Near-Optimal Power Control in Wireless Networks: A Potential Game Approach

Utku Ozan Candogan, Ishai Menache,
Asu Ozdaglar and Pablo A. Parrilo

Laboratory for Information and Decision Systems
Electrical Engineering and Computer Science
Massachusetts Institute of Technology

ITMANET Meeting
September 2009

Motivation

- Game-theoretic analysis has been used extensively in the study of networks in general and in wireless networks in particular for two major reasons:
  - Game-theoretic tools enable a flexible control paradigm where agents autonomously control their resource usage to optimize their own selfish objectives
  - Even when selfish incentives are not present, game-theoretic models and tools provide potentially tractable decentralized algorithms for network control
- Most work on network games has focused on:
  - Static equilibrium analysis without establishing how an equilibrium can be reached dynamically
  - Properties of equilibria without systematically considering incentive mechanisms that can implement general system-wide objectives
- Natural distributed user dynamics converge to an equilibrium in very restrictive classes of games; potential games is an example

This Work

- These considerations motivate two important questions:
  - Can we extend the class of games with desirable dynamic properties beyond potential games?
  - Can we develop simple pricing schemes that will steer the limit point of these dynamics to a desirable operating point on the performance region?
- In this project, we introduce the potential game approach:
  - Approximate the original game with a potential game that has an (additively) separable structure in the individual resources
  - Enables design of a simple pricing scheme that induces the equilibrium of the potentialized game to align with the optimum of any system objective
  - We use the proximity of the two games to establish through Lyapunov-based analysis that natural user dynamics (applied to the original game) converge “within a neighborhood" of the system-wide optimum

Our Contributions

- We apply the potential game approach to study power control in a CDMA wireless system.
- We provide a general distributed power control scheme that would (approximately) achieve any system objective despite the selfishness of the mobiles.
    - Our approach can be used for network regulation under any SINR regime with explicit performance guarantees.
- More generally, we introduce a general framework that shows that any game has a canonical decomposition that has 3 components: potential, harmonic, and nonstrategic components
    - Enables a new approach for studying dynamics in arbitrary games by considering their potential components
    - More details in the poster!

The Network Model

- A set of mobiles (users) $\mathcal{M} = \{1, \dots, M\}$ share the same wireless spectrum (single channel).
- We denote by $\mathbf{p} = (p_1, \dots, p_M)$ the power allocation (vector) of the mobiles.
- Power constraints: $\mathcal{P}_m = \{p_m \mid \underline{P}_m \le p_m \le \bar{P}_m\}$, with $\underline{P}_m > 0$.
  - Upper bound represents a constraint on the maximum power usage
  - Lower bound represents a minimum QoS constraint for the mobile
- The rate (throughput) of user m is given by

$$r_m(\mathbf{p}) = \log\left(1 + \gamma \mathrm{SINR}_m(\mathbf{p})\right),$$

where, $\gamma > 0$ is the spreading gain of the CDMA system and

$$\mathrm{SINR}_m(\mathbf{p}) = \frac{h_{mm} p_m}{N_0 + \sum_{k \neq m} h_{km} p_k}.$$

Here, h_{km} is the channel gain between user k's transmitter and user m's receiver.

User Utilities and Equilibrium

- Each user is interested in maximizing a net rate-utility, which captures a tradeoff between the obtained rate and power cost:

$$u_m(\mathbf{p}) = r_m(\mathbf{p}) - \lambda_m p_m,$$

  where λ_m is a user-specific price per unit power.
- We refer to the induced game among the users as the power game and denote it by $\mathcal{G}$.
- Existence of a pure Nash equilibrium follows because the underlying game is a *concave game*.
- We are also interested in “approximate equilibria" of the power game, for which we use the concept of ϵ-(Nash) equilibria.
  - For a given ϵ, we denote by $\mathcal{I}_\epsilon$ the set of ϵ-equilibria of the power game $\mathcal{G}$, i.e.,

$$\mathcal{I}_\epsilon = \{\mathbf{p} \mid u_m(p_m, \mathbf{p}_{-m}) \geq u_m(q_m, \mathbf{p}_{-m}) - \epsilon, \quad \text{for all } m \in \mathcal{M}, \ q_m \in \mathcal{P}_m\}$$

System Utility

- Assume that a central planner wishes to impose a general performance objective over the network formulated as

$$\max_{\mathbf{p} \in \mathcal{P}} U_0(\mathbf{p}),$$

  where $\mathcal{P} = \mathcal{P}_1 \times \cdots \times \mathcal{P}_m$ is the joint feasible power set.
- We refer to $U_0(\cdot)$ as the system utility-function.
- We denote the optimal solution of this system optimization problem by p^* and refer to it as the desired operating point.
- Our goal is to set the prices such that the equilibrium of the power game can approximate the desired operating point p^*.

Potential Game Approximation

- We approximate the power game with a potential game.
- A game $\mathcal{G} = \langle \mathcal{M}, \{u_m\}, \{\mathcal{P}_m\} \rangle$ is a potential game if there exists a function$\Phi : \mathcal{P} \to \mathbb{R}$ such that

$$\Phi(p^m, \mathbf{p}^{-m}) - \Phi(q^m, \mathbf{p}^{-m}) = u^m(p^m, \mathbf{p}^{-m}) - u^m(q^m, \mathbf{p}^{-m}),$$

  for all $m \in \mathcal{M}$, $p_m, q_m \in \mathcal{P}_m$, and $\mathbf{p}_{-m} \in \mathcal{P}_{-m}$.
- The potential function Φ aggregates the preferences of all players.
  - Every finite potential game has a pure equilibrium.
  - Many learning dynamics (such as better-reply dynamics, fictitious play, spatial adaptive play) “converge” to a pure Nash equilibrium [Monderer and Shapley 96], [Young 98], [Marden, Arslan, Shamma 06, 07].

Potentialized Game

- We consider a slightly modified game with player utility functions given by
$$\tilde{u}_m(\mathbf{p}) = \tilde{r}_m(\mathbf{p}) - \lambda_m p_m$$
where $\tilde{r}_m(\mathbf{p}) = \log\left(\gamma \mathrm{SINR}_m(\mathbf{p})\right)$.
- We refer to this game as the potentialized game and denote it by $\tilde{\mathcal{G}} = \langle \mathcal{M}, \{\tilde{u}_m\}, \{\mathcal{P}_m\} \rangle$.
- For high-SINR regime (γ satisfies $\gamma \gg 1$ or $h_{mm} \gg h_{km}$ for all $k \neq m$),the modified rate formula $\tilde{r}_m(\mathbf{p}) \approx r_m(\mathbf{p})$ serves as a good approximation for the true rate, and thus $\tilde{u}_m(\mathbf{p}) \approx u_m(\mathbf{p})$.

Pricing in the Modified Game

Theorem

The modified game $\tilde{\mathcal{G}}$ is a potential game. The corresponding potential function is given by

$$\phi(\mathbf{p}) = \sum_m \log(p_m) - \lambda_m p_m.$$

- $\tilde{\mathcal{G}}$ has a unique equilibrium.
- The potential function suggests a simple linear pricing scheme.

Theorem

Let $\mathbf{p}^*$ be the desired operating point. Assume that the prices λ^* are given by

$$\lambda_m^* = \frac{1}{p_m^*}, \qquad \text{for all } m \in \mathcal{M}.$$

Then the unique equilibrium of the potentialized game coincides with p^*.

Near-Optimal Dynamics

- We will study the dynamic properties of the power game $\mathcal{G}$ when the prices are set equal to λ^*.
- A natural class of dynamics is the best-response dynamics, in which each user updates his strategy to maximize its utility, given the strategies of other users.
- Let $\beta_m : \mathcal{P}_{-m} \to \mathcal{P}_m$ denote the best-response mapping of user m, i.e.,

$$\beta_m(\mathbf{p}_{-m}) = \arg\max_{p_m \in \mathcal{P}_m} u_m(p_m, \mathbf{p}_{-m}).$$

- Discrete time BR dynamics:

$$p_m \leftarrow p_m + \alpha\left(\beta_m(\mathbf{p}_{-m}) - p_m\right) \quad \text{for all } m \in \mathcal{M},$$

- Continuous time BR dynamics:

$$\dot{p}_m = \beta_m(\mathbf{p}_{-m}) - p_m \quad \text{for all } m \in \mathcal{M}.$$

- The continuous-time BR dynamics is similar to continuous time fictitious play dynamics and gradient-play dynamics [Flam, 2002], [Shamma and Arslan, 2005], [Fudenberg and Levine, 1998].

Convergence Analysis – 1

- If users use BR dynamics in the potentialized game $\tilde{\mathcal{G}}$, their strategies converge to the desired operating point p^*.
  - This can be shown through a Lyapunov analysis using the potential function of $\tilde{\mathcal{G}}$, [Hofbauer and Sandholm, 2000]
  - Our interest is in studying the convergence properties of BR dynamics when used in the power game $\mathcal{G}$.
- Idea: Use perturbation analysis from system theory
  - The difference between the utilities of the original and the potentialized game can be viewed as a perturbation.
  - Lyapunov function of the potentialized game can be used to characterize the set to which the BR dynamics for the original power game converges.

Convergence Analysis – 2

- Our first result shows BR dynamics applied to game $\mathcal{G}$ converges to the set of ϵ-equilibria of the potentialized game $\tilde{\mathcal{G}}$, denoted by $\tilde{\mathcal{I}}_\epsilon$.
- We define the minimum SINR:

$$\underline{SINR}_m = \frac{\underline{P}_m h_{mm}}{N_0 + \sum_{k \neq m} h_{km} \overline{P}_k}$$

- We say that the dynamics *converges uniformly* to a set S if there exists some $T \in (0, \infty)$ such that $\mathbf{p}^t \in S$ for every $t \geq T$ and any initial operating point $\mathbf{p}^0 \in \mathcal{P}$.

Lemma

The BR dynamics applied to the original power game $\mathcal{G}$ converges uniformly to the set $\tilde{\mathcal{I}}_\epsilon$, where ϵ satisfies

$$\epsilon \leq \frac{1}{\gamma} \sum_{m \in \mathcal{M}} \frac{1}{\underline{SINR}_m}.$$

- The error bound provides the explicit dependence on γ and $\underline{SINR}_m$.

Convergence Analysis – 3

- We next establish how “far" the power allocations in $\tilde{\mathcal{I}}_\epsilon$ can be from the desired operating point $\mathbf{p}^*$.

Theorem

For all ϵ, $\mathbf{p} \in \tilde{\mathcal{I}}_\epsilon$ satisfies

$$|\tilde{p}_m - p_m^*| \leq \overline{P}_m\sqrt{2\epsilon} \quad \textit{for every } \tilde{p} \in \tilde{\mathcal{I}}_\epsilon \textit{ and every } m \in \mathcal{M}$$

- Idea: Using the strict concavity and the additively separable structure of the potential function, we characterize $\tilde{\mathcal{I}}_\epsilon$.

Convergence and the System Utility

- Under some smoothness assumptions, the error bound enables us to characterize the performance loss in terms of system utility.

Theorem

Let $\epsilon > 0$ be given. (i) Assume that U_0 is a Lipschitz continuous function, with a Lipschitz constant given by L. Then

$$|U_0(\mathbf{p}^*) - U_0(\mathbf{p})| \leq \sqrt{2\epsilon} L \sqrt{\sum_{m \in \mathcal{M}} \overline{P}_m^2}, \quad \text{for every } \mathbf{p} \in \tilde{\mathcal{I}}_\epsilon.$$

(ii) Assume that U_0 is a continuously differentiable function so that $|\frac{\partial U_0}{\partial p_m}| \leq L_m$, $m \in \mathcal{M}$. Then

$$|U_0(\mathbf{p}^*) - U_0(\mathbf{p})| \leq \sqrt{2\epsilon} \sum_{m \in \mathcal{M}} \overline{P}_m L_m, \quad \text{for every } \mathbf{p} \in \tilde{\mathcal{I}}_\epsilon.$$

Numerical Example – 1

- Consider a system with 3 users and let the desired operating point be given by $\mathbf{p}^* = [5, 5, 5]$.
- We choose the prices as $\lambda_m^* = \frac{1}{p_M^*}$ and pick the channel gain coefficients uniformly at random.
- We consider three different values of $\gamma \in \{5, 10, 50\}$.

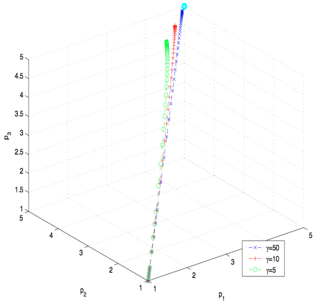

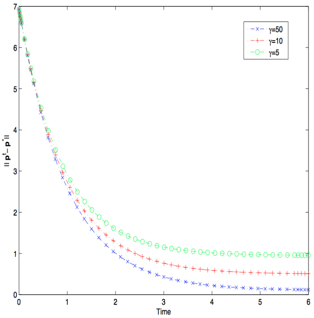

Sum-rate Objective

- We next consider the natural system objective of maximizing the sum-rate in the network.

$$U_0(\mathbf{p}) = \sum_m r_m(\mathbf{p}).$$

- The performance loss in our pricing scheme can be quantified as follows.

Theorem

Let $\mathbf{p}^*$ be the operating point that maximizes sum-rate objective, and let $\tilde{\mathcal{I}}_\epsilon$ be the set of ϵ-equilibria of the modified game to which the BR dynamics converges. Then

$$|U_0(\mathbf{p}^*) - U_0(\mathbf{p})| \leq \sqrt{2\epsilon}(M-1)\sum_{m\in\mathcal{M}} \frac{\overline{P}_m}{\underline{P}_m}, \qquad \textit{for every } \mathbf{p} \in \tilde{\mathcal{I}}_\epsilon.$$

Numerical Example – 2

- Consider $M = 10$ users and assume that the power bounds are given by $\underline{P}_m = 10^{-2}$, $\overline{P}_m = 10$ for all $m \in \mathcal{M}$.

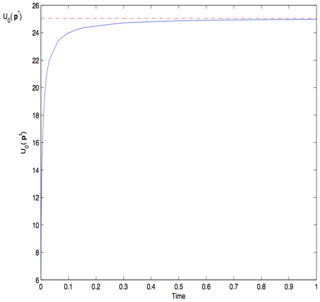

(c) The change in sum-rate as a function of time for $\gamma = 10$.

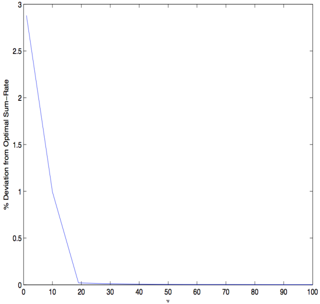

(d) The effect of γ on performance loss.

Summary and Future Work

- We have introduced the potential-game approach for distributed power allocation, which (approximately) enforces any power-dependent system-objective.
- By exploiting the relation between the power game and its approximation (with a potential game), the prices in the potential game induce near-optimal performance in the underlying system.
- We provide bounds on deviation from the maximum system utility in a dynamical sense.
- Future Work:
  - Distributed implementation of optimal desired operating point.
  - Potential game approach for other wireless resource allocation problems.