UM Translog: A Planning Domain for the Development and Benchmarking of Planning Systems [*]

Scott Andrews, Brian Kettler, Kutluhan Erol, and
James Hendler

Department of Computer Science,
Institute for Advanced Computer Studies,
and Institute for Systems Research
University of Maryland
College Park, MD 20742
E-mail: hendler@cs.umd.edu
June 1995

[*]This research was supported in part by grants from NSF (IRI-8907890), ONR (N00014-J-91-1451), AFOSR (F49620-93-1-0065), the ARPA/Rome Laboratory Planning Initiative (F30602-93-C-0039) and by ARI (MDA-903-92-R-0035), subcontract through Microelectronics and Design, Inc. Dr. Hendler is also affiliated with the Institute for Systems Research and the Institute for Advanced Computer Studies.

Abstract

The last twenty years of AI planning research has discovered a wide variety of planning techniques such as state-space search, hierarchical planning, case-based planning and reactive planning. These techniques have been implemented in numerous planning systems (e.g., [12, 8, 9, 10, 11]). Initially, a number of simple toy domains have been devised to assist in the analysis and evaluation of planning systems and techniques. The most well known examples are "Blocks World" and "Towers of Hanoi." As planning systems grow in sophistication and capabilities, however, there is a clear need for planning benchmarks with matching complexity to evaluate those new features and capabilities. UM Translog is a planning domain designed specifically for this purpose.

UM Translog was inspired by the CMU Transport Logistics domain developed by Manuela Veloso. UM Translog is an order of magnitude larger in size (41 actions versus 6), number of features and types interactions. It provides a rich set of entities, attributes, actions and conditions, which can be used to specify rather complex planning problems with a variety of plan interactions. The detailed set of operators provides long plans ($\tilde{4}0$ steps) with many possible solutions to the same problem, and thus this domain can also be used to evaluate the solution quality of planning systems. The UM Translog domain has been used with the UMCP, UM Nonlin, and CaPER planning systems thus far.

Contents

1 Background and Motivation

The last twenty years of AI planning research has discovered a wide variety of planning techniques such as state-space search, hierarchical planning, case-based planning and reactive planning. These techniques have been implemented in numerous planning systems (e.g., [12, 8, 9, 10, 11]). Initially, a number of simple toy domains have been devised to assist in the analysis and evaluation of planning systems and techniques. The most well known examples are "Blocks World" and "Towers of Hanoi." As planning systems grow in sophistication and capabilities, however, there is a clear need for planning benchmarks with matching complexity to evaluate those new features and capabilities. UM Translog is a planning domain designed specifically for this purpose.

UM Translog was inspired by the CMU Transport Logistics domain developed by Manuela Veloso[7]. UM Translog is an order of magnitude larger in size (41 actions versus 6), number of features and types interactions. It provides a rich set of entities, attributes, actions and conditions, which can be used to specify rather complex planning problems with a variety of plan interactions. The detailed set of operators provides long plans ($\tilde{4}0$ steps) with many possible solutions to the same problem, and thus this domain can also be used to evaluate the solution quality of planning systems.

UM Translog is currently being used in the evaluation of a case-based planning system, CaPER [6, 5], and a hierarchical task network planning system, UMCP [3, 2]. It is also being used by UM Nonlin [4], a common lisp implementation of Austin Tate's Nonlin HTN planning system, to generate a plan library for plan reuse by CaPER.

Due to the complexity and size of the domain, it is not easy to define UM Translog problems manually. In order to facilitate the use of UM Translog, a system has been implemented for generating random problems or assisting the user in defining his/her own problems. This system can be adapted to the input format requirements of other planning systems with reasonable effort. We hope to make the Problem Specifier/Generator code available soon.

Documentation and code related to UM Translog, UM Nonlin, UMCP and CaPER can be accessed through Worldwide Web (WWW) at URL http://www.cs.umd.edu/projects/plus/[1] or by anonymous FTP.[2]

Because several of these planning systems are ongoing research projects, the domain definition may change. We are distributing this domain definition and the affiliated code for using it with UM Nonlin and UMCP free of charge but without any implied guarantees or promises of support.

We hope that other researchers will contribute their planning domains so that a library of benchmark planning domains can be established, similar to the benchmark library used by the machine learning community.

1.1 A Guide to This Document

Section 2 presents briefly the our goals in defining the UM Translog domain and some of its general features. The domain is then described in a planner-independent manner in the sections that follow.

[1]The UM Translog homepage is at URL http://www.cs.umd.edu/projects/plus/UMT/.

[2]Instructions for FTPing UM Translog can be obtained by anonymous FTP to ftp.cs.umd.edu (file is /pub/nonlin/um-translog/README).

Entities and predicates are described in Sections 3 and 4, respectively. Operators are described in Section 5 (Actions) and Section 6 (Tasks/Methods). Some planner-specific information about this domain as defined for the UM Nonlin and UMCP systems is given in the appendices (details can be found in the actual domain definitions available from the WWW site or via FTP).

2 Overview

In this domain, the planner is given one or more goals, where a goal is typically the delivery of a particular package from an origin to a destination. Our goal for UM Translog domain was to create a domain more complex than toy domains such as blocks world. To do this we modelled additional aspects of transport logistics not in the CMU Transport Logistics domain, and which we believed were somewhat realistic. These include the following features and restrictions:

- transport is by air, rail, or road
- transport is intracity or intercity
- several basic methods of transport are available including local transport via road, direct transport via a specific direct route, and "indirect" transport via a transportation hub
- customer locations and transportation centers (airports and train stations) are grouped into cities which are grouped into regions
- intercity transport uses specific routes
- transport via air or rail uses specific transportation centers (airports and train stations)
- particular vehicles and packages have special (un)loading methods and actions
- packages and vehicles have (sub)types which must be compatible
- vehicles, equipment, routes, and transportation centers may be temporarily unavailable
- certain cities may not allow hazardous packages to be transported through them

Our goal, however, was *not* to model "real" transport planning as described by a domain expert, for example. Obviously the kind of transport logistic planning done in the "real world" by UPS, the military, etc. is considerably more complicated, particularly as the world state is highly dynamic and often not fully known by the planner. Some major aspects of transport logistics planning we are *not* modelling in order to keep the domain "manageable" and to stay with in the capabilities of the planning systems we are testing include:

- package sizes and vehicle capacities (We assume unlimited vehicle capacity.)
- numeric distances between locations[3]
- temporal aspects including delivery deadlines, action durations (travel time, etc.), action concurrency, etc.

[3] Distances are "simulated" to some extent by the grouping of locations into cities and cities into regions and the explicit specification of intercity routes.