

# Swarm AI: A General-Purpose Swarm Intelligence Design Technique

Keywords: Swarm Intelligence, Intelligent Systems Design, Multiagent systems, Soccer, Emergence

## Abstract

This paper introduces Swarm AI, the first general framework for designing Swarm Intelligence approaches to problems. We outline principles of Swarm AI, discuss its connection to previous work, and analyze the advantages and disadvantages of this method. Finally, we describe a case study of applying Swarm AI to a new problem. Specifically, we discuss the design of a soccer team, using the principles of Swarm AI to create agents capable of playing soccer effectively.

## 1 Introduction

Swarm Intelligence (SI) is a growing new discipline that views intelligence as a function of social interactions between individuals. SI is based on the study of social insects, which individually are quite simple but have complex group behavior. Work in SI began when scientists noticed how intellectually primitive organisms like ants and bees work together in colonies to accomplish very complex tasks. Swarm Intelligence algorithms attempt to solve problems in the world of computer science by creating teams of simple agents that are guided by nature-influenced rules. [1] defines swarm intelligence as the “collective behavior that emerges from a group of social insects. [2], a recent book on Swarm Intelligence, defines it as a phenomenon that emerges from the behavior of a large number of rule-based particles. The former definition arises from its author’s work on the Ant System, while the latter definition is derived from that book’s concentration on Particle Swarm Optimization.

Unfortunately, one barrier to the spread of Swarm Intelligence applications is the lack of a common definition, as researchers in this field tend to bias their explanations of SI toward the specific work they are doing. Furthermore, no common design framework has been put forth to relate the different techniques labeled as Swarm Intelligence [3,4,5]. Certainly, one can read about previous approaches like Dorigo’s Ant System, which has been extended and improved greatly since its proposal 14 years ago [6]. However, there is no methodology for designing a new SI approach for problems that SI has never been applied to.

It is our intent in this paper to propose a new problem-solving technique based on Swarm Intelligence, which we believe will fill this gap. Our technique can be used to

create new SI approaches to problems, providing a general framework for their design. We shall refer to this technique as Swarm Artificial Intelligence (Swarm AI) and will define it as the design of a multiagent system of simple communicating agents that cooperate to solve problems.

Thus, this paper is an attempt to demystify Swarm Intelligence by presenting a technique that inherits the advantages of SI while casting off the vague terminology that makes SI so hard to apply to new problems. First, we will discuss the defining principles of Swarm AI and how this new framework is different from other SI research. Second, we will list the advantages and disadvantages of Swarm AI. Finally, we will present a case study we conducted in applying Swarm AI to the problem of controlling a team of soccer players.

## 2 Swarm AI Principles

According to [7], the multiagent systems discipline examines social processes in groups of agents where intelligence is a property of a system of agents. By viewing social insect societies as multiagent systems and the individual insects as agents, we can attempt to formalize the rules that govern those societies and apply them toward the development of solutions to engineering problems that may or may not resemble the problems these natural systems face. To accomplish this objective, through the analysis of previous work we have developed the following three principles of Swarm AI. These are by no means comprehensive because we still have much to learn about Swarm AI, but they establish the main ideas behind this new technique.

### 2.1 Principle 1: Create a system of agents that work individually on a common problem

The first principle of Swarm AI is that a multiagent system is used to solve a particular problem. This problem is separated into parts that are given to the autonomous agents of the system. By autonomous agents, we mean agents that cannot be directly controlled by any other agent or central component. The designer of the system must partition the problem, so the agents themselves are not responsible for task distribution. Often, the nature of the problem suggests a logical way to divide the problem, which the designer may take advantage of. The individual

agents work on completing the sub-problems assigned to them, contributing to the solving of the global problem.

## **2.2 Principle 2: Agents are simple, fast, and have a limited perspective**

In Swarm AI, each agent is given a limited or local perspective of its environment. There is no explicit limit to what kind of sensory information an agent can receive, but it is important to keep the amount of information a Swarm AI agent processes small. The less information a Swarm AI agent receives, the simpler it is to design that agent and the faster this agent will carry out the processing of that information to reach a decision. We want an agent to maximize its speed in performing simple tasks without worrying about the big picture. Thus, it is often useful to discard all the information that is not directly related to the task a Swarm AI agent has been assigned. In addition, agents do not use symbolic representation and rely on reactive or utility-based behaviors assisted by fast heuristics to make decisions. These agents may not think ahead too much or use heuristics that take a significant amount of time to run, since this would take away from the speed of a Swarm AI agent. [8] provides a more detailed discussion of the advantages of using simple agents in multiagent systems.

## **2.3 Principle 3: Indirect Simple Inter-Agent Communication**

Based on the first two principles alone, it may seem that Swarm AI agents are rather limited in their use because their focus on local data may cause them to neglect potentially useful “big picture” information. Fortunately, the third principle of Swarm AI provides a way for the system as a whole to both perceive and act globally. The third and final requirement of a Swarm AI system is that its agents have a simple and indirect method of communicating with each other. The method has to be simple since we know from Principle 2 that the Swarm AI agents are themselves simple and will not be able to decipher complex messages.

More specifically, Swarm AI communication can be performed in one of two ways. First, it may be done through the environment in which the agents exist. This method is called ‘stigmergy’ and occurs when one agent alters the environment such that another agent may notice the change and potentially alter its own behavior as a result [9]. From this, we can say that the first agent indirectly communicated with the second by affecting the environment in which they coexist. The second method of communication that can be exhibited by Swarm AI agents is changing one’s own state such as location or velocity to make other agents who observe this change in state act differently.

Direct communication between agents is not permitted in Swarm AI because it complicates agent design, requiring agents that have message sending and receiving capability. Also, direct communication uses more space and time resources than the other two methods, which makes it undesirable for problem solving if it can be avoided. [8] provides a philosophical argument based on natural systems for why it is beneficial to keep agent interactions local by disallowing direct communication between the system’s agents.

## **2.4 Swarm AI: The Main Idea**

At this point, it is important to briefly describe Swarm AI once again, in light of the principles introduced above. Swarm AI is a multiagent system design technique. It works by creating a new SI approach without forcing the designer to model the behavior of some particular social insect. The main idea is to create a system of simple local agents that would individually work on parts of the problem the designer is trying to solve. As they work, these agents must have a simple indirect method of communicating with each other, which will allow them to exchange information and coordinate efforts globally. The agents in a Swarm AI system must operate quickly on small amounts of data and should not use symbolic reasoning which would slow them down, violating the second principle of Swarm AI. In creating these agents, the designer must use some domain knowledge of the problem being solved. Once a multiagent system of such agents is created, it will have a high level behavior that should be successful in solving the global problem.

## **3 Swarm AI and Previous Work**

Swarm AI is a general design method that is meant to solve various computer science and AI problems, regardless of whether they resemble social insect environments. Unlike much of SI work, Swarm AI is not concerned with modeling a particular natural phenomenon and then matching it to a specific kind of problem as was done in [4,10]. Rather, the three Swarm AI principles provide a foundation for building multiagent systems to solve problems without being required to study ants, bees, spiders, or other ‘inspirational’ social animals.

The local interactions of the Swarm AI agents indirectly contribute to a high level group behavior; this process is commonly referred to as “emergence” in SI literature. Therefore, Swarm AI can be seen as a particular method for engineering emergence such that it is useful for problem solving. This idea of deliberately creating emergent behavior to solve AI problems was originally proposed by Alexis Drogoul. He did this with what he called a “distributive approach” for various problems, including Pengo and chess [11,12]. Additionally, [2]

discusses Particle Swarm Optimization, a technique for using large groups of simple particles that exhibit flocking behavior to solve complex optimization problems. Swarm AI builds on the work of Drogoul, Kennedy, Dorigo, and other SI designers by organizing the common features of their approaches into a single framework.

Swarm AI is progressive because it deals with creating new SI techniques for problems. Since Dorigo's discovery of the Ant System, his approach has been improved and applied to various optimization domains from the classic Traveling Salesman and Vehicle Routing problems to scheduling and network routing [13]. While this work is important, Swarm AI is a tool for facilitating the creation of new systems that we believe can have as much impact as Dorigo's classic design.

As far as we know, we are the first to put forth an all-purpose design technique based on Swarm Intelligence. For example, while there has been much Swarm Intelligence research in Robotics, it almost exclusively deals with 3D physical domain problems like motion [2]. We must note that [3] attempts to build a "unifying model of Swarm Intelligence," but their co-field framework is restricted to approaches that concentrate on agent movement and not on other kinds of decision-making (see Section 5.1 for two examples of such problems in the soccer domain). Furthermore, their model does not incorporate communication by changing one's state as discussed in Section 2.3. Consequently, while their approach is insightful, it is ultimately limited in scope. Swarm AI is much more general because it is not restricted to working in a particular domain such as optimization or movement and allows for a variety of agent and communication method designs.

Though it may appear ambitious, Swarm AI is a natural step in the development of Swarm Intelligence as a discipline. Swarm AI is meant to be a generalization of all SI techniques, recognizing that they are based on simple communicating agents working together. At this point, we do not know much about the workings of such systems of agents because they have for a long time been dismissed as either Ad Hoc or incapable of solving complex problems. However, SI research has in the last decade demonstrated the power of such systems, and it is now time to give them a common name and begin understanding exactly how they function.

## 4 Advantages and Disadvantages of Swarm AI

The primary advantage of Swarm AI as a problem-solving technique is its high performance efficiency to design complexity ratio. The inherent simplicity of Swarm AI agents and their means of communication makes them easy to understand, which in turn, results in a fast design process of a Swarm AI system. At the same time, Swarm AI systems are very efficient because their agents are necessarily fast. These agents don't plan about the future

and don't use symbolic reasoning, both of which are computationally expensive. Furthermore, Swarm AI approaches have limited memory requirements because they feature simple reactive or utility-based agents that don't need to store a lot of previous information. In addition, our experiments (see Section 5) suggest that Swarm AI systems can be robust, maintaining good performance in varied and rapidly changing environments. Overall, Swarm AI provides all the benefits of Swarm Intelligence without the necessity of problem solvers studying life sciences phenomenon alongside biologists or zoologists, which is still common in Swarm Intelligence research as exemplified by [4].

The major disadvantage of Swarm AI is that it provides approximated solutions, so it is not applicable to problems where exact results are required. Also, the design of Swarm AI systems requires the tweaking of parameters to increase performance, much like in genetic algorithms. Furthermore, at this point the Swarm AI method does not assist the designer in distributing tasks among agents. However, we feel that future work in applying Swarm AI to new problems will alleviate this particular difficulty.

## 5 Case Study: Applying Swarm AI to a New Problem

### 5.1 The Soccer Problem and Our Swarm AI Approach

To test the Swarm AI methodology, we decided to apply it to the problem of creating a team of soccer players, which operates in a complex dynamic real-time strategic environment. We wanted to see if we could create a viable team, and the results of applying Swarm AI to this problem exceeded our expectations as we obtained a team that can perform sophisticated multi-player soccer behaviors in addition to the basic ones. Also, using the Swarm AI principles we developed an interesting new SI approach we call Swarm Painting for handling decisions that do not deal with movement.

First, we created an abstract 2D soccer simulation where two teams can play five on five games, using simplified ball and player physics. Our goal was then to create a team of Swarm AI agents that could play soccer against another implemented team of reactive agents that don't communicate with each other. This opponent team has a centralized "defensive coach agent" that tells each player whom it should cover on defense.

We identified 3 global problems that have to be solved by a soccer team:

- 1) Defensive assignments: on defense, whom each friendly player should guard
- 2) Passing: deciding when the ball carrier should pass the ball.
- 3) Offensive spacing: on offense, how should all the

players without the ball move to be open for a pass while being ready for an enemy counterattack.

Because there are three high level problems that need to be solved, we created three separate Swarm AI approaches for solving them. We took advantage of shared design elements and allowed the Swarm AI players to switch from one problem solving behavior to another as the game situation changes.

To create an intelligent team of soccer players, we created a multiagent system that satisfies the Swarm AI principles. Fulfilling the first Swarm AI principle, we divided up the thinking for the team among all the friendly players. This seemed natural since our problems deal with figuring out how each player on a single team should act, so we initially wanted to involve them and nothing else (like the ball, the goals, or enemy players) in the decision-making process. Therefore, we assigned an agent to control each player on the friendly team, which proved sufficient for solving the first two of the above-mentioned global problems.

The soccer field is large, and there are too many events occurring at the same time for a single player agent to keep track of them all in real time. The second Swarm AI principle suggests we limit the perception area of each agent and how much thinking it does. Consequently, our Swarm AI player-agents focus on the immediate area around them and the locations of players in that area, ignoring information like player velocities and not thinking beyond what to do at the current point in time. In addition to the physically local area around them, Swarm AI players keep in mind who has control of the ball and the locations of the ball, the goals, and enemy players they may be guarding. The set of actions available to each player is defined by the rules of soccer, so they can move, kick the ball, and take it away from enemies by running into them. Using these agent requirements, we designed a team of utility-based player agents that can carry out primitive behaviors like chasing, shooting on goal, passing, and others. Each global problem required a separate solution that we reached by satisfying the principle of communication in different ways.

## 5.2 Defensive Assignments

In real soccer, human players performing defensive assignments try to let others know whom they intend to guard. We used stigmergy, as discussed in the third Swarm AI principle, to allow friendly players to exchange this information. Specifically, each friendly player has a unique color that its teammates can recognize as that player's color. When a player A decides to guard an enemy player B, A "paints" B with its own color. Later, when A's teammates are thinking about guarding the painted B, they will see A's paint on B and understand that A has committed to guarding B. Players have a simple utility function that tells them it's better to guard nearby

dangerous enemy players (ones with the ball or those who are close to the friendly goal) who are not already covered by a teammate. In addition, players can overwrite others' paint with their own when they want to guard an already covered player, and paint degrades over time as discussed in the next section.

This way of relaying information allows defenders to have a sophisticated real-time discourse about who is guarding whom using the enemy players as the environment through which the communication is performed. The resulting team behavior ensures that the most dangerous enemy players are usually covered and sometimes double-teamed when two players happen to decide that an enemy player is important enough to warrant two defenders. The coordinated double-teaming behavior occurs without direct communication as a product of this Swarm AI-created approach, which we call Swarm Painting.

## 5.3 Passing

For the passing problem, the first Swarm AI principle suggests we involve friendly players without the ball in helping the ball carrier decide when to pass. This is a good idea because the ball carrier is already trying to avoid enemies, move the ball up-field, and shoot on goal if it is possible. Therefore, it has plenty of other things to think about, and it would be beneficial to decrease the amount of information processing that particular agent has to perform.

To accomplish this distributed thinking, we have the players without the ball change their own state to communicate with the observing ball carrier about their openness status (using the second type of communication discussed in Section 2.3). Again, we use painting where players without the ball decide if they are open and paint themselves a certain color to let the passer know they are currently not covered. In this way, a player's observable state is changed, which communicates a message to the passer when it sees the "self-paint." As a result, the ball carrier only has to add up the number of open or "self-painted" teammates and has a likelihood of passing the ball that is proportional to the number of painted friends it sees at the moment. This self-painting variant of the Swarm Painting technique we discussed earlier allows the Swarm AI players to make timely coordinated decisions about when to pass the ball.

In the two Swarm Painting approaches, we used paint that degrades over time like the Ant System's pheromones [6], finding it effective in relating information about being open or defensive responsibilities in the dynamic soccer environment. For example, the paint on a player that has self-painted himself to show that it is open disappears over time, which is appropriate because that player is not guaranteed to remain open for long.

## 5.4 Offensive Spacing

The offensive spacing problem was solved by introducing agents for the enemy players, satisfying the first Swarm AI principle in a way similar to how Drogoul used enemy chess pieces in [11]. The behavior of agents for this problem is governed by simple rules that determine attraction and repulsion vectors toward the ball, the goals, and other players. We arrived at these rules by observing that human soccer players tend to move toward their own goal on defense and toward the enemy's goal on offense, while also moving toward the ball. Furthermore, real players can communicate without words just by changing their own location. For example, if a player without the ball who is trying to become open comes closer to a teammate who is also trying to become open, the second player will move away from the first because of a general desire to be spread out. In this way, the first player can influence the second by changing its own location. We implemented this behavior by again using the "change own state to influence observers" method of communication. In this case, a player agent's state is simply its location, a change in which causes nearby teammates to spread out when they observe it or move away if the communicating agent is an enemy that is coming closer. This results in the advantageous global behavior of players without the ball spreading out across the field while making runs to become open.

## 5.5 Results of Applying Swarm AI to Soccer

One result of applying Swarm AI to the soccer problem was the creation of Swarm Painting, a new SI approach that we think may be applied to other problems that deal with decisions unrelated to movement. Note that Swarm Painting was designed using the Swarm AI methodology combined with domain knowledge about real soccer – not by mimicking some social insect phenomenon.

The design process took us less than a month, and we quickly had a working prototype. In testing, the Swarm AI team performed better than its opponent, scoring an average of 1.64 times more goals in the series of experiments we conducted, which is a measure of both its offensive and defensive capabilities in relation to the non-communicating reactive agents team. Also of note is that the Swarm AI team behaved realistically, performing complex multi-player soccer maneuvers like give-and-gos, passing into space, and "total football" ball possession. These behaviors were not pre-programmed by us; they are products of the simple communication between the Swarm AI soccer agents.

## 6 Future Directions

This paper is only the beginning of the road towards the development of Swarm AI as a problem-solving method. The design framework proposed in this paper is far from comprehensive since we have so far identified a number of principles and are not ready to present a step-by-step process that one may follow. More specifically, it is still unclear how the designer of a multiagent system should distribute the problem solving, and we intend to provide some additional theory to make this step less Ad Hoc. In addition, we have much to learn about deciding which communication technique to use and what subtypes are available within the two categories described by the third Swarm AI principle. We also think it would be beneficial to perform studies on how simple and/or fast Swarm AI agents have to be to operate effectively. Thus, we feel that there a lot more we can learn about the relationship between agent speed and complexity with regard to performance, taking into account the nature of the communication method used. Work in this direction will improve the Swarm AI framework by giving designers more specific guidelines to follow and help answer questions like, "Exactly how simple do Swarm AI agents have to be?" As a result, we plan to apply Swarm AI to additional problems to study the effects of various design decisions on the behavior of Swarm AI systems.

While we value the current work done on extending previous SI approaches, we feel it is imperative that more researchers try to come up with new Swarm Intelligence approaches now that a design framework is available. As more methods are created using Swarm AI, we will gain a better understanding of how to apply it and what kinds of problems it can handle. We also believe that it would be advantageous to apply evolutionary methods in tweaking Swarm AI systems, which could result in an even more powerful problem-solving method. Finally, [9] discusses the potential of using hierarchies of Swarm systems, and we think Swarms of Swarms may very well be the future of problem solving. However, our current goal is to understand how to create useful single Swarm systems, which we may later learn to combine in order to produce even more intelligent behaviors for solving harder tasks.

## 7 Conclusion

We introduced Swarm AI, a design framework for creating Swarm Intelligence approaches to problems. This framework is defined by three principles:

- 1) Create a system of agents that work individually on a common problem
- 2) Agents are simple, fast, and have a limited perspective
- 3) Indirect simple inter-agent communication

We have demonstrated how this approach is a step forward from previous SI work in several ways, including its lack of emphasis on studying social insects and relative generality. In fact, Swarm AI can be seen as a bridge

between Swarm Intelligence and engineering problems, since this method takes advantage of things learned from nature without forcing whoever uses it to actually model natural phenomenon. Furthermore, Swarm AI is an excellent choice as a problem-solving technique because it offers good performance at a small design time and complexity cost. Our own experiment with applying Swarm AI to soccer problems led to successful results. The Swarm AI soccer team not only soundly defeated a non-Swarm opponent team, but also exhibited sophisticated globally complex soccer behavior in a dynamic real-time environment. Ultimately, Swarm AI is a promising new technique that generalizes Swarm Intelligence into a useful design framework.

## References

- [1] Bonabeau, E. and Theraulaz, G. (Mar. 2000) "Swarm smarts," *Scientific American*, pp. 72-79.
- [2] Kennedy, J., Eberhart, R. C., Shi, Y. (2001). "Swarm Intelligence," San Francisco: Morgan Kaufmann Publishers.
- [3] Mamei, M., Zambonelli, F., Leonardi, L. (2002) "Co-Fields: A Unifying Approach to Swarm Intelligence", *3rd International Workshop on Engineering Societies in the Agents' World*, Madrid (E), LNAI Vol. 2577.
- [4] Bourjot, C., Chevrier, V., Vincent, T. (2003) "A New Swarm Mechanism based on Social Spiders Colonies: from Web Weaving to Region Detection," *Web Intelligent and Agent Systems: An International Journal, WIAS*.
- [5] Bonabeau, E., Dorigo, M., Theraulaz, G. (1999) "Swarm Intelligence: From Natural to Artificial Systems," New York: Oxford University Press.
- [6] Maniezzo, V., Gambardella, L. M., De Luigi, F. (2004) "Ant Colony Optimization," *New Optimization Techniques in Engineering*, by Onwubolu, G. C., and B. V. Babu, Springer-Verlag Berlin Heidelberg, pp. 101-117.
- [7] Wooldridge, M. (1999) "Intelligent Agents," in G. Weiss, editor: *Multiagent Systems*, The MIT Press.
- [8] Parunak, H. D. V. (1997) "Go to the Ant: Engineering Principles from Natural Agent Systems," *Annals of Operation Research*,(75) pp. 69-101.
- [9] Hoffmeyer, J. (1994) "The swarming body," *Proc. 5th Congress of the International Association for Semiotic Studies*, Berkeley.
- [10] Dorigo, M., Maniezzo, V., Colorni, A. (1991) "Positive feedback as a search strategy," *Technical Report 91-016*, Dipartimento di Elettronica, Politecnico di Milano, Italy.
- [11] Drogoul A. (1993) "When Ants Play Chess (Or Can Strategies Emerge from Tactical Behaviors?)" *From Reaction to Cognition: Select Papers, Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '93)*, pp. 13-27. Springer Verlag.
- [12] Drogoul, A., Ferber, J., and Jacopin, E. (1991) "Pengi: Applying Eco-Problem-Solving for Behavior Modelling in an Abstract Eco-System," *Modelling and Simulation: Proc. of ESM'91*, Simulation Councils, Copenhagen, pp. 337-342.
- [13] Dorigo, M., Stützle, T. (2002) "The ant colony optimization metaheuristic: Algorithms, Applications and Advances," *Glover, F. and Kochenberger, G., editors: Handbook of Metaheuristics*, pp. 251-285, Kluwer Academic Publishers.