Title: Automated Travel Itinerary Planning - An Operations Research Perspective

Presentation Authors: Simon Dunstall, Andreas Ernst, Mark Horn, Phil Kilby, Mohan Krishnamoorthy and David Sier. (All authors affiliated with CSIRO Mathematical and Information Sciences, Clayton, Victoria, Australia).

Abstract: In this presentation we describe key aspects of the *Electronic Travel Planner* (ETP), the pilot version of which has recently been developed by CSIRO Australia. ETP prepares and presents travel itineraries for tourists. These itineraries are created by combining tourism product data, information relating to the traveller's requirements, preferences and desires for their trip, and the itinerary generation capabilities of the ETP *Activity Planner*.

ETP captures the preferences, interests and values of a traveller. This information about the traveller profile is assembled into a *user request*. The user request incorporates information such as the specified starting and finishing location of the trip, the locations that the traveller is and is not interested in visiting, the types of activity the traveller likes and dislikes, the traveller's accommodation requirements, the duration of the trip and the budget available for the trip.

The role of the Activity Planner in ETP is to generate an itinerary. It builds this itinerary by matching the traveller profile with available travel products (i.e. airline flights, bus and rail journeys, accommodation, tours and activities). The data-gathering components of ETP act as a filter for identifying the travel products that are available at the time of the trip and which can be considered for an itinerary. The data is drawn from a database of many accommodation options and activities, and transport options are drawn directly from transport schedules.

ETP also provides the user with a natural-language explanation (or justification) of each proposed itinerary. These explanations are formulated relative to the traveller profile expressed by the user request, and are constructed using information gathered during the execution of the itinerary-generating algorithms.

We concentrate our attention on the Activity Planner component of ETP, and in doing so adopt an OR-oriented perspective of the prototype ETP technology. We present models for traveller desires and preferences, the criteria applied for measuring itinerary "goodness", and the scheduling models used in itinerary generation. From this basis we give an overview of the Activity Planner optimisation problems, and outline our approach to solving these problems using a combination of ideas gleaned from the areas of scheduling, routing, assignment and constraint programming.

The Activity Planner generates answers to the fundamental travel-related questions of "where, what and when". It selects the locations visited by the traveller ("where"), decides which tours and attractions are to be taken in by the traveller ("what"), and determines a detailed timetable for the chosen travel and tourism activities ("when"). In doing so it tackles a series of difficult and highly-constrained optimisation problems, each of which undertakes a combination of selection, assignment, scheduling and multi-criteria decision-making functions.

The Activity Planner translates the user request into a set of constraints and objective-function components. Each tourism product is associated with a set of *attributes*. The set of available attributes varies between travel products, and there are four broad categories of product: transport, locations and localities, activities and accommodation. For example, the set of transport attributes includes the elements "business class" and "rail journey", and the set of accommodation attributes includes "swimming pool", "disabled access" and "four star".

The user request explicitly assigns a *preference level* to a selection of attributes, and default values are adopted for the attributes that are not part of the user request. There are six preference levels, specified on a scale of increasing desirability, ranging from "forbidden" through to "mandatory". For many attributes, the default preference level is "permitted". Locations and localities, for example "Island of Hokkaido", "Tokyo" and "Central Australia", are not marked with attributes. Rather, preferences for these are specified directly by the user request, either implicitly or explicitly.

The extreme preference levels translate into hard constraints that strictly force travel products either in- or out- of the itinerary-generation process. The intermediate preference levels give rise to objective function components (*attribute response functions*) that either reward or penalise an itinerary for possessing quantities of that attribute. The attribute response functions can be non-linear. The values of the attribute response functions are summed to give a quantity termed the *gathered attribute value (GAV)*. The GAV reflects the "enjoyment value" of the itinerary.

The budget-related information in the user request is translated into a utility function that expresses, for a given total itinerary cost, the GAV value that the traveller would "expect". The overall objective function value is the difference between the GAV and the utility function value, and we seek to maximise this value. The utility function is parameterised by the user request and is the sum of a linear function and an exponential "barrier function" that becomes increasingly significant as the itinerary cost approaches the traveller's upper budget limit.

Itinerary construction is governed by a series of hard rules that capture the "system constraints" of travel. For example, a traveller must have booked accommodation for every night that they are not travelling, they must spend time resting before and after travelling between locations, and a traveller may only undertake one activity at a time.

An itinerary is composed of three scheduling layers. The locational layer specifies the location occupied by the traveller, or travel on a transport service, for each instant of the trip. The accommodation layer specifies the accommodation for each night that is not spent on a transport service. The activities layer schedules day tours and other activities between inter-location travel intervals.

The "where" question is addressed by an algorithm that combines TSP tour-scheduling with decision-tree searching. The root-node sequence consists of the user-specified first and last locations. Constraint programming approaches are used to weight and rank candidate locations, and the locations are inserted one- or two- at a time into the sequence of locations. The TSP algorithm seeks to minimise expected travel cost.

Each trial sequence is then used to create a sub-problem, the solution of which involves selecting transport services between locations and determining the time of travel (the "when" question), and scheduling touring activities and selecting accommodation (the "what" question). There is a strong interaction between these problems, both through the constraints of space and time, and the behaviour of the objective function. For example, an itinerary that specifies more time in a location with less-expensive accommodation and activities will save cost, yet allocating more time to an interesting location may result in the gathering of greater attribute value over the same time interval.

The "when" and "what" questions are addressed by a nested series of algorithms that operate on specific parts of the problem. The outermost algorithm selects transport services and touring activities. It forms a priority list of the options for each, where the priority depends on several factors. Ideas established within the constraint programming literature are used within a module that computes an influential factor in the priority value. In this module a selection of the problem constraints are analysed and the criticality of the inclusion or exclusion of a particular activity is determined. The priorities are dynamically updated during the execution of the outermost algorithm.

Elements are selected from the priority list and introduced into the itinerary. The introduction of an element spawns a process termed *location scheduling*, where the activities assigned to each location are scheduled according to their availability timetables. The OR methods used during location scheduling include bisection search, auction algorithms and network optimisation algorithms. Accommodation is selected as part of the location scheduling process, and "free time" in the itinerary is allocated to locations in order to maximise the objective function value in the neighbourhood of the current solution.

Computational testing of the prototype ETP has shown that further development iterations will be required in order for achieve computation-time performance that is suitable for, say, an on-line web-based application of the technology. Such an application is one of a series of potential business opportunities for the technology. For modestly sized input data sets, computation times can extend to several minutes, although many trials yielded results in a matter of seconds.

Our experience with the ETP itinerary-planning problem has clearly highlighted to us the need for thoughtful modelling of the needs, preferences, interests and values of travellers: it is the travellers who are the ultimate judges of a solution. Our personal observations of the quality of the itineraries that can be generated using ETP allow us to conclude that automated itinerary planning is an achievable goal for real world applications. Furthermore, the development of ETP has led to the identification and solution of a sizeable number of itinerary-planning sub-problems, each of which offer both research and commercial opportunities.