

A-TEAMS

How Distributed, Autonomous Agents Can Cooperate to Solve Both Off-Line and Real-Time-Control Problems

Institute for Complex Engineered Systems

Carnegie Mellon University
College of Engineering
1201 Hamburg Hall
Pittsburgh, PA 15213-3890

ph: 412.268.3372
fax: 412.268.5229
<http://www.ices.cmu.edu>

For More Information:

Sarosh Talukdar
ph: 412.268.8778
email: talukdar@cmu.edu
<http://www.ece.cmu.edu/~talukdar>

WHAT IS AN ASYNCHRONOUS TEAM?

Asynchronous teams (A-Teams) are open, high-performance organizations for solving difficult problems, particularly, problems from the areas of planning, design, scheduling and real-time control. A-Teams get their openness (ease of assembly) from the use of autonomous agents that can be added or removed without any supervisory system getting in the way. A-Teams get their high-performance (quick convergence to good solutions) by exploiting multiplicity (large numbers of agents, computers, and other resources). Indeed, what an A-Team does well, perhaps better than any other type of organization, is exploit multiplicity.

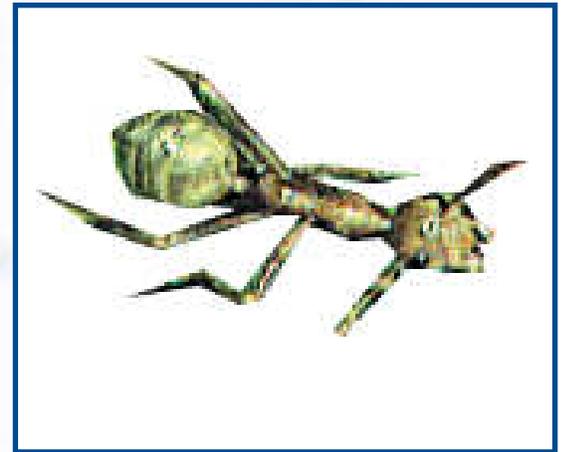
WHERE DID THE IDEAS COME FROM?

An A-Team is a combination of features taken from a variety of natural and synthetic systems, including insect societies, cellular communities, brainstorming protocols, blackboards, genetic algorithms, simulated annealing, and tabu search.

HOW DOES AN A-TEAM WORK?

Structurally, an A-Team is a strongly cyclic network of autonomous agents and shared memories. Each agent contains some problem-solving skills; each memory contains a population of trial-solutions to a part of the problem-to-be-solved. The trial-solutions circulate continually. Agents cooperate by selecting and modifying these solutions as they circulate. The modifications can be constructive (an improvement to a part of the selected solution), or destructive (the erasure of the selected solution). All the agents can work asynchronously (each at its own speed) and in parallel. Some of the agents are simple, others can be complex, some of the agents can be computer-based, others can be human.

An A-Team's structure allows for multiplicity to the point of excess, along five dimensions: 1) multiplicity of skills, so no relevant skill, no matter how specialized or redundant it may seem, need be excluded; 2) multiplicity of agent-types, so a wide variety of agents, including software and human agents, can be accommodated; 3) multiplicity of representations, so the different representational needs of all the agents can be met; 4) multiplicity of computers, so all the agents can work in parallel, whether they are co-located or distributed; and 5) multiplicity of trial-solutions, so a suitably wide range of promising alternatives can be simultaneously considered.



Both experiment and theory confirm two key properties. First, construction and destruction are complementary processes in an A-Team: adept destruction can compensate for inept construction, and vice-versa. Second, there are profound benefits-of-scale, meaning that the performance of an A-Team can invariably be improved by expanding it along one or more of its five "multiplicity-dimensions."

APPLICATIONS

A-Teams can be applied to both off-line problems, like planning, design and scheduling, and real-time control problems, like the operation of electric power or packet switching networks. At CMU, A-Teams have been developed for traveling salesman problems, high-rise building design, reconfigurable robot design, fault diagnosis, operation planning, and job shop scheduling. Work is now underway to develop an A-Team whose agents will be distributed over an electric power network and will continually learn how to do their control jobs better. At IBM, A-Teams have been used for steel and paper mill scheduling.

FOR FURTHER INFORMATION

Contact Sarosh Talukdar at 412 268 8778 or talukdar@cmu.edu. Links to papers on A-Teams can be found at: www.ece.cmu.edu/~talukdar/