Coalition Formation for Large-Scale Electronic Markets *

K. Lerman[1] and O. Shehory[2]
1. Information Sciences Institute
Univ. of Southern California
Marina del Rey, CA 90292-6695
2. IBM Research Lab in Haifa, the Tel-Aviv site,
Tel-Aviv, 61336 Israel

Abstract

Coalition formation is a desirable behavior in a multi-agent system, when a group of agents can perform a task more efficiently than any single agent can. Computational and communications complexity of traditional approaches to coalition formation, e.g., through negotiation, make them impractical for large systems. We propose an alternative, physics-motivated mechanism for coalition formation that treats agents as randomly moving, locally interacting entities. A new coalition may form when two agents encounter one another, and it may grow when a single agent encounters it. Such agent-level behavior leads to a macroscopic model that describes how the number and distribution of coalitions change with time. We increase the generality and complexity of the model by letting the agents leave coalitions with some probability. The model is expressed mathematically as a series of differential equations. These equations have steady state solutions that describe the equilibrium distribution of coalitions. Within a context of a specific multi-agent application, we analyze and discuss the connection between the global system utility and the parameters of the model.

1. Introduction

In the last few years, the electronic marketplace has witnessed an exponential growth in worth and size, and projections are for this trend to intensify in coming years. Yet, the tools available to market players are very limited, thus imposing restrictions on their ability to exploit market opportunities. Implemented electronic markets do not offer means for buyers and sellers to group together and exploit the benefits of such grouping. A buyer in today's electronic market can buy through two broad types of mechanisms: (1) fixed prices are given by a seller at the seller's site, and the buyer has to make a "take it or leave it" decision; (2) the buyer submits his/her price to an auction server and, according to the rules of the auction, the server determines the winner. No explicit negotiation or grouping are supported, and no tools or protocols for these are provided. In addition, all of the burden of buyer's interaction with the sellers' servers and auctions is on the buyer.

Research on agent coordination and cooperation (e.g., [14, 1]) shows that, at the conceptual level, such activities may increase the agents' benefits, even when the agents are self-interested, as common in the marketplace. Researchers have shown, via theoretical analysis and simulations, that forming coalitions can serve as a feasible means for such cooperation [13, 14, 9]. However, the coalition formation mechanisms proposed to date apply to *deliberative* agents. While they might be suitable for dozens of agent, they will not scale up to thousands of agents, that are expected to participate in the electronic marketplace, due to their computational and communications complexity. We are interested in a new paradigm of MAS design that can accomodate a large number of agents and still provide a good enough performance (in terms of agent benefits and consumption of computational resources). We draw inspiration from natural systems in which complex global structures and behaviors result from local interactions among many simple elements [17]. The model presented in this paper provides means for coalition formation among *simple* buyer agents in an electronic marketplace. The low computational overhead of the model makes it appropriate for even very large systems. In particular, we address the wholesale market, where sellers benefit from selling large quantities of goods in bulk, because it reduces their manufacturing, advertising and distribution costs. They usually choose to pass some of the savings to the buyers. If the buyers do not individually need large quantities of the goods, it is still beneficial for

*This material is based on work supported in part by MURI contract N00014-96-1222, DARPA contract F30602-98-2-0138, NSF grant IRI-9612131 and an ISI-ISD Research Fund Award.

them to form buyers' coalitions, allowing them entry into the wholesale market and reducing the price per unit. To the best of our knowledge, there is no implementation of a wholesale virtual market to date where agents collaborate and perform many-to-many multi-attribute negotiations on behalf of buyers and sellers.[1]

2. The goal

The main contributions of this paper are a low complexity coalition formation mechanism that is applicable to large scale MAS and an analytical study of this mechanism. Our research has several goals:

- *Propose a feasible mechanism* in which interactions between agents with simple local strategies lead to desirable group behavior, and in particular to beneficial coalition formation.
- Show that, unlike other current models of coalition formation, our approach incurs *minimal communications and computation costs*. This low complexity approach provides a unique solution to the case of coalition formation among *rational self-interested* agents in large-scale MAS, with which we deal.
- *Decompose* the proposed coalition formation mechanism into its component agent strategies. We show that even the simplest set of agent actions leads to non-trivial group behavior. We also show how to relate the details of the mechanism to the agent's goals and actions.
- *Create a mathematical description* of the proposed model. We construct a macroscopic (coalition-level) model and discuss how model parameters are related to the details of microscopic (agent-level) interactions.
- *Study the model quantitatively* for different parameter values. We relate these results to predictions about group formation behavior, *e. g.*, the number and size of coalitions in a large scale MAS.

3. The approach

Some previous attempts to mathematically study a large-scale MAS used a physics-based [15] approach. Several advantages were shown to stem from such an approach. First, a MAS based on local interactions can promise very low, sometimes even constant, communication complexity. The physics-based models also provide simple agent decision making mechanisms. This simplicity results in reduced computational complexity. The simplicity of interactions makes the global properties of the system amenable to mathematical analysis. In other models, this is usually done experimentally or via simulations, and not analytically.

To date, a physics-based approach has been used to study coordination and task allocation in a cooperative MAS. In our work we present a new MAS model in which cooperative behavior, such as coalition formation, arises out of the interactions among many *simple self-interested* agents. We study quantitatively the bulk properties of a large scale MAS, and draw connection between agent properties and global system behavior.

We refer to a system of multiple mobile purchasing agents, each of which is given a task to obtain goods, with the goal of minimizing the price paid for the goods. We assume that agents have, or can acquire, contact information of vendors which supply the requested goods and the retail (base) price for the product. Such information can be provided via middle agents [4] and other agent location mechanisms [16]. We also assume that, given no additional information, agents have no *a priori* preference among the vendors, and each agent makes a random selection from its list of vendors and moves to the vendor site. If it encounters other agents or a coalition of agents, it can join the coalition or, in the case of single agents, form a new coalition with them. We assume that the agents are self-interested and will choose to join coalition when such a choice is available and beneficial. However, since further exploration of the coalition formation space may result in joining a more beneficial coalition, our model, unlike previous ones, allows agents to leave coalitions. The agents join coalitions by placing an order to purchase a product, and they leave coalitions by withdrawing an order for the product. The orders remain open for some period of time to allow new orders to come in. At the end of the specified time, the orders are filled, and the price each agent pays for the product is based on the size of the final purchasing coalition.

The coalition formation mechanism outlined above requires minimal communication between agents, and because their decision depends solely on local conditions, it also requires no global knowledge. Agents learn indirectly about the presence and size of the coalition at a particular vendor site by querying the vendor for the current price of the product.[2] The price is lower if a buying coalition exists

[1] Note that some e-stores provide a service which has some similarities. For instance, Mobshop (http://www.mobshop.com) provides increasing savings according to a sale volume, and one can join a sale to receive this discount. Not surprisingly, they provide the service only for products *they* sell. For small bulk orders for products, they offer a slight discount from the manufacturer suggested price, which is usually high compared to discount sales (with no aggregation) on the web. Participation in real wholesale markets as our mechanism should allow may provide much better deals by letting buyers bypass middlemen such as Mobshop.

[2] Vendors may lie about these prices, however since agents can ask other agents as well, such lying may result in bad reputation and harms the vendor's business. Although manipulative behavior of all sides is possible, there are mechanisms that can vastly reduce the expected utility from, and the probability of, such behavior (e.g., in [9]).

at that vendor site.[3] In general, there is an inverse relationship between the size of the coalition and the price of the product. In the simplest case, this relationship is linear – the larger the coalition, the lower the price. We will specify the coalition price in a later section, when we use it to compute savings by all agents. Such indirect communication is analogous to stigmergetic communication in insect societies that has been shown to lead to interesting collective behaviors [3].

4. The model

We decompose the coalition formation mechanism to its basic components. This approach is general—it is applicable to other systems in which aggregation occurs. The components are then used to build a mathematical model of the MAS. As models are in general, our model is an idealized representation of a process. To be useful, our model explicitly takes into account the salient details of the coalition formation process it describes. Thus, it is an indispensable tool for understanding the process. Similar to others, our model can become a better approximation of reality by including more details. These additions can correct the quantitative results of the simpler model, but not invalidate them. Therefore, we focus on the simplest model of the mechanism that leads to coalition formation in a MAS. This initial model may be refined and adapted to different MAS environments. We leave the extensions of the model for future work.

The following axioms capture the most important features of the proposed coalition formation process:

- Agents are homogeneous in a sense that each agent has the same goal (to purchase a specific product at the lowest price) and follows the same strategy for coalition formation.
- Agents are free to choose among vendors; coalitions are not.
- Agents encounter other agents and coalitions randomly.
- Each agent's strategy is determined entirely by local conditions, such as the size of the coalition present at a particular vendor site.
- It is beneficial for agents to join a coalition, because as a member of a coalition, it will pay less for the product.
- Some constraints may limit the maximum size of the coalition. E.g., Manufacturing and distribution constraints limit the bulk order to a maximum size; therefore, vendors will not accept bulk orders greater than this maximum.
- Agents are self-interested; therefore, given several alternatives, they will prefer and select the more beneficial ones.
- The agents are spatially uniformly distributed, apart from non-uniformities inherent in the coalitions. Even if this is not true initially, over time agents will tend to become uniformly distributed, because they have no preference among vendors.
- An agent may leave an existing coalition with some probability. This is a reasonable strategy, because in many cases, there may exist a better coalition for it to join.

This model excludes explicit negotiation from the coalition formation process. In this paper we show that, for the problem domain we study, the model as described above still leads the agents to form robust coalitions.

4.1. The microscopic *vs* macroscopic descriptions

Using the above axioms we may construct a microscopic theory of the coalition formation process, that treats the individual agent as a fundamental unit in the model. This model would describe how agents make decisions to join coalitions. Simulating a system composed of many such agents, modeled as cellular automata for instance, would give us an understanding of the global behavior of the system.

Alternatively, we may construct a macroscopic model that treats coalitions as the fundamental units. A macroscopic description offers several advantages, the most important of which is that such a model directly describes the global properties of the system we are interested in studying, namely the number and size of coalitions, and how these quantities change with time. This approach is more computationally efficient, because it uses many fewer variables than the microscopic model. The macroscopic theories tend to be more universal, and therefore, more powerful. The same mathematical description can be applied to other systems governed by the same abstract principles, which are outlined in the axioms above. At the heart of this argument is the concept of a separation of scales, which holds that the details of the microscopic interactions (e.g, among agents) are only relevant to computing the values of the parameters of the macroscopic model (e.g., how fast coalitions grow). This idea has been used by physicists to construct a single model that describes the behavior of seemingly disparate systems, e.g., pattern formation in convecting fluids and chemical reaction-diffusion systems [17]. Of course, the two descriptive levels are related, and it may be possible in some cases to exactly derive the parameters of the macroscopic model from the microscopic theory.

In this paper, we construct a macroscopic model that captures the dynamics of the coalition formation process. This

[3] Conceptually, a buyers' coalition can be located at sites other than the vendor's site, e.g., at some mediator's or coordinator's site.

is a phenomenological model, since we do not derive it from the microscopic theory. The model is expressed mathematically as a set of first order differential equations that describe how the number of coalitions of each size evolves in time. As we argued above, this description is not intrinsic to the e-commerce application we chose to study—it can describe a number of systems where aggregation occurs. In our application, however, there is a simple connection between the microscopic behavior of the agents and the parameters of the model. For other applications this relationship may be more complex, and it might be necessary to derive the parameters.

In the following sections we present the mathematical model and study the behavior of solutions for different parameter values. We show that solutions reach a steady state in which the distribution of coalitions no longer changes. We define a utility gain function, the measure of savings achieved by all agents in the system, and calculate its value for each steady state solution. We find that the steady state distribution and utility gain depend strongly on the rate at which agents leave coalitions. We discuss the implications of the behavior of these solutions on an agent's design.

5. The macroscopic model

Let $r_1(t)$ denote the number of unaffiliated agents (monomers) in the system at time t, $r_2(t)$ the number of coalitions of size two (dimers), *etc.*; $r_n(t)$ the number of coalitions of size n at time t, up to a maximum coalition size m. We assume that there is no net change in the number of agents, and therefore expect a realistic dynamic process to conserve the total number of agents in the system (*i.e.*, $\sum_{n=1}^{m} n r_n = N$, where N is the total number of agents in the system).

5.1. Global utility gain

The global utility gain measures the efficiency of the system—the price discount all agents receive by being members of coalitions. The value of this metric is expected to be high when there are many large coalitions, and conversely, it is low, meaning the system is less efficient, when there is a large number of unaffiliated agents. Note that global benefit is achieved even while each agent is selfishly maximizing its individual gain [8]. The retail price that an unaffiliated agent pays to the vendor for the product is p, and the coalition price that each member pays is $p_n < p$, which depends on the size of the coalition. In the simplest model we let $p_n = p - \Delta p(n-1)$, where Δp is some price decrement. The total discount for all agents is:

$$G = Np - \sum_{n=1}^{m} p_n n r_n$$

Expanding and using conservation of the number of agents to eliminate terms, yields the following expression for the utility gain per agent:

$$G/N = \Delta p \left(\sum_{n=1}^{m} \frac{n^2 r_n}{N} - 1 \right) . \qquad (1)$$

5.2. Dynamic equations

Initially (at $t = 0$) the system consists of N agents and no coalitions. We assume that there is no spatial dependence in the agent distribution, apart from coalition-based aggregation. A series of coupled rate equations [2] describe how the number of coalitions of different size changes in time. The solutions of these rate equations yield the coalition distribution at any given time. The equations are written as follows:

$$\begin{aligned}
\frac{dr_1}{dt} &= -2D_1 {r_1}^2(t) - r_1(t) \sum_{n=2}^{m-1} D_n r_n(t) + \\
&\quad 2B_2 r_2(t) + \sum_{n=3}^{m} B_n r_n(t) , \\
\frac{dr_n}{dt} &= r_1(t) \left(D_{n-1} r_{n-1}(t) - D_n r_n(t) \right) - \\
&\quad B_n r_n(t) + B_{n+1} r_{n+1}(t) , \\
\frac{dr_m}{dt} &= D_{m-1} r_1(t) r_{m-1}(t) - B_m r_m(t)
\end{aligned}$$

Here $r_n(t)$ is the number of coalitions of size n at time t, and $\frac{dr_n}{dt}$ is the rate of change of this number. Parameter D_n, the attachment rate, controls the rate at which unaffiliated agents join coalitions of size n. This parameter includes contributions from two factors: the rate at which agents encounter n-mers ($\propto r_1 r_n$, where the proportionality factor determines how many vendor sites an agent visits in a given period of time), and the probability of joining the coalition of size n. B_n, the detachment rate, gives the rate at which agents leave coalitions of size n. The solutions are subject to the initial conditions: $r_1(t = 0) = N$ and $r_n(t = 0) = 0$ for all $n > 1$.

The rate equations can be interpreted in the following way: when two unaffiliated agents find themselves on the same site at the same time, they form a coalition of size two, decreasing the number of unaffiliated agents in the system by two (hence the factor 2 in the first equation) and increasing the number of size-two coalitions by one. The number of single agents also decreases when an agent joins a coalition of size n. However, the number of single agents can also increase when an agent leaves a coalition of any size. Likewise, the number of n-mers increases when an unaffiliated agent encounters a group of size $n-1$, or an agent leaves a coalition of size $n+1$, but it decreases when n-mers

themselves absorb unaffiliated agents or have the agents detach from them. The last equation states that agents cannot join coalitions of maximum size, m.

We restrict our attention to the uniform attachment–uniform detachment case: $D_n = D$, $B_n = B$ for all n. To simplify the analysis, we rewrite the equations in dimensionless form by making the following variable transformations: $\tilde{r}_n = r_n/N$, $\tilde{t} = DNt$, and $\tilde{B} = B/DN$. $\tilde{r}_n$ is the density of coalitions of size n. The rate equations in dimensionless form are:

$$\frac{d\tilde{r}_1}{d\tilde{t}} = -2\tilde{r}_1^2(t) - \tilde{r}_1(t)\sum_{n=2}^{m-1}\tilde{r}_n(t) + 2\tilde{B}\tilde{r}_2(t) + \sum_{n=3}^{m}\tilde{B}\tilde{r}_n(t)\,, \quad (2)$$

$$\frac{d\tilde{r}_n}{d\tilde{t}} = \tilde{r}_1(t)\left(\tilde{r}_{n-1}(t) - \tilde{r}_n(t)\right) - \tilde{B}\tilde{r}_n(t) + \tilde{B}\tilde{r}_{n+1}(t)\,, \quad (3)$$

$$\frac{d\tilde{r}_m}{d\tilde{t}} = \tilde{r}_1(t)\tilde{r}_{m-1}(t) - \tilde{B}\tilde{r}_m(t). \quad (4)$$

Note that the attachment rate no longer explicitly appears in the equations. There is now a single variable parameter in the equations, the dimensionless detachment rate $\tilde{B}$, which measures the relative strength of detachment *vs* the rate at which agents join coalitions. We investigate the behavior of solutions of the equations as this parameter is varied.

Numerical integration of these equations[4] shows that solutions reach a steady state, after which the coalition densities no longer change, but the time it takes for them to do so depends sensitively on $\tilde{B}$. The steady state coalition densities and the global utility gain (calculated from Eq. 1) also depend on this parameter.

Figure 1 shows the time evolution of the solutions for three different values of the dimensionless detachment rate: $\tilde{B} = 0$ (no-detachment case), $\tilde{B} = 10^{-5}$ and $\tilde{B} = 10^{-2}$. Maximum coalition size is six in all cases. For the no-detachment case, the density of unaffiliated agents quickly drops to zero while the system reaches its final configuration where coalitions of size two and three predominate. In contrast, it takes much longer for solutions to reach their final values for $\tilde{B} = 10^{-5}$ than either for $\tilde{B} = 0$ or $\tilde{B} = 10^{-2}$. The density of unaffiliated agents, $\tilde{r}_1$, reaches a small but finite value for $\tilde{B} = 10^{-2}$, indicating that, unlike the no-detachment scenario, there are some agents left at late times who are not part of any coalition. Notice that this density is larger for $\tilde{B} = 10^{-2}$ than for $\tilde{B} = 10^{-5}$. In general, we expect the number of free, or unaffiliated, agents to increase as the detachment rate is increased. We note that for $\tilde{B} \neq 0$, the steady state is an equilibrium state: even though agents are continuously joining and leaving coalitions, the overall distribution of coalitions does not change. For $\tilde{B} = 0$, the system gets trapped in an intermediate non-equilibrium state before it is able to form larger coalitions.

The equations were integrated numerically for $m = 6$ and different values of $\tilde{B}$. Figure 2 shows how the steady state coalition densities change as $\tilde{B}$ is increased. The data are plotted on a logarithmic scale to facilitate the display of variations that occur over many orders of magnitude. The left-most set of points are the results for the no-detachment case, $\tilde{B} = 0$, where the steady state consists mostly of coalitions of size two and three, a quickly decreasing number of larger coalitions, and no unaffiliated agents. When $\tilde{B}$ is small and finite, the number of unaffiliated agents is small and largest coalitions dominate. $\tilde{r}_1$ grows linearly with $\tilde{B}$ (on a log-log scale) over several decades, until $\tilde{B} \approx 10$. At that point coalitions start to "evaporate" quickly, and as a result, the number of coalitions of larger size drops precipitously.

The global utility gain (per agent), calculated according to Eq. 1, is shown as a solid line in Figure 2, with the scale displayed on the right-hand side. The utility gain is largest for small finite $\tilde{B}$. Its value for $\tilde{B} = 10^{-6}$ is $G/N = 4.87\Delta p$ — a substantial increase over the no-detachment case value of $G/N = 2.00\Delta p$. The utility gain roughly follows the number of coalitions of maximum size: it decreases slowly as the detachment rate grows to $\tilde{B} \approx 10$, thereafter it drops quickly to zero. For large detachment rates, there is virtually no utility gain, as the system is composed mainly of unaffiliated agents. The large increase in the utility gain for small $\tilde{B}$ comes at a price, namely the time required to reach the steady state. While it takes $\tilde{t} \approx 10$ for solutions to reach the final state for $\tilde{B} = 10$, it takes $\tilde{t} \approx 10^9$ for the solutions to equilibrate for $\tilde{B} = 10^{-6}$. Note that there is a discontinuity at $\tilde{B} = 0$: the steady state solutions are qualitatively different for $\tilde{B} \to 0$ than at $\tilde{B} = 0$, because in the latter case, the system gets trapped in a non-equilibrium state before it gets a chance to reach a true equilibrium.

We can obtain analytic expressions for the steady state densities in terms of $\tilde{r}_1$ by setting the left-hand side of Eqs. 2–4 to zero. We find that at late times the densities obey a simple relationship:

$$\tilde{r}_n = \tilde{B}^{-(n-1)}\tilde{r}_1^n\,. \quad (5)$$

By studying the behavior of solutions for different values of m, we empirically obtain a scaling law for the steady state monomer density,

$$\tilde{r}_1 \propto \tilde{B}^{\frac{m-2}{m-1}}\,. \quad (6)$$

This result is valid in the parameter range that we are interested in, namely where the utility gain is large and slowly varying. Equations Eq. 5 and Eq. 6, together, allow us to

[4] Numerical integration was carried out using Mathematica 4.0.

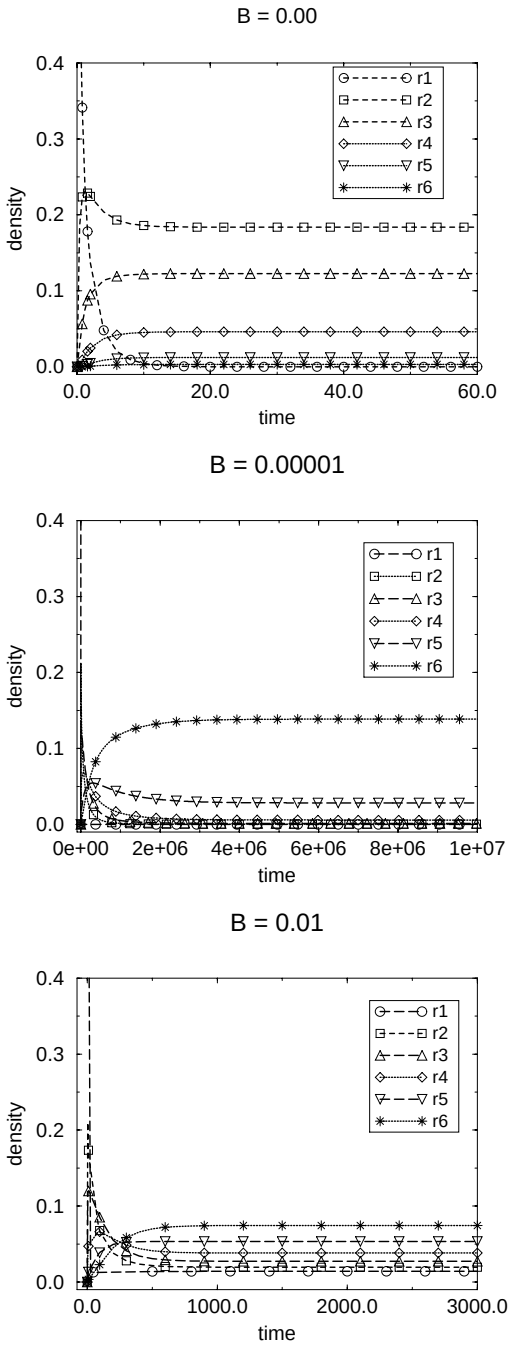

Figure 1. Time evolution of coalition densities for $m = 6$ and three detachment rates: $\tilde{B} = 0$, $\tilde{B} = 10^{-5}$, and $\tilde{B} = 10^{-2}$.

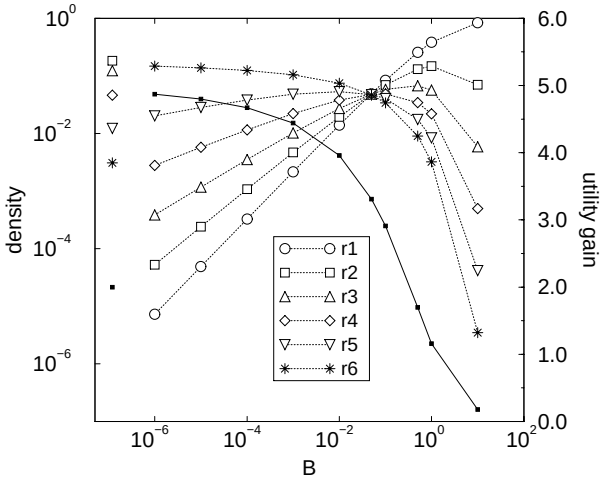

Figure 2. Steady state distribution of coalition densities and the global utility gain *vs.* the dimensionless detachment rate. Solid line is the global utility gain per agent.

predict how the steady state density of coalitions of any size changes as the detachment rate, $\tilde{B}$, and the maximum coalition size, m, are changed. In particular, as m becomes large, the exponent of $\tilde{B}$ approaches 1. In this case $\tilde{r}_n \propto \tilde{B}$, that is the number of coalitions of every size grows linearly with $\tilde{B}$.

5.3. Lessons learned

We conclude that when the agents are not allowed to leave coalitions, there is some utility gain in the steady state, reflecting the presence of small coalitions. However, introducing even a very small detachment rate to the basic coalition formation process allows the system to reach an equilibrium steady state. The increase in the global utility gain is more than twice that for the no-detachment case. The price for higher utility gain is that the time required to reach the steady state solution grows very large as $\tilde{B}$ becomes small. The system is composed mainly of coalitions of the largest size. As the relative strength of the detachment rate increases, the utility decreases, because coalitions become smaller, and the number of unaffiliated agents grows until there is virtually no utility gain. However, utility gain remains large and decreases slowly over many orders of magnitude of $\tilde{B}$. The agents designer has much leeway in choosing parameter values that result in a substantial global benefit, while not requiring to wait for too long a time for this benefit to be achieved. The agent designer can also pre-

dict the final distribution of coalitions, even for very large systems.

6. Related work

Cooperation among multiple agents was discussed previously in several publications. For example, Ephrati *et al.* [5] suggest a coordination mechanism for agents in the tile-world, where minimal communication is required, though, the suggested mechanism does not provide means for forming groups or coalitions. Moreover, that mechanism was implemented and tested only with very small number of agents, and its scaling up is yet to be checked. Even if it scales up, no guarantees or tools for prediction and analysis of large-scale behavior are provided. The lack of such tools is also apparent in work by Sen *et al.* [11]. However, in difference from Ephrati, Sen provides means for the formation of groups of agents. Nevertheless, the groups his mechanism supports are usually pairs of agents. Yet, an important difference between Sen's work and ours is that, while we address the more complex case of self-interested agents, Sen addresses the case of cooperative agents. Note that in another work Sen addresses the case of self-interested agents [12], however there no coalition formation is discussed. Another work, by Shehory *et al.* [15] suggests a low communication complexity coordination mechanism for large scale MAS. It also suggests a model similar to a physical model, and that the tools provided by physics can be used to analyze and predict large-scale behavior. Yet, that work does not support the formation of groups, and refers to cooperative agents rather than self-interested ones as we do in this research.

Cooperation among self-interested agents was addressed using other approaches. For example, Axelrod *et al.* [1] show how cooperative behavior can arise among selfish autonomous agents. They use game dynamics to *simulate* interactions between two agents, in which each agent has to make a decision, each with a different payoff to the agent. The agent's decision depends on choices made by other agents. Some strategies lead to stable cooperation due to mutual payoff increase. Yet, coalition formation is not addressed, and results are arrived at via simulations and not analytically, in contrast to our work. Others [10, 6] have applied game dynamics formalism to distributed control, where many agents adjust their strategies (a decision to compete or to cooperate) to increase their share of a finite resource. Though, the focus of that work was on adaptation in a distributed system, i.e., how a group of agents can learn to cooperate to achieve a common goal without central control. Some of the presented game dynamics systems (see, for example, Huberman and Hogg's work on computational ecologies [7, 6]) are amenable to mathematical analysis, though results about the stability of the system were usually achieved via simulation. In studying the global dynamics of a system of locally interacting agents, that research is similar to ours. The aim of mathematical analysis is to demonstrate the existence of evolutionary stable strategies that drive the system to the steady optimal solution. However, there are differences between that work and ours. There, the cooperation mechanisms are usually only abstractly specified, as an increase in the payoff matrix. Yet, an agent's behavior is more complex—it can change in response to the actions of other agents—and as a result, the global stability is often sacrificed to adaptability [7]. We do not address adaptability and strategic diversity; however, though we have not carried out a detailed stability analysis, our work suggests that the steady state solutions are stable.

7. Future direction

There are many issues in our work that remain unexplored. Our model can be extended to provide a wider range of solutions by incorporating a varying likelihood of agents leaving coalitions (e.g., agents may be less likely to leave large coalitions since the probability of finding a better coalition to join may be small). We can increase the flexibility of coalition-vendor relations (e.g., allow a coalition to approach multiple vendors). It should be of interest to examine stability of the steady state solutions to perturbations (e.g., agent failure): is stability maintained? This question is similar to the effect of noise on the system, which we have also not explored. Another promising direction that can follow from our approach is the interplay between agent complexity and system complexity. In other words, how much adaptability and stability can one achieve with simple agents that have a fixed strategy compared to the more complex agents that can change their strategies. It would be interesting to apply rigorous mathematical analysis to these problems.

8. Acknowledgements

KL wishes to acknowledge R. S. Ross for illuminating discussions concerning physical phenomena that are analogous to coalition formation, and A. Bugacov for helping with numeric integration.

References

[1] R. Axelrod and W. D. Hamilton. The Evolution of Cooperation. *Science*, 211:1390–1396, 1981.

[2] A.-L. Barabasi and H.E. Stanley. *Fractal Concepts in Surface Growth.* Cambridge University Press, Cambridge, 1995.

[3] E. Bonabeau, M. Dorigo and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems.* Oxford University Press, New York, 1999.

[4] K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proceeding of IJCAI-97*, pages 578–583, Nagoya, Japan, 1997.

[5] E. Ephrati, M. Pollack, and S. Ur. Deriving multi-agent coordination through filtering strategies. In *IJCAI95*, pages 679–685, Montreal, Quebec, Canada, 1995.

[6] O. Guenter, T. Hogg and B. A. Huberman. Learning in Multiagent Control of Smart Matter. *AAAI Workshop on Multiagent Learning*, 1997.

[7] B. A. Huberman and T. Hogg. The Behavior of Computational Ecologies. in *The Ecology of Computation*, B. A. Huberman Ed., Elsevier, North-Holland,1988.

[8] N. R. Jennings and J. R. Campos. Towards a Social Level Characterization of Socially Responsible Agents. *IEEE Proceedings on Software Engineering*, 144(1):11–25, 1997.

[9] T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst-case guarantees. *Artificial Intelligence*, 1999.

[10] A. Schaerf, Y. Shoham and M. Tennenholtz. Adaptive Load Balancing: A Study in Multi-Agent Learning. *Journal of Artificial Intelligence Research*, 2:475–500, 1995.

[11] S. Sen and M. Sekaran and J. Hale. Learning to Coordinate without sharing information. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi–Agent Systems*. MIT Press, 1995.

[12] S. Sen. Reciprocity: A foundational principle for promoting cooperative behavior among self-interested agents. In Victor Lesser, editor, *Proceedings of the First International Conference on Multi–Agent Systems*. MIT Press, 1995.

[13] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.

[14] O. Shehory and S. Kraus. Feasible formation of coalitions among autonomous agents in non-super-additive environments. *Computational Intelligence*, 15(3):218–251, 1999.

[15] O. Shehory, S. Kraus, and O. Yadgar. Emergent cooperative goal satisfaction in large-scale automated-agent systems. *Artificial Intelligence*, 1999.

[16] O. Shehory. A scalable agent location mechanism. In N.R. Jennings and Y. Lespérance, eds., *Intelligent Agents VI—Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, 2000.

[17] D. Walgraef. *Spatio-Temporal Pattern Formation, with Examples from Physics, Chemistry, and Materials Science.* Springer, New York, 1997.