

Behaviosites: A Novel Paradigm for Affecting Distributed Behavior

From a Healthy Society to a Wealthy Society

Amit Shabtay, Zinovi Rabinovich, and Jeffrey S. Rosenschein

School of Engineering and Computer Science
The Hebrew University, Jerusalem, Israel
ashabtay@gmail.com
{nomad, jeff}@cs.huji.ac.il

Abstract. In this paper we present the Behaviosite paradigm, a new approach to affecting the behavior of distributed agents in a multiagent system, which is inspired by biological parasites with behavior manipulation properties. Behaviosites are special kinds of agents that “infect” a system composed of agents operating in that environment. The behaviosites facilitate behavioral changes in agents to achieve altered, potentially improved, performance of the overall system. Behaviosites need to be designed so that they are intimately familiar with the internal workings of the environment and of the agents operating within it, and behaviosites apply this knowledge for their manipulation, using various infection and manipulation strategies.

To demonstrate and test this paradigm, we implemented a version of the El Farol problem, where agents want to go to a bar of limited capacity, and cannot use communication to coordinate their activity. Several solutions to this problem exist, but most yield near-zero utility for the agents. We added behaviosites to the El Farol problem, which manipulate the decision making process of some of the agents by making them believe that bar capacity is lower than it really is. We show that behaviosites overcome the learning ability of the agents, and increase social utility and social fairness significantly, with little actual damage to the overall system, and none to the agents.

1 Introduction

Biology-based technology (biomimetics) often provides insight into ways of improving technology based on biological metaphors [20]; nature and evolution have worked hard at finding solutions to many problems that are also present in artificial environments. Parasites have been examined in the biomimetic context, in particular for their special abilities (such as safe navigation inside the human body, needed for microendoscopy [12]).

However, one of the most interesting abilities of biological parasites has received little attention in technological contexts — their ability to manipulate and alter their host’s behavior. Usually, parasites alter host behavior so as to transmit themselves to the next host, as in the case of rabies (by generating aggressive host behavior). In rare cases, biological parasites can actually benefit the host or the host’s society. The word

“parasite” [13] derives from the Greek word *parasitos*, which means “beside the grain”, and originally had a positive connotation; a parasite was a fellow guest who ate beside you at the dinner table. Only later did it receive the meaning of someone eating at the expense of another.

In the computer science literature, the term “parasite” is used in three different contexts, none dealing with behavior manipulation: parasites in simulations of evolution, parasites as a driving force in genetic algorithms, and parasites as a common name for malware (computer viruses, trojan horses, etc.).

1.1 The Behaviosite Paradigm

In this paper, we present a novel paradigm that employs a special kind of agent (called a *behaviosite*) that manipulates the behavior of other agents so as to achieve altered, possibly improved, performance of the entire system. The behaviosite (by definition) is not itself necessary for the normal conduct of the system; thus, it is termed a kind of “parasite”.

Within the field of MultiAgent Systems (MAS), the behaviosite is closely related to the idea of adjustable autonomy (AA) [16], although it approaches issues of autonomy in a novel way. AA is about agents varying the level of their own autonomy based upon the situation (thus leveraging three alternative modes of operation, fully autonomous, semi-autonomous, and teleoperated). AA often deals with agents transferring control to human users [19] in cooperative settings. In the context of the Behaviosite paradigm, an agent transfers some of its autonomy to the behaviosite (usually unwittingly, and in any case the agent is better described as *ceding autonomy*). The behaviosite manipulates the agent’s behavior to achieve altered performance of the system. In this way, the behaviosite may create improved performance of the overall system (as we will demonstrate in the parasitized El Farol problem [1]), or it may facilitate new behaviors of the system.

Introducing behaviosites into a system may be either planned as part of the overall system design, or added (after the fact) to an already working system. When planned as part of the overall system, it is possible to choose between applying *internal behaviosites* (design hooks inside agents, so that behaviosites can attach to the agent and manipulate them) or *external behaviosites* (which manipulate only the input/output of the agent, vis-à-vis the environment). When applied to an already working system, external behaviosites are usually the only option. Both internal and external behaviosites are discussed in Section 3.

Internal behaviosites require cooperation at the level of agent designer(s), so that the appropriate hooks are in place for them to run (and alter agent behavior). Centrally-designed MAS systems, of course, may include behaviosite hooks because overall system performance can be improved. However, even in systems of self-interested, heterogeneous agents, inclusion of a behaviosite hook may be mandated as a requirement for operation within a given environment — thus giving the environment designer an additional tool for improving overall system behavior.

1.2 A Distributed Approach to Altered Behavior

One of the major strengths of the Behaviosite paradigm is that it is a distributed solution to issues raised in a distributed environment. Consider, for example, the El Farol problem [1] as an example of such a distributed environment. All agents want to go to a bar called El Farol, but it has limited (comfortable) capacity. If there are more attendees at the bar than that capacity, all attendees suffer from the crowdedness. With no option for communication or collusion, an agent must learn the behavior of other agents *en masse*, in order to reach a decision: go to the bar, or stay at home.

With very simple agent behavior, the system reaches an equilibrium around the given capacity. Unfortunately, personal and social utilities are extremely low, since about half the time the bar is overcrowded. To this setting, we introduce behaviosites, and the problem becomes the “parasitized El Farol” problem. We show that with simple infection and manipulation strategies of these special agents, it is possible to dramatically increase personal and social utilities, and even improve social fairness.

Although the El Farol problem is artificial, it has implications for many fields, including game theory, congestion problems in networking, logic, and economics [8]. In most, if not all, of these areas, the Behaviosite paradigm is applicable. There are, in addition, other possible applications for behaviosites, as will be discussed in Section 5.

The rest of the paper is organized as follows. In Section 2 we present a formalization of the Behaviosite paradigm. In Section 3, we discuss the El Farol problem, and how behaviosites could be employed in its solution. In Section 4 we briefly overview the concept of “parasite” in biology and computer science. We conclude in Section 5 with an overall discussion of our approach and of future work.

2 Behaviosite Formalization

2.1 Overall Structure

The Behaviosite paradigm is composed of three parts: the environment, agents (also referred to as hosts), and behaviosites. Environment, agents, and behaviosites will be referred to collectively as “the system” (as in ecosystem). See Figure 1 for an illustration of the environment/agent relationship.




Fig. 1. The environment/agent relationship

Environment: encapsulates all of the external factors, conditions, and influences that affect a community of agents and behaviosites. For example, during the course of a

program run, it stores the runtime state of the system, and it is discarded when the program ends. In some cases, the environment can be a degenerate instance, as in the case where agents only influence one another. In other cases, it may have an important role in the conduct of the system, as in the parasitized El Farol problem.

Agent society: A society of agents is a system composed of multiple, interacting, possibly cooperating agents (see [21]). As there are many definitions of agency, we use the one suggested by Franklin and Graesser [6]: *An autonomous agent is a system situated within and as part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect [sic] what it senses in the future.*

Behaviosite: Most basically, a specialized type of agent. A behaviosite is an additional property/information added to a system with a society of agents, and is not (and should not be) a property of the agent or the environment. The behaviosite is not required for the system, but should be beneficial in some sense; too many behaviosites can degrade some aspects of the system. Below, we present specific definitions regarding behaviosite elements. These definitions are further clarified in Section 3, which presents (in the context of the El Farol problem) what behaviosites are, and their advantages within a system.

2.2 Behaviosite Specifics

Required Traits *Benefiting the system:* Behaviosites come with costs to the system, as will be discussed below. The design and use of behaviosites is unnecessary, unless they are beneficial to the system in some respect. Behaviosites may influence a system in two desired ways: they may increase social utility, or they may create new features in a working system. Both are accomplished by altering the behavior of some of the agents. Note that in the behaviosite paradigm, unlike with parasites in nature, the utility of the behaviosite itself is not important when declaring them successful.

System knowledge: A behaviosite must be designed with deep understanding of how the system works: agent-agent interactions, agent-environment interactions, and also internal workings of the agents in some cases. Such knowledge is essential for the success of the behaviosites in improving the system. One type of system in which behaviosites can be beneficial is a system that works in a suboptimal equilibrium. When building a system, usually suboptimality is caused by the need to make it robust against failure. However, there are examples in which sub-optimality results from the inherent structure of the system. The internet is one example of a system in suboptimal equilibrium, with regard to the congestion problem caused by packet routing. The El Farol problem presents another, similar example.

Not a property of the agent nor of the environment: Behaviosites exist in the middle ground between agents and the environment, and should not be a property of either of them.

Optional Traits *Hidden or apparent infection:* We usually will not want agents to know who is infected by the behaviosite. Such knowledge may help individual agents exploit the situation, harnessing its own behaviosite or some other parasitized agent to

its own needs and thereby damaging the designed mechanism of environment-agents-behaviorsites. However, there are some settings in which such knowledge can benefit the system (as also occurs in nature), for example, the elimination of ill-functioning agents (a process known as “apoptosis” in cells).

Finding the host: In many cases in nature, finding the host is essential for the parasite. In the Behaviosite paradigm, it may also be an issue. A behaviosite designer may endow the environment with the responsibility of infecting agents in the system using some strategy (like in the parasitized El Farol problem), or it may leave the task of infection to the behaviosites themselves, thus making the behaviosites more autonomous. In certain contexts, hosts may also be equipped with defense mechanisms against infection.

Behaviosite communication: Behaviosites may communicate with one another within a host (the host may be infected with more than one behaviosite), or across hosts. The latter enables formation of a network of parasitized hosts, which may act in some kind of parasite-induced coalition. This may enable the behaviosites to be a catalyst for the creation of norms or social laws.

2.3 Placing the Behaviosite

To alter host behavior, the behaviosite designer must decide where to place the behaviosite in the flow of the system, as can be seen in Figure 2.




Fig. 2. (a) marks where an external behaviosite can alter the host’s behavior, by altering its input/output. (b) marks where an internal behaviosite can alter the host’s behavior, by altering internal data or by partial or full replacement of behavior modules.

In general, behaviosites can be divided into two main groups, external behaviosites, and internal behaviosites. External behaviosites can alter the input or output of the agent vis-à-vis the environment (Figure 2a). In this way, the host designer(s) need not know of the possible existence of a behaviosite now or in the future. Moreover, the behaviosite designer may introduce them into an existing system.¹

¹ Of course, these external behaviosites affect inter-agent communication only if it, too, travels via the environment.

On the other hand, internal behaviorsites usually require the system to be designed “plug and play” for the behaviorsite (Figure 2b). The host designer(s) will leave a (sufficiently protected) hook, into which behaviorsites can be plugged, possibly created by another designer. Allowing the behaviorsite internal access may seem dangerous, but it holds many advantages.

Internal behaviorsites can have two main manipulation methods: changing the host’s internal data (Figure 2b, interaction between see-behavior-act modules) or replacing some or all of its behavioral modules (Figure 2b, behavior module). In the first form of manipulation, only relevant data will be changed, and in some scenarios, it can actually solve the problem of dealing with noise resulting from the transduction of input from the environment.²

The hook for the second manipulation method is very easy to program, by using the behavioral design patterns suggested by the “gang of four” [7]. Introducing a new behaviorsite to a well-designed system may at times be much like programming a new, special behavior for an agent (though it should not be a regular property of the host). In specific systems, if needed, security measures will have to be taken to protect the host from malicious parasites, which could exploit the existence of the hook. Such security measures are abundant; one easy-to-implement example is public/private key encryption (further consideration of these security issues is beyond the scope of this paper).

2.4 Cost of Manipulation

In different scenarios, behaviorsites have different types of costs. The first cost that should be considered is the cost to the system designer, for there is the immediate issue of designing behaviorsites and testing them. Analyzing system behavior can be very complex, and analysis of a system in which behaviorsites are integrated is even more complex. This is because their impact on the system is usually hard to predict without extensive testing. Other costs of behaviorsites are system specific, such as balancing behaviorsite complexity and number with the benefit they give to the system, cost of possibly increased variation in the society, oscillations, run-time costs, and so forth. Despite all of the above, there are many scenarios in which behaviorsites prove themselves very useful.

3 Parasitized El Farol Problem

3.1 The El Farol Problem and Related Work

The El Farol problem (or the Santa Fe bar problem) was introduced by Brian Arthur [1] as a toy problem in economics for using inductive reasoning and bounded rationality. In this problem, N (e.g., 100) agents decide independently each week whether to go to the El Farol bar or not. Comfortable capacity is limited, and the evening is enjoyable

² This is somewhat reminiscent of Brooks’ subsumption architecture [3], in which various levels can be designed to suppress input/output on other levels.

only if the bar is not overcrowded (specifically, fewer than 60 agents out of the possible 100 attend).

In the original problem, no communication or collusion is possible, and the only information available to the agents is the number of attendees in the past. An agent will go if it predicts less than 60 agents will go, and will stay home otherwise. Arthur suggested that bounded rationality, together with learning, can yield solutions to problems of resource allocation in decentralized environments, using the El Farol problem as an example [9]. At this point, the problem in the research literature diverged into two branches, differentiated by their utility functions. The first branch was as in Arthur’s paper:

$$Util(ag[i]) = \begin{cases} x & \text{attended and undercrowded} \\ 0 & \text{did not attend} \\ -y & \text{attended and overcrowded} \end{cases}$$

The second branch used another kind of utility function:

$$Util(ag[i]) = \begin{cases} x & \text{part of minority group} \\ -y & \text{not part of minority group} \end{cases}$$

In this paper, we use the first utility function, where $x = y = 0.5$.³ Arthur showed [1] that if each agent uses a set of personal basic deterministic strategies (such as going if more than 55 agents went last time), combined with a simple learning algorithm, then the system converges to the capacity after some initial learning time. Moreover, membership kept changing, and some degree of fairness was maintained. Arthur described the emergent ecology as almost organic in nature. In his work, Edmonds [5] took this analogy a step further by using genetic algorithms for learning, and allowing communication. Edmonds showed that these were sufficient for the development of different, interesting social roles.

However, both solutions produced systems that fluctuated chaotically around the capacity, resulting in low personal and social utility. High variation in the system was one reason for suboptimality. The second was that at many times the bar was overcrowded (above capacity), thus giving negative utility to attending agents.

Bell [2] tried to deal with the overcrowdedness problem by a simple adaptive strategy used by all agents, which led to an outcome close to the socially optimal attendance. Bell’s solution holds two pitfalls. The first is that the same strategy must be used by all agents, perhaps too strong a requirement in a distributed system with different designers. The second is that in his simulations most of the attendees were “regulars” who came all of the rounds, thus making for social injustice with the “casuals”.

Greenwald [8] also tackled this suboptimality problem, but using a different method. Agents that attended El Farol were required to “pay” some of their utility, and this payoff was distributed among all agents that did not attend. Not surprisingly, the optimal payoff for optimal social utility converged to a 40% attendance fee, for the scenario of 100 agents and bar capacity of 60. Again, this solution may also be problematic in many distributed systems, since it requires a central “utility distributor” that can reach all agents.

³ For a review of the minority problem, see [14].

The learning algorithm used in the parasitized El Farol simulation was the additive updating learning algorithm, introduced in [9]. Basically, each agent has a pool of simple, personal deterministic strategies. Each such behavior has a weight, initially distributed uniformly. The weights of all these basic strategies are updated in an additive manner once a round, according to the number of attendees in that round, above or below capacity (0.5 or -0.5 utility to attendees, respectively, or zero if it chooses not to attend).

3.2 Using Behaviosites in the Parasitized El Farol Problem

The El Farol bar problem presents a distributed system, with suboptimal social and personal utility. The main idea of the parasitized El Farol problem is to increase social utility with as few side effects as possible, using behaviosites. The main problem was that agents learn and adapt themselves to new situations. If the behaviosites are not carefully crafted, then either their effect will soon vanish, and the system will return to the equilibrium around the capacity with suboptimal social utility, or in each round behaviosites will need to make a stronger impact to achieve the same effect.




Fig. 3. An illustration of the environment/agent/behaviosite relationship

The system is composed of environment, agents (hosts) and behaviosites (Figure 3). Each agent has an internal behaviosite field that is always occupied, usually with the null behaviosite (see [7]) which has no effect over its host’s behavior. Each round, the environment infects some of the agents according to an infection strategy (as will be discussed below), and infection lasts only one round. Agents are asked to provide a decision (whether they will go to the bar or not). Then, the agent gives its behaviosite private information before making the decision. If the behaviosite is not the null behaviosite, it alters the relevant host’s data (the behaviosite is placed between “see” and “behavior” of Figure 2). In this way, the agent comes to a decision not using the data received from the environment, but rather by using manipulated data.

Environment infection strategies: In this simulation, the environment had three strategies of infection, which were not mixed. In the first strategy, all agents were candidates for infection (*infect all*). In the second, only attending agents in the given round were candidates for infection (*infect attending*). In the third, candidates for infection were all attending agents at the given round, but only when the bar was overcrowded (*infect overcrowded*). In each strategy, only a percentage of the candidates for infection actu-

ally got infected (responsibility of the behavior site design), depending on the infection rate (0%-100%).⁴

Behavior site manipulation strategies: The chosen manipulation strategy was quite simple. The behavior site replaced the parasitized agents' belief regarding the current capacity with a lower one, common to all behavior sites. Since capacity information was kept as the private history of the agent, this decrease was also considered in future rounds by the agent. If a parasitized agent decided to attend the bar and the number of attending agents was higher than the capacity the agent currently believed, it suffered a utility decrease, even if the bar was actually undercrowded. This was intended to enforce a stricter approach; agents were affected by the world according to their personal beliefs, and not according to some global truth.

3.3 Parasitized El Farol Simulation Results

The parasitized El Farol problem was simulated for 2000 rounds, assuming 100 agents in the system. Three different capacities were tested: 50, 60, and 80, where the behavior site manipulation strategy was to decrease capacity for parasitized agents (50 decreased to 40, 60 also decreased to 40, and 80 decreased to 60). For each such capacity, three different environment infection strategies were tested, as mentioned above: *infect all*, *infect attending*, and *infect overcrowded*, giving a total of 9 different simulations. For each of the 9 simulations, a different percentage of the agents who were candidates for infection were actually infected, with infection rates ranging from 0% to 100% with jumps of 10% (total of 11 different simulations for each). Each of these 99 different simulations was repeated 50 times.

Mean Attendance As was mentioned above, agents' mean attendance in the El Farol problem converges to the capacity of the bar. In the parasitized El Farol problem, we would like to increase the social utility, without lowering by too much the mean attendance. From the bar owner's point of view, this is one of the most important parameters. From a social utility point of view, decreased mean attendance takes us further away from an optimal solution.

Figure 4 shows how different infection strategies affected the mean attendance. For all three different capacities, the *infect all* strategy decreased the mean attendance severely, as a function of the behavior sites' infection rate. For 100% infection rate of potential agents (in this case, all agents), the system converged to the capacity induced by the behavior sites. For example, for capacity of 60, infection rate of 100% of course resulted in convergence to 40, since the behavior sites' manipulation strategy is to make the host believe that the capacity is 40.

The *infect overcrowded* strategy resulted in a relatively low effect on the mean attendance, since behavior sites infect agents only in a most particular situation — only attendees, and only when the bar is overcrowded. The outcome of the *infect attending* strategy depended on the bar's capacity. For capacity of 80, the *infect all* and *infect attending* strategies yielded very similar results, since most of the agents are candidates

⁴ Another possible design is that the entire infection strategy is a property of the behavior site, but this was not implemented.




Fig. 4. Mean attendance for capacities of 50, 60, 80, respectively. Legend: :Infect All, —:Infect Attending, - - - :Infect Overcrowded

for infection. For the capacities of 50 and 60, *infect attending* strategy's effect resembles the *infect overcrowded* strategy, since it affects only a relatively small portion of the agent society.

Mean Utility The main objective in the parasitized El Farol problem simulation was to show that it is possible to increase social utility using behaviorsites. This can be seen clearly in Figure 5, which compares the effect of different capacities for each strategy. Another goal, with roughly the same importance, was not to cause too much harm to the system, namely social injustice, which will be discussed later.

infect all: Infecting all agents resulted in a similar graph shape for all three capacities, differing in the position of the peak, with impressive improvement of mean utility. For capacities of 60 and 80 the improvement was 7–9 times the minimal received in a system without behaviorsites, and relatively close to the optimal (0.12 utility out of possible 0.3 for capacity of 60, and 0.23 utility out of possible 0.4 for capacity of 80). For

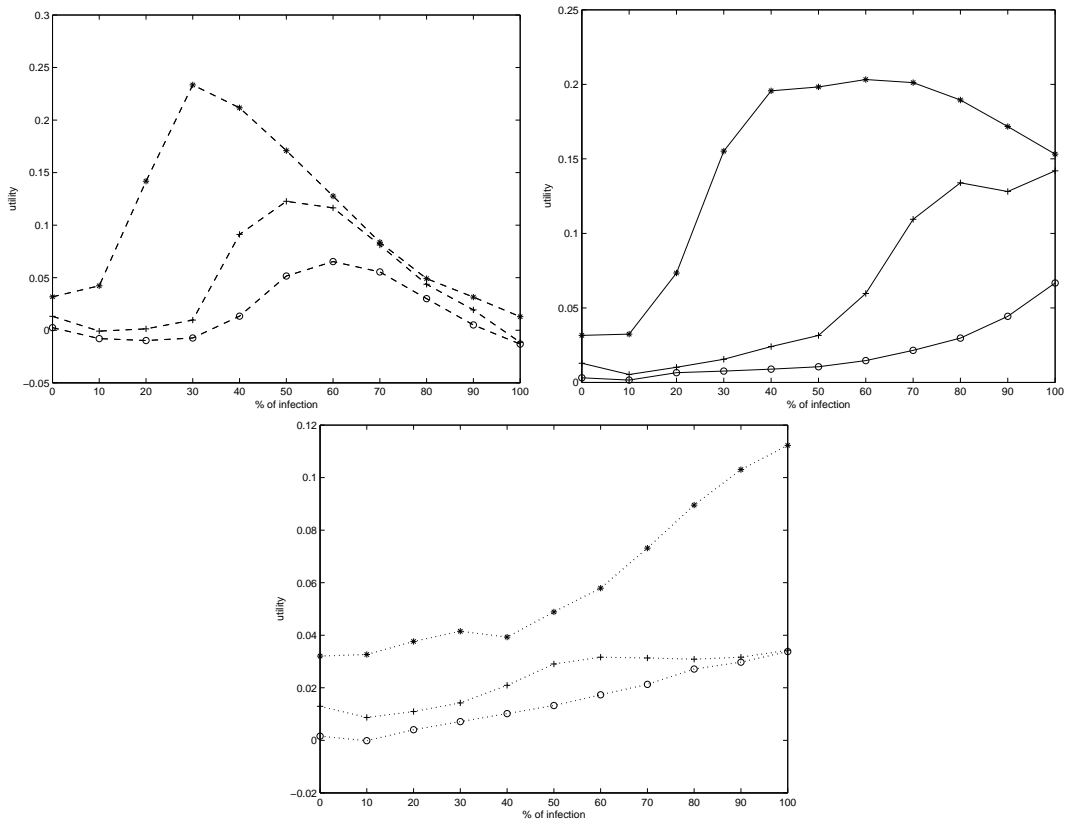


Fig. 5. Mean utility for Infect All, Infect Attending, Infect Overcrowded strategies, respectively. Capacity legend: o-o: 50, ++:60, *-*:80

capacity of 50, improvement was 27 times the minimal, though still near zero (0.065 out of possible 0.25).

infect overcrowded: For all three different capacities, the mean utility increase was moderate, with respect to other strategies. For capacities of 60 and 80, the maximal mean utility was about 2.5–3.5 times the minimal mean utility, received by a system with no behaviorsites. For the capacity of 50, the maximal mean utility was 22 times better than the minimal mean utility received for a system with no behaviorsites, but still near zero (0.034). For capacity of 80, the maximal mean utility (0.112) was not significantly lower than the optimal possible (0.4), especially when compared to a system without behaviorsites (utility of 0.032).

infect attending: The outcome of this strategy depended heavily on the level of the used capacity. For high capacities (specifically 80), the strategy's behavior was very much like *infect all*, because infecting all attending agents is very much like infecting all agents. For low capacities (specifically 50), the strategy's outcome was very much like *infect overcrowded*, because it affected only a relatively small portion of the society, and

the actual effect of the behavior site was mild. In between (specifically 60), the behavior of this strategy combined features of both.




Fig. 6. System’s fairness for capacities of 50, 60, 80, respectively. Legend: - - - :Infect All, — :Infect Attending, :Infect Overcrowded

Forced Justice As was described earlier, two previous attempts to increase social utility resulted in two proposed solutions, both with undesirable features in certain settings. One feature was the need to charge utility from attending agents [8], which basically kept the social utility within the system. Another led to the creation of bullies [2] — regulars and casuals. In the original El Farol problem, the membership of the agents kept changing, thus keeping some level of social fairness. Social fairness is an important parameter when characterizing the effect of behavior sites on the system, since we do not wish to create regulars and casuals. This is especially true if agents are self-interested, and have no prior incentive to allow behavior sites to infect and affect them. Social fairness was calculated by the following formula:

$$1 - \frac{1}{\#trials} \sum_{t \in trials} \frac{Personal\ Attendance\ SD[t]}{Mean\ Attendance[t]}$$

Social fairness is captured by the mean of standard deviation, normalized according to the mean of a specific iteration. In Figure 6, we can see that for capacities of 50 and 60, *infect attending* and *infect overcrowded* strategies constantly increase the social fairness as a function of infection rate. For capacity of 80, they improved social fairness only a small amount, since already the vast majority of agents go to the bar. However, the *infect all* strategy increased social fairness as a function of infection rate until a certain maximum was reached, and then decreased again. For capacity of 80, it reached exactly the level of social fairness of capacity 60, as would be expected.




Fig. 7. Mean utility with and without harming behaviorsites for capacities of 50, 60, 80, respectively. Legend: - - - :Infect All, —:Infect Attending, :Infect Overcrowded. + is for behaviorsites that do not cause harm.

Cost of Manipulation In the specific design of this parasitized El Farol problem, behaviorsites caused actual damage to the received utility of agents. If an agent came to an undercrowded bar, but believed it was overcrowded, it would suffer utility loss that unparasitized agents would not — thus imposing a cost upon agents. From the Designer Perspective, behaviorsites in the parasitized El Farol problem have several types of costs

— run time, integration time of behaviorsites, and testing time of various infection and manipulation strategies.

Reflections about Design The behaviorsite in the parasitized El Farol problem presented above is an example of an internal behaviorsite, in a system designed to accommodate behaviorsites. However, it is also possible to create an external behaviorsite, one that the agents need know nothing about. The behaviorsite can reside in the interface between the environment and the agent (Figure 2a). Parasitized agents can receive manipulated information from the environment regarding the current capacity, or manipulated information of number of attendees in a specific round. The effect of such altered data should be the same as presented above, and social utility should, again, increase.

4 Related Work

4.1 Related Work in Biology

In nature, most common recorded behavior manipulations are done by parasites with a complex life cycle, as it ensures host-to-host transmission. For example, in the case of *Dicrocoelium* [17], parasitized ants' behavior (the upstream hosts) is modified so that they crawl up grass stalks, thus improving chances of being ingested by a passing cow (the downstream host). Although most parasite manipulations are harmful to the host, there are examples in which parasites are beneficial, either to the host or to the society.

Increasing the fitness of the hosts' society can usually be seen in the case of close kin, such as ants or bees. Bumblebees have successfully mastered the use of parasite-induced altered behavior to their own advantage [15]. A parasitized worker bee prolongs its life span by changing its behavior and staying in the cold of night at the field (the cold retarded parasitoid development). This way, the entire colony benefits from this infection, due to increased foraging.

4.2 Related Work in Computer Science

The term "parasite" is used in various contexts in the computer science literature. However, in each context its meaning is substantially different from its use in the Behaviorsite paradigm. Host-parasite paradigm in computer science previously existed in three main roles: for simulations of evolution such as Ray's Tierra computer program [18], as a driving force in genetic algorithms (usually in the context of predator-prey and co-evolution [10]), and as malware in the electronic world. None talked about using parasites to change the behavior of individual agents, to shift the equilibrium of the entire society.

In their work, Mamei, Roli and Zambonelli [11] presented an idea similar to that of the Behaviorsite Paradigm, within the context of cellular automata. They showed that via an external perturbation that dynamically influenced the state transition of several cellular automata, they can achieve a global effect within the system (i.e., control of macro-spatial structures, as they termed it). Large-scale spatial structures emerged as a result of those perturbations. The Behaviorsite Paradigm presented in this paper posits

a more highly developed description of the perturbation mechanism — special agents that infect and manipulate other agents. Mamei et al. maintain, as do we, that in the near future there will exist more systems characterized by distribution and autonomy, operating within highly dynamic environments. These systems will by their nature require decentralized control mechanisms, such as those presented in [11] and in this paper.

Computer Agents' Behavior Behaviosites affect behavior, but what constitutes “behavior”, both the agent’s and the system’s, is quite difficult to define; there are numerous ways of approaching the issue. Two of the most common are logic-based behavior and Brooks’ [3,4] emergent behavior and subsumption architecture. Behavior is considered as an aspect of the individual agent, of an agent in a society, and of the whole society. Much research explores the creation of agents with defined behavior, or analyzes an existing system’s behavior. For our purposes, we do not attempt to define exactly what behavior is, but rather use Brooks’ view, that behavior results from the agent’s interaction with the environment.

5 Discussion and Future Work

In this paper, we presented the Behaviosite paradigm in the spirit of biomimetics — parasites manipulating their hosts’ behavior. Although the term “parasite” is not new in computer science, it has been used in different contexts than the Behaviosite paradigm describes. We specified the behaviosite concept, which in essence is a special type of agent, which infects and manipulates other agents to achieve altered performance of the system. We described the parasitized El Farol problem, and showed that social and personal utilities and social fairness can be increased using behaviosites; in some cases, mean attendance deviates from the capacity by only a small amount.

The behaviosites used were internal behaviosites with a fixed, unchanging influence over the agents (though it is equally possible to use external behaviosites, or those with dynamic influence). The parasitized El Farol problem simulated self-interested agents. However, behaviosites could also be integrated into a cooperative society to form new norms or social laws, or to eliminate ill-functioning agents.

In our view, the major contribution of the Behaviosite paradigm is conceptual — thinking of a system composed of environment and agents as whole, something that can be manipulated using external forces (behaviosites) that employ local changes (affecting agents’ individual behaviors) to bring about altered behavior of the entire system. However, as users of this paradigm, we do not consider ourselves omnipotent — we just use existing capabilities of the (possibly self-interested) agents and environment. Hence this is not simply mechanism design, but rather augmenting the system, or parasitizing it. It also differs from existing research in adjustable autonomy, because of this holistic view, which AA is lacking.

After understanding what the Behaviosite paradigm is, one might wonder if the term “parasite” is appropriate. There are, indeed, significant differences. The behaviosite is not an autonomous agent in and of itself. Moreover, in nature the utility gain of behavior manipulation is of the parasite, and not the host; is it not an important part of being a parasite? However, we intend parasitic activity to be a metaphor, and need not

require an exact match between nature and behaviorsites. The key analogy is the fact that behaviorsite actions constitute behavior manipulation of the host.

Behaviorsites are not just a way of propagating false information within a system. The intimacy of agent-behaviorsite induces far more powerful effects. Lies can be disregarded or overcome by agents. However, when a host is parasitized, the behaviorsite is considered an almost integral part of the agent. Agents can doubt external information, but rely on their internal beliefs — as was shown in the parasitized El Farol problem, not acting on the basis of personal belief did not help “shake” the influence of the behaviorsite.

As with the agent metaphor itself, the Behaviorsite paradigm can be very appealing, but could also conceivably be overused. It is certainly the case that it is possible to use solutions other than behaviorsites — behavioral change can be entirely inside the agent (e.g., random coin flips inside agents in the El Farol problem, with strategies similar to the ones applied by the behaviorsites), or it can be totally a property of the environment. The analogy with agents is clear, for although sometimes an agent-like approach is inappropriate (and solutions can theoretically always be engineered in other ways), the key question is “what is the best solution?”. In some cases, the Behaviorsite paradigm presents conceptual advantages.

In the future, we would like to further strengthen the Behaviorsite paradigm by showing it is applicable and desirable in other scenarios. One example is using external behaviorsites in an already-built real-time environment, namely the internet, for dealing with the packet congestion problem. The solution should be similar to that of the parasitized El Farol problem with some modifications. Another application is the automatic generation of stories; it is a rapidly growing field, with high potential in the gamer community. Behaviorsites are an excellent solution for altering stories by changing some of the characters’ behavior and adding new, unpredictable system behavior in a distributed manner.

6 Acknowledgment

This work was partially supported by Israel Science Foundation grant #039-7582.

References

1. W. B. Arthur. Inductive reasoning and bounded rationality (the El Farol problem). *Amer. Econ. Review*, 84(406), 1994.
2. A. M. Bell and W. A. Sethares. The El Farol problem and the internet: Congestion and coordination failure. *Computing in Economics and Finance 1999* 812, Society for Computational Economics, March 1999.
3. R. A. Brooks. Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 569–595, Sydney, Australia, 1991.
4. R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
5. B. Edmunds. Gossip, sexual recombination and the El Farol Bar: Modelling the emergence of heterogeneity. *Journal of Artificial Societies and Social Simulation*, October 1999.

6. S. Franklin and A. Graesser, editors. *Is it an agent, or just a program?: A taxonomy for autonomous agents*, 1996. In Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages.
7. E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional Computing Series, 1994.
8. Grace Hopper Celebration of Women in Computing. *Fair and Efficient Solutions to the Santa Fe Bar Problem*. A. Greenwald, J. Farago and K. Hall, October 2002.
9. A. Greenwald, B. Mishra, and R. Parikh. The Santa Fe Bar problem revisited: Theoretical and practical implications. *Festival on Game Theory: Interactive Dynamics and Learning, SUNY Stony Brook*, July 1998.
10. W. D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. *Artificial Life II*, 1992.
11. M. Mamei, A. Roli, and F. Zambonelli. Emergence and control of macro-spatial structures in perturbed cellular automata, and implications for pervasive computing systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 35(3):337–348, May 2005.
12. A. Menciassi and P. Dario. Bio-inspired solutions for locomotion in the gastrointestinal tract: background and perspectives. *Phil. Trans. R. Soc. Lond. A*, 361(1811):2287–2298, October 2003.
13. J. H. Miller. The critic as host. *Critical Inquiry*, 3(3):439–447, 1977.
14. E. Moro. The minority game: an introductory guide. *Advances in Condensed Matter and Statistical Physics*, 2004.
15. C. B. Muller and P. Schmid-Hempel. Exploitation of cold temperature as defense against parasitoids in bumblebees. *Nature*, 363:65–67, 1993.
16. Matthias Nickles, Michael Rovatsos, and Gerhard Weiss, editors. *Post-proceedings of the First International Workshop on Computational Autonomy — Potential, Risks, Solutions*, volume 2969 of *Lecture Notes in Computer Sciences*. Springer-Verlag, 2003.
17. R. Poulin. *Evolutionary Ecology of Parasites — From individuals to communities*. Springer, 1997.
18. T. Ray. Tierra. <http://www.his.atr.jp/ray/tierra>, 1990.
19. P. Scerri, D. V. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 17:171–228, September 2002.
20. J. F. V. Vincent and D. L. Mann. Systematic technology transfer from biology to engineering. *Phil. Trans. R. Soc. Lond. A*, 360(1791):159–174, February 2002.
21. G. Weiss. *Multiagent systems: A modern introduction to distributed artificial intelligence*. MIT Press, 1999.