Distributed Wireless Resource Allocation Game

Jianwei Huang

Department of Electrical Engineering & Computer Science
Northwestern University

Joint work with R. Berry and M. Honig

Motorola Labs Seminar
Sept. 9, 2005

Wireless Resource Allocation

- Increasing number of wireless technologies
  - Cellular, wireless LANs, WiMax, sensor networks, etc.
- Limited radio resources:
  - Frequency bands, time slots, orthogonal codes, transmission power, etc.
- How to allocate limited resources to support users' increasing demands and heterogeneous Quality of Service (QoS) requirements?

Wireless Ad Hoc Networks

- Formed by a group of wireless nodes without a centralized controller
  - Applications: battle field, disaster relief, sports events, ...
- Fundamental problem:
  - Conflicts: individual QoS requirements, energy constraints.
  - Cooperation: interferences mitigation, data relay.
  - Decisions have to be made distributively.
- Game theory can be used to analyze the conflicts and cooperations of distributed decision makers.

Talk Outline

- Network Model and Problem Formulation
- Game Theory Background
- Analysis and Solutions

Network Model and Problem Formulation

Ad Hoc Network Model

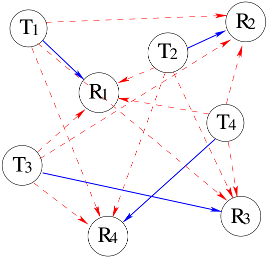

- Single-hop transmission.
    - No multi-hop relay traffic.
- One-to-one half duplex transmission.
    - Each active node can only transmit to (or receive from) one other node.
- Multiple access based on spread spectrum.
    - Concurrent transmissions in the same frequency band (channel).
    - Mutual interference is the performance bottleneck.

User Representation

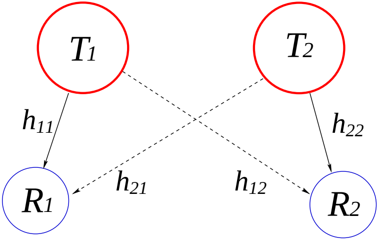

- A network with two transmitter-receiver node pairs.

User Representation

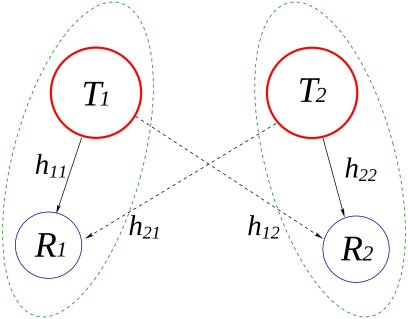

- A network with two transmitter-receiver node pairs.

User Representation

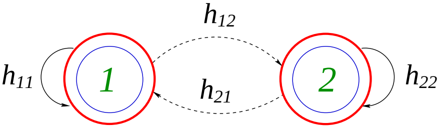

- A network with two transmitter-receiver node pairs.
- Can be represented as two users.

Signal-to-Interference Plus Noise Ratio (SINR)

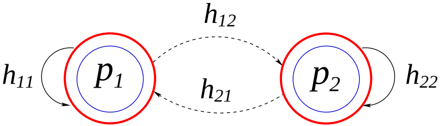

- User i's transmitter transmits at power $p_i \in [P_i^{\min}, P_i^{\max}]$ over the entire bandwidth of BHz.
- User i's QoS depends on SINR at its receiver:

$$\gamma_i = \frac{p_i h_{ii}}{n_0 + \frac{1}{B}\sum_{j\neq i} p_j h_{ji}}$$

Signal-to-Interference Plus Noise Ratio (SINR)

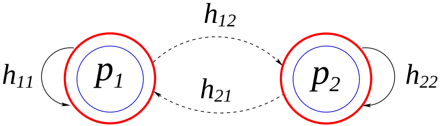

- User i's transmitter transmits at power $p_i \in [P_i^{\min}, P_i^{\max}]$ over the entire bandwidth of BHz.
- User i's QoS depends on SINR at its receiver:

$$\gamma_i = \frac{p_i h_{ii}}{n_0 + \frac{1}{B}\sum_{j \neq i} p_j h_{ji}}$$

- Interference Mitigation = Power Allocation.

User Preferences

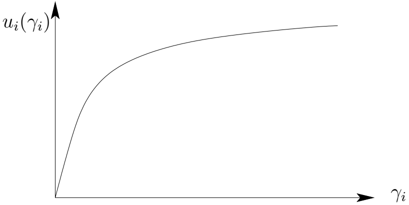

- User's QoS preferences given by utility $U_i(\gamma_i)$.
  - Increasing and strictly concave in γ_i.
  - Rate-adaptive applications with elastic demands.
  - Private information, only known by the user himself.
- **Goal**: find distributed power allocation to maximize users' total utility.

Total Utility Maximization Problem

- Related work
  - Power control with target SINRs: [Holiday, Goldsmith, Glynn'03]
  - Gradient-based search: [Chiang'05]
- **Our approach**:
  - Construct a game theoretic model for distributed power control.
  - Develop algorithms that converge fast and perform well under practical conditions.

Game Theory Background

What is Game Theory?

- The study of conflicts and cooperations between rational decision-makers.
  - Interactive decisions affecting each other.
  - Each decision maker wants to maximize one's own benefit.

Basic Elements of a Game

- Three elements
  - A set of two or more players.
  - Each player has a set of action choices.
  - Each player has a payoff that is a function of the action profile (all players' actions).

Basic Elements of a Game

- Three elements
  - A set of two or more players.
  - Each player has a set of action choices.
  - Each player has a payoff that is a function of the action profile (all players' actions).

Example

- Player 1's choices: $\{Silent, Transmit\}$
  Player 2's choices: $\{Silent, Transmit\}$
- (Player 1's payoff,Player 2's payoff)

		Player 2	
		Silent	*Transmit*
Player 1	*Silent*	$(0,0)$	$(0,8)$
	Transmit	$(8,0)$	$(3,3)$

Basic Elements of a Game

- Three elements
  - A set of two or more players.
  - Each player has a set of action choices.
  - Each player has a payoff that is a function of the action profile (all players' actions).

Example

- Player 1's choices:
  $\{Silent, Transmit\}$
  Player 2's choices:
  $\{Silent, Transmit\}$
- (Player 1's payoff,Player 2's payoff)

		Player 2	
		Silent	*Transmit*
Player 1	*Silent*	$(0,0)$	$(0,8)$
	Transmit	$(8,0)$	$(3,3)$

Basic Elements of a Game

- Three elements
  - A set of two or more players.
  - Each player has a set of action choices.
  - Each player has a payoff that is a function of the action profile (all players' actions).

Example

- Player 1's choices:
  {*Silent*, *Transmit*}
  Player 2's choices:
  {*Silent*, *Transmit*}
- (Player 1's payoff,Player 2's payoff)

		Player 2	
		Silent	*Transmit*
Player 1	*Silent*	$(0,0)$	$(0,8)$
	Transmit	$(8,0)$	$(3,3)$

Basic Elements of a Game

- Three elements
  - A set of two or more players.
  - Each player has a set of action choices.
  - Each player has a payoff that is a function of the action profile (all players' actions).

Example

- Player 1's choices: {*Silent*, *Transmit*}
  Player 2's choices: {*Silent*, *Transmit*}
- (Player 1's payoff,Player 2's payoff)
- Game result: Nash Equilibrium

		Player 2	
		Silent	*Transmit*
Player 1	*Silent*	(0,0)	(0,8)
	Transmit	(8,0)	(3,3)

Nash Equilibrium (NE)

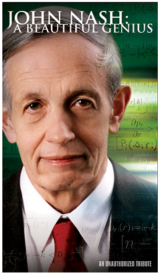

- NE is the most commonly used concept to describe the result of a game.
- NE is a vector of all players' actions, at which no player can increase his payoff by deviate unilaterally.
  - If a player thinks others will act according to NE, then it is to his best interest to do that as well.
- NE may not be socially optimal - the cost of selfishness and distributed decision making.

How to Find NE?

- NE is a state where every player is doing his best.
- **Best Response**: the action that maximizes player *i*'s payoff as a function of other players' actions.

$$BR_i(a_{-i}) = \arg\max_{a_i \in \mathcal{A}_i} Payoff_i(a_i, a_{-i})$$

  - Obtained by a single-player optimization (often relatively easy).

How to Find NE?

- NE is a state where every player is doing his best.
- **Best Response**: the action that maximizes player i's payoff as a function of other players' actions.

$$BR_i(a_{-i}) = \arg\max_{a_i \in \mathcal{A}_i} Payoff_i(a_i, a_{-i})$$

  - Obtained by a single-player optimization (often relatively easy).
- NE is a fixed point of all players' best responses
  - Difficult to find in general.
  - Easier for games with special structures.

Analysis and Solutions

Game Model I

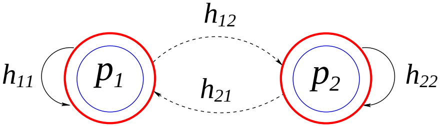

- Basic elements
  - Player = User
  - Action = Transmission Power $p_i \in [P_i^{\text{min}}, P_i^{\text{max}}]$
  - Payoff = Utility U_i

Game Model I

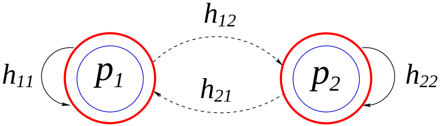

- Basic elements
  - Player = User
  - Action = Transmission Power $p_i \in [P_i^{\text{min}}, P_i^{\text{max}}]$
  - Payoff = Utility U_i
- **NE**: each user transmits at the maximum power.

Game Model I

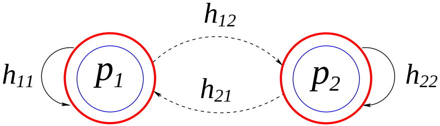

- Basic elements
  - Player = User
  - Action = Transmission Power $p_i \in [P_i^{\text{min}}, P_i^{\text{max}}]$
  - Payoff = Utility U_i
- **NE**: each user transmits at the maximum power.
- **Why**: no penalty for transmitting at high power levels.

Game Model I

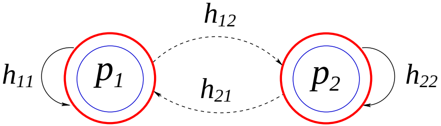

- Basic elements
  - Player = User
  - Action = Transmission Power $p_i \in [P_i^{\min}, P_i^{\max}]$
  - Payoff = Utility U_i
- **NE**: each user transmits at the maximum power.
- **Why**: no penalty for transmitting at high power levels.
- **Solution**: add penalty for large power (interference).

Pricing-based Power Control Algorithm

- **Price Announcing**: user i announces "price" π_i (per unit interference)
- **Power Updating**: user i updates power p_i to maximize surplus:

$$s_i = U_i(\gamma_i) - p_i \sum_{j \neq i} \pi_j h_{ij}.$$

- Can we formalize it as a game?

Game Model II

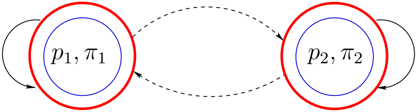

- Basic elements
    - Player = User
    - Action = Transmission Power and Interference Price (p_i, π_i)
    - Payoff = Surplus s_i

Game Model II

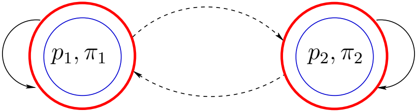

- Basic elements
  - Player = User
  - Action = Transmission Power and Interference Price (p_i, π_i)
  - Payoff = Surplus s_i
- **NE**: each user announces an infinite large price, and chooses $p_i = 0$ (or very small value).

Game Model II

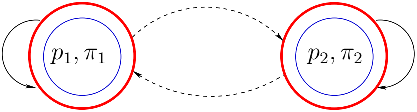

- Basic elements
  - Player = User
  - Action = Transmission Power and Interference Price (p_i, π_i)
  - Payoff = Surplus s_i
- **NE**: each user announces an infinite large price, and chooses $p_i = 0$ (or very small value).
- **Why**: no penalty for announcing a large price.

Game Model II

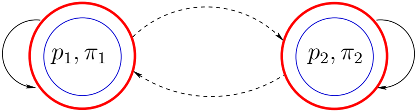

- Basic elements
  - Player = User
  - Action = Transmission Power and Interference Price (p_i, π_i)
  - Payoff = Surplus s_i
- **NE**: each user announces an infinite large price, and chooses $p_i = 0$ (or very small value).
- **Why**: no penalty for announcing a large price.
- **Solution**: prices should reflect users' sensitivities to the interference.

Game Model III: Fictitious Game

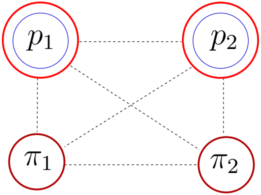

Fictitious Power Player (FP)

Fictitious Price Player (Fπ)

- Split each user into two fictitious players (FP and Fπ).

Game Model III: Fictitious Game

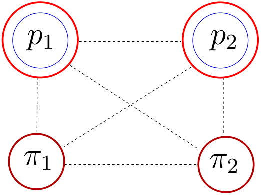

Fictitious Power Player (FP)

Fictitious Price Player (Fπ)

- Split each user into two fictitious players (FP and Fπ).
- Basic elements

Player	FP	Fπ
Action	Transmission Power p_i	Interference Price π_i
Payoff	Surplus s_i	$-\left(\pi_i - \frac{\partial u_i}{\partial(\sum_{j\neq i} p_j h_{ji})}\right)^2$
Best Response	$p_i = \arg\max_{\hat{p}_i} s_i(\hat{p}_i)$	$\pi_i = \frac{\partial U_i}{\partial(\sum_{j\neq i} p_j h_{ji})}$

Game Model III: Fictitious Game

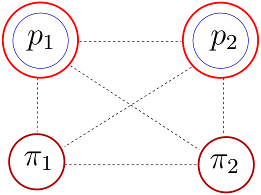

Fictitious Power Player (FP)

Fictitious Price Player (Fπ)

- Split each user into two fictitious players (FP and Fπ).
- Basic elements

Player	FP	Fπ
Action	Transmission Power p_i	Interference Price π_i
Payoff	Surplus s_i	$-\left(\pi_i - \frac{\partial u_i}{\partial(\sum_{j\neq i} p_j h_{ji})}\right)^2$
Best Response	$p_i = \arg\max_{\hat{p}_i} s_i(\hat{p}_i)$	$\pi_i = \frac{\partial U_i}{\partial(\sum_{j\neq i} p_j h_{ji})}$

- **NE**: optimal solution for the total utility maximization problem (for certain utility functions).

Asynchronous Distributed Pricing (ADP) Algorithm

- **Best Response Dynamics** of the fictitious game:
  - **Price Announcing**: user *i* announces "price" to reflect marginal utility sensitivities (best response of player Fπ):

$$\pi_i = -\frac{\partial u_i(\gamma_i)}{\partial \left(\sum_{j \neq i} p_j h_{ji}\right)}.$$

  - **Power Updating**: user *i* updates power p_i to maximize surplus (best response of player FP):

$$p_i = \arg\max_{\hat{p}_i} \left(U_i(\gamma_i(\hat{p}_i)) - \hat{p}_i \sum_{j \neq i} \pi_j h_{ij} \right).$$

- Repeat two phases asynchronously across users.

ADP Algorithm

- Scalable and distributed:
  - Each user only announces one price.
  - Each user only need to limited channel gains to maximize surplus.
- Marginal price can be interpreted as Pigovian taxation
  - Users pay for their behaviors (transmission power) that cause negative effects (interferences) to other users.
  - Typically centralized imposed by government.
- ADP algorithm discovers the right Pigovian taxes in a distributed fashion.
- **Question**: can ADP algorithm converge?

Convergence

- **Answer**: converges under mild conditions of the utility functions.

Convergence

- **Answer**: converges under mild conditions of the utility functions.
- Define the coefficient of relative risk aversion of a utility $U(\gamma)$ to be

$$Q(\gamma) = -\frac{\gamma U''(\gamma)}{U'(\gamma)}.$$

  - larger $Q(\gamma) \Rightarrow$ "more concave" U.

Convergence

- **Answer**: converges under mild conditions of the utility functions.
- Define the coefficient of relative risk aversion of a utility $U(\gamma)$ to be

$$Q(\gamma) = -\frac{\gamma U''(\gamma)}{U'(\gamma)}.$$

  - larger $Q(\gamma) \Rightarrow$ "more concave" U.

- **Prop:** If for all user i:

  (a) $P_i^{\min} > 0$, and

  (b) $Q(\gamma_i) \in [1, 2]$ for all feasible SINR γ_i;

  then there is a unique optimal allocation and the ADP algorithm globally converges to it.

  - e.g. above condition is always satisfied with log utilities $U(\gamma) = \log(\gamma)$.
  - proof: show the fictitious game is a supermodular game.

Supermodular Game

- Players' actions have positive feedback on each other
  - e.g., if one user increases its transmission power, then the other users would like to increase as well.
- How to identify:
  - Each player's best response is non-decreasing in other players' actions.
- Properties:
  - NEs generally exist and have nice structures.
  - Best response dynamics converges fast.
  - If there is a single NE, best response dynamics globally converge to it.

Convergence

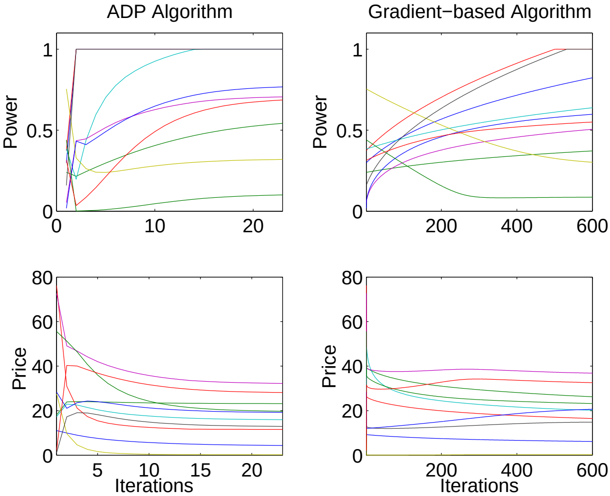

log utilities, 10 users

Limited Price Information Exchange

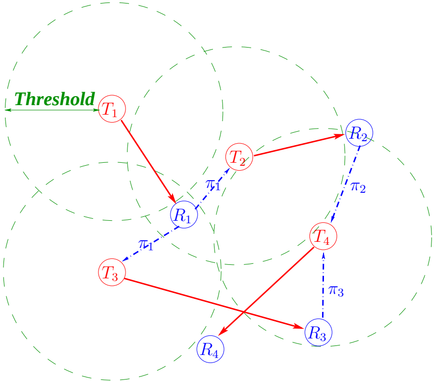

- Previously we assume that each user can decode all the prices.
- In practice, users (receivers) may only decode price messages within threshold distance.

ADP with Limited Pricing

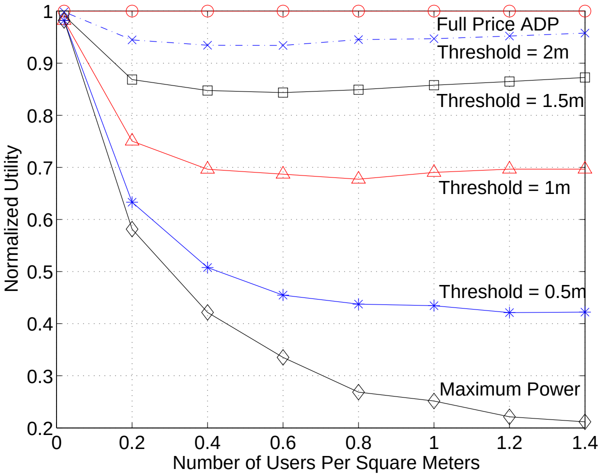

utility $\log(1+\gamma_i)$, 10m × 10m area

Comparison of ADP with 802.11 (RTS/CTS)

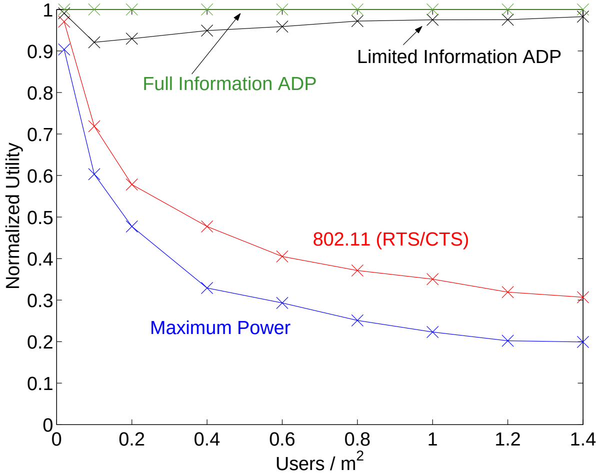

rate utility $\log(1+\gamma_i)$, 10m × 10m area

Comparison of ADP with 802.11 (RTS/CTS)

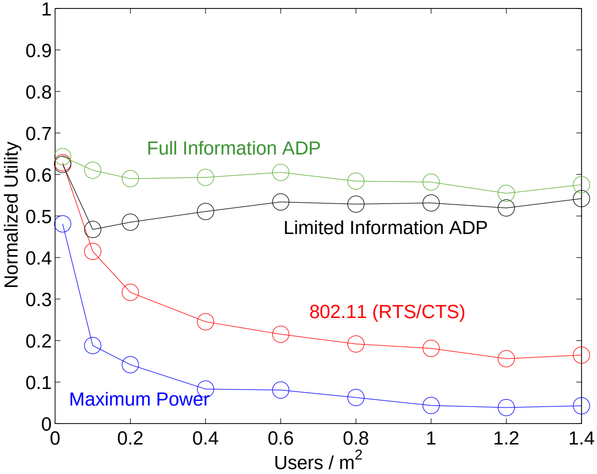

quantized utility $\log(1+\gamma_i)$ ($\{0,5,10,15,20\}bits/Hz$), 10m × 10m area

Multi-channel Model

- Assume each user can transmit over K independent channels (e.g. multi-carrier system).
- Received SINR in channel k for user i

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

- Can allocate power across channels subject to total power constraint:

$$\sum_k p_i^k \leq P_i^{max}.$$

- User's utility is "carrier separable"

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

Modified ADP algorithm

- Each user *i* keeps a dual price, μ_i, to represent total power constraint.
  - Decouple the power control problem across channels.

Modified ADP algorithm

- Each user i keeps a dual price, μ_i, to represent total power constraint.
    - Decouple the power control problem across channels.
- In each channel k, users perform a modified ADP algorithm:
    - **Interference price updates**: each user i announces an interference price π_i^k to reflect marginal utility loss.
    - **Power updates**: user chooses power p_i^k to maximize:

$$U_i^k(\gamma_i^k) - p_i^k \left(\sum_{j \neq i} h_{ij} \pi_j^k + \mu_i \right).$$

Modified ADP algorithm

- Each user i keeps a dual price, μ_i, to represent total power constraint.
  - Decouple the power control problem across channels.
- In each channel k, users perform a modified ADP algorithm:
  - **Interference price updates**: each user i announces an interference price π_i^k to reflect marginal utility loss.
  - **Power updates**: user chooses power p_i^k to maximize:

$$U_i^k(\gamma_i^k) - p_i^k \left(\sum_{j \neq i} h_{ij} \pi_j^k + \mu_i \right).$$

- **Dual price updates**:

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

  - Solves the dual of the total utility maximization problem.

Convergence

- Under similar restrictions on the utility functions to single channel case, this globally converges to unique optimal power allocation.
- The proof needs "separation of time-scales":
  - **Primal Updates**: the modified ADP in each channel under fixed dual prices.
  - **Dual Iterations**: the updates of dual prices.
  - The Primal Updates need to converge between two adjacent Dual Iterations.
- Simulation shows that algorithm converges without separation of time-scale.

Multi-channel Example

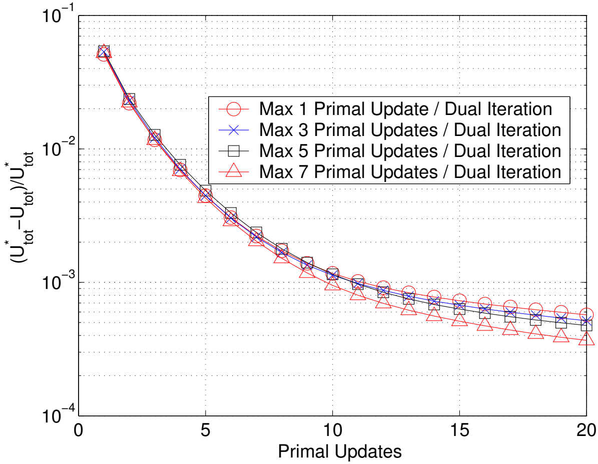

log utilities, 16 channels, 50 users

Conclusions and Outlook

- Distributed resource allocation can be well studied using game theory.
- Presented an asynchronous distributed pricing (ADP) algorithm for power control in wireless ad hoc networks.
- Analyzed the performance and convergence in single- and multi-channel networks under both idealistic and practical conditions.
- **Ongoing and future work**:
    - Joint scheduling and power control in multi-hop networks.
    - Competition and pricing for multiple service providers.
    - Incentive design in relay networks.

References

- J. Huang, R. Berry, and M. Honig, "Distributed Interference Compensation for Wireless Networks," submitted to *IEEE Journal on Selected Areas of Communications*, 2005.
- J. Huang, R. Berry, and M. Honig, "Spectrum Sharing with Distributed Interference Compensation, " to appear in *IEEE DySPAN (Dynamic Spectrum Access Networks)*, Baltimore, MD, Nov. 2005
- J. Huang, R. Berry, and M. Honig, "A Price-based Distributed Power Control Algorithm for Wireless Ad Hoc Networks, " to appear in *Proceedings of IEEE MILCOM*, Atlantic City, NJ, Oct. 2005
- J. Huang, R. Berry, and M. Honig, "Distributed Interference Compensation for Multi-channel Wireless Networks, " to appear in *Proceedings of Annual Allerton Conference on Communication, Control, and Computing*, Allerton, IL, Sept. 2005
- J. Huang, R. Berry, and M. Honig, "A Game Theoretic Analysis of Distributed Power Control for Spread Spectrum Ad Hoc Networks, " to appear in *Proceedings of IEEE International Symposium on Information Theory*, Adelaide, Australia, Sept. 2005