Swarm Engineering

Sanza T. Kazadi
Jisan Research Institute
28 N. Oak Avenue
Pasadena, CA 91107
sanza@jisan.org

Abstract: This paper surveys the research on swarm engineering, a recent offshoot of swarm intelligence. In particular, we examine the motivation and approach of swarm engineers in developing swarm-based applications and propose a method that is essentially top-down.

1. The Genesis of Swarm Intelligence

It has almost been two decades now since Rodney Brooks (Brooks, 1986) published his seminal work on the design of micro robots that kicked off the minimalist design rediscovery in robotics design and artificial intelligence research. This work proposed a new way of viewing intelligence as a result of both the individual's interaction with the environment and thought processes. It was about the same time that Deneubourg (Denebourg, 1983) was examining ant systems and discovering the remarkable dynamics of ant hills that had heretofore been poorly understood. This also marked the first year that the artificial life community became organized. These three contemporaneous occurrences combined to pave the way for the exploration of systems of simple agents.

Historically, it was the decision-making capability of swarms that captivated the imagination of scientists. This capability seemed to have the potential to lead to intelligent swarms of agents capable in groups of accomplishing the kind of decisions that artificial intelligence had been exploring with individual agents. We have been able to illustrate this kind of decision making in artificial ant simulations, and to demonstrate how little it takes to get the kind of dynamics one finds in ant hills. However, whether or not this will lead to swarms of agents whose aggregate behavior leads to swarms capable of making more sophisticated decisions remains to be seen.

2. Swarm Engineering

For the vast majority of researchers working with swarms, the actual capability of the swarm is much more important than the details of whether or not the swarm is performing an intelligent behavior. Dorigo and associates (Dorigo, 1996) have developed methodologies, for instance, which have centered around the ability of agents whose global interactions are mediated in ways similar to that used by ants.

Recently, the vast distributed nature of the Internet, several researchers are exploring the use of distributed agent systems in commerce, computation, communication, etc (Foster, 2001). Alongside these computational algorithms arose several algorithms based on robotic platforms (Holland, 1999; Kazadi, 2000). The earliest of these were centered around the movement of cylindrical objects called "pucks", but later work centered around swarm-based coordination, cooperative movement, etc. These various methodologies, again, are no longer centered around the development of intelligent behaviors, but rather, around predictable and controllable behavior.

Thus, it is safe to say that in the current climate the focus of the large majority of research on swarms has moved away from the swarm intelligence paradigm. Rather, the focus seems to be on the design of predictable, controllable swarms with well-defined global goals and provable minimal conditions. This seems much more like an engineering paradigm than a study of swarm-based artificial intelligence. As a result, the term *swarm engineering* seems much more fitting for the current research activities of many researchers around the world. To the swarm engineer, the important points in the design of a swarm are that the swarm will do precisely what it is designed to do, and that it will do so in reliably and on time. The benefits of using a swarm over using single devices are that the redundancy of the swarm allows for a number of things that single devices cannot - automatic recruitment of more effort on hard or large tasks, specialization on multiple simultaneous tasks, and robustness to failure. Swarm engineering then consists of the application of swarm-based methods to problems that will benefit from these properties of swarms in ways that are predictable.

Much of what was done with swarms in the initial years was not particularly generalizable. That is, it wasn't easy to see how the methods that had been developed in solving a particular task could be applied to other tasks outside of the immediate scope of the task at hand. To the swarm engineer, particularly when carrying out research on any particular topic, it is more important to know to which classes of problems a particular technique will work and then to gauge its properties. The initial work, then had to be bolstered by a more complete understanding of the class of problem to which the technique could be applied. This work has begun happening, particularly in the fields of distributed cluster-based construction, collaborative robotics, and ant optimization algorithms (Kazadi, 2002; Kazadi, 2004; Lerman, 2003; Olfati-Saber 2002). These studies have begun grounding the work done in a theoretical framework which has begun to provide measurable and provable predictions about the behaviors and characteristics of swarms.

In particular, swarm engineers have begun moving towards a more top-down procedure for generating swarms of agents. This top-down mechanism is probably more effective than the initial bottom-up methodology of the early artificial intelligence studies because it is very easy to generate complex dynamics from simple systems in ways that cannot be immediately predicted from the system design. As a result, the hallmarks of swarm engineering studies are approximating the following:

Step 1: Generate a provable condition that each agent will have to satisfy in order to accomplish the global task.
Step 2: Build agents that satisfy the condition.

Of course, the majority of the effort of this methodology is in Step 1, and the methods for proving that the minimal requirements lead to the global outcome are varied. The accomplishment of Step 1 requires engineers to understand clearly what the class of problem is to which the swarm methodology is being applied. This allows the generalization of the task under study to other tasks potentially far removed from the original task.

3. Future of Swarm Engineering

At present, Swarm Engineering is an emerging field, with a growing number of people interested in the techniques. There are many questions to be answered, particularly in the area of classification of swarm-based tasks and general methods. To date, there have been several different task areas identified, including deliberate and stigmergic communication, distributed recruitment, sensor fusion, distributed computation, and medical uses of swarms. It is currently unknown if the number of different areas in which swarms have been applied is exhaustive. A catalog of methods for each of these areas needs to be built up so that it may be possible to choose from one of many different methods when designing an application. More effort must be focused on work which provably describes swarm behaviors. This is likely to lead the field in directions that are currently out of the reach of our current methodologies and expectations.

References

Brooks R. *A robust layered control system for a mobile robot*. **IEEE J. Robotics and Automation**, RA-2 (1): 14-23, 1986.

Dorigo M., Maniezzo V. and Colorni A. *The Ant System: Optimization by a Colony of Cooperating Agents*. **IEEE Transactions on Systems, Man, and Cybernetics-Part B**, 26(1):29-41.

Deneubourg, J. et. al. *Probabilistic behaviour in ants: a strategy of errors?* **Journal of Theoretical Biology**, 105: 259-271, 1983.

Foster I., Kesselman C., and Tuecke S. *The anatomy of the grid. Enabling scalable virtual organizations*. **International Journal of Supercomputer Applications**, 2001.

Holland O. and Melhuish C. *Stigmergy, self organization, and sorting in collective robotics*. **Artificial Life**, 5:173-202, 1999.

Kazadi, S. **Swarm Engineering**. Ph.D. Thesis, California Institute of Technology. 2000.

Kazadi, S., Abdul-Khaliq, A., and Goodman, R. *On the convergence of puck clustering systems*. **Robotics and Autonomous Systems**, 38(2): 93-117, 2002.

Kazadi, S. et. al. *On the dynamics of puck clustering systems*. **Robotics and Autonomous Systems**, 46(1): 1-27, 2004.

Lerman, K., Galstyan, A. and Hogg, T. *Mathematical Analysis of Multi-Agent Systems*. cs. RO/0404002, 2003.

Martinoli A., Easton K., and Agassounon W.. *Modeling of Swarm Robotic Systems: A*

Case Study in Collaborative Distributed Manipulation. **Special Issue on Experimental Robotics**, P. Dario and B. Siciliano, editors, **Int. Journal of Robotics Research**, 23(4), 2004.

Olfati-Saber R. and Murray R. *Graph rigidity and distributed formation stabilization of multi-vehicle systems*. **Proceedings of the 41st Conference on Decision and Control**, Las Vegas, NV, Dec. 2002.

Schoonderwoerd R. and Holland O. *Ant-based load balancing in telecommunications networks*. **Adaptive Behavior**, 5(2): 169-207, 1997.