Distributed Interference Compensation for Multi-channel Wireless Networks

Jianwei Huang, Randall Berry, Michael L. Honig

Department of Electrical Engineering & Computer Science
Northwestern University

Allerton Conference
Sept. 29, 2005

Motivation

- Communication networks with multiple interfering channels:
    - Wireless ad hoc networks with frequency selective channels.
    - Multi-cell OFDM networks with inter-cell interference.
    - DSL with crosstalk.
- How to mitigate interference and improve network performance?

Multi-channel Wireless Network Model

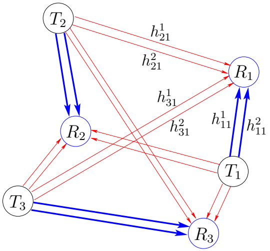

- Several transmitter-receiver pairs (users).
  - Single-hop and half-duplex transmissions.
- Channels gains are fixed (slow fading).
- Concurrent transmissions cause mutual interference.
- No centralized controller.
- How can interference mitigation be done efficiently in a distributed way with limited (scalable) information exchange?

Related Work and Our Contributions

- Most previous work has focused on maximization of achievable rate region:
  - Partially characterized capacity regime [Costa'85].
  - Centralized optimal power allocation for interference channel. [Cendrillon et. al.'05].
  - Decentralized algorithms: iterative water-filling, [Yu et. al.'02], signal space partitioning [Popescu et. al.'04].

Related Work and Our Contributions

- Most previous work has focused on maximization of achievable rate region:
  - Partially characterized capacity regime [Costa'85].
  - Centralized optimal power allocation for interference channel. [Cendrillon et. al.'05].
  - Decentralized algorithms: iterative water-filling, [Yu et. al.'02], signal space partitioning [Popescu et. al.'04].
- Our contributions:
  - Maximize total network utility (total rate is a special case).
  - Distributed price-based algorithm that achieves global optimality in certain cases.
  - Fast and robust convergence.

Talk Outline

- Network model and performance metric.
- **Asynchronous distributed pricing (ADP)** algorithm for power control.
- Optimality and convergence analysis:
  - Primal ADP algorithm
  - Dual ADP algorithm

Network Model

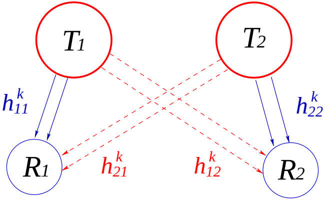

- I transmitter-receiver pairs (users).
- K parallel channels for all users.

Network Model

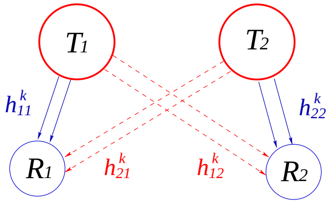

- I transmitter-receiver pairs (users).
- K parallel channels for all users.
- User i transmits in channel k with power $p_i^k \geq P_i^{\min}$ such that

$$\sum_{k=1}^{K} p_i^k \leq P_i^{\max} \quad (\textit{individual total power constraint})$$

- Received Signal-to-Interference plus Noise Ratio (SINR) in channel k for user i:

$$\gamma_i^k = \frac{p_i^k h_{ii}^k}{n_0^k + \sum_{j \neq i} p_j^k h_{ji}^k}$$

Utility Functions

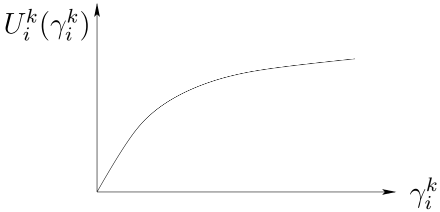

- User i's QoS preference is given by utility

$$U_i = \sum_k U_i^k(\gamma_i^k).$$

  - Separable across channels.
  - $U_i^k(\gamma_i^k)$ is increasing and strictly concave in γ_i^k.
  - Rate-adaptive applications with elastic demands.
  - Private information, only known to the user.
- Network performance = total network utility

Total Utility Maximization Problem

- **Goal**: allocate power in a distributed way to maximize total utility.
- **Challenges**:
  - Power assignment across channels is coupled due to individual total power constraints.
  - Power assignment across users is coupled due to mutual interferences.
  - Objective function may not be concave in power.
- **Our approach**: distributed cooperation by exchange of interference prices.

Asynchronous Distributed Pricing (ADP) Algorithm

- Repeat two steps asynchronously across users.
  - **Price announcement**: users announce prices in each channel, which reflect sensitivities to interferences.
  - **Power updates**: users update powers to maximize their surplus (utility - payment).

Asynchronous Distributed Pricing (ADP) Algorithm

- Repeat two steps asynchronously across users.
  - **Price announcement**: users announce prices in each channel, which reflect sensitivities to interferences.
  - **Power updates**: users update powers to maximize their surplus (utility - payment).
- Prices can be interpreted as a Pigovian tax that internalizes the negative externality caused by interference.

Asynchronous Distributed Pricing (ADP) Algorithm

- Repeat two steps asynchronously across users.
  - **Price announcement**: users announce prices in each channel, which reflect sensitivities to interferences.
  - **Power updates**: users update powers to maximize their surplus (utility - payment).
- Prices can be interpreted as a Pigovian tax that internalizes the negative externality caused by interference.
- **Primal** and **Dual** ADP differ in how to deal with the coupling across channels due to individual total power constraints.

Primal ADP Algorithm

- Solve total utility maximum problem directly.

Primal ADP Algorithm

- Solve total utility maximum problem directly.
- **Price Announcement**: user i announces an "interference price" π_i^k in channel k (one price per user per channel)

$$\pi_i^k = \left| \frac{\partial U_i^k(\gamma_i^k)}{\partial \left(\sum_{j \neq i} p_j^k h_{ji}^k \right)} \right|.$$

Primal ADP Algorithm

- Solve total utility maximum problem directly.
- **Price Announcement**: user i announces an "interference price" π_i^k in channel k (one price per user per channel)

$$\pi_i^k = \left| \frac{\partial U_i^k(\gamma_i^k)}{\partial \left(\sum_{j \neq i} p_j^k h_{ji}^k \right)} \right| .$$

- **Power Update**: user i updates the power in all channels to maximize surplus (need local channel measurements)

$$\max \sum_{k=1}^{K} \left(\underbrace{\overbrace{U_i^k(\gamma_i^k(p_i^k))}^{\text{Utility}} - \overbrace{p_i^k \sum_{j \neq i} \pi_j^k h_{ij}^k}^{\text{Payment}}}_{\text{Surplus in channel } k} \right)$$

$$s.t. \sum_{k=1}^{K} p_i^k \leq P_i^{\max}$$

Properties of Primal ADP

- **Optimality**: a power and price profile $(\mathbf{p}^*, \boldsymbol{\pi}^*)$ is a fixed point of the Primal ADP algorithm if and only if $\mathbf{p}^*$ satisfies the Kuhn-Tucker condition of the total utility maximization problem.
  - A unique fixed point corresponds to the global optimal power allocation.

Properties of Primal ADP

- **Optimality**: a power and price profile $(\mathbf{p}^*, \boldsymbol{\pi}^*)$ is a fixed point of the Primal ADP algorithm if and only if $\mathbf{p}^*$ satisfies the Kuhn-Tucker condition of the total utility maximization problem.
  - A unique fixed point corresponds to the global optimal power allocation.
- **Convergence**: depends on system parameters
  - Game theoretic approach: converges with unique or multiple fixed points.
  - Contraction mapping approach: converges to a unique fixed point.

Fictitious Power-Price Game

- Model each user as two fictitious players in a game:
  - Price player chooses prices to maximize payoff

$$S^{price}(\boldsymbol{\pi}_i) = -\sum_{k=1}^{K} \left(\pi_i^k - \left| \frac{\partial U_i^k(\gamma_i^k)}{\partial \left(\sum_{j \neq i} p_j^k h_{ji}^k \right)} \right| \right)^2$$

  - Power player chooses power to maximize payoff

$$S^{power}(\mathbf{p}_i) = \sum_{k=1}^{K} \left(U_i^k(\gamma_i^k(p_i^k)) - p_i^k \sum_{j \neq i} \pi_j^k h_{ij}^k \right).$$

Fictitious Power-Price Game

- Model each user as two fictitious players in a game:
  - Price player chooses prices to maximize payoff

$$S^{price}(\boldsymbol{\pi}_i) = -\sum_{k=1}^{K}\left(\pi_i^k - \left|\frac{\partial U_i^k(\gamma_i^k)}{\partial\left(\sum_{j\neq i} p_j^k h_{ji}^k\right)}\right|\right)^2$$

  - Power player chooses power to maximize payoff

$$S^{power}(\mathbf{p}_i) = \sum_{k=1}^{K}\left(U_i^k(\gamma_i^k(p_i^k)) - p_i^k \sum_{j\neq i} \pi_j^k h_{ij}^k\right).$$

- Primal ADP algorithm corresponds to the best response updates of the fictitious game.
  - A fixed point of the algorithm is a Nash Equilibrium (NE) of the game.
- Convergence of algorithm = convergence of best response updates
  - Fast convergence when game is supermodular.

Convergence as a Supermodular Game

- **Supermodular game**:
  - A class of games with strategic complementaries.
  - Each player's payoff S_i has increasing differences:

$$\frac{\partial^2 S_i}{\partial x_i \partial x_j} \geq 0$$

Convergence as a Supermodular Game

- **Supermodular game**:
  - A class of games with strategic complementaries.
  - Each player's payoff S_i has increasing differences:

$$\frac{\partial^2 S_i}{\partial x_i \partial x_j} \geq 0$$

  - Key properties:
    (1) An NE exists.
    (2) The best response updates converge with proper initializations.
    (3) If the NE is unique, then the asynchronous best response updates globally converge to it.

Convergence as a Supermodular Game

- **Supermodular game**:
  - A class of games with strategic complementaries.
  - Each player's payoff S_i has increasing differences:

$$\frac{\partial^2 S_i}{\partial x_i \partial x_j} \geq 0$$

  - Key properties:
    - **(1)** An NE exists.
    - **(2)** The best response updates converge with proper initializations.
    - **(3)** If the NE is unique, then the asynchronous best response updates globally converge to it.
- **Prop**: when $K = 1, 2$, with suitable restrictions on utilities and users, the fictitious game is supermodular, and
  - The Primal ADP algorithm converges with proper initializations.
  - If there is a unique fixed point, the algorithm globally converges to it.

Convergence of Primal ADP

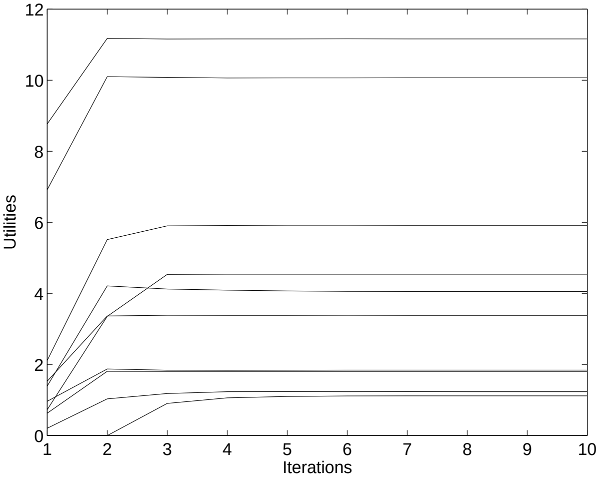

$U_i = \sum_{k=1}^{K} \log\left(1+\gamma_i^k\right)$, 20 channels, 10 users

Convergence Analysis: Contraction Mapping

- Primal ADP converges globally to the unique optimal solution for:
  - Two users and multiple channels.
  - $U_i = \sum_{k=1}^{K} \log\left(1 + \gamma_i^k\right)$.
  - Small normalized interference: $h_{ji}^k / h_{ii}^k \le \zeta$ for all users and channels.

Convergence Analysis: Contraction Mapping

- Primal ADP converges globally to the unique optimal solution for:
  - Two users and multiple channels.
  - $U_i = \sum_{k=1}^{K} \log\left(1 + \gamma_i^k\right)$.
  - Small normalized interference: $h_{ji}^k / h_{ii}^k \leq \zeta$ for all users and channels.
- Compare: **Iterative Water-Filling** (IWF) [Yu et. al.'02]
  - No exchange of price information.
  - Each user chooses power to maximize $\sum_{k=1}^{K} \log\left(1 + \gamma_i^k\right)$.
  - Converges under similar conditions (not necessarily optimal).
- Primal ADP achieves better performance than IWF.

Comparison of PADP and IWF

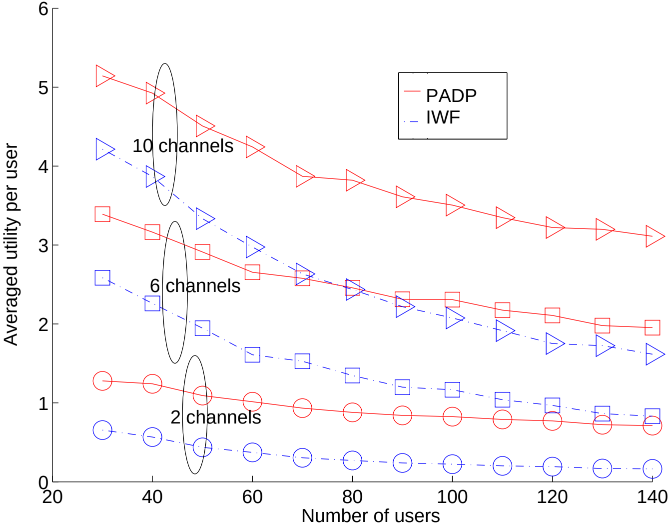

Dual ADP Algorithm

- User i relaxes the total power constraint by using a dual price μ_i.

Dual ADP Algorithm

- User i relaxes the total power constraint by using a dual price μ_i.
- **Primal Updates**: for fixed dual prices, in each channel k,
  - User announces interference price π_i^k as before.
  - User chooses power p_i^k to maximize surplus in channel k

$$\max_{p_i^k \in \left[P_i^{\min}, P_i^{\max}\right]} U_i^k(\gamma_i^k(p_i^k)) - p_i^k \left(\sum_{j \neq i} h_{ij} \pi_j^k + \mu_i \right).$$

Dual ADP Algorithm

- User i relaxes the total power constraint by using a dual price μ_i.
- **Primal Updates**: for fixed dual prices, in each channel k,
  - User announces interference price π_i^k as before.
  - User chooses power p_i^k to maximize surplus in channel k

$$\max_{p_i^k \in \left[P_i^{\min}, P_i^{\max}\right]} U_i^k(\gamma_i^k(p_i^k)) - p_i^k \left(\sum_{j \neq i} h_{ij} \pi_j^k + \mu_i \right).$$

- **Dual Iterations**: update the dual prices as

$$\mu_i(t) = \left[\mu_i(t^-) + \kappa \left(\sum_{k \in \mathcal{K}} p_i^k(t^-) - P_i^{max} \right) \right]^+.$$

- Solves the dual of the total utility maximization problem.

Convergence of Dual ADP

- Converges with “sufficiently” concave utilities.

Convergence of Dual ADP

- Converges with "sufficiently" concave utilities.
- Define the coefficient of relative risk aversion of a utility $U_i^k(\gamma_i^k)$ to be

$$Q_i^k(\gamma_i^k) = -\frac{\gamma_i^k U_i^{k''}(\gamma_i^k)}{U_i^{k'}(\gamma_i^k)}.$$

  - larger $Q_i^k(\gamma_i^k) \Rightarrow$ "more concave" U_i^k.

Convergence of Dual ADP

- Converges with "sufficiently" concave utilities.
- Define the coefficient of relative risk aversion of a utility $U_i^k(\gamma_i^k)$ to be

$$Q_i^k(\gamma_i^k) = -\frac{\gamma_i^k U_i^{k''}(\gamma_i^k)}{U_i^{k'}(\gamma_i^k)}.$$

  - larger $Q_i^k(\gamma_i^k) \Rightarrow$ "more concave" U_i^k.

- **Prop:** If for each user i:
  (1) $P_i^{\min} > 0$, and
  (2) $Q_i^k(\gamma_i^k) \in [1, 2]$ for all feasible SINR γ_i^k;
  then there is a unique optimal power allocation, and the Dual ADP algorithm globally and geometrically converges to it with small enough stepsize κ.

Convergence of Dual ADP

- The constraint $Q_i^k(\gamma_i^k) \in [1, 2]$ ensures the problem is concave in transformed variables $log(p_i^k)$ [Chiang'05].

Convergence of Dual ADP

- The constraint $Q_i^k(\gamma_i^k) \in [1, 2]$ ensures the problem is concave in transformed variables $log(p_i^k)$ [Chiang'05].
- Convergence with constant step-size κ depends on the Lipschitz condition for the gradient of the dual function.
  - The dual function is not separable across users (different from dual decomposition in wireline networks [Low and Lapsley'99])

Convergence of Dual ADP

- The constraint $Q_i^k(\gamma_i^k) \in [1,2]$ ensures the problem is concave in transformed variables $log(p_i^k)$ [Chiang'05].
- Convergence with constant step-size κ depends on the Lipschitz condition for the gradient of the dual function.
  - The dual function is not separable across users (different from dual decomposition in wireline networks [Low and Lapsley'99])
- The proof needs "separation of time-scales":
  - The Primal Updates need to converge between two successive dual iterations.
  - Simulations show that the algorithm converges without separation of time-scales.

Convergence of Dual ADP

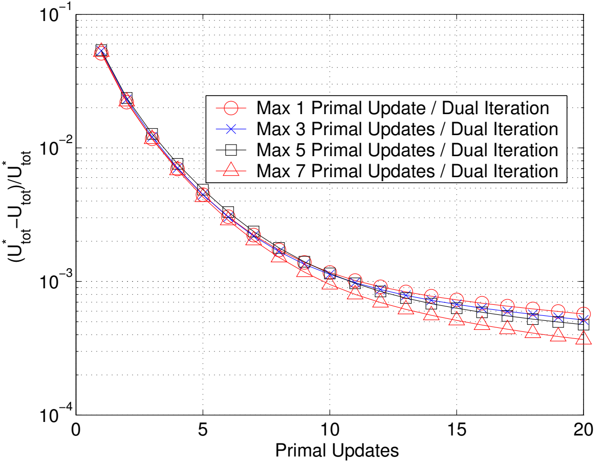

$U_i = \sum_{k=1}^{K} \log\left(\gamma_i^k\right)$, 16 channels, 50 users

Conclusions

- Presented distributed power control algorithm for multi-channel wireless networks.
- Users exchange price information to reflect their sensitivities to interference (Pigovian tax).
- Proved optimality and convergence of **Primal** and **Dual** ADP algorithms in various cases.
- Observed convergence in more general settings.
- Currently extending to multi-hop networks with joint scheduling and power control.