Local decision procedures for avoiding the Tragedy of Commons

Sabyasachi Saha and Sandip Sen

Math & CS Dept, University of Tulsa
600 South College Avenue
Tulsa, OK 74104-3189
{sahasa, sandip}@ens.utulsa.edu

Abstract. The social sciences literature abound in problems of providing and maintaining a public good in a society composed of self-interested individuals [6]. Public goods are social benefits that can be accessed by individuals irrespective of their personal contributions. Such problems are also addressed in the domain of agent based systems [15]. In this paper we address the problem of the Tragedy of the Commons [9], a particularly common social dilemma that leads to inefficient usage of shared resources. We present a decision procedure following which rational agents can optimally use a shared resource using only local information. Our experimental results confirm that the tragedy of the commons is successfully avoided and the shared resource is utilized to its capacity when agents following our prescribed decision procedure.

1 Introduction

The viability of an individual in a society depends critically on the behavior of other members. Interesting computational problems in agent societies include paradoxes that involve reduction of system throughput when more resources are added to an existing system. Other social dilemmas arise when myopic, local-utility-maximizing decision-making by individual members of the society lead to a loss of utility for everyone. Such problematic scenarios appear frequently in natural and artificial societies.

In a society, the common infrastructures, goods and services are typically shared between members. For example, if we consider the problem of city traffic, we find that congestion problems arises out of self-interested drivers having to share common resources like roads, bridges etc. It often happens that the shared resource has a capacity and if the load is more than its capacity the resource performance or its perceived utility to the users decrease sharply. In a society of self-interested rational agents, or humans, each individual will try to maximize their utility from the shared resources. From the local perspective of a given agent, the more extensive use of a resource produces greater utility. If decision-making is predicated only on this local perspective, each user in the system can myopically try to maximize its load on the common resource. As a result, the combined load is likely to exceed the capacity of the common resource and adversely affect everyone and result in a decrease in everyone's utility from the resource. This is the well-known *Social dilemma* problem of the *Tragedy of the commons*.

The examples of the *Tragedy of the commons* are now seen from the problem of global warming, congestion of traffic to the problem of sharing communication channel bandwidth.

An example of Tragedy of the commons lie in the example of network congestion if every packet is sent with highest possible priority. Suppose there are some routes of different quality. If everybody wants to route through the best possible route then it leads to a congestion which worsen every routing through that route.

More recently, attention has been drawn to the tragedy of the commons in the context of autonomous agent systems [15]. These and other problems arise in multiagent societies as multiple, distributed decision-makers try to maximize local utility by taking decisions based only on limited global knowledge. Correspondingly, multiagent system researchers have developed various approaches to resolve resource conflicts between distributed agents. For example, some researches have addressed the problem of effectively sharing common resources [2]. They proposed an agent as a planner who will make all resource allocation decisions. But this central planning approach requires nearly perfect global knowledge of all agents and the environment which is not very reasonable in complex, distributed and dynamic domains. Durfee and Lesser proposed a distributed partial-global planning [3] approach for coherent coordination between distributed problem solvers through the exchange of partial local plans. Approaches that emphasize economic mechanisms like contracting and auctions, allocate resources based on perceived utility [12]. While the economic approaches are interesting, we believe that they do not provide a satisfactory resolution to social dilemma problems without an adequate discussion of varying individual wealth and interpersonal utility comparisons. The *COIN* approach to solving social dilemmas allows distributed computation but requires an "omniscient" agent to set up the utility functions to be optimized locally [14].

Glance and Hogg [4] make the important observation that computational social dilemmas can produce situations where it is impossible to arrive at globally optimal system configurations based only on distributed, rational decision-making with local knwoledge. They contrast such computational problems with traditional complexity analysis in algorithm theory where solutions are hard, but not impossible to find.

The motivation of our work on computational social dilemma has been to investigate mechanisms to resolve conflicts while requiring minimal global knowledge or imposing minimal behavioral restrictions on the agents. For example, in [1] it is shown that a genetic algorithm based optimization framework can solve a well-known social dilemma problem, the Braess' Paradox [8]. The GA-based function optimization approach is a centralized mechanism. Munde *et. al.* used a more decentralized, adaptive systems approach using GAs, to address both the Braess' paradox and the Tragedy of the Commons [11]. Though decision making is decentralized in this approach, the survival of individuals, as determined by fitness-proportionate selection scheme, is a centralized procedure. Though the latter procedure can be approximated in a decentralized manner, a further criticism of the approach, the somewhat altruistic decision-procedure used by the distributed agents, is difficult to address.

In this paper we concentrate on *Tragedy of the commons* problem with the goal of designing defensible decentralized procedures relying on only minimal local information that can still solve this dilemma. In the following we first review the problem of social dilemmas and discuss the tragedy of the commons in more detail. Then we present a local decision procedure for addressing the tragedy of the commons. We assume that all agents in the system use our suggested decision procedure. We then experimenatlly demonstrate that our suggested local decision procedure produces optimal global utilization of the shared resource.

2 Social dilemmas

A social dilemma arises when agents have to decide between contributing or not contributing towards a public good without the enforcement mechanism of a central authority [5]. Individual agents have to tradeoff local and global interests while choosing their actions. A selfish individual will prefer not to contribute towards the public good, but utilize the benefits once the service is in place. If a sufficient number of agents make the selfish choice, the public good may not survive, and then everybody suffers. In general, social laws, taxes, etc. are enforced to guarantee the preservation of necessary public goods. Consider a scenario where a public good is to be initiated provided enough contribution is received from the populace. Let us assume that the public good, $\mathcal{G}$, costs C, and the benefit received by individual members of the populace is B. Let us also assume that in a society of N agents, $P < N$ individuals decided to contribute to the public good. Assuming that the cost is uniformly shared by the contributors, each contributing agent incurs a personal cost of $\frac{C}{P}$. If enough agents contribute, we can have $\frac{C}{P} < B$, that is even the contributors will benefit from the public good. Since we do not preclude non-contributors from enjoying the public good in this model, the non-contributors will benefit more than the contributors. If we introduce a ceiling, M, on the cost that any individual can bear, then the public good will not be offered if $\frac{C}{P} > M$. In this case, everybody is denied the benefit from the public good.

Similarly in a resource sharing problem, where the cost of utilizing a resource increases with the number of agents sharing it (for example, congestion on traffic lanes). Assume that initially the agents are randomly assigned to one of two identical resources. Now, if every agent opts for the resource with the least current usage, the overall system cost (cost incurred per person) increases [7]. So, the dilemma for each agent is whether or not to make the greedy choice.

2.1 Tragedy of the Commons

In his book, *The Wealth of Nations* (1776), Adam Smith conjectured that an individual for his own gain is prompted by an "invisible hand" to benefit the group [13]. As a rebuttal to this theory, William Forster Lloyd presented the *tragedy of the commons* scenario in 1833 [9]. Lloyd's scenario consisted of a pasture shared by a number of herdsmen for grazing cattles. This pasture has a capacity, say C, i.e., each time a cattle added by a herdsman result in a gain as long as the total number of cattles in the pasture, x, is less than or equal to C. When $x > C$, each addition of a cattle result in a decrease in the quality of grazing for all. Lloyd showed that when the utilization of the pasture gets close to its capacity, overgrazing is guaranteed to doom the pastureland. For each herdsman, the incentive is to add more cattles to his herd as he receives the full proceeds from the sale of additional cattle, but shares the cost of overgrazing with all herdsmen. Whereas the common resource could have been reasonably shared by the herdsman exhibiting some restraint, greedy local choices made by the herdsmen quickly leads to overgrazing and destruction of the pasture. The question the herdsman will face is "What is the utility of adding one more animal to my herd?" [6]. He observes that "Freedom in a commons brings ruin to all." and convincingly argues that enforced laws, and not appeals to conscience, is necessary to avoid the *tragedy of the commons*.

Muhsam [10] has shown that if some or all other herdsmen add cattle when $x > C$, an individual must add a head if he or she wishes to reduce the loss suffered as a result. A rational, utility-maximizing agent will have no choice but to add to the herd, and hence, to the overall deterioration of the resource performance. This means that it is only possible to reach a *co-operative equillibrium*.

In our paper, we now define an abstract version of the Tragedy of the Commons problem, to be used in the rest of the paper, as follows: a shared resource can effectively support C units of load, but if the jointly applied load, x, exceeds C, the quality of the service received from the resource deteriorates. We call C the *critical load* of the resource. The above constraint is modeled by a utility per unit load function as follows:

$$\begin{aligned} U(x) &= K, \quad \text{when } x \leq C, \\ &= K * \frac{C}{x}, \quad \text{otherwise}, \end{aligned} \tag{1}$$

where K is a constant and $U(x)$ denotes the utility per unit load when a total of x units of load is applied on the system. It can be shown here that for any rational, self interested, utility maximizing agents it is always a better option to add more loads when the other agents are adding more load to the system. And also it is clear that when every agent will go on adding load the utilization of the system will be deteriorating as a result of decreased per unit utility. In such a situation, intelligent agents will try to reach a co-operative equillibrium to optimize the resource utilization. In this paper, we have presented such a mechanism.

3 A Local Decision Procedure for the Tragedy of the Commons Problem

In this section, we will provide an probabilistic distributed algorithm to solve the problem of the *Tragedy of the commons* and also discuss the convergence of the algorithm.

3.1 Algorithm

We assume the following. $A[i], i = 1..N$ are the agents in the society. C is the critical load of the shared resource, and x is the current combined load on the shared resource. $upper_i$ and p_i^u are private fields for agent i.

Step 1 Each agent apply a random load, l_i^0, on the shared resource. (We assume that the initial combined load is less than the capacity of the shared resource, i.e., $\sum_i l_i^0 < C$). Then the following steps are independently followed by each agent.

Step 2 Each agent i increments its load, l_i, on the shared resource. Note that, every time an agent increases (or decreases) its load implies increases (or decreases) it's load to the system by one unit.

Step 3 An agent i recieve its utility from the resource based both on the load it applied and the total load on the resource(i.e. x). The resource derives each agent's per unit utility from the equation 1 and send it to each agent. If the per unit utility received by this agent is not less than the best per unit utility it has received in the past, go to Step 2.

Step 4 Agent i decrements its load and sets its increasing probability, $p_i^u = 1$, and $upper_i$ to $false$.

Step 5 If $upper_i$ is $false$, agent i increments its load by one on the shared resource with probability p_i^u. Otherwise, agent i maintains its previously applied load.

Step 6 Agent i receives its updated utility based on current system load x and load applied by this agent. If an agent had increased its load in Step 5 and the new per unit utility is worse than the best per unit utility that agent has ever received, it decrements its load by one and sets p_i^u to half of it's previous value. Otherwise, an agent i who has increased its load in Step 5 sets $upper_i$ to $true$.

Step 7 If $upper_i$ is $false$ and $p_i^u > p_t$ (where p_t is a small thresold probability) go to Step 5.

Step 8 Agent i maintains its current load l_i. Repeat Step 8.

3.2 Convergence to equillibrium

The system reaches equilibrium when all agents reach Step 8 of the algorithm. At this state, each agent feels that any increment/decrement of its load will reduce its utility. Hence the load on the system do not change.

After an agent has passed through Step 3 it realizes that increasing its load may decrease its per unit utility which in turn decreases the utility obtained from the resource. So, every agent at Step 4 removes one load it added last time. Observe that every agent will execute Step 4 in the same iteration. Here we assume that in *Step*1 after all the agents add a random load, the total load administered to the resource did not exceed the resource capacity, C[1]. So, after Step 4 there may be three possibilites from the perspective of each autonomous agent:

- It is the only agent who is using this resource and adding one more load will decrease the per unit utility.
- More than one agent is using this resource, and the resource has reached it's critical load *i.e.* if any of the agents adds one more load everybody's per unit utility will be decreased.
- More than one agent is using this resource and some but not all of the agents may add one more load without crossing the critical load.

A rational agent can reason after Step 4 that to prevent over-utilization of the resource it should add one more load. For the first two possibilities above, it should not increment load. But as it is not sure which of the possibilities correspond to the current situation, it can use a probabilistic exploration scheme outlined in Steps 5 through 7 to reach its optimum load. It starts with initial increment probabilty of 1. At Step 5 it adds one load with its *increment probabilty*. If the load increment produces increased per unit utility then it does not change it's load any more. Otherwise, if the addition of one more load reduces it's per unit utility, it halves it's *increment probability* and tries later to add one more load with this reduced *increment probability*. It keeps on probing in this manner until it's probability falls below a threshold. The motivation behind this halving exploration process is similar to exponential backoff used for conflict resolution in shared communication channels like a token ring. The realization is that there must be other agents in the system who are trying to increase their loads as well, and unless every agent back off a little from their eagerness to increase system load, no one can benefit. Such exponentially decaying probabilities make it more likely for the system to converge to an equilibrium.

In our algorithm, an agent uses only local feedback to determine the load it applies on the system. Global information about individual loads used by other agents is never used. Our claim is that when *equilibrium* is reached, the combined load on the resource is exactly the *critical load* or capacity of the resource, i.e., the agents are using the resource optimally. They have reached this optimality through a distributed decision procedure using only local knowledge and without the directive of any central authority.

Now, we present some arguments for our decision procedure producing convergence to the resource capacity, C. Suppose after Step 4 the load reaches L

[1] If this assumption is violated a minor modification to the algorithm is required to reduce the loads and bring the total load near C. Our implementation includes this modification.

and the increment probability to 1. So, it is clear that $L \leq C < L + N$ has to be satisfied, i.e., there must be some y for which $L + y = C$, where $0 \leq y < N$. After Step 5 of the first iteration of the loop Step 5 through Step 7 the expected number of load added to the system is N. So, the total load on the system exceeds critical load and every agent reaches to Step 5 with system load L and increment probability 0.5. Now in the second iteration of the loop, the expected load after Step 5 is $L + \lfloor \frac{N}{2} \rfloor$. C may be greater than, equal to or less than this load. In the first two possibilities the load on the resource is expected to increase to $L + \lfloor \frac{N}{2} \rfloor$ and the number of agents who will be further willing to increase their load by one is remaining $N - \lfloor \frac{N}{2} \rfloor$ with their increment probability set to 0.5 in the next iteration of the loop. When $C < L + \lfloor \frac{N}{2} \rfloor$, all of the N agents will withdraw their added load resulting in the load on the resource be still at L and each agent's increment probability reduced to 0.25 in the next iteration of the loop. So, after each subsequent loop through Steps 5 through 7, the system progresses towards convergence because either the number of agents willing to increment their load is reduced to half compared to that of the last iteration of the loop with the same individual increment probability, or the number of agents willing to increment their load remains the same with the corresponding increment probability reduced to half compared to that of the last iteration of the loop.

So, after $\lfloor \log_2 N \rfloor + 1$ iterations of the loop Step 5 through Step 7, the load on the resource will be C and $C - L$ agents will feel content with their maximum possible load and $N - C + L$ number of agents will be still trying to add one more load with increment probability lies between 0.5 to $2^{-\log_2 N}$.

Now, we can say that after sufficient steps (with a loose upper bound of $\lceil \frac{(-\log_2(p_t))^2}{2} \rceil$, where p_t is the small thresold probability) of the loop Step 5 through Step 7 these $N - C + L$ agents satisfy the criterion of equillibrium as its increment probability will be less than a threshold. So, in any of the three possibilities of perception, every agent will not change its current load. This is also the optimal procedure for reaching the *critical load* of the resource as can be inferred from the theory of binary search.

4 Experimental Section

We have experimented with different scenarios considering different number of agents in the system and different critical loads. We now show results from some of the experiments with 10 and 100 agents. Here we take the value of K as 1. We set the critical load to different values with each agent following the decision procedure outlined above. We have chosen resource capacities such that equal distribution of load will not produce optimal resource utilization. These scenarios provide more difficult challenges compared to the case whre system capacity can be uniformly shared by the agents.

In Figure 1 (left), the load capacity is set to 62 and there are 10 agents in the society. The figure shows how the autonomous agents reaches equilibrium with the total load applied equalling the resource capacity. In this figure

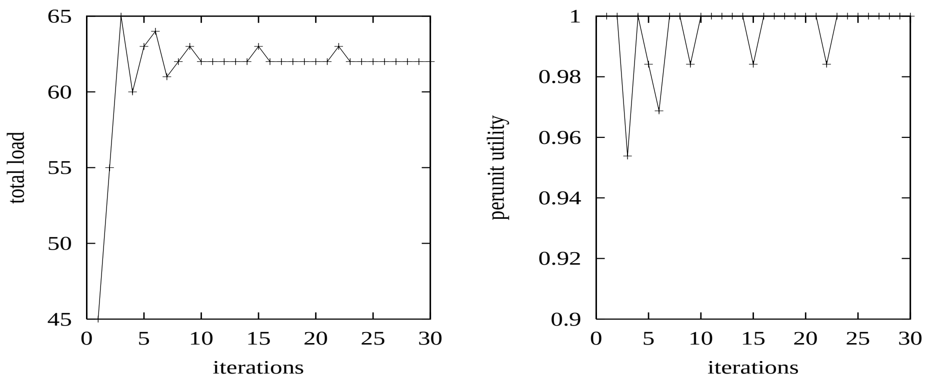

Fig. 1. (left)Variation of total load and (right) variation in average per unit utility of an agent: in the system with 10 agents and a load capacity of 62.

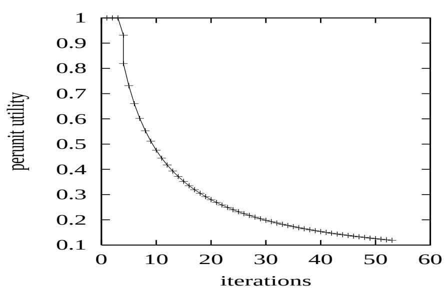

Fig. 2. Variation of average per unit utility of an agent in a system with 10 agents and a load capacity of 62 where the agents are not using local decision procedure and go on adding load in each iteration

we have shown that after initial random allocation of the load to the resource, agents steadily increment the load and then when the load exceeds the resource capacity, agents decrease their loads to reach equilibrium. The convergence phenomena is similar to the overshooting and undershooting typically observed in control systems where the control variable overshoots and undershoots the desired set point before settling. In Figure 1 (right), we present the variation of the average per unit utility of an agent over the course of a run. We can observe that initially there are a lot of deviations in the average per unit utility per agent as the system overshoots the critical load. Finally, however, optimal capacity is used at equilibrium and average per unit utility of an agent is reached to the maximum(which is 1 here as the value K) .

In this framework of the experiment, in Figure 2 we show how the average per unit utility of an agent goes down, if all agents go on adding load instead of using this decision mechanism.

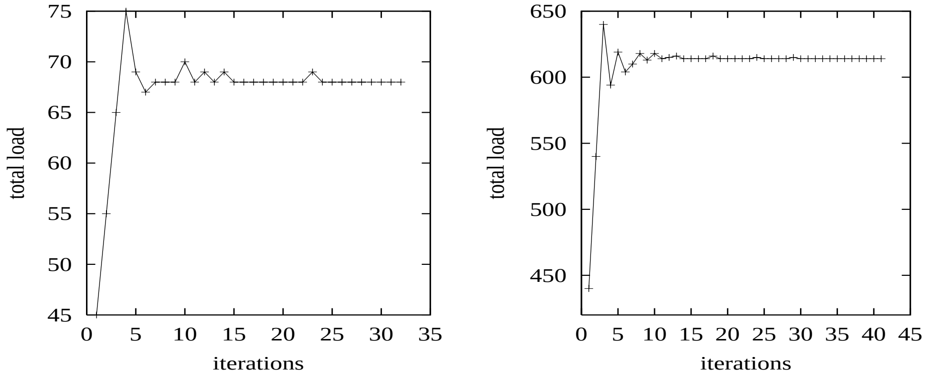

Fig. 3. (Left) Variation of total load on a system with 10 agents and a load capacity of 68 and (Right) Variation of total load on a system with 100 agents and a load capacity of 614.

In Figure 3 (Left), the critical load is changed to 68. This to ensure that the algorithm works for all types of critical loads. Here we have shown what happens after Step 4 of the algorithm. In the right figure of Figure 3, we use a larger society of autonomous agents where the size of the society is 100. Here the critical load is set to 614.

We have noted that the system reaches equilibrium with combined load equal to resource's critical capacity in all the scenarios we have experimented with. We also verify this claim with running each experiment for 100 times and observe no deviation from the convergence.

5 Conclusions

In this paper we have presented an algorithm to avoid the problem of Tragedy of the commons in a society of rational agents based only on local feedback in the form of utility received for the current load applied on a shared resource. We have shown that our proposed procedure results in the shared resource is used at its capacity load starting from arbitrary initial loads. As the tragedy of the commons is an important, and common problem which can lead to inefficiencies in the usage of shared resources, our procedure can have wide applicability.

Our procedure results in an equilibrium where some agents have higher utility than others even when everyone starts with the same load. This is because the equilibrium nature is static in the sense that no agent change their load after reaching equillibrium. We plan to work towards a more "fair", dynamic euqilibrium where agents increase/decrease their load around so that the capacity load is maintained while individuals with higher than average utility change over time. We also plan to augment our procedure such that equilibrium is reached in fewer iterations.

One of the drawback of this approach is that this considers only integral load. It will be interesting to study how it can be improved to work with real-valued loads.

Acknowledgments This work has been supported in part by NSF grants IIS-9724672 and IIS-0209208.

References

1. Arora, N., and Sen, S.: Resolving social dilemmas using genetic algorithms: Initial results. In *Proceedings of the 7th International Conference on Genetic Algorithms*, pages 689–695, San Mateo, CA, 1997. Morgan Kaufman.
2. Cammarata, S., McArthur, D., and Steeb, R.: Strategies of cooperation in distributed problem solving. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 767–770, Karlsruhe, Federal Republic of Germany, August 1983.
3. Durfee, E. H., and Lesser, V. R.: Using partial global plans to coordinate distributed problem solvers. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 875–883, Milan, Italy, August 1987.
4. Glance, N. S., and Hogg, T.: Dilemmas in computational societies. In *First International Conference on Multiagent Systems*, pages 117–124, Menlo Park, CA, 1995. AAAI Press/MIT Press.
5. Glance, N. S., and Huberman, B. A.: The dynamics of social dilemmas. *Scientific American*, 270(3): pages 76–81, March 1994.
6. Hardin G.: The tragedy of the commons. *Science*, 162: pages 1243–1248, 1968.
7. Hogg, T., and Huberman, B. A.: Controlling chaos in distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6): pages 1325–1332, December 1991. (Special Issue on Distributed AI).
8. Irvine A. D.: How Braess' paradox solves Newcomb's problem. *International Studies in the Philosophy of Science*, 7(2): pages 141–160, 1993.
9. Lloyd. W. F.: *Two Lectures on the Checks to Population*. Oxford University Press, Oxford, England, 1833.
10. Muhsam H. V.: A world population policy for the World Population Year. In *Jouranl of Peace Research*, 1(2): pages 97–99, 1973.
11. Mundhe, M. and Sen, S.: Evolving agent societies that avoid social dilemmas. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2000*, pages 809–816, 2000.
12. Sandholm, T. W. and Lesser, V. R.: Equilibrium analysis of the possibilities of unenforced exchange in multiagent systems. In *14th International Joint Conference on Artificial Intelligence*, pages 694–701, San Francisco, CA, 1995. Morgan Kaufman.
13. Smith A.: *The Wealth of Nations*. A. Strahan, Printer-stree; for T. Cadell Jun. and W. Davies, in the Strand, Boston, MA, 10 edition, 1802.
14. Tumer, K. and Wolpert, D. H.: Collective intelligence and Braess' paradox. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 104–109, Menlo Park, CA, 2000. AAAI Press.
15. Turner R. M.: The tragedy of the commons and distributed ai systems. In *Working Papers of the 12th International Workshop on Distributed Artificial Intelligence*, pages 379–390, May 1993.