A Knowledge-Based Methodology for Designing Robust Multi-Agent Systems

Mark Klein
Massachusetts Institute of Technology
(617) 253-6796
m_klein@mit.edu

Chrysanthos Dellarocas
Massachusetts Institute of Technology
(617) 258-8115
dell@mit.edu

Juan Antonio Rodriguez-Aguilar
iSoco S.A.
jar@isoco.com

ABSTRACT
This paper describes a methodology that multi-agent system designers can use to identify, and find suitable responses for, potential failure modes (henceforth called 'exceptions') in multi-agent systems.

Categories and Subject Descriptors
I.2.11 Distributed Artificial Intelligence

General Terms: Design

Keywords: Exception handling

1. INTRODUCTION
Multi-agent systems, especially open ones that must operate with agents from diverse sources on the infrastructures at hand, must be able to operate robustly despite many possible failure modes ('exceptions'). This paper describes a methodology that human multi-agent system designers can use to identify, and find suitable responses for, these potential failures (henceforth called 'exceptions').

2. METHODOLOGY
Our exception analysis methodology is based on the insight that coordination fundamentally involves the making of commitments, and that exceptions (i.e. coordination failures) can as a result be viewed as *violations* of the commitments agents require of one another. Exception analysis thus consists of the following three steps:

Identify Commitments: Commitments can be identified by finding all the places in the protocol where one agent depends on some other agent in order to discharge its own responsibilities successfully. There are two general kinds: *design*-time commitments that are part of the mechanism definition (e.g. that a bidding agent will send only one bid at a given price level), as well as *run*-time commitments created as a result of the operation of the mechanism (e.g. that an agent will in fact perform the task it was allocated with the contracted cost quality and delivery time).

Identify Commitment Violations (Exceptions): The next step is to identify, for each commitment, how that commitment can be violated, i.e. what its characteristic exceptions are. The basic set of possible exceptions can be identified simply as the possible negations of the commitment itself.

Identifying Exception Handlers: Once we have identified the exceptions that potentially characterize a given coordination mechanism, we are ready to identify possible techniques for *handling* these exceptions. If the coordination mechanism does not currently provide any handlers for an important exception type, then the mechanism designer needs to select one or more handlers and incorporate them into the coordination mechanism being analyzed. There is a vast range of possible exception handling techniques. One can argue in fact that a very significant proportion of human organizational innovations, including the police, law courts, disaster relief agencies, insurance, contracts and so on, can all be viewed as being exception handling techniques. As with exception *causes,* however, the range of possible *handlers* appears to be limited only by human creativity and there appears to be no way to identify all potentially useful exception handlers from first principles.

The fact that exception and handler identification can be time-consuming as well as at least partly experiential in nature has led us to explore whether it is possible to systematically accumulate this knowledge so that we can analyze new mechanisms more quickly and completely. We have been able to do so using a substantively extended version of the MIT Process Handbook , a process knowledge repository which has been under development at the Center for Coordination Science (CCS) for the past six years.

3. FOR MORE INFORMATION
For additional information about this and related work, see http://ccs.mit.edu/klein/.

AAMAS'02, July 15-19, 2002, Bologna, Italy.