

The Use of Optimistic and Pessimistic Resource Profiles to Inform Search in an Activity Based Planner

Brian Drabble & Austin Tate
Artificial Intelligence Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN
United Kingdom

Tel: (+44) 31 650 2732 Fax: (44) 31 650 6513

Email: B.Drabble@ed.ac.uk & A.Tate@ed.ac.uk

Abstract

Resource reasoning has been at the heart of many of the successful AI based scheduling systems – yet no attempt has been made to integrate the best techniques from scheduling with the best techniques from AI activity based planning. This paper presents a set of incremental algorithms which create two separate profiles to represent the optimistic and pessimistic use of resources within an activity plan. These allow the planner to ensure that there is a feasible assignment of resources available within any plan state being considered. The paper demonstrates how these profiles can be used to track the usage of a variety of different resource types and how they can be used to provide detailed and relevant information when a resource constraint conflict is detected.

1 Introduction

Resource reasoning has been at the heart of many of the successful AI based scheduling systems – yet no attempt has been made to integrate the best techniques from scheduling with the best techniques from AI activity based planning. The reason for wishing to reason about resources in an activity based planner is clear. One of the prime motivations for not considering a particular course of action is that you have insufficient resources with which to carry it out. These resources can vary from people, to money, to space in a car park. Resource reasoning provides a powerful way of pruning the search space and guiding the planner towards a successful plan.

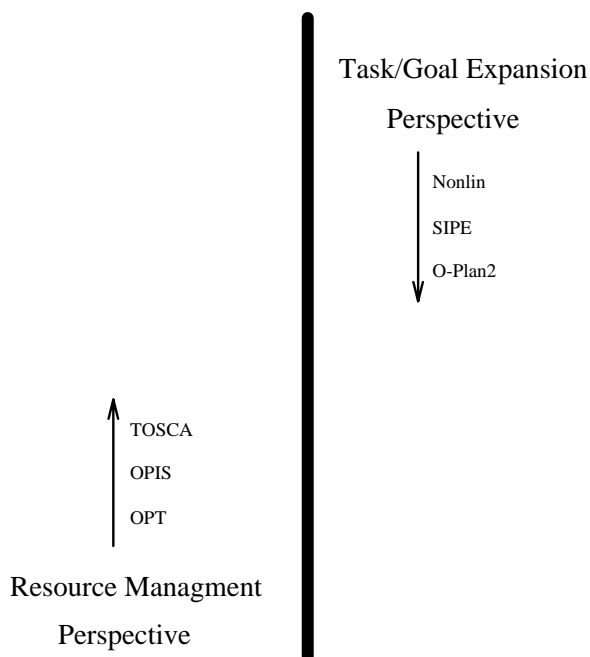
Scheduling problems have tended to be dominated by complex resource contentions and relatively simple process plans whereas activity plans have tended to have complex process options with simple resource uses. However, this view of planning and scheduling as being two separate problems has been enforced by the different approaches of AI researchers and not by the nature of the problem itself. Activity based planning and resource based scheduling can be viewed as opposite ends of a continuum with the middle area being of particular interest in real world applications. This middle area of the continuum contains problems such as:

- project management for product introduction, systems engineering, construction, process flow for assembly, integration and verification, etc.
- planning and control of supply and distribution logistics.
- mission sequencing and control of spacecraft such as Voyager and ERS-1.

Activity based planning systems have attempted to address some of the problems in reasoning about resources – NOAH [6], Nonlin+ [9],[12] and SIPE-2 [14] are examples, but they have had limited success. At the same time resource based schedulers such as OPIS [8], Micro-Boss [7] and TOSCA [1] have attempted to use more complex process plans than those used in earlier scheduling systems. Again these attempts have met with limited success. There have also been attempts to handle a richer model of resources primarily within a constraint management framework such as in CAMPS [2, 13], AMPS [3, 5] and EMPRESS [4]. Figure 1 shows how these differing systems can be related in the maturity of their approaches to resource reasoning and handling process plans.

This paper describes an approach to resource reasoning which takes the idea of a rich resource model as developed in AI based scheduling systems and presents a series of incremental algorithms which allow such a resource model to be used in an activity based planner framework. The techniques described in this paper

Complex Process Options with
Simple Resources Uses



Large Resource Contention Dominated
Problems with Simple Process Plans

Figure 1: The Continuum of Planning and Scheduling Systems

are currently being integrated into the Resource Utilisation Manager (RUM) of the O-Plan2 planner [11].

O-Plan2 is aimed to be relevant to the types of problems which were outlined above. O-Plan2 uses a number of *Constraint Managers* to maintain information about a plan while it is being generated [10]. The information can then be used to prune search (where plans are found to be invalid as a result of propagating the constraints managed by these managers) or to order search alternatives according to some heuristic priority. Constraint Managers are intended to provide efficient support to a higher level of the planner where decisions are taken. They do not take any decision themselves. They are intended to provide complete information about the constraints they are managing or to respond to questions being asked of them by the decision making level.

2 Resource Management in an Activity Planner

Resource constraint management within the O-Plan2 system is carried out by a Resource Utilisation Manager (RUM). It is the function of the RUM to check on the levels of resources being used at certain points in the plan. The RUM is informed of resources level changes in a plan by means of Resource Utilisation Entries (RUE's). A RUE can change resource levels in one of five different ways:

1. **Set** a resource level to be a particular value (or within a particular range)
For example, top up a fuel tank to its maximum capacity.
2. **Allocate** a certain amount of resource i.e. reduce the amount of resource remaining as available from that point within the plan. Semantically, an allocation must be paired with a subsequent deallocation.
3. **Deallocate** a certain amount of resource back to the resource pool, i.e. increase the amount of resource available from that point in the plan.
4. **Consume** a certain amount of resource.
5. **Produce** a certain amount of “new” resource.

The initial declaration of resource types (e.g. fuel, food, money, plumbers, etc.) is accomplished by using a `resource_type` definition in the O-Plan2 domain description language (Task Formalism – TF), together with the information required to define that resource. For example,

```
types
  fuel_loc = (port1 port2 port3 ship1),
  fuel_storage = (tank1 tank2 tank3 tank4 tank5),
  prov_type = (frozen chilled fresh),
  prov_loc = (port1 port2 port3 port4),
  prov_storage = (warehouse1 warehouse2 warehouse3);

resource_types
  consumable_producible_by_agent
    {resource fuel ?{type fuel_loc}
      ?{type fuel_storage}} = gallons,
  consumable_producible_by_agent
    {resource provisions ?{type prov_type}
      ?{type prov_loc}
      ?{type prov_storage}} = kilos;
```

The actual usage and setting of resource levels in the plan is achieved by RUE's which are derived from **resource** statements in TF action schemas. These provide changes to resources levels i.e. increments (**produces, deallocates**), decrements (**consumes, allocates**) and **sets**.

The RUM maintains resource usage profiles that reflect the changes of resource levels indicated by the RUEs. There can be uncertainty in two dimensions: in the actual level of resource changes and in the time at which such a change occurs. The RUM manages resource usage profiles in order to provide the following functionality for the planner:

- **Adding a new resource utilisation into a resource profile**
As actions are expanded in the plan new resource utilisations will need to be added to the resource profile. The RUM will need to constrain the resources affected and monitor for resource violations.
- **Modifying an existing resource utilisation entry**
Existing resource entries will be modified during the plan as their time and resource windows are constrained by other activities. The RUM will propagate the effects of these changes through only those resource entries affected.
- **Providing feedback when a constraint violation occurs**
The RUM is able to provide specific advice relevant to the particular problem which has arisen. By using the type of a resource to restrict the options proposed, the RUM can suggest altering the resource levels in other related resource entries and/or modifying the time constraint of related resource entries.

3 Management of Resource Specification and Aggregate Resource Usage

Resource information in O-Plan2 action schemas is used to to restrict search and to ensure that resource usage in a plan stays within the bounds indicated. There are two types of resource statements. One gives a *specification* of the **overall** limitation on resource usage for a schema (over the total time that the schema's expansion can span). The other type of statement describes actual resource *utilisation* **at** points in the expansion of a schema. It must be possible (within the flexibility admitted by the actual resource utilisation statements) for a point in the range of the aggregate of the resource utilisation statements to be within the overall resource specification given.

1. the **specification** of the limits on resources used within a schema and all its possible expansions. For example, a schema to move a ship from one port to another may specify that it may consume between 100 and 1000 gallons of fuel depending on which ship and which pair of a specified set of ports is chosen.

```
resource consumes (resource fuel) = 100:1000 gallons overall;
```

2. the **utilisation** of resources on a particular action or at a particular time point within a schema. In this example ship1 receives 5000 gallons of fuel into its single fuel tank at the end of action 5 from tank1 at port1.

```
resource produces (resource fuel ship1 tank1) = 5000 gallons
                    at end_of node-5,
consumes (resource fuel port1 tank1) = 5000 gallons
                    at end_of node-5;
```

All resource specifications and utilisations are maintained as **min/max** pairs, specifying the upper and lower bounds known at the time. Resource declarations which describe resource specifications and utilisation statements (perhaps still only partially specified) are held in the plan being developed by O-Plan2. The current best numerical bounds on resource utilisation statements are also converted to RUE's which are stored in a Resource Utilisation Table (RUT) with (notionally) one table per specific resource available. The entries of the RUT are held in ascending time point order. The following table (Table 1) is a fragment from a RUT for a specific fuel tank.

The entries within the RUT are fully qualified entries and as such represent actual resource utilisation. A schema which states that it produces 500 gallons of fuel from port1 is viewed as a *specification* as the actual change in a resource cannot be specified relative to a port – but only for a specific fuel tank at a location (a port or a ship).

No	Type	Resource	Quantity	Time Point	Min	Max
1	*	(resource fuel port1 tank1)	20:20	tp1	0	0
2	-	(resource fuel port1 tank1)	20:30	tp2	4	8
3	+	(resource fuel port1 tank1)	15:15	tp19	6	6
4	*	(resource fuel port1 tank1)	10:15	tp36	7	7

Table 1: Example Resource Utilisation Table

4 Optimistic and Pessimistic Resource Profile Management

The algorithm used to track resource levels uses two distinct measures:

1. Optimistic Resource Profile (ORP)

This describes the maximum resource that could be available with optimistic assumptions and is calculated from:

- (a) the **set** resource statements
- (b) the minimal resource usage at the maximum time value of a time point for an RUE with negative influences, i.e. **allocates, consumes**.
- (c) the maximal resource usage at the minimum time value of a time point of an RUE with positive influences, i.e. **deallocates, produces**.

For example, if action 1 **allocates** between 20 and 30 resource units between time 4 and time 8 then the ORP normally decreases by 20 at time 8 (unless a **set** is given at the same time point).

2. Pessimistic Resource Profile (PRP)

This describes the minimum resource that would be available with pessimistic assumptions and is calculated from:

- (a) the **set** resource statements
- (b) the maximal resource usage at the minimum time value of a time point of an RUE with negative influences, i.e. **allocates, consumes**.
- (c) the minimal resource usage at the maximum time value of a time point of an RUE with positive influences, i.e. **deallocates, produces**.

For the above ORP example, the PRP normally decreases by 30 units at time 4 (again unless a **set** is given)

By calculating the changes in anticipated resource levels at specified points along a time line, a profile can be generated for the ORP and the PRP. Using the first three entries of the RUT described in Table 1, the following graph (Figure 2) of resource levels against time can be generated.

To generate the profile the RUM needs to keep track of various pieces of information and to be sensitive to the type of change which is being carried out for each RUE in the RUT. The changes which the RUM must deal with are the addition of a new RUE or the modification of an existing RUE. The information which is maintained is as follows:




Figure 2: Optimistic and Pessimistic Profiles of Resource Utilisation

1. Optimistic Increment (**OptInc**) which is defined as the incremental change in the level of resource at a particular time point ignoring **sets**. It is calculated from summing:
 - (a) If the time point in question is the maximum time point of an **allocates** or **consumes** then add in the minimum change in resource for each record
 - (b) If the time point in question is the minimum time point of a **deallocates** or **produces** then add in the maximum change in resource for each record.

2. Pessimistic Increment (**PesInc**) which is defined as the incremental change in the level of resource at a time point ignoring **sets**. It is calculated from summing:
 - (a) If the time point in question is the maximum time point of a **deallocates** or **produces** then add in the minimum change in resource for each record
 - (b) If the time point in question is the minimum time point of an **allocates** or **consumes** then add in the maximum change in resource for each record.

3. A *base* value of PRP to assist in incrementally computing the actual PRP.
4. A *base* value of ORP to assist in incrementally computing the actual ORP.
5. Whether a **set** is involved as one of the RUE's at the time point.
6. Dependency records containing lists of RUE's affected by resource information at time point.

Formulae to Maintain PRP and ORP

The *base* value for PRP and ORP and the ORP and PRP themselves can be calculated for any time point using the following formulae:

FORMULA 1 - PESSIMISTIC PROFILE:

```

IF one or more set entries are present THEN
  IF there are over lapping sets THEN
    PRP = minimum of all overlapping sets
  ELSE
    IF there are overlapping deallocates or produces then
      base_PRP = base_PRP + the maximum resource value for all
        overlapping deallocate or produce RUE's
    ELSE
      base_PRP = minimum of the minimum of the set value
    ENDIF
  ENDIF
  PRP = base_PRP
ELSE
  base_PRP = PRP at a previous time point in the RUT or
    0 if none available
  PRP = base_PRP + PesInc
ENDIF

```

FORMULA 2 - OPTIMISTIC PROFILE:

```

IF one or more set entries are present THEN
  IF there are overlapping sets THEN
    ORP = maximum of all overlapping sets
  ELSE
    IF there are overlapping allocates or consumes then
      base_ORP = base_ORP + the maximum resource value for all
        overlapping allocate or consume RUE's
    ENDIF
  ENDIF

```

```

ELSE
    base_ORP = maximum of the maximum set value
ENDIF
ENDIF
ORP = base_ORP
ELSE
    base_ORP = ORP at a previous time point in the RUT or
                0 if none available
    ORP = base_ORP + OptInc
ENDIF

```

Example when a Resource “Set” is Involved

The above formulae will now be demonstrated on an example in which there are positive or negative resource changes which may occur within the time range of a set (e.g. see Figure 3).




Figure 3: Set Spanning a Resource Change

In this example the allocation of the resource will take place some time between $t=2$ and $t=8$ and will allocate between 4 and 6 units. However, there are two sets which may occur in parallel with this allocation. This is further complicated for illustrative purposes in this example by including a deallocate. By using the algorithms described above it is possible for the RUM to construct the PRP and ORP profiles and to find that there is at least some possible allocation of resources which is valid.

5 Detection of Resource Utilisation Failures

The failure of the addition of a RUE or during propagation of RUE entries represents an attempt by the plan to use more of the resource than there is available. The failures types which have been identified so far are as follows:

1. **ORP less than ZERO** This failure means the even with that most optimistic assumptions there is insufficient resource available.
2. **ORP less than PRP** This failure means the resource utilisation has been declared incorrectly within the domain description (TF) definitions.

The RUM informs the planner of the RUE which has a fault and the possible tactics available to resolve the conflict. These are as follows:

1. increase the earliest start time of a failing action i.e. start it later.
2. alter the latest finish or earliest start time of possible actions which contribute to the problem. This will depend on whether the actions are taking or giving back a resource. It may make more sense to give some resource back earlier or take a resource later (if time constraints allows) rather than reduce your own resource utilisation.
3. increase the maximum resource level available at a point by adding a **set** of a particular resource. For example, if the authority can be found an extra shift of workers or an extension to the working day may resolve the resource problem.

The actual tactics proposed are sensitive to the resource type for the RUE involved.

6 Summary

This paper has described a mechanism for the incremental management of optimistic and pessimistic resource usage profiles in an activity planning framework. A rich resource model can be handled which can manage uncertainty in the time at which resources are used and the absolute resource levels involved in any resource level change.

The technique allows for an AI planner to check the feasibility of resource availability for plans being considered in the search for a solution. The techniques allows for the maintenance of resource usage profiles within which a specific resource allocation should be possible.

Acknowledgements

Current O-Plan2 work is supported by the US Advanced Research Projects Agency (ARPA) and the US Air Force Rome Laboratory acting through the Air Force Office of Scientific Research (AFSC) under contract F49620-92-C-0042. The United States Government is authorised to reproduce and distribute reprints for government purposes notwithstanding any copyright notation hereon. The project is monitored by Dr. Northrup Fowler III at the USAF Rome Laboratory.

All views expressed are those of the authors only. The information contained in the paper has benefited from discussions with many talented researchers who have worked on the O-Plan and related projects. Thanks are due for productive discussions to our co-workers on the O-Plan2 project: Jeff Dalton and Glen Reece.

References

- [1] Beck, H. *TOSCA: A Novel Approach to the Management of Job-shop Scheduling Constraints*, in Realising CIM's Industrial Potential: Proceedings of the Ninth CIM-Europe Annual Conference, Amsterdam, 12-14 May 1993, (eds, C. Kooij and P.A. MacConaill and J. Bastos), pp138-149.
- [2] Brown, R., *Knowledge-based Scheduling and Resource Allocation in the CAMPS Architecture*, in Proceedings from the International Conference on expert Systems and the Leading Edge in Production Planning and Control (ed.. M. Oliff.), Benjamin/Cummings, Menlo Park, CA 1987.
- [3] Dawson, B., Day, D.S. and Mulvehill, A., *The AMPS Final Report*, Final Report, USAF Rome Laboratory Technical Report RADC-TR-90-131. July 1990.
- [4] Hankins, G.B., Jordan, J.W., Katz, J.L., Mulvehill, A.M., Dumoulin, J.N. and Ragusa, J., *EMPRESS: Expert Mission Planning and REplanning Scheduling System*, in Expert Systems in Government Symposium, 1985.
- [5] Mulvehill, A., *CAMPS/AMPS FY88 Year End Report*, The MITRE Corporation, Technical Report M89-28, May 1989.
- [6] Sacerdoti, E., *A Structure for Plans and Behaviours*, Artificial Intelligence Series, North Holland, 1977.
- [7] Sadeh, N., *Look-ahead Techniques for Micro-opportunistic Job Shop Scheduling*, Ph.D., CMU-CS-91-102, School of Computer Science, Carnegie Mellon University, 1991,

- [8] Smith, S.F., *A Constraint-Based Framework for Reactive Management of Factory Schedules*, in Proceedings International Conference of Expert Systems and the Leading Edge in Production Planning and Control, Charleston, South Carolina, May, 1987
- [9] Tate, A., *Generating Project Networks*, in Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-77), Cambridge, Mass., USA, 1977.
- [10] Tate, A., *The Emergence of "Standard" Planning and Scheduling System Components*, in Proceedings of the Second European Workshop on Planning (EWSP-93), IOS Press, Sweden.
- [11] Tate, A., Drabble, B. & Kirby, R., *O-Plan2: An Open Architecture for Command, Planning and Control*, in *Intelligent Scheduling* (Eds. M. Fox and M. Zweben), Morgan Kaufmann.
- [12] Tate, A. & Whiter, A., *Planning with Multiple Resource Constraints and an Application to a Naval Planning Problem*, in Proceedings of the First Conference on Artificial Intelligence Applications, pp. 410-416, AAAI, Denver, Colorado, USA, December 1984.
- [13] Zweben, M., *CAMPS: A Dynamic Replanning System*, The MITRE Corporation, Technical Report M87-50, December 1986.
- [14] Wilkins, D., *Practical Planning*, Morgan Kaufmann, 1988.