A Framework for Truth Maintenance in Multi-Agent Systems

Brett Bojduj[1], Ben Weber[1], and Dennis Taylor[2]

[1] California Polytechnic State University, San Luis Obispo, U.S.A.
bbojduj@calpoly.edu
bgweber@calpoly.edu
[2] CDM Technologies, Inc., San Luis Obispo, U.S.A.
djtaylor@cdmtech.com

Summary. Maintaining logical consistency in a knowledge base that receives asynchronous events from collaborative agents poses a challenge, due to multiple agent perspectives of the knowledge base. Facades and filters further complicate this problem by distorting the definition of consistent knowledge. Therefore, maintaining logical consistency of the knowledge base requires an infrastructure for handling truth maintenance. This paper presents a generic, object-oriented framework for truth maintenance in collaborative multi-agent systems. The core of the framework is an agent that autonomously reasons on system events, thus guaranteeing the integrity of the knowledge base independent of external agents. Specialization to a particular domain is achieved through the description of tests that verify the consistency of the knowledge base. This paper shows an example of this approach in a real-world, multi-agent system and discusses performance and maintainability in such a system.

1 Introduction

A challenge for large software systems is maintaining logical consistency between collaborative agents operating asynchronously to solve a problem. Multi-agent systems may use a divide-and-conquer approach to solve a problem [4]. They can do this by assigning different aspects of the problem to different agents. Each agent in the system may have a different perspective of the problem and the definition of consistency may vary between agents. Therefore, multi-agent systems require that solutions are consistent across all agent perspectives. This paper presents an external truth maintenance agent as a technique for maintaining consistency between collaborative agents.

Multi-agent systems have been used to solve constraint satisfaction problems [7]. For large, highly-constrained problems, implementation of an agent that simultaneously solves all constraints becomes impractical. A common solution is to construct collaborative agents that subdivide the problem into

smaller problems and share solutions. Decomposing the problem requires that solutions shared between two agents are consistent with both agents' perspectives of the problem. Consistency can be achieved if agents communicate through a shared knowledge base that is monitored by an external truth maintenance agent. The truth maintenance agent verifies that solutions in the knowledge base are consistent with respect to all constraints in the problem. This alleviates individual agents from the burden of maintaining consistency for the entire solution.

This paper explores the use of a truth maintenance agent to maintain logical consistency of a knowledge base in a global logistics and scheduling application. The task of planning is distributed between several collaborative agents and the consistency of the knowledge base is verified by an external truth maintenance agent. The presented technique enables metaheuristic-based agents to generate consistent solutions in a near real-time planning environment.

2 Truth Maintenance Systems

Most agents require logically consistent information in order to perform their reasoning functions. In dynamic environments, additional information is added to the problem definition while the system is running, which may conflict with current solutions. If logically inconsistent data enters a system, existing solutions must be marked as invalid or revised in order to allow agents to reason accurately.

Doyle proposed a truth maintenance system to handle cases when logically inconsistent data enters a system [2]. Each set of data in the system is annotated with a node that specifies the consistency of the data. When new information enters the system, every node is updated to reflect the new problem definition. Huhns and Bridgeland extended the truth maintenance system to enable consistency in multi-agent systems [5]. When new information enters the system, every agent verifies that its shared data is consistent. This validation process becomes a bottleneck, but is required for blackboard architectures. If the knowledge base becomes compromised, then agents in the system communicating through the blackboard will be unable to properly perform their reasoning tasks.

Traditional truth maintenance systems guarantee consistency of a knowledge base, but are unsuitable for real-time systems. Elkan proved that truth maintenance systems that deal with logically inconsistent information are NP-complete [3]. Therefore, systems must compromise global consistency in order to operate in real-time environments. The amount of computation required to enforce consistency can be reduced if agents consider different aspects of the problem. Bertel et al. presented the concept of "aspect filters" which can be applied to truth maintenance systems [1]. When new data enters the system, consistency is maintained within only the affected aspects of the knowledge