Solving The Tragedy of the Commons by Adapting Aspiration Levels

Onkur Sen[1] and Sandip Sen[2]

[1] Oklahoma School of Science and Mathematics
onkursen@gmail.com
[2] Department of Computer Science
University of Tulsa
sandip@utulsa.edu

Abstract. The Tragedy of the Commons involves a community utilizing a shared resource (the "commons") which can sustain a maximum load beyond which its performance degrades. If utility received is proportional to the load applied on the system, individuals will maximize their applied load. Such greedy behavior will eventually lead to the total load exceeding the threshold. Thereafter, individuals will get less for adding more load on the system, which signifies a social dilemma. We develop computational models of such situations to study the effects of individual aspiration levels. An aspiration level corresponds to the satisficing return for an individual, which can be adjusted based on experience. We develop a computational model where individuals choose the load applied on the system based on their aspiration levels, thereby affecting the stability and performance of the "commons". The significance of our research is two-fold. First, our model can be used to evaluate the effects of using aspiration levels in computational environments, e.g., distributed computing networks, peer-to-peer systems, computational grids, etc. Secondly, our approach can model the effects of use of aspirations by individuals and organizations in our world, with effects ranging from local, e.g., traffic patterns, to global, e.g., environmental effects.

1 Introduction

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* that represents critical problems for large-scale systems with multiple independent actors. The problem involves a group of individuals (henceforth referred to as *agents*) who are using a shared resource which can only sustain a maximum threshold load. When the system's total load is below the threshold, utility is proportional to load. However, once the threshold has been passed, the system performance suffers, which is reflected in a rapid decline in yield per unit load. The goal of an agent is to maximize its utility, or the yield returned from its load on the system. Hence, an agent has an incentive to increase its load. However, if all agents behave in such myopic manner, the system will face more load than it can handle, i.e., total load will cross the threshold, and thereafter, all individuals will suffer from reduced utility. Therefore, the tragedy of the commons is a classic instance of a *social dilemma* because behavior based on short-term gains leads to losses in the long run for everyone. The classic example given is a group of herders, the agents, who are utilizing a shared pasture, the commons [9]. The herders try to maximize their herd size (load), but this leads to a depletion of the pasture's available resources, which is an indicator that the herders' utilities are decreasing. Real world examples of the *Tragedy of the commons* are increasingly observed in different scales and societal contexts, e.g., unsustainable agricultural practices and habitat destruction, traffic congestion, etc. 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. From a reverse viewpoint, the tragedy of the commons reappears in problems of pollution. Here it is not a question of taking something out of the commons, but of putting sewage, or chemical, radioactive, and heat wastes etc into the environment to be specific into the water. The utility is what they achieved by putting this in environment instead of purifying this. Since this is true for everyone as an independent, rational individuals. But this collectively harm everyone of the individuals with a greater impact when the pollution in the environment becomes high. Thus we are locked into a system of "fouling our own nest". More recently, attention has been drawn to the tragedy of the commons in the context of autonomous agent systems [21]. These and other problems arise in multiagent societies [16] as multiple, distributed decision-makers try to maximize local utility by taking decisions based only on limited global knowledge. We are particularly interested in social dilemmas and, in particular, tragedy of the commons situations, in distributed computational frameworks, including computer networks, peer-to-peer (P2P) and grid systems, modem pools, shared databases, printers, and servers. The tragedy of the commons often leads to reduced throughput system inefficiency and failure and undermines user satisfaction and trust in such distributed systems. For example, when the number of users connecting to a fixed-size modem increases, the time required to connect increases, which leads individuals to hog available resources (users tend to maintain their connection rather than logging off after finishing their work). This, in turn, leads to further delays in establishing a connection. We propose and evaluate a distributed computational approach to solve social dilemmas which assigns each agent an aspiration level that is used to suppress greedy behavior and choose an effective load that will provide satisfactory performance [5]. An aspiration level is a measurement of the utility an agent needs to be satisfied, which is adjusted up or down over time based on the utility received from the environment. In our model, we assume agents have very limited information about the system. In particular, the agent does not know the number of agents, the load each applies, or the total load. Therefore, the agent adjusts its behavior based on a history of its past loads and corresponding utilities. Our research goal is to develop such local decision-making procedures that will allow the system to work at near-optimum capacity without going significantly over the threshold and, hence, solve the tragedy of the commons. The rest of the paper is organized as follows: in Section 3 we present our problem formulation; in Section 4 we discuss our proposal for adapting load to be applied and aspiration levels; in Section 5 outline our experimental framework; in Section 6 we analyze our experimental results and their implications; in Section 7, we form a conclusion about our results and propose future work.

2 Related Work

2.1 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 [7]. 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. In the following, we present a few quintessential social dilemmas:

– Consider a scenario where a public good is to be initiated provided enough contribution is received from the populace. Public goods are benefits produced by the society and available to all of its members regardless of individual contribution. Examples of public goods in human societies include provision of parks, roads, a clean environment and national defense. A selfish individual will prefer not to contribute towards the public good, but utilize the benefits once the service is in place. Any public good needs the contribution of a certain percentage of the populace to be either initiated or maintained. Therefore, if a sufficient percentage of agents in the populace makes the selfish choice, the public good may not be produced. Every agent then has to face the dilemma of whether to contribute or to exploit. 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 [10]. So, the dilemma for each agent is whether or not to make the greedy choice.

2.2 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 [19]. As a rebuttal to this theory, William Forster Lloyd presented the *tragedy of the commons* scenario in 1833 [14]. Lloyd's scenario consisted of a pasture shared by a number of herdsmen for grazing cattle. This pasture has a capacity, say L, i.e., each time a cattle added by a herdsman result in a gain as long as the total number of cattle in the pasture, x, is less than or equal to L. When $x > L$, 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 cattle 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?" [9]. 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*. Literature about the tragedy of the commons is extensive [9]. Diecidue and van de Ven show how aspiration levels are linked to expected utility [5]. Gilboa and Schmeidler extend on this and provide mathematical procedures to adjust aspiration levels based on utility returned [8]. Macy and Flache indicate how agents learn over time in social dilemmas [16] Researchers in computational systems have discussed the existence of the tragedy in the commons in computational systems [3, 12, 21]. Muhsam [15] has shown that if some or all other herdsmen add cattle when $x > L$, 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 equilibrium*. Our contribution is to successfully adapt the aspiration level mechanism to solve the tragedy of the commons.

2.3 Computational Approaches

Multiagent systems 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 [4] 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 [18]. 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 [20]. Glance and Hogg [6] 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 knowledge. 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 [11]. 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 [17]. 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.

3 Problem Formulation

For an agent $i \in \mathbf{N}$, the set of all agents, we denote U_i^t, h_i^t, and L_i^t as the **utility**, **aspiration level**, and **load**, respectively, at time t. In addition, let the **total load** of the system at time t be $\mathcal{L}^t$, the **threshold load** be ϕ, and $n = |\mathbf{N}|$ be the **population size**.

Utility Functions U_i^t can be expressed under one of the following circumstances:

$$\begin{aligned} U_i^t &= L_i^t \times \delta, \text{ where } \mathcal{L}^t < \phi \\ &= L_i^t \times \delta e^{-k(\mathcal{L}^t - \phi)}, \text{ otherwise.} \end{aligned}$$

where δ is the utility per unit load under normal operating conditions and k is the **environmental factor**, which determines the rate at which system performance decreases after the threshold is crossed. Hence, after the load on the resource crosses the threshold the utility per unit load diminishes exponentially.

4 Adjusting Aspiration Levels and Load

If $\mathcal{L}^t < \phi$, we update the aspiration level [8] as follows:

$$h_i^t = \alpha U_i^t + (1 - \alpha) U_i^{t-1},$$

where α is a **learning rate** which weights the current utility received with the previous aspiration level. The initial load applied by an agent is selected randomly in the range $(0, l_{init})$. The agent subsequently increases its applied load by a constant Δ_i as long as the utility received is higher than its aspiration level, i.e., $U_i^t > h_i^t$. However, once the utility decreases, our agents choose the next load as halfway between the current load and L_i^τ, where τ is the last time that caused an increase in its aspiration level. Therefore,

$$\begin{aligned} L_i^t &= L_i^{t-1} + \Delta_i, \text{ where } U_i^t > h_i^t \\ &= \frac{L_i^{t-1} - L_i^\tau}{2}, \text{ otherwise.} \end{aligned}$$

We now present the algorithm used by an agent at each time step t:

Data: Aspiration level, load, utility
$U_i^t = \text{getUtility}()$;
if $U_i^t < h_i^t$ **then**
  $L_i^t = \frac{L_i^{t-1} - L_i^\tau}{2}$;
else
  $h_i^t = \alpha U_i^t + (1 - \alpha) U_i^{t-1}$;
  $L_i^\tau = L_i^t$;
  $L_i^t = L_i^t + \Delta_i$;
end

5 Experimental Framework

We introduce the performance metrics used and the different scenarios we use to evaluate our approach to solving social dilemmas.

5.1 Performance Metrics

To evaluate our framework, we collected and analyzed data from various system metrics:

Social Welfare: Social welfare is the total utility received by all agents and captures the overall system performance: $\mathcal{U}^t = \sum_{i=1}^{n} U_i^t$. This was our primary metric.

Total Load: To evaluate system efficiency, we see how close $\mathcal{L}^t$ is to ϕ. A truly efficient system will have a load close to its threshold without crossing it.

Average and Standard Deviation of Agent Utility: Average agent utility, $\mu^t = \frac{\mathcal{U}^t}{n}$, is proportional to social welfare. The standard deviation of agent utility, $\sigma^t = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (U_i^t - \mu^t)^2}$, is used to analyze the deviations in load between agents. These can be used to check for *free riders*, i.e., agents who benefit at the expense of others.

5.2 Base Case Scenario

To help better understand the underlying system dynamics and the working of our approach, we created and experimented with a base case scenario. We chose parameters that would allow us to perform representative experiments while still allowing us to inspect individual behavior and environment modules. For the base case, we choose $l_{init} = 0.2$, $\alpha = 1$ and $\forall i, \Delta_i = 0.01$ and $h_i = 0$. In addition, we used the following default system parameter values: $\phi = 7.5$, $k = 3$, and $n = 10$.

5.3 Parameters to be Varied During Experimentation

Moving beyond the base case, we carefully generated a large variety of environments to critically evaluate the strengths and weaknesses of our proposed approach. To generate these environments, we vary the following three parameters over the corresponding stated ranges (the parameters for the base case are in bold).

Environmental Factor: We vary k over the set $\{0.25, 0.5, 1, \mathbf{3}, 5\}$.
Population Size: We vary n over the set $\{\mathbf{10}, 100, 1000\}$.
Learning Rate: We vary α over the set $\{0.5, 0.75, \mathbf{1}\}$.

For statistical verification, we averaged data for the metrics mentioned over 10 independent runs.

6 Results and Discussion

6.1 Base Case

Figure 1 shows graphs describing the base case. Initially agents receive higher utility with increased load, which results in increased aspirations. With increasingly higher individual loads applied by agents on the system, $\mathcal{L}^t$ surpasses ϕ between iterations 65 and 70, and the social welfare drops immediately. Correspondingly, individual agent utilities drop below their aspiration levels, and they revert back to applying lower loads. Hence, the total load drops under ϕ, then back over by a smaller margin between iterations 70 and 85. There are further overshoots and undershoots until the system stabilizes, with both $\mathcal{L}$ and $\mathcal{U}$ approaching ϕ. This shows that the social dilemma was successfully countered by adapting aspiration levels in the base case scenario.

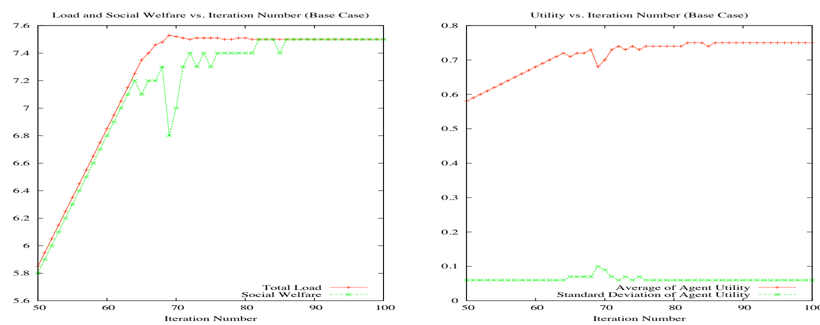

Fig. 1. Total Load and Social Welfare (left) and Average and Standard Deviation of Agent Utility (right) for the base case.

6.2 Varying Environmental Factor (k)

Figure 2 shows graphs for total load and social welfare over different values of k. For all values of k except 0.25, the system stabilizes at ϕ. However, when $k = 0.25$, some agents acting as free riders are able to increase their load and utility while others maintain their load at a constant level. This is because with small increments, the utility of the free riders continue to increase whereas the utilities of the other agents plummet. To study this problem analytically, let $j \in N$ be a lone free rider and N_{-j} be the remaining agents, L_j be the current load of j and L the total load. For the free rider, the utility for incrementing its load by Δ_j is positive, i.e., $L_j^t \times \delta e^{-k(\mathcal{L}^t-\phi)} > (L_j^t + \Delta_j) \times \delta e^{-k((\mathcal{L}^t+\Delta_j)-\phi)}$. For any other other agent, $z \in N_{-j}$, the utility decreases as they hold their load constant, i.e., $L_z^t \times \delta e^{-k(\mathcal{L}^t-\phi)} < L_z \times \delta e^{-k((\mathcal{L}^t+\Delta_j)-\phi)}$. Essentially, for $k = 0.25$, the utility of free riders is not decreasing rapidly enough as they apply more load. Hence, such environments sustain a minority of free riders at the expense of the majority. This result is an intriguing outcome of the dynamics of aspiration level adaptation and environmental characteristics that deserve further investigation. We performed significance tests using the average and standard deviations of each of the evaluation metrics discussed above. The converged values of social welfare and total load on the system were statistically different for $k = 0.25$ compared to the other values for k.

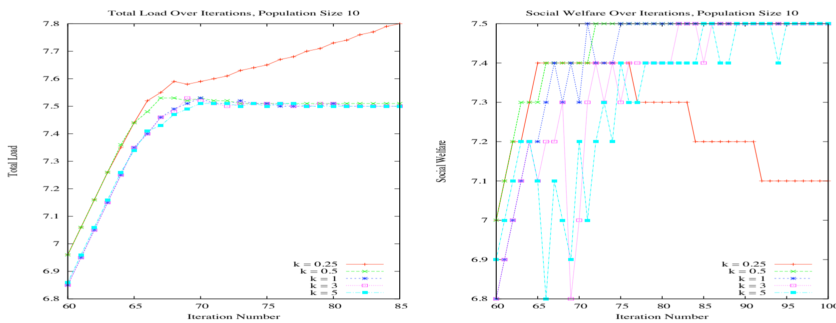

Fig. 2. Total Load (left) and Social Welfare (right) for various k ($n = 10$).

6.3 Varying Population Size (n)

Figures 3 and 4 show graphs for total load and social welfare for different population sizes with different k values. For $n = 100$, we see a similar converging pattern as for $n = 10$. For all k values but 0.25, the system stabilizes. However, with the presence of free riders at $k = 0.25$, the load on the system continues to increase, and the utility decreases. When $n = 1000$, however, we do not observe runaway free-riding and social welfare do not continue to plummet. Rather, the social welfare of the system stabilizes, albeit at suboptimal values. This intriguing result can be explained by observing that as there are more free riders, eventually the system reaches a state where further increase of load is detrimental for all. Hence, though free riding is not eliminated, its effect is curtailed. This self-stabilizing nature of the system for larger population sizes is a very interesting, and unexpected, property of the use of aspiration levels. It would be interesting to further study this phenomena to see if the drop in system performance can be minimized while avoiding free riding. The system, though suffers from suboptimal performance for $k \in \{0.25, 0.5, 1\}$. We conclude that for the current approach, for large populations, larger values of k are required to produce optimal system performance, i.e., system performance must degrade steeply to completely eliminate free riding.

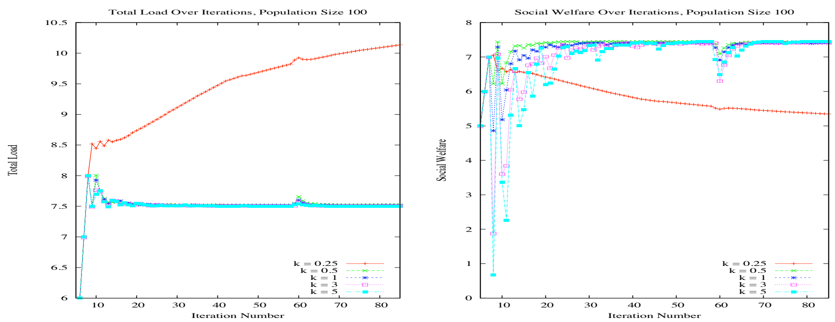

Fig. 3. Total Load (left) and Social Welfare (right) for various k ($n = 100$).

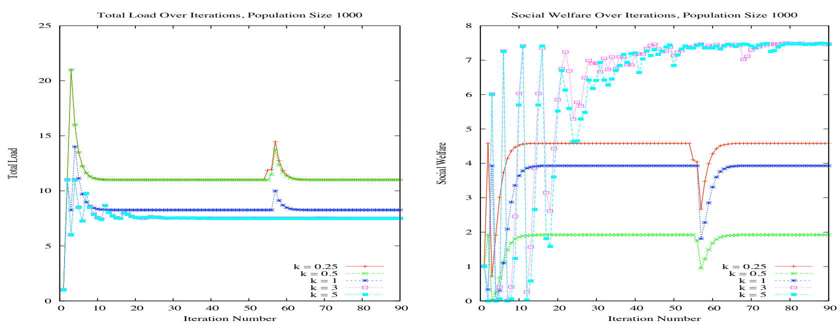

Fig. 4. Total Load (left) and Social Welfare (right) for various k ($n = 1000$).

6.4 Varying Learning Rate (α)

The graphs in Figure 5 show a substantial difference in the system when α is changed. The system stabilizes when $\alpha = 1$. However, the system performs a little worse when $\alpha = 0.75$ ($\mathcal{L}$ increases by 0.005) and degrades further when $\alpha = 0.5$. This is because the aspiration level is not being adjusted fast enough towards the current utility level. This suggest that fast learners will be able to avoid social dilemmas more consistently.

7 Conclusions and Future Work

Our research goal is to develop a distributed computational approach to solve the tragedy of the commons, a social dilemmas. We investigate a solution approach where agents adapt their aspiration levels to counter such social dilemmas with limited information about the system. The aspiration levels were adjusted based on utility returned at the given load being applied. We systematically varied environmental factors as well as an agent's behavioral parameters to observe

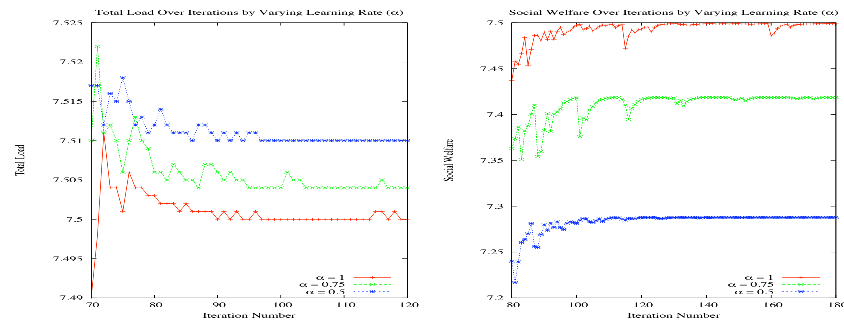

Fig. 5. Total Load (left) and Social Welfare (right) for Various $\alpha, k = 3, n = 10$.

and analyze the scope and the effectiveness of our approach. An interesting result was that in benign environments, where the system degrades slowly above its threshold load, a minority of free riders was able to benefit at the expense of the community. For large systems with many agents, however, free riders limit each others' exploitation and mitigate the adversarial effect on the system. We also observed that faster learners were successful in avoiding the social dilemma in a wider range of scenarios. We further plan to run experiments to evaluate our proposed mechanism in the following scenarios:

Asynchronous Updating: In our current work, all agents update their loads at the same time. We want to investigate situations where agents update their loads asynchronously.

Dynamically Changing Population: We plan to investigate open environments, i.e., situations where agents may enter and leave the population, thereby changing the population size.

Irreversible Systems: In the environments we studied, the performance of the system returns to its previous level if the load is reduced after crossing the threshold. We plan to study real-world environments where crossing the threshold causes irreversible damage, i.e., the optimum utility can not be regained after crossing the threshold. Maintaining effective performance in such environments will be especially challenging.

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. de Cote, Enrique Munoz et al.: Learning to cooperate in multi-agent social dilemmas. *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems,* pages 783-785, 2006.
4. 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.
5. Diecidue, Enrico and van de Ven, Jeroen.: Aspiration Level, Probability of Success and Failure, and Expected Utility. *International Economic Review*, 49(2):683-700, May 2008.
6. 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.
7. Glance, N. S., and Huberman, B. A.: The dynamics of social dilemmas. *Scientific American*, 270(3): pages 76–81, March 1994.
8. Gilboa, Itzhak and Schmeidler, David.: Reaction to price changes and aspiration level adjustments. *Review of Economic Design*, 6:215-223, 2001.
9. Hardin G.: The tragedy of the commons. *Science*, 162: pages 1243–1248, 1968.
10. 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).
11. Irvine A. D.: How Braess' paradox solves Newcomb's problem. *International Studies in the Philosophy of Science*, 7(2): pages 141–160, 1993.
12. Ito, Akira.: How do autonomous agents solve social dilemmas? *Intelligent Agent Systems: Theoretical and Practical Issues,* LNCS vol. 1209, 1209:177-188, 1997.
13. Kollock, Peter.: Social Dilemmas: The Anatomy of Cooperation. *Annual Review of Sociology*, 24:183-214, 1998.
14. Lloyd. W. F.: *Two Lectures on the Checks to Population*. Oxford University Press, Oxford, England, 1833.
15. Muhsam H. V.: A world population policy for the World Population Year. *Journal of Peace Research*, 1(2): pages 97–99, 1973.
16. Macy, Michael W. and Flache, Andreas.: Learning Dynamics in Social Dilemmas. *Proceedings of the National Academy of Sciences of the United States of America*, pages 7229-7236, 14 May, 2002.
17. 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.
18. 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.
19. 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.
20. 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.
21. 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.