

SELFISH OPTIMIZATION IN COMPUTER NETWORKS

Yechiam Yemini
Department of Computer Science
Columbia University

New York, New York 10027

August 1981

SELFISH OPTIMIZATION IN COMPUTER NETWORKS**

YECHIAM YEMINI,
COMPUTER SCIENCE DEPARTMENT
COLUMBIA UNIVERSITY
NY, NY 10027

ABSTRACT

This paper describes two applications of decentralized (Pareto) optimization to problems of computer communication networks. The first application is to develop a generalized principle for optimality of multi-hop broadcast channel access schemes. The second application is to decentralized flow-control in fixed virtual-circuit networks (e.g., SNA) using power maximization as the performance index. The decentralized approach to optimum network behavior yields, among other results, characterization of fair global objective functions, and optimal decentralized greedy network control algorithms. The main conclusion of this paper is that Pareto-optimality methods can be successfully used to develop optimal decentralized behavior algorithms where a centralized approach is (sometimes provably) not applicable.

*This research has been supported in part by an NSF grant No. MCS 8110319 and by the Defense Advanced Research Project Agency of the Department of Defense.

*This paper will be presented at the 20-th IEEE conference on Decision and Control, December 1981

Table of Contents

1. INTRODUCTION	1
2. SELFISH PACKET BROADCASTING	2
3. SELFISH POWER CONTROL	5
3.1. Fair Global Performance Measures	8
3.2. Greedy Algorithms For Selfish Optimization	9
3.3. Open Problems	11
4. CONCLUSIONS	11

1. INTRODUCTION

How should one derive optimal behavior algorithms for computer communication networks? The classic approach to the problem views the network as a single entity to which a global performance objective is assigned. This leads to a centralized optimization problem. The major shortcoming of this approach is that when one has overcome the complexity of deriving an optimal solution the network problem is still not solved, since the centralized objective usually leads to a centralized behavior policy. This centralized optimal behavior needs to be decentralized to serve as an adequate solution. The process of decentralization is usually more difficult than that of solving the original optimization problem. Therefore decentralization is usually an ad-hoc approximation process with little formal methodological support.

An alternative approach is to view the network as a loose collection of interfering agents (i.e., nodes, processes), each of which is assigned a selfish utility function which it seeks to optimize. The problem then becomes that of finding an adequate compromise among the selfish needs of the agents. One usually adopts Pareto optimality as the norm for rational behavior. That is, the agents should select a policy which is not dominated by any other policy*. The major advantage of this selfish approach to optimal network behavior, compared with the global approach, is that it generates policies that are immediately decentralized. The major disadvantage is that global optimization is much better understood, formally speaking, than selfish optimization.

The objective of this paper is to demonstrate the value of selfish optimization in providing useful solutions to decentralized network control problems. We present two successful applications of the selfish approach to two computer network problems: broadcast channel sharing and flow control. The presentation style is intentionally semi-formal, and the models selected are as simple as possible to avoid unnecessary mathematical complexity that would hide the forest behind the trees. (We do, however, point to some "trees" of further research interest).

*In the sense that no subset of agents can improve their performance without a performance degradation of some other agents.

2. SELFISH PACKET BROADCASTING

This section briefly summarizes results published in [YEMI 79] and is included for the sake of completeness. We give only a rudimentary description of the results: the interested reader may find precise derivations in [YEMI 79].

Consider the problem of channel sharing, (i.e., designing a multiaccess scheme), in a network of packet switched broadcast units. The single-hop access scheme problem, when all the broadcast units can hear each other, has been thoroughly explored (see [TOBA 80] for a recent survey). However, the multi-hop problem remains a terra-incognita. Therefore, let us consider the general case of a network where the broadcast units are not necessarily within hearing distance of each other. Let us also assume that packets are routed using a fixed routing scheme*. Finally, the communication channel will be assumed to be time-slotted to packet-size slots. A transmission may only take place within a certain slot, and will be successfully received if it does not collide with another transmission at its destination.

An access scheme is an algorithm to decide which busy units* should be selected to transmit in any given time slot, that is, an algorithm to schedule channel access rights.

Let us establish a mathematical model of the problem. We use numbers $\{1, 2, \dots, N\}$ to denote the broadcast units. Consider an access algorithm: at any given slot unit i may be assigned a transmission right with probability p_i ; it will use this right and transmit if it is busy. Therefore the behavior of an access algorithm may be described by its choice of transmission policy vectors $\underline{p}^e (p_1, p_2, \dots, p_N)$.

*This assumption may be easily relaxed to allow general models of routing; it is only used to simplify the discussion.

*A unit is said to be busy if it has packets ready for transmission

Let $S_i(\underline{p})$ denote the thruput (i.e., probability of successful transmission) obtained by unit i when the transmission policy is \underline{p} . The performance of an access algorithm is completely described by specifying the transformation $\underline{S}(\underline{p}) \triangleq (s_1, s_2, \dots, s_N)$, of transmission policies to attainable thruputs.

The global approach might seek to optimize some global function of the thruput vector $\underline{S}(\underline{p})$ (e.g., the average thruput). Clearly an optimal solution to the global problem is to select a maximal set of non interfering units and let them transmit with probability 1 while the others are kept quiet. Unfortunately, this policy cannot be effectively decentralized.

The selfish approach considers the thruput $S_i(\underline{p})$ as the utility of unit i ; the different units seek to jointly maximize their individual thruputs. A thruput vector \underline{S} is said to dominate the vector \underline{S}' if $S_i > S'_i$ for all i , with at least one strict inequality. A thruput vector is Pareto optimal if it is not dominated by any other attainable thruput vector. A transmission policy which attains a Pareto optimal thruput is said to be a Pareto optimal policy. An access scheme would clearly prefer Pareto-optimal transmission policies. The selfish optimization problem is to characterize Pareto optimal transmission policies.

Let \underline{p}_0 denote a Pareto optimal policy whose thruput is \underline{S}_0 . A small perturbation in the policy $\Delta \underline{p}$ results in a perturbation $\Delta \underline{S}$ in the thruput; the two perturbations being related through:

$$\Delta \underline{S} = \partial \underline{S} \Delta \underline{p}$$

Where $\partial \underline{S}$ is the Jacobian matrix of the transformation $\underline{S}(\underline{p})$. It is easily demonstrable that a necessary condition for Pareto optimality of a policy \underline{p} is that the Jacobian matrix $\partial \underline{S}(\underline{p})$ be singular at \underline{p} .

Define E_i to be the expected number of slots that are empty at the destination of unit i given that unit i is busy, and $S_{i/j}$ to be the thruput of unit i , given that unit j is busy and interferes with unit i . It can be shown [YEMI 79] that the necessary condition for Pareto optimality is that there exist multipliers $\underline{\alpha} \triangleq (\alpha_1, \alpha_2, \dots, \alpha_N)$ such that:

$$\alpha_i E_i = \sum_{j \in I(i)} \alpha_j S_{j/i} \quad (1)$$

where $I(i)$ is the set of units with which unit i interferes.

These optimality conditions may be interpreted as follows. Each unit is given a multiplier α which indicates its relative "dollar" value. The left hand side of the equation represents the dollar value of channel loss by unit i due to silence at its destination; the right hand side represents the dollar value of thruput that unit i might interfere with; we thus use silence and thruput to denote the two sides of the equation (1) respectively. The optimality rule may now be restated: if a transmission policy is Pareto optimal than it must equate "silence" and "thruput" of each unit. This optimality principle is intuitively plausible; a broadcast unit should only waste its silence "dollars" if it can expect other units to use this silence to gain an equal amount of thruput "dollars".

Let us briefly apply these optimality conditions to the classic multiaccess problem of a single hop network. Suppose the broadcast units use the Slotted-ALOHA transmission policy, i.e., a busy unit tosses a coin with probability of transmission p and decides whether or not it should transmit accordingly. The problem is to find an optimal transmission policy p . The classic solution is to maximize the overall thruput $S=np(1-p)^n$, where n is the number of busy units; this is maximized when $p=1/n$. Let us apply the Pareto optimality conditions to this model; equating silence $= (1-p)^n$ and thruput $= (n-1)p(1-p)^{n-1}$ yields that the optimal choice of p is $p=1/n$, in agreement with the global approach.

Let us apply the Pareto optimality condition to the Urn scheme [YEMI 78]; again, we consider a single hop network with n busy units. The Urn scheme selects at each slot k random units out of the N units and gives them access rights. The problem is to find a k which optimizes the performance. The global approach yields [YEMI 78] $k=(N-n+1)/n$ as the value which optimizes the overall thruput. Now, one can easily verify that:

$$\underline{\text{silence}} = \frac{\binom{N-k-1}{n-1}}{\binom{N-1}{n-1}}$$

$$\underline{\text{thruput}} = \frac{\binom{k}{1} \binom{N-k-1}{n-2}}{\binom{N-1}{n-1}}$$

If silence and thruput are equated the conclusion is that the optimum policy is to select $k=(N-n+1)/n$, again in agreement with the global optimization result.

There is more, however, to these results than the mere reassurance that global optimization results may be rederived using selfish optimization. Both Slotted-ALOHA and the Urn scheme require knowledge of the channel load n for their control. This knowledge is not readily available. The characterization of Pareto optimality in terms of silence and thruput suggests new access control algorithms for both schemes. Namely, rather than estimating n , which is not directly observable, the control algorithm should estimate silence and thruput and adjust the control parameters (p or k) to equate the two quantities at each unit. While these algorithms require further study (e.g., how do we guarantee convergence), they are inherently decentralized and do not depend on information which is not observable. In addition, the characterization of Pareto optimal policies is parameterized by the cost multipliers α . By assigning different values to the broadcast units, one obtains a priority mechanism for access schemes. Finally, the selfish approach was successfully applied to a very general multi-hop network, a problem that has so far resisted the global approach.

3. SELFISH POWER CONTROL

A major objective of flow-control mechanisms in computer communication networks is to regulate the use of shared communication resources to achieve an adequate delay-thruput response. As with other queueing systems, one has two conflicting objectives: to maximize thruput and to minimize delay. It was suggested by Giessler et al. [GIES 78] that the dilemma may be resolved by using a single performance measure power, defined as the ratio of thruput to delay. In [GIES 78] and then [KLIE 79] the properties of power and its

generalizations were explored. It was shown that maximization of power offers a desirable network objective.

Bharath-Kumar and Jaffe [KUJA 81] considered power-maximization as a mechanism to control flow over virtual circuits (VCs). They show that certain notions of global network power, when optimized, lead to unfair flow control mechanisms.

In another paper [JAFF 81] it is shown that different notions of power cannot be optimized by decentralized algorithms based on local observations of thruputs and delay. This negative result epitomizes the shortcomings of using the global optimization approach and supports our application of the selfish optimization approach to derive decentralized optimal power control algorithms.

To fix the ideas, consider a packet-switched computer communication network which provides virtual circuit communication between nodes. We assume that a VC, once established, uses a fixed path through the network. This is the view on which the SNA and TYMNET architectures are based, for instance.

The VCs share the communication resources over which their traffic is multiplexed. The problem is to derive a flow control mechanism for the VCs to adjust their mutual thruputs in order to maximize their power. In this paper we consider a simple example to illustrate the issues and derive optimal selfish policies. (The generalization of our results to a network is discussed in a forthcoming paper).

Consider the case of two VCs sharing a single link. This is illustrated in figure 3-1 below.

Let x and y denote the thruput rates of VC-1 and VC-2 respectively and let μ be the capacity of the physical link shared by the two virtual circuits. We assume that arrivals to the virtual circuits are Poisson and that transmission time is exponentially distributed with rate 1. The expected delay over the VCs is given by $D_i(x,y)=1/(\mu-x-y)$.




Figure 3-1: Two interfering virtual circuits

The utility of each VC is its power, given by

$$\begin{cases} P^1(x,y) \triangleq x/D_1(x,y) = x(\mu-x-y) \\ P^2(x,y) \triangleq y/D_2(x,y) = y(\mu-x-y) \end{cases}$$

respectively. The objective of a selfish power control policy is to have each node select a Pareto optimal thruput.

A necessary condition for a thruput pair (x,y) to be Pareto optimal is that the Jacobian matrix of the transformation $\underline{P}(x,y) \triangleq (P^1, P^2)$ is singular. This is equivalent to the existence of a multiplier α such that:

$$\begin{cases} 0 = \alpha P_x^1 + P_y^1 = \alpha(\mu-2x-y) - x \\ 0 = \alpha P_x^2 + P_y^2 = -\alpha y + (\mu-x-2y) \end{cases} \quad (2)$$

Where P_x^i and P_y^i denote the partial derivatives of $P^i(x,y)$.

It is easy to see that if $\alpha = -1$ then these optimality conditions are equivalent to $x+y=\mu$ while for $\alpha \neq -1$ these conditions are equivalent to $x+y=\mu/2$. Clearly the first case i.e., $\alpha = -1$ is not an adequate policy (the physical link will be saturated with the flows causing the power of both VCs to be 0). Therefore, we conclude that the set of Pareto optimal thruput pairs is the line $x+y=\mu/2$.

3.1. Fair Global Performance Measures

In [KUJA 81] a number of global performance measures, based on the functions of the individual powers, are examined. The main concern of that paper is with fairness of global objectives. A global performance objective is fair if it does not lead to a policy which provides zero power to any of the VCs.

Given a global objective measure, one may consider its level curves in the (x,y) thrupt plane, i.e., the curves on which the performance measure is constant. A sound global objective should not select an optimal thrupt pair which is not Pareto optimal. Therefore, optimal thrupt pairs lie at the meeting points of the line $x+y=\mu/2$ and the highest attainable level curve of the global objective.

One can see immediately that any linear combination of individual powers is not a fair global objective,* since its level curves in the (P^1, P^2) plane are straight lines and would select the policy yielding the power pair $(0, (\mu/2)^2)$ (i.e., corresponding to the thrupt pair $(0, \mu/2)$) if VC-2 is given more weight and the policy $(\mu/2, 0)$ if VC-1 is given more weight. This lack of fairness may be easily generalized to any global objective function, the level curves of which are concave with respect to the Pareto optimal line when observed in the direction of the origin.

On the other hand, consider the global objective function $g(x,y) \hat{=} P^1(x,y)P^2(x,y)$. The level curves of this function in the (P^1, P^2) plane are strictly convex with respect to the line of Pareto-optimal powers. The optimum global policy thus intersects the Pareto-optimal line in its interior. Moreover, $g(x,y)$ is symmetrical so that the optimum thrupt pair allocates identical thrupts to both virtual circuits. These properties of the product measure clearly render it fair [KUJA 81].* To summarize, the geometry of Pareto-optimal solutions provides an easy explanation of fairness.

*That is, except for the trivial case when all powers are taken with equal weights and thus any Pareto optimal policy is globally optimal.

*If we choose a product measure with the P^i 's raised to different powers, other fair optimum behaviors are obtained.

3.2. Greedy Algorithms For Selfish Optimization

How can individual VCs adjust their thruputs on the basis of local observations, to reach a Pareto optimal policy? We shall consider greedy algorithms i.e., algorithms where a VC adjusts its thrupt rate according to the gradient of its utility (power). Formally, let ∂P^i be the gradient of the power of VC-i, a greedy algorithm is one where

$$\begin{cases} \frac{dx}{dt} = \partial P^1 \underline{a}^1 \\ \frac{dy}{dt} = \partial P^2 \underline{a}^2 \end{cases} \quad (3)$$

Here the vectors \underline{a}^i represent directions along which the gradient is projected. The idea behind greedy algorithms is that a VC should change its thrupt (the left hand sides of the above equations) proportionally to the increase in utility incurred to it. The increase in utility is represented by a projection of the respective gradient, given by the right hand side of the equations 3. Note that we use a continuous-time approximation of a discrete-time process. This helps simplify the computations while not influencing the results since our model is essentially stationary and the model of time has no intrinsic significance.

Before proceeding to analyze the dynamics of greedy algorithms let us consider the steady state limit of the process described by equations 3. During a steady state the thrupt pair (x,y) satisfies:

$$\begin{cases} 0 = \partial P^1 \underline{a}^1 \\ 0 = \partial P^2 \underline{a}^2 \end{cases} \quad (4)$$

Let us note in passing that the greedy algorithm presented in [KUJA 81] is obtained by selecting $\underline{a}^1 \triangleq (1,0)$ and $\underline{a}^2 \triangleq (0,1)$. This selection represents a process where the VCs take turns adjusting their thrupts; each VC, in its turn, maximizes its power for the given thrupt of the other VC. Therefore, the thrupt adjustment process is such that each VC only considers its own

influence on its power (i.e., the respective component of the gradient) ignoring the changes of power resulting from the choices of the other VC. The solution of the steady state equations for this choice of \underline{a}^i is easily computed to be $(\mu/3, \mu/3)$, which is not Pareto optimal. Clearly, the reason why this greedy algorithm does not lead to an optimal solution is the lack of coordination in the choice of direction vectors \underline{a}^i (the two nodes are pulling the cart in orthogonal directions).

It may be easily demonstrated that the steady state equation (4) is equivalent to the necessary conditions for Pareto optimality (equation (2)) if and only if the direction vectors \underline{a}^i are equal. Therefore, in order to ensure that the greedy algorithm will converge to the Pareto optimal line $x+y=\mu/2$, all that is required is to select the direction vector \underline{a} in the linear space spanned by $(\alpha, 1)$ where $\alpha \neq -1$.

Having seen the significance of the projection vectors \underline{a}^i for proper coordination, let us study their role further. Returning to equation 3, we may interpret the right hand side as the directional derivatives of P^i in the direction specified by \underline{a}^i . Therefore, one possible interpretation of the greedy algorithms is that the two VCs alternate synchronously adjusting their thruputs to maximize their power in the direction indicated by the respective \underline{a}^i . When the two direction vectors are identical, another useful interpretation of the greedy algorithm arises. Namely, suppose each VC iterates adjusting its thruptut proportionally to the observed changes of its power. However, let us assume that the two VCs iterate at different speeds. The coordinates of the common direction vector \underline{a} represent the relative speeds of iteration of the two VCs.

Finally, let us consider the stability of convergence of the greedy algorithm when both VCs use the same direction vector $\underline{a}=(\alpha, 1)$. The equations describing the evolution of the greedy algorithm (equation 3) are linear with a matrix whose eigenvalues are $\alpha+1$ and $2(\alpha+1)$. Therefore when $\alpha < -1$ the algorithm is stable (note again the singular role of the value $\alpha = -1$).

3.3. Open Problems

When one tries to add more realistic details to the simple model of interaction described above, a few major mathematical difficulties arise, requiring the development of adequate tools. The first problem is that of the interaction between the control algorithms and the underlying stochastic processes. In our model it is assumed that convergence to a steady state is much faster than the speed of iteration of the control algorithm. This renders the algorithms quasi-static; that is, the time between any two iterations of the algorithm must be greater than the time constant of the steady state convergence. What if we wish the control algorithm to proceed at faster speeds? The simple mathematics above is no longer applicable. What is the dynamics of the control algorithm when it is no longer a quasi-static process?

Another problem is that of incorporating delayed and partial observations into the model. In the model above the information required for control (e.g., the change of power incurred) is available instantaneously. This makes sense for quasi-static control algorithms, but what if the control algorithm proceeds at a speed comparable to the time it takes to obtain observations?

Finally, modelling the stochastic asynchronous operation of the distributed agents is important. The model above assumed that the control algorithms executed by the different agents are synchronous. The physical significance of the relative speeds in which the agents execute their algorithms was discussed. In reality, however, one can expect asynchronous operation. How should the models account for this?

4. CONCLUSIONS

The two examples discussed here demonstrate the advantages of the selfish optimization approach in providing an explanation of optimal decentralized behavior and offering useful decentralized optimality algorithms. In a recent work [BROO 82] similar ideas have been successfully applied to generate optimal decentralized traffic light control algorithms. Other applications of selfish optimization to network problems are currently being pursued.

- [BROO 82] Brooks, R., Phd Thesis, Computer Science Department, University of Massachusetts at Amherst, to appear.
- [GIES 78] Giessler, J. H., Konig, A., and Pade, F.
Free Buffer Allocation -- An investigation by Simulation.
Computer Networks 2, 1978.
- [JAFF 81] Jaffe, J. M.
Flow Control Power is Nondecentralizable.
IEEE Transactions on Communications (9), September, 1981.
- [KLIE 79] Kleinrock, L.
On Flow Control in Computer Networks.
In Proceedings, International Conference on Communications.
1979.
- [KUJA 81] Bharat-Kumar, K., and Jaffe, J. M.
A New Approach to Performance Oriented Flow Control.
IEEE Transactions on Communications COM-29, April, 1981.
- [TOBA 80] Tobagi, F. A.
Multiaccess Protocols in Packet Communication Systems.
IEEE, Transactions on Communications , April, 1980.
- [YEMI 78] Yemini, Y. and Kleinrock, L.
An Optimal Adaptive Scheme for Multiple-Access Broadcast Communication.
In Proceedings of ICC, pages 7.2.1-7.2.5. June, 1978.
Toronto.
- [YEMI 79] Yemini, Y., and Kleinrock, L.
On a General Rule for Access Control, or Silence is Golden...
In J. Lagrange (editor), Flow Control in Computer Networks, .
North-Holland, 1979.