colorado State University

April 30, 2008

Progress Toward Ph.D.

- Contributions:
  - Algebraic equivalence of matrix conjugate direction and matrix multistage filters. (Joint work with Prof. Louis Scharf.)
  - Downlink opportunistic scheduling for OFDM systems.
  - Approximate stochastic dynamic programming for opportunistic fair scheduling in wireless networks. (Joint work with Sudhir Moola.)
  - Opportunistic scheduling with mixed constraints. (Ongoing)
  - Average reward MDP with temporal fairness constraints. (Ongoing)
- Publications:
  - 4 published in conference proceedings
  - 1 published in journal
  - 1 submitted to conference
- Graduation: Dec. 2008 (tentative)

Outline

Outline

Challenges of Broadband Wireless Networks

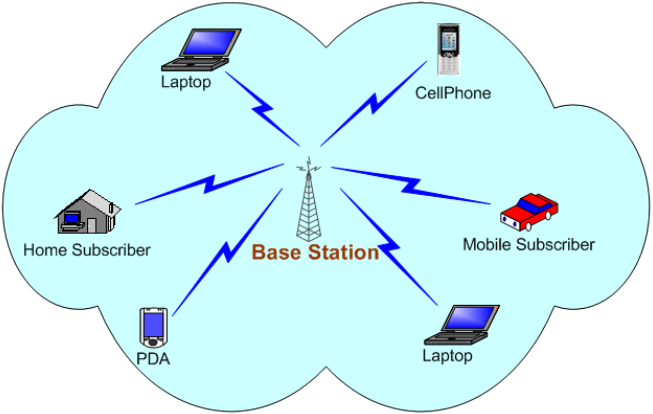

Challenges of Broadband Wireless Networks

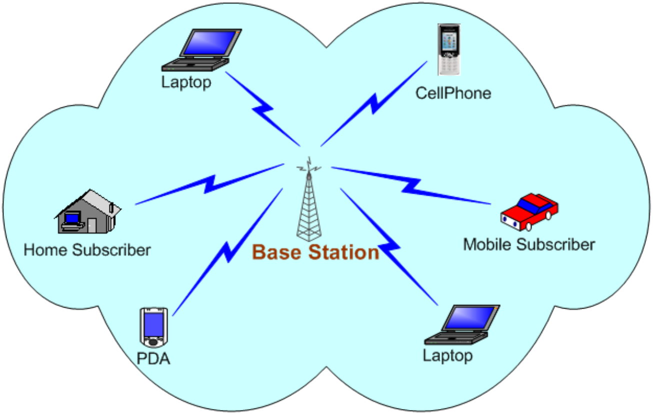

- High-speed data rate: $\sim$100Mbps to 1Gbps
- Heterogeneous Quality of Service (QoS) provisioning

Challenges of Broadband Wireless Networks

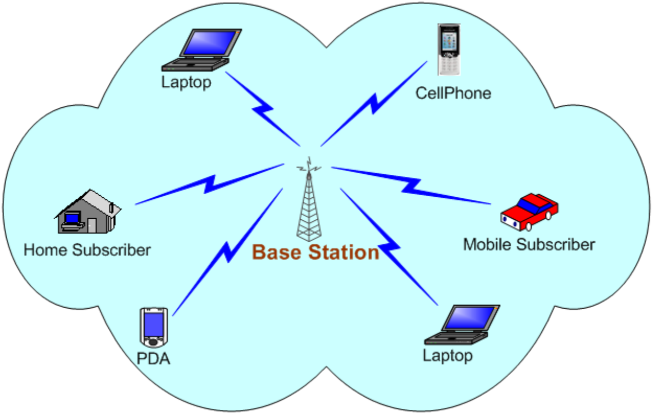

- High-speed data rate: ∼100Mbps to 1Gbps
- Heterogeneous Quality of Service (QoS) provisioning

⇒ Flexible and efficient resource management: scheduling, admission control, power control, etc.

Characteristics of Wireless Channels

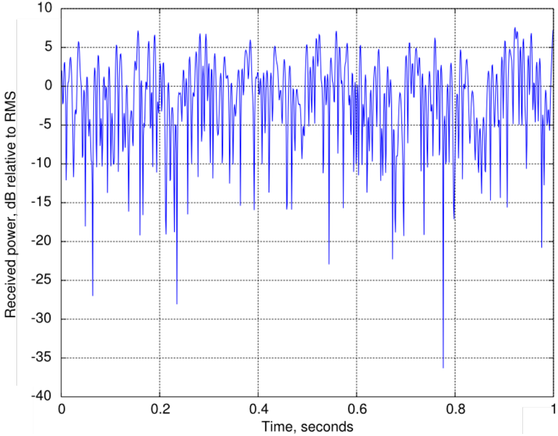

Figure: Rayleigh fading with maximum 100Hz Doppler shift

Characteristics of Wireless Channels

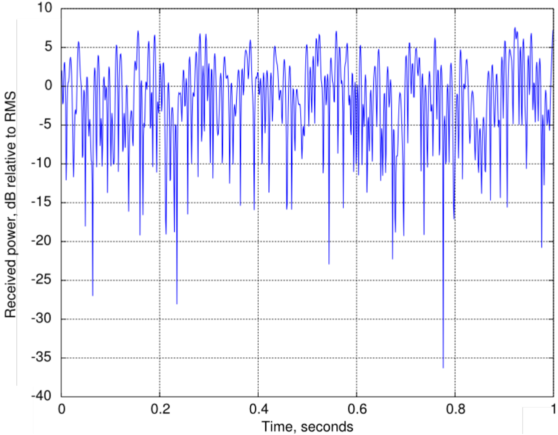

Figure: Rayleigh fading with maximum 100Hz Doppler shift

- Radio propagation: path loss, shadowing, and multipath fading
- *Time-varying* and *location-dependent* channel conditions

Outline

1. Background
2. Opportunistic Scheduling for OFDM Systems
3. Stochastic Dynamic Programming for Opportunistic Scheduling
4. Ongoing and Future Work
   - Opportunistic Scheduling with Mixed Constraints
   - Average Reward MDP with Temporal Fairness Constraints

Motivation

- Channel conditions vary significantly over users and time.
- Multiuser diversity: gain in performance with many users experiencing independently fluctuating channel conditions.
- OFDM provides efficient frequency utilization by exploiting frequency diversity.
- Combine multiuser diversity and frequency diversity opportunistically.

Opportunistic Scheduling: An example

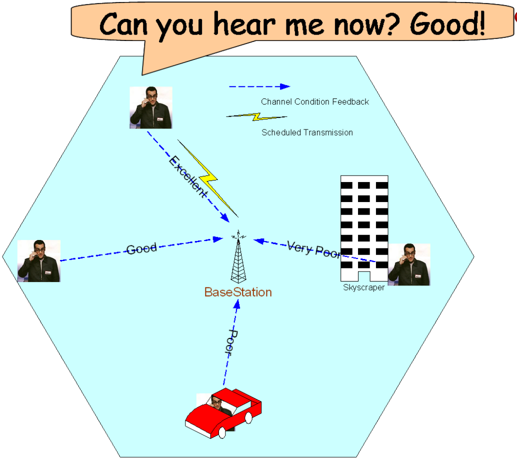

- One base station and four active users in a cell.

Opportunistic Scheduling: An example

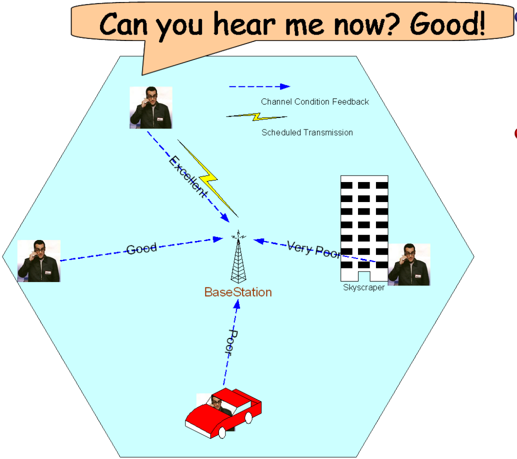

- One base station and four active users in a cell.
- At each time, scheduler in the base station picks the "relatively best" user to transmit in a given channel.

Opportunistic Scheduling: An example

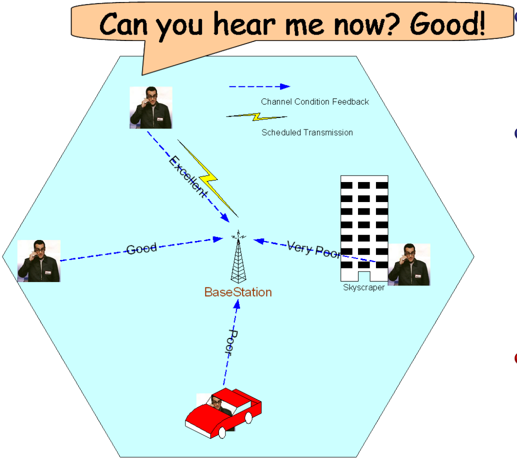

- One base station and four active users in a cell.
- At each time, scheduler in the base station picks the "relatively best" user to transmit in a given channel.
- Scheduling decision is based on channel condition feedback and QoS constraints.

Fairness and QoS Requirements

- Scheduler must allocate resources fairly among users under specific fairness/QoS constraints.
- Examples of (long-term) constraints:
  - Temporal fairness: user i is scheduled at least r_i of the time.
  - Utilitarian fairness: user i receives at least a_i of the overall system utility.
  - Minimum-performance guarantee: user i receives at least a utility of C_i.
  - Proportional fairness: aggregate of proportional change in utility is non-positive.

Downlink Scheduling over Multiuser OFDM

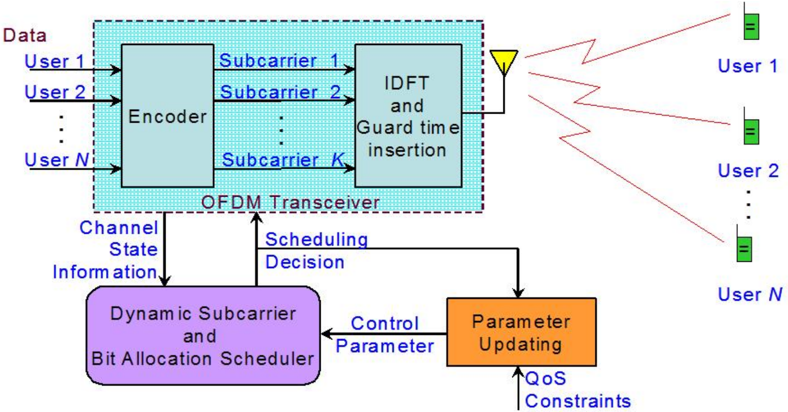

- Basic idea: opportunistically schedule users to available subcarriers and time slots while maintaining certain fairness/QoS constraints.

System Model

- $\omega_{i,k}^t$: channel utility for user i over subcarrier k at time t. Example: instantaneous throughput.
- Channel utility matrix: $\omega^t = (\omega_{i,k}^t; i = 1, \ldots, N, k = 1, \ldots, K)$.
- A_k^t: index of user scheduled over subcarrier k at time t. $\bar{A}^t = (A_1^t, A_2^t, \ldots, A_K^t)$: action.
- Scheduling policy: a rule that specifies the action at each time.
- At time t, policy π selects user $A_k^t = \pi_k(\omega^t)$ for subcarrier k, and receives "reward" $\omega_{i,k}^t$.
- A *feasible* policy satisfies specific fairness/QoS constraints.

Utility of Policy

- Given policy π, define

$$R_i(\pi) = \liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T}\sum_{k=1}^{K} \mathbf{1}_{\{\pi_k(\omega^t)=i\}}, \quad i = 1, \dots, N,$$

$$U_i^T(\pi) = \frac{1}{T}\sum_{t=1}^{T}\sum_{k=1}^{K} \omega_{i,k}^t \mathbf{1}_{\{\pi_k(\omega^t)=i\}}, \quad i = 1, \dots, N,$$

$$U(\pi) = \limsup_{T\to\infty} \sum_{i=1}^{N} U_i^T(\pi).$$

- $R_i(\pi)$: worst-case fraction of time allocated to user i.
- $U(\pi)$: best-case system performance of policy π.
- Temporal fairness requirement: each user i requires at least r_i of time over all subcarriers, where $r_i \geq 0$ and $\sum_{i=1}^{N} r_i \leq 1$.
- $R_i(\pi) \geq r_i$.

Temporal Fairness Scheduling

Optimal Scheduling Problem

$$\max_{\pi \in \Pi} U(\pi)$$
$$\text{subject to } R_i^T(\pi) \geq r_i, \quad i = 1, \ldots, N$$

Optimal Scheduling Policy

$$\pi^*(\omega^t) = \operatorname*{argmax}_{\bar{A}^t} \left\{ \sum_{i=1}^{N} \sum_{k=1}^{K} \left(\omega_{i,k}^t + v_i^* \right) \mathbf{1}_{\{\pi_k(\omega^t)=i\}} \right\},$$

where, for all i, v_i^* satisfies:

1. $v_i^* \geq 0$;
2. $R_i(\pi^*) \geq r_i$;
3. If $R_i(\pi^*) > r_i$, then $v_i^* = 0$.

Remarks on Optimal Policy

- π^* is a stationary, memoryless, index policy.
- $\vec{v}^* = (v_1^*, \ldots, v_N^*)$: an "offset" or "threshold" to satisfy the fairness constraints.
- $\vec{v}^*$ can be estimated online in practice (e.g., via a stochastic approximation algorithm).
- Similar results apply for other fairness/QoS criteria:
  - utilitarian fairness,
  - minimum-performance guarantees,
  - proportional fairness.

Assignment Problem

Recall our optimal policy

$$\pi^*(\omega^t) = \underset{\bar{A}^t}{\operatorname{argmax}} \left\{ \sum_{i=1}^{N} \sum_{k=1}^{K} \left(\omega_{i,k}^t + v_i^*\right) \mathbf{1}_{\{\pi_k(\omega^t)=i\}} \right\}$$

Assignment Problem

Recall our optimal policy

$$\pi^*(\omega^t) = \underset{\bar{A}^t}{\text{argmax}} \left\{ \sum_{i=1}^{N} \sum_{k=1}^{K} \left(\omega_{i,k}^t + v_i^*\right) \mathbf{1}_{\{\pi_k(\omega^t)=i\}} \right\}$$

Assignment Problem

$$\begin{aligned} \text{maximize} & \sum_{i=1}^{N} \sum_{k=1}^{K} c_{ik} x_{ik} \\ \text{subject to} & \sum_{i=1}^{N} x_{ik} = 1, \quad k = 1, \ldots, K, \\ & \sum_{k=1}^{K} x_{ik} \leq 1, \quad i = 1, \ldots, N, \\ & x_{ik} \in \{0,1\},\ c_{ik} \geq 0,\ N \geq K. \end{aligned}$$

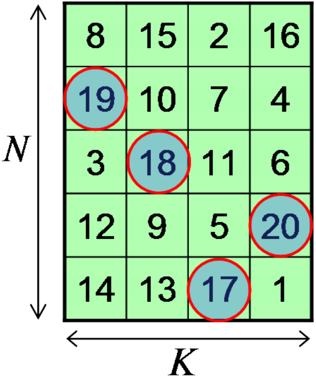

Algorithms for Assignment Problem

- Brute-force approach: exhaustively search over $\binom{N}{K}$ possible choices.
- Hungarian algorithm and our modified Hungarian algorithm are optimal solutions to this problem with $O(N^3)$ $(N = K)$ computation.
- Max-max algorithm:
  - a variation of "min-min" method.
  - a heuristic, suboptimal algorithm with lower computation complexity.

System Gain over Round-Robin

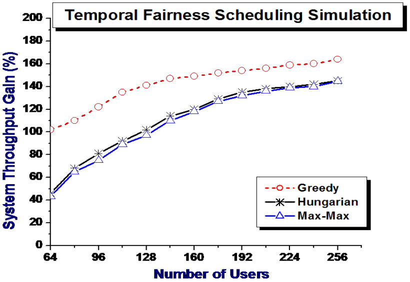

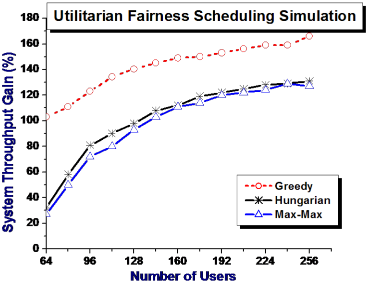

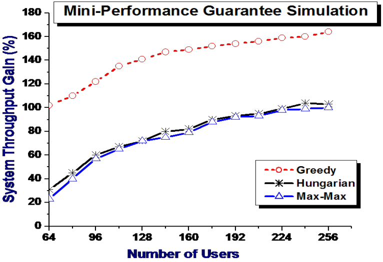

Temporal Fairness Among Users

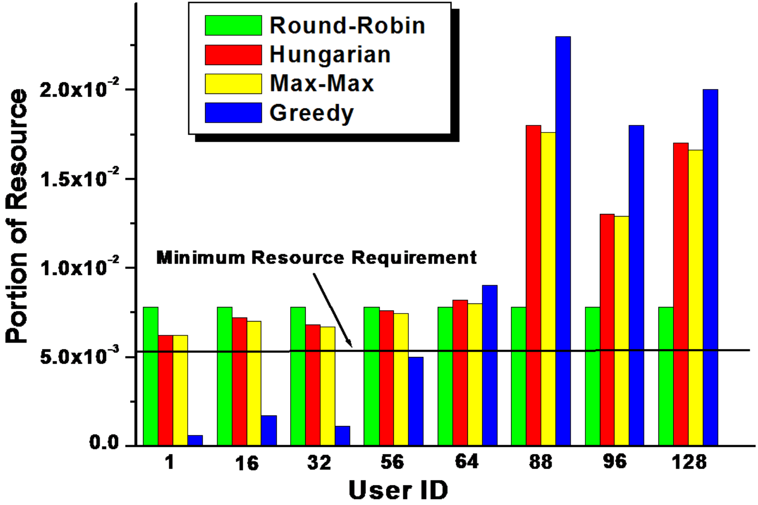

Outline

Motivation

- Opportunistic scheduling: schedule a single user with the best instantaneous channel condition to transmit.
- Multipath Rayleigh fading channels: Finite-State Markov Channel (FSMC) model.
- Elastic traffic vs. non-elastic traffic.
- Markov decision processes (MDPs) and dynamic programming: sequential decision making under uncertainty.

A TDM Uplink Queueing Model

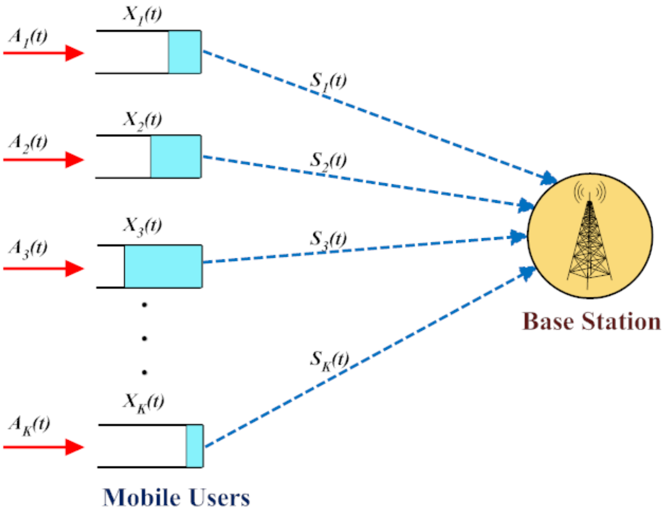

Fair Scheduling Problem

- Goal: maximize the system performance under certain QoS requirements of users.
  - Throughput maximization.
  - Delay minimization: average system queue length.
- Scheduler must allocate resources fairly among users under specific fairness/QoS constraints.
- Scheduling decision depends on
  - Instantaneous channel conditions $S(t)$.
  - Packet queue lengths $X(t)$.
  - Exogenous packet arrivals $A(t)$.
  - Specific fairness/QoS requirements.

Problem Formulation as a MDP

- **State space** S: the set of all vectors $s \in \mathbb{R}^{2K}$

$$s = (x_1, x_2, \dots, x_K, s_1, s_2, \dots, s_K)$$

where x_k and s_k are the queue length and the channel-state information of user k. The state of the system at time slot t is

$$X_t = (X_1(t), X_2(t), \dots, X_K(t), S_1(t), S_2(t), \dots, S_K(t)).$$

- **Action space** A: $\{1, 2, \dots, K\}$.
- **Transition probability function**: determined by
  - queue-length evolution.
  - dynamics of the channels.
- **Reward**:
  - Throughput maximization:

$$r(X_t, \pi_t) = \sum_{k=1}^{K} \mathbf{1}_{\{\pi_t = k\}} \min\left(X_k(t), S_k(t)\right).$$

  - Delay minimization:

$$r(X_t, \pi_t) = -\sum_{k=1}^{K} X_k(t).$$

Temporal Fair Scheduling Problem

Constrained MDP Formulation

$$\max_{\pi \in \Pi} J_\pi(s) = E_\pi \left[\sum_{t=0}^{\infty} \alpha^t r(X_t, \pi_t) \middle| X_0 = s \right], \quad s \in S,$$

subject to

$$\lim_{T \to \infty} \frac{1}{1/(1-\alpha)} E_\pi \left[\sum_{t=0}^{T} \alpha^t \mathbf{1}_{\{\pi_t = a\}} \middle| X_0 = s \right] \geq C(a), \ \forall a \in A. \qquad (1)$$

- Eqn. (1): the expected discounted temporal fairness constraint.
- $C(a)$: the minimum discounted time-fraction in which action (user) a should be chosen.
- $0 \leq C(a) \leq 1$ and $\sum_{a \in A} C(a) \leq 1$.

Optimal and Suboptimal Scheduling Policies

Optimal Scheduling Policy

- Explicit optimal Bellman equation for the above constrained MDP.
- Trade-off between immediate and future rewards.
- Optimal scheduling policy based on the equation.

Temporal Fair Rollout

- An efficient suboptimal approximation method: one-step lookahead.
- Optimal value function is approximated by some base policy: opportunistic fair scheduling.
- Achieve much better performance than the base policy.

$\Rightarrow$ *rollout improvement property*.

Performance for Delay Minimization Problem

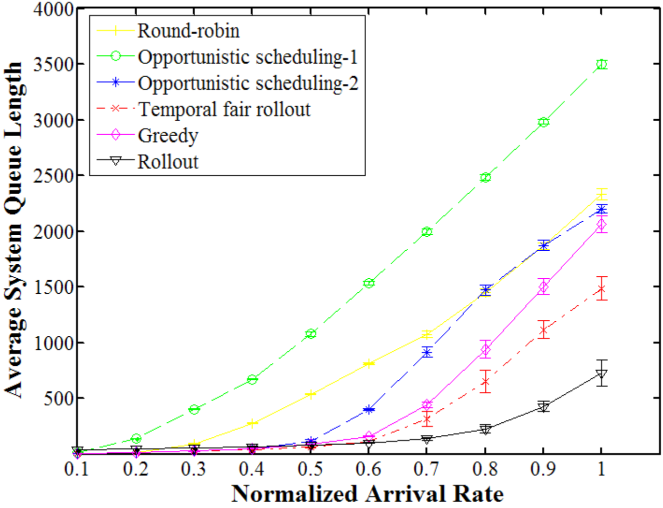

Figure: Average system queue length vs. normalized arrival rate.

Fairness in Delay Minimization Problem

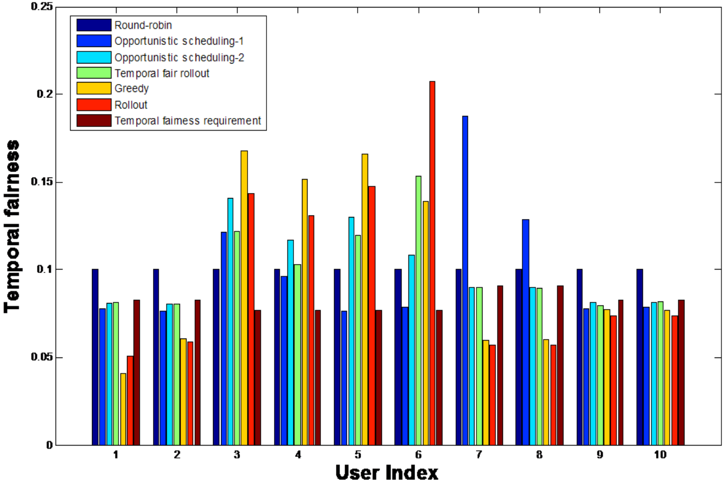

Figure: Time fraction allocation for delay minimization problem.

Outline

Outline

Motivation

- Future wireless multimedia networks require heterogeneous QoS support.
- A single user could require multiple different fairness/QoS constraints.
- Maximal constraints are useful for implementing multi-tiered services.
- Maximal constraints can decrease the subscribers' QoS sensitivity to the number of subscribers in the network.

Scheduling with Maximal and Minimum Bounds

Scheduling with Maximal and Minimum Data Rates

$$\max_{\pi \in \Pi} E\left(U(\pi)\right) = \sum_{i=1}^{N} E\left(U_i \mathbf{1}_{\{\pi(\vec{U})=i\}}\right) \tag{2}$$

$$\text{subject to } E\left(U_i \mathbf{1}_{\{\pi(\vec{U})=i\}}\right) \geq C_i, \qquad i = 1, 2, \ldots, N,$$

$$E\left(U_i \mathbf{1}_{\{\pi(\vec{U})=i\}}\right) \leq D_i, \qquad i = 1, 2, \ldots, N.$$

- $\vec{C} = (C_1, C_2, \ldots, C_N)$: a feasible predetermined *minimum data rate requirement vector*.
- $\vec{D} = (D_1, D_2, \ldots, D_N)$: a feasible predetermined *maximal data rate requirement vector*.
- $D_i \geq C_i \geq 0, \ \forall i$.
- For simplicity, here assume that channels are *stationary* and *ergodic*.

Optimal Scheduling Policy

Define the policy π^* as follows:

$$\pi^*(\vec{U}) = \operatorname*{argmax}_i \left((\theta_i - \mu_i) U_i \right), \tag{3}$$

where the control parameters θ_i and μ_i are chosen such that:

1. $\theta_i \geq 1$, $0 \leq \mu_i \leq 1$, $\forall i$;
2. $E\left(U_i \mathbf{1}_{\{\pi^*(\vec{U})=i\}}\right) \geq C_i$, $\forall i$;
3. If $E\left(U_i \mathbf{1}_{\{\pi^*(\vec{U})=i\}}\right) > C_i$, then $\theta_i = 1$, $\forall i$;
4. $E\left(U_i \mathbf{1}_{\{\pi^*(\vec{U})=i\}}\right) \leq D_i$, $\forall i$;
5. If $E\left(U_i \mathbf{1}_{\{\pi^*(\vec{U})=i\}}\right) < D_i$, then $\mu_i = 0$, $\forall i$.

Theorem

The policy π^* defined in (3) is an optimal solution to the problem defined in (2), i.e., it maximizes the system performance while satisfying the maximum and minimum data rate requirements for individual users.

Scheduling with Mixed Constraints

Scheduling with Minimum Temporal Fraction and Data Rate

$$\max_{\pi \in \Pi} E\left(U(\pi)\right) \tag{4}$$

$$\text{subject to } P\{\pi(\vec{U}) = i\} \geq r_i, \qquad i = 1, 2, \ldots, N,$$

$$E\left(U_i \mathbf{1}_{\{\pi(\vec{U})=i\}}\right) \geq C_i, \qquad i = 1, 2, \ldots, N.$$

- r_i: the minimum time-fraction that should be assigned to user i.
- $\vec{C} = (C_1, C_2, \ldots, C_N)$: a feasible predetermined *minimum data rate requirement vector*.
- $r_i \geq 0$ and $\sum_{i=1}^{N} r_i \leq 1$.

Optimal Scheduling Policy

Define the policy π^* as follows:

$$\pi^*(\vec{U}) = \underset{i}{\operatorname{argmax}} \left(\theta_i U_i + \alpha_i\right), \tag{5}$$

where the control parameters α_i and θ_i are chosen such that:

1. $\alpha_i \geq 0$, $\theta_i \geq 1$, $\forall i$;
2. $P\{\pi^*(\vec{U}) = i\} \geq r_i$, $\forall i$;
3. If $P\{\pi^*(\vec{U}) = i\} > r_i$, then $\alpha_i = 0$, $\forall i$;
4. $E\left(U_i \mathbf{1}_{\{\pi^*(\vec{U})=i\}}\right) \geq C_i$, $\forall i$;
5. If $E\left(U_i \mathbf{1}_{\{\pi^*(\vec{U})=i\}}\right) > C_i$, then $\theta_i = 1$, $\forall i$.

Theorem

The policy π^* defined in (5) is an optimal solution to the problem defined in (4), i.e., it maximizes the system performance while satisfying the minimum time fraction and data rate requirements for individual users.

Future Work

- Derive the scheduling schemes for different combinations of QoS/fairness constraints (including maximal and minimum).
- Find out a generalized QoS-constrained opportunistic scheduling framework to model and solve the above scheduling problems.
- Compare the related results with ours and evaluate the performance of our work by simulations.
- Summer 2008.

Outline

Problem Definition

- For any policy π, we define

$$J_\pi(s) = \lim_{T\to\infty} \frac{1}{T} E_\pi \left[\sum_{t=0}^{T-1} r(X_t, \pi(X_t)) \middle| X_0 = s \right], s \in \mathcal{S}.$$

- The policy π^* is *average-reward-optimal* if

$$J_{\pi^*}(s) = \max_{\pi\in\Pi} J_\pi(s), \qquad \forall s \in \mathcal{S}.$$

- Goal: to find an average-reward-optimal policy π^*, while maintaining certain constraints on actions.

Average Reward MDP with Temporal Fairness

$$\max_{\pi\in\Pi} J_\pi(s) \qquad s \in \mathcal{S} \tag{6}$$

$$\text{subject to } \lim_{T\to\infty} \frac{1}{T} E_\pi \left[\sum_{t=0}^{T-1} \mathbf{1}_{\{\pi(X_t)=a\}} \middle| X_0 = s \right] \geq C(a), \quad \forall a \in \mathcal{A}.$$

$C(a)$: the minimum relative frequency that should take action a, $C(a) \geq 0$ and $\sum_{a\in\mathcal{A}} C(a) \leq 1$.

Optimal Scheduling Policy

Conjecture

Suppose we have a bounded function $h : S \rightarrow \mathbb{R}$, a function $u : A \rightarrow \mathbb{R}$, a constant g, and a stationary unichain policy π^* such that for $s \in S$,

1. $\forall a \in A$, $u(a) \geq 0$;
2. $\forall a \in A$, $\lim_{T\rightarrow\infty} \frac{1}{T} E_{\pi^*} \left[\sum_{t=0}^{T-1} \mathbf{1}_{\{\pi^*(X_t)=a\}} \middle| X_0 = s \right] \geq C(a)$;
3. $\forall a \in A$, if $\lim_{T\rightarrow\infty} \frac{1}{T} E_{\pi^*} \left[\sum_{t=0}^{T-1} \mathbf{1}_{\{\pi^*(X_t)=a\}} \middle| X_0 = s \right] > C(a)$, then $u(a) = 0$;
4. $$g + h(s) = \max_a \{r(s,a) + u(a) + \sum_{j \in S} p(j|s,a)h(j)\};$$
5. $$\pi^*(s) = \operatorname*{argmax}_a \{r(s,a) + u(a) + \sum_{j \in S} p(j|s,a)h(j)\}.$$

.

Then π^* is an average-reward-optimal policy for problem (6).

Future Work

- Derive the optimal scheduling policy for average reward MDP with temporal fairness constraints: start from our conjecture.
- Investigate the corresponding policy iteration and value iteration method.
- Develop some efficient approximation methods (such as rollout).
- Possible application in wireless sensor networks: network lifetime.
- Summer-fall 2008.

The End

THANK YOU!

OFDM and OFDMA