

Active Coordination of Distributed Human Planners

Karen L. Myers Peter A. Jarvis Thomas J. Lee

Artificial Intelligence Center
SRI International
333 Ravenswood Ave.,
Menlo Park, CA 94025
{myers, jarvis, tomlee}@ai.sri.com

Abstract

Effective coordination of distributed human planners requires timely communication of relevant information to ensure the overall coherence of activities and the compatibility of assumptions. This paper presents a framework called CODA that provides targeted information dissemination among distributed human planners as a way of improving coordination. Within CODA, each planner declares interest in different types of plan change that could impact his or her local plan development. As individuals develop plans using a plan authoring tool, their activities are monitored; changes that match declared interests trigger automatic notification of appropriate planners. In this way, distributed planners can receive focused, real-time updates of plan changes that are relevant to their local planning efforts.

Introduction

The scope and complexity of large-scale planning tasks generally requires cooperation among multiple planners with differing areas of expertise, each of whom contributes portions of the overall plan. These planners may be distributed both geographically and temporally, further complicating coordination.

As a concrete illustration, special operations forces (SOF) mission planning involves numerous people working on separate but interconnected facets (e.g., strategic, logistical, medical, intel) of an overall plan. The SOF planning process itself is time constrained, concurrent, and iterative. Individual planners construct subplans based on their expectations for the operating environment and requirements. As the overall plan develops, these expectations change and modifications must be made to reflect new information. Currently, such changes are communicated informally by word of mouth, or transmitted in batch mode at regularly scheduled coordination sessions. This approach can lead to omissions and delays that reduce the effectiveness of the overall planning process and the quality of the resulting plans.

The SOF planning domain lies well beyond the range of current automated planning technologies. Moreover, fully automated solutions are unlikely ever to succeed, due to two

main factors. First, effective planning for this domain involves a huge strategic component that requires evaluating options with respect to expected plan quality. This type of knowledge is extremely difficult to capture and model, being grounded in a combination of intuition, vast experience, and common sense. Planning knowledge bases constructed to date mostly ignore strategic issues, focusing instead on modeling preconditions/postconditions for action applicability and effects. Second, SOF planning tasks (e.g., disaster relief, counterterrorism) tend to be unique and distinctive, making it difficult to formulate reusable background knowledge with adequate coverage. As an additional consideration, human planners in this area are reluctant to cede full control to automated planning systems, even in situations where automation is possible.

The above characteristics apply to a range of planning domains, including more general military operations planning, workflow management, and certain types of space missions. Techniques from the AI planning community can still contribute to complex problem solving in these domains. In particular, *plan authoring* tools that build on AI planning concepts are being introduced to improve the quality and process of plan development (Valente *et al.* 1999; GTE 2000; Muñoz-Avila *et al.* 1999; Knoblock *et al.* 2001; desJardins *et al.* 2002). Plan authoring tools provide a set of plan editing operations and manipulation capabilities that support users in exploring and developing plans. These tools may provide some automated capabilities; however, their main role is to augment rather than replace human planning skills. Plan authoring tools introduce a degree of structure to the planning process, yielding principled representations of plans with well-defined semantics. Planning aids that reason over these structures can be defined, including tools to support interplanner coordination.

Three main considerations drive the development of coordination techniques within this type of setting.

Selectivity Coordination strategies must strike a balance between disseminating too little and too much information. Failure to communicate enough information could be disastrous, if planners do not receive notification of critical changes. Conversely, flooding planners with much extraneous information can lead to cognitive overload, with critical changes being lost in a sea of irrelevant updates. Overcommunication can be a further problem for

situations where bandwidth is limited (e.g., wireless devices) or communication is expensive.

Timeliness Information about plan changes must be received in a timely manner to ensure adequate time for human planners to assess impact on their local plans and develop appropriate repairs.

Nonintrusiveness Users are generally reluctant to embrace technologies that require changes to traditional work habits (Conklin & Yakemovic 1991). In particular, a user is unlikely to perform activities beyond his normal sphere of responsibility unless there are immediate and substantial benefits for him. Thus, coordination tools that impose demands on users must incentivize their participation.

This paper describes the CODA (Coordination of Distributed Activities) framework, which provides automated support for focused information sharing during collaborative plan development by a team of humans. While motivated by the SOF planning problem, CODA more generally targets applications where distributed human planners are assigned responsibility for developing subportions of a global plan. These subplans are expected to have a moderate degree of coupling due to the need for coherent strategy, coordinated actions, and sharing of limited resources.

Within CODA, each planner declares the kinds of plan change that are of interest to him or her; we call these declarations *plan awareness requirements* (or *PARs*). As users develop plans with a plan authoring tool, their activities are monitored. Changes that match plan awareness requirements are forwarded automatically to the person who declared interest in them. In this way, distributed planners can receive focused, real-time updates of plan changes that are relevant to their local planning efforts.

Because local planners declare precisely the information in which they are interested, CODA satisfies the criteria for selectivity stated above. The declaration of plan awareness requirements constitutes the only additional effort required by the human planners above and beyond their normal operation. Because this declaration process will produce direct benefits to the planner (namely, notification of changes in which they are interested), the motivation for this specification phase is strong. Finally, the real-time nature of the plan updates within CODA addresses the timeliness concern.

The paper begins with an overview of the planning model within CODA, followed by a description of the CODA architecture. Next, we define our language for expressing plan awareness requirements along with a formal semantics and algorithms for matching them to evolving plans. We then describe an implementation of the CODA framework, focusing on different matching modes that it supports and tools for creating and registering plan awareness requirements. Finally, we compare our approach to related work in distributed AI, active databases, and concurrent engineering.

Plan Model

The generative planning community typically models a plan as a partial order of actions, and planning as an action selection problem. More generally, planning can be viewed

Objective a goal to be achieved within a plan

Action an activity that can be performed to achieve objectives

Effect a world-state condition denoting an expected result of plan activities

Event a world-state condition denoting an expected externality

Decision Point a condition for branching among subplans

Role a required capacity within a plan

Relation a semantic connection among plan objects (e.g., an Action *supports* an Objective or *enables* an effect)

Figure 1: Plan Ontology for CODA

as a *design task* that involves creating, refining, and linking objects to produce a composite structure that satisfies stated design objectives. The plan authoring model underlying CODA embodies this more general model of planning.

Plan Ontology and Structure A CODA plan is composed of a collection of objects drawn from the plan ontology presented in Figure 1. Each plan object is characterized by a unique *id* and may optionally have a (not necessarily unique) *name*. A collection of *domain objects* augments the plan ontology, representing the basic entities within a specific domain.

While most of the ontology elements correspond to standard concepts in AI planning, *roles* are somewhat different. A role defines an abstract capacity within a plan, independent of the specific object that will eventually fill that capacity. For example, *place* roles are used frequently within the SOF domain: in a disaster relief plan, there may be one or more *assembly-point* roles that denote places where evacuees should assemble to be evacuated. A planner often knows early on that assembly-points with certain capacities are required, but the exact physical locations for them and the actions that make use of them may not be chosen until later in the planning process.

Planning Process The planning process for CODA involves specifying roles required within a plan, selecting objects to instantiate those roles, selecting particular actions to be executed, declaring expected effects and events, and asserting relationships among objects and actions. The planning process is incremental in nature, involving the selection, extension, and refinement of selected objects and actions until the plan meets the human planner's criteria for adequacy with respect to stated objectives.

Plan Query Language The language for expressing plan awareness requirements (described in a subsequent section) builds on a general-purpose plan query language for the CODA plan ontology. The query language consists of a typed first-order language specialized to the Generic Frame Protocol (GFP) model of frame representation systems (i.e., it includes the full range of GFP-defined functions and predicates for manipulating classes, instances, and relationships

