Mechanism Design for Efficient Decentralized Network Control: The Case of Power Allocation in Wireless Networks

Virgilio Rodriguez
ECE Department
Polytechnic University
Brooklyn, NY 11201
E-mail: vr@ieee.org

August 17, 2003

Abstract

A “mechanism” is a set of rules governing the interaction of selfish entities, which attempts to lead these entities to a desirable outcome. This work applies a relatively simple mechanism, available in the economics literature, to achieve an efficient decentralized allocation of power among data-transmitting terminals. The resulting operating point is “efficient”, because terminals end up “fairly” compensating each other for the interference each one causes. The same ideas can be fruitfully applied in more general networks, and also outside the engineering context.

1 Introduction

It has long been recognized that decentralized control algorithms offer many advantages over their centralized counterparts. Some of the reasons often cited include complexity, signaling overhead, and unavailability of local information to a central controller. And certain modern communication and/or computing paradigms, such as ah-hoc wireless networks, and peer-to-peer applications over communication networks, make central controllers highly impractical, if not outright impossible to implement.

Along these lines, several recent scholarly publications, recognize that algorithms useful for engineering applications can be obtained via the formulation of radio resource management issues, in particular power control in wireless data applications, on the foundations of microeconomic theory (See [1], for an introductory discussion). This approach is centered around the *decentralized* maximization, under appropriate rules and constraints, of a quality-of-service (QoS) index, referred to as a “utility function”. This maximization *may* or may *not* involve a human user choosing resources during transmission. The choices may be made by “software agents” inside transmitting terminals. These agents may be entirely programmed by the network administrator, so that they behave in the way the network owner wants. Or these agents may be controlled and/or tuned or trained by the actual human operator.

In any case, decentralized QoS maximization can be modeled as a “game”: a situation in which each of several selfish agents choose a “strategy” in order to maximize its own “payoff”. Generally, the payoff to a given player depends on the chosen strategies by all players. For instance, in a wireless network, the transmission power chosen by a terminal becomes interference for others. And this interference affects the payoff/utility (QoS) of all terminals.

A key solution concept is a Nash equilibrium; i.e., an allocation (a strategy per player) such that no player would be better off by *unilaterally* “deviating” (changing strategy). For instance, in a wireless network in which each of several mutually-interfering terminals choose transmission power to maximize a bits/Joule index (“utility”), a Nash equilibrium would be a power level per terminal, such that no terminal could increase its utility by unilaterally adjusting its power. In this context, it is well understood that, if transmission power is limited, a Nash equilibrium does exist. And even if power is unlimited, a Nash equilibrium may exist under certain circumstances. This is amply discussed in [3].

However, it is well understood that Nash equilibria are generally "inefficient". For instance, when data-transmitting wireless terminals choose power levels to maximize a sensible QoS index in bits/Joule, they settle on equilibrium power levels that are “too high”. That is, they would all be better off if they all agreed to simultaneously lower their respective transmission power. The challenge is to get selfish terminals to move toward a more efficient operating point “on their own”. This is usually attempted through some form of pricing.

An approach employed in [5] to induce the terminals toward a lower-power equilibrium is to introduce a “tax” on transmission power. That is, terminals are programmed to maximize an expression of the form $u(p;I) - cp$, where $u(p;I)$ denotes the utility of the terminal when its transmission power is p, and its interfering power (caused by noise and the other terminals) equals I; and c is a “tax” on power. This leads to lower power levels at equilibrium, and an increase in the utility of each terminal. Unfortunately, this approach has a significant drawback: while the original utility function is a nice “single-peaked” (quasi-concave) function of the transmission power, the modified function $u(p;I) - cp$ is *not*. With a *non*-quasi-concave utility function, the wireless data power control game may have *NO* equilibrium (even with limited power). To work around this issue, [5] goes into a more complex analysis involving “super-modularity”. This requires certain unnatural impositions and assumptions which are best avoided. In any case, this approach leads to a “better” but still inefficient operating point.

Another approach to guide competing selfish entities toward a “socially optimal” outcome is to design an appropriate “mechanism”; i.e., a set of procedures, penalties and rewards designed to guide the entities toward a desired outcome. In order to achieve an efficient decentralized allocation of power among mutually interfering terminals, this work applies the relatively simple mechanism introduced in [7]. In order for this mechanism to work, there must exist one “transferable good” with which terminals can compensate each other. This good could be money, or some form of service credits, such as time of usage.

The intuition of the mechanism being proposed can best be captured by considering a 2-terminal situation in which only terminal 1 interferes with terminal 2 (but *not* vice-versa). (This can actually happen with successive interference cancellation decoding.) Terminal 2 must declare the amount of the transferable good it wishes to *charge* terminal 1 as compensation for each unit of interference. Likewise, terminal 1 must quote the price it *offers* to pay terminal 2 as compensation. But terminal 1 faces a penalty proportional to the absolute value of any difference between its offered price and the price demanded by terminal 2 as compensation. The fair compensation terminal 2 should receive is the true dis-utility or “cost” induced on terminal 2 by the interference. Under the proposed mechanism, at equilibrium, terminal 1 ends up paying precisely the fair amount to terminal 2.

To see why the equilibrium compensation is fair, first notice that, at equilibrium, neither terminal must want to change its price (by definition). This implies that, at equilibrium, the interfering terminal will NOT pay more than the true cost caused on the other terminal by its interference. If the amount paid by terminal 1 exceeds the cost its interference causes on terminal 2, then terminal 2 is in fact “making a profit” on each unit of interference produced by terminal 1. But if each unit of interference leaves a net “profit” on terminal 2, it is optimal for this terminal to induce terminal 1 to *increase* its interference even further. But in order to accomplish this, terminal 2 must *decrease* its charge. Thus, as long as terminal 2 asking price is higher than the true cost it suffers as result of the interference, it is optimal for this terminal to lower its own price. Therefore, at equilibrium, the price charged by terminal 2 is exactly equal to the true “cost” that the interference causes.

When 2 terminals interfere each other, each terminal must quote two prices. One price is the amount of the transferable good it *offers to pay* the other terminal as compensation for each unit of interference it creates. The second price is the amount this terminal wishes to *be paid* as compensation for each unit of interference caused by the other terminal. But each terminal faces a penalty if its prices differ from those quoted by the other. At equilibrium, both terminals quote the same prices, and no penalty is necessary. Terminals end up fairly compensating each other for the interference each one creates. It can be shown that the resulting allocation is “socially optimal” in a reasonable sense.

This framework can be extended to accommodate many mutually-interfering terminals. Additionally, the basic mechanism can be applied in many situations unrelated to power, or even communications.

Below, the system model is built, and the pricing mechanism is specified more formally. Then, the maximization problem faced by each of two terminals is technically studied, reaction (best-response) functions are derived, and equilibrium allocation found. Subsequently, the efficiency of the equilibrium allocation is shown. Afterward, the extension of this analysis to a many terminal situation is discussed. Finally, some concluding remarks are made.

References

[1] Goodman, D.J., N. Mandayam, “Power Control for Wireless Data”, IEEE Pers. Comm., Vol. 7, No. 2, pp. 48 -54 , Apr 2000

[2] Ji, H., "Resource Management in Communication Networks via Economic Models", Ph.D. Thesis, Rutgers Univ., New Jersey, 1997

[3] Rodriguez, V., “Robust Modeling and Analysis for Wireless Data Resource Management”, IEEE WCNC, 2003

[4] Rodriguez, V., “Maximizing the Ratio of a Generalized Sigmoid to its Argument”, WICAT Tech. Rep. 02-010, Polytechnic University, 2002 `http://wicat.poly.edu/reports`

[5] Saraydar, C.U., N.B. Mandayam and D.J. Goodman, “Efficient power control via pricing in wireless data networks”, IEEE Trans. on Comm., Vol. 50, No. 2 , pp. 291 -303, Feb. 2002

[6] Varian, H., , “Microeconomic Analysis”(3rd ed), New York : W.W. Norton & Co., 1992

[7] Varian, H., “A solution to the problem of externalities when agents are well-informed”, The American Economic Review, Vol. 84, No. 5, pp. 1278-93, Dec. 1994