New Planning Benchmarks: Vehicle Routing Problem Variants

Peter Gregory
University of Strathclyde
Glasgow, UK
peter.gregory@cis.strath.ac.uk

Abstract

The Vehicle Routing Problem (VRP) is a classical problem in Operations Research. It is a generalisation of the Travelling Salesman Problem in which a set of locations must be visted by one of a number of vehicles. All vehicles start and end at a depot location. The aim of the problem is to minimise the total distance covered by the trucks.

There are many different variants of the VRP. For example, Capacitated VRP, VRP with Time Windows, VRP with Multiple Depots and VRP with Pickup and Delivery. All of these variants are interesting planning problems, as they deal with resource management, time windows, routing, etc.

In this work, we model some of the variants of VRP in PDDL. We provide this as an intersting new benchmark set for planning. Already, many planning benchmarks are based in some way on transportation. Modelling VRP directly allows access to a great quantity of existing benchmark instances. This has the benefit of posing difficult and realistic challenges to existing planners.

Introduction

The Vehicle Routing Problem (Dantzig and Ramser 1959; Clarke and Wright 1964; Laporte 2009) is a problem from Operations Research that has many practical applications in industry. In the planning community, each International Planning Competition has included one, or more, transportation-style (Helmert 2003) domain. Examples include in IPC 2, Logistics; in IPC 3, Driverlog, Zeno-Travel and arguably even Rovers. More recent competitions have seen the TPP, Trucks and Transport domains. In essence, transportation-style problems form natural planning problems.

We investigate translating several variants of the Vehicle Routing Problem into PDDL. We find that there is a very compact model that can be created that requires only minimal changes in order to represent the different VRP variants. The planning model we create is very flexible and can be modified to represent new constraints very quickly.

We show differences in the modelling assumptions in the Operations Research and planning communities. We show that the performance of variaus planners on the benchmarks.

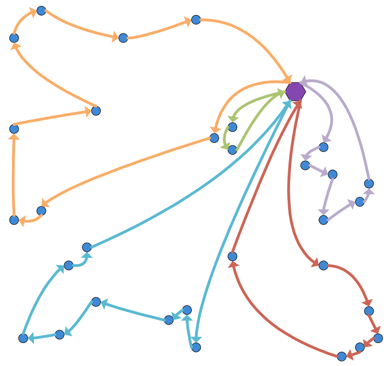

Figure 1: A Solution to a Vehicle Routing Problem instance

The VRP and its Variants

The basic vehicle routing problem can be defined in the following way:

Definition 1 (Vehicle Routing Problem) *An instance of the Vehicle Routing Problem has the following components:*
V*: a set of vehicles*
L*: a set of locations, where* L_0 *is a depot location where all vehicles start*
$d : L \times L \mapsto \mathbb{R}$*: the distance between each pair of locations*
A solution to a VRP instance is a tour for each vehicle, such that all locations are visited, and each vehicle returns to the depot. An optimal solution is a solution in which the total distance of all tours is minimised.

In effect, the Vehicle Routing Problem is the Travelling Salesman Problem with multiple vehicles. Although industrial interest in the VRP is as old as the problem (Clarke and Wright 1964), the VRP is too simplistic an abstraction to correctly model many real transportation problems. Because of this, many variations of the VRP have been studied over time, building on the basic VRP with extra features.

The first of these we discuss is the **Capacitated VRP** in which each vehicle has finite capacity and each location has a finite demand. A valid solution to the Capacitated VRP is one in which, for each vehicle, the sum of the demands of the tour for that vehicle does not exceed the vehicle's capacity.

The optimal solution is simply the shortest solution that does not break the capacity constraints.

Another variation of the VRP is the **VRP with Time Windows** in which there is a specified temporal window of opportunity in which to visit each location. Each location has a service time, in which the vehicle has to wait, in order to unload goods. This variation can naturally be combined with the Capacitated VRP.

Yet another variation, **VRP with Multiple Depots**, generalises the idea of a depot, in such a way that there are several depots from which each customer can be served. Similarly, the **VRP with Satellite Facilities** models the situation when vehicles can be replenished with goods.

In the **Multi-Commodity VRP**, each location has a demand for different commodities. There is more than one definition of Multi-Commodity VRP. In one of these (Repoussis, Tarantilis, and Ioannou 2006) definitions, each vehicle has a set of compartments in which only one commodity can be loaded. The problem then becomes that of deciding which commodities to place in which compartments in order to minimise distance travelled.

There are many other variants described in the literature. From the earliest works on the VRP, until the most recent, researchers have been solving a real-world problem. The different variants have been developed when the simpler methods have proved insufficient to model the underlying problem to be solved.

PDDL Models

Presently, we have two of the VRP variants (CVRP and CVRPTW) encoded in PDDL, with a translator that converts between the standard formats to PDDL problems. These will be made available via the author's webpage, along with encodings of the other VRP variants.

Encodings

Figure 2 shows our model of the CVRP. We believe the model to be intuitive and neat. There are two actions, one that visits a location with a truck, another that drives the truck back to the depot. The level of resource in the vehicle is maintained by the 'space' value, which is decreased by the demand of a location when the vehicle visits that location.

In the initial state, the space in the trucks is assigned to the maximum capacity. The goal in the problem is to have visited all of the locations, and have all trucks back at the depot ('home' in the model).

We omit the most of the CVRPTW domain definition, however, we do include the visit action in Figure **??** for an idea of the differences from CVRP. The key difference between the two models is the addition of an 'open' predicate, which is a precondition of visiting a location. Each location has an 'open' predicate, which is activated (and later deactivated) as a timed initial literal. In fact, CVRP problems can be modelled in the CVRPTW domain, simply by asserting all locations open in the initial state, corresponding to time-windows that are open forever. We intend to provide a unifying domain in which all VRP variants can be modelled. However, for the purposes of carrying out experiments, we

```
(:types place vehicle)

(:predicates
 (home ?v - vehicle)
 (not-visited ?p - place)
 (visited ?p - place)
 (at ?v - vehicle ?p - place)
 )

(:functions
 (distance ?p1 - place ?p2 - place)
 (quantity ?v - vehicle)
 (demand ?p - place)
 (cost)
 )

(:constants l0 - place)

(:action go-home
 :parameters (?v - vehicle ?pfrom - place)
 :precondition (and
                (at ?v ?pfrom)
               )
 :effect       (and
                (not (at ?v ?pfrom))
                (home ?v)
                (increase (cost) (distance ?pfrom l0))
               )
)

(:action visit
 :parameters (?v - vehicle
              ?pfrom - place
              ?pto - place)
 :precondition (and
                (at ?v ?pfrom)
                (>= (space ?v) (demand ?pto))
                (not-visited ?pto)
               )
 :effect       (and
                (not (at ?v ?pfrom))
                (at ?v ?pto)
                (not (not-visited ?pto))
                (visited ?p - pto)
                (decrease (quantity ?v) (demand ?pto))
                (increase (cost) (distance ?pfrom ?pto))
               )
)
```

Figure 2: PDDL Description of the CVRP.

```
(:durative-action visit
 :parameters (?v - vehicle ?pfrom - place ?pto - place)
 :duration (= ?duration (+ (distance ?pfrom ?pto)
                           (servicetime ?pto)))
 :condition (and (at start (at ?v ?pfrom))
                 (at start (>= (quantity ?v) (demand ?pto)))
                 (at end (open ?pto))
                 (at start (not-visited ?pto))
            )
 :effect (and (at start (not (at ?v ?pfrom)))
              (at end   (at ?v ?pto))
              (at start (not (not-visited ?pto)))
              (at end   (visited ?pto))
              (at end   (decrease (quantity ?v) (demand ?pto)))
              (at end   (increase (cost) (distance ?pfrom ?pto)))
              ))
```

Figure 3: PDDL Description of the visit action in the CVRPTW domain.

have produced seperate domain files. This is because we wish to test on the largest number of planners, and not all planners support temporal domains, for example.

Comparisons with Existing Domains

In constrast to typical transportation domains, our model has no explicit load and unload actions. This is because each vehicle can be assumed to be fully loaded when it leaves the depot, as there is an implicit assumption that there are enough resources at the depot to service all orders.

In all current transportation domains, the underlying map is typically represented as a connected undirected graph. In the VRP benchmark instances, each location is modelled as a point on a two-dimensional plane, and the distance between the two locations is simply the Euclidian distance between the points. This means enforcing the quadratic number of distances between points in the initial state in the planning model.

In the VRP variants, each location is visited exactly once. Therefore, our model maintains a 'not-visited' predicate for each location that is deleted once the location has been serviced. This disallows other trucks from visiting that location. The assumptions made in the VRP variants (that vehicles start and end in a depot, that each location is visited once, that trucks do not deliver partial orders) are made for two reasons: firstly, they are assumptions that simplify the problems enough so that large real-world instances can be tackled. And secondly, they accuaratelly reflect reality enough to be useful.

It might be argued that the VRP is not really an idiomatic planning problem. One reason for believing this argument is that all goals in the VRP are very loosely coupled, and there are very limited causal relationships between the actions. Whilst this case can be argued for the VRP, for the more complex cases of CVRP, CVRPTW and beyond, it is much more difficult to argue. Vehicle capacities and varying demand at locations compicates the decision of which routes to take. The ability to model time windows is a relatively recent innovation in PDDL, and was intended to model problems just like the CVRPTW. The more structure that is added to the VRP, the more it looks like an idiomatic planning problem.

Empirical Performance

Any planner that can solve the CVRP must satisfy the PDDL requirement 'fluents', for the CVRPTW, planners must also satisfy the 'durative-actions' 'timed-initial-literals'. The planners we tried were LPG-TD (Gerevini, Saetti, and Serina 2005), SGPlan521 (Chen, Wah, and Hsu 2006), LPRPG (Coles et al. 2008), POP-F (Coles et al. 2010) and MIPS-XXL (Edelkamp and Jabbar 2006). Of these planners, SGPlan521 and LPG-TD reliably solve CVRP problems, whilst LPG-TD and POP-F reliably solve CVRPTW problems.

Figure **??** shows the performance of planners on CVRP and CVRPTW benchmark instances. The CVRP instances used (Augerat et al. 1995) range from between 32 to 100 locations, and from five to 15 vehicles. We show the optimal solution costs as a comparison, and see that LPG performs

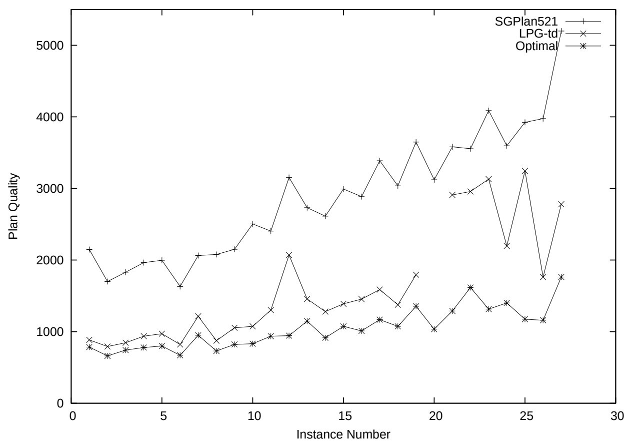

(a) Randomly Distributed Locations: Rural CVRP Benchmark

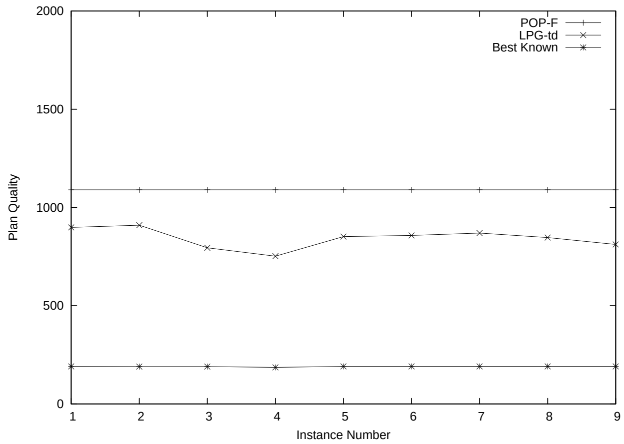

(b) CVRPTW Benchmarks

Figure 4: Plan Quality. We show the quality of the solutions produced by state of the art planners compared to the highest quality results found in the Operations Research literature (Augerat et al. 1995; Solomon 1995).

reasonably well for small instances, however scaling quite badly. SGPlan521, on the other hand, produces low quality solutions more reliably. Indeed, VRP variants have been solved successfully with decomposition-based approaches (Shaw 1998) in the OR community.

In the CVRPTW case, we use the instances from (Solomon 1995). The actual problem set has instances with 25, 50 and 100 locations. Each instance has 25 vehicles available. We only consider the smallest of these, as the larger ones are currently out of reach for planners. Again, LPG performs the best of the planners, although it must be remembered that POP-F is not an any-time planner.

Discussion

It would be unreasonable to expect domain-independent planners to solve as large instances of vehicle routing problems as specific solvers. The instances studied in this work are very small in relation to those currently solved by state of the art vehicle routing solvers. Despite this, in the CVRP problem class, LPG performs reasonably well and for small instances performs close to optimal.

Studying benchmark vehicle routing problems can be valuable to planning in several ways. It provides an almost open-ended challenge as there are standard benchmark instances orders of magnitude larger than those reported here. It challenges the way that we typically model planning problems. For example, it is typical in planning to represent maps as sparsely connected graphs. This makes the problem of finding shortest paths between the locations part of the planning problem. One possible explanation of why we overburden our planning systems with this path-planning is that a sparsely connected graph produces fewer grounded actions than a fully-connected graph. However, the two problems are equivalent, and planners should have solutions to the problem of dealing with large numbers of grounded actions.

Another challenge to PDDL modelling philosophy is that in transportation domains, load and unload actions tend to be modelled explicitly. However (except in some special cases), in vehicle routing, this is not the norm. The PDDL transportation domains purport to model logistics-style problems, but lift what may appear to be assumptions (that all trucks start and end at a depot, that all locations are visited once, etc) that are often, in fact, real-world problem constraints.

Conclusions

We have presented the VRP, and its variants, as an interesting new benchmark class for planning. Studying the VRP is useful for several reasons:

- It is a difficult real-world problem in which planners are still relatively weak.
- There are existing benchmarks for which we can assess the performance of planners against existing techniques.
- Technological developments that have made VRP solvers effective could also improve the performance of planners.

It is unrealistic to expect planners to perform at the same level as problem-specific solvers. It is realistic, however, to expect planners to solve problems that are structured in the same way as real-world problems. We believe that this new set of benchmarks can help in achieving this goal.

References

Augerat, P.; Belenguer, J.; Benavent, E.; Corbern, A.; Naddef, D.; and Rinaldi, G. 1995. Computational results with a branch and cut code for the capacitated vehicle routing problem. Technical Report 949-M, Universite Joseph Fourier, Grenoble.

Chen, Y.; Wah, B. W.; and Hsu, C. 2006. Temporal Planning using Subgoal Partitioning and Resolution in SGPlan. *Journal of Artificial Intelligence Research* 26:323–369.

Clarke, G., and Wright, J. 1964. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research* 12(4):568–581.

Coles, A. I.; Fox, M.; Long, D.; and Smith, A. J. 2008. A hybrid relaxed planning graph-lp heuristic for numeric planning domains. In *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS 08).*

Coles, A. J.; Coles, A. I.; Fox, M.; and Long, D. 2010. Forward-chaining partial-order planning. In *Twentieth International Conference on Automated Planning and Scheduling (ICAPS 10).* AAAI Press.

Dantzig, G., and Ramser, J. H. 1959. The Truck Dispatching Problem. *Management Science* 6(1):80–91.

Edelkamp, S., and Jabbar, S. 2006. Cost-Optimal External Planning. In *AAAI*, 821–826.

Gerevini, A.; Saetti, A.; and Serina, I. 2005. Integrating Planning and Temporal Reasoning for Domains with Durations and Time Windows. In *IJCAI*.

Gerevini, A.; Saetti, A.; and Serina, I. 2006. An approach to temporal planning and scheduling in domains with predictable exogenous events. *J. Artif. Intell. Res. (JAIR)* 25:187–231.

Helmert, M. 2003. Complexity results for standard benchmark domains in planning. *Artif. Intell.* 143(2):219–262.

Laporte, G. 2009. Fifty years of vehicle routing. *Transportation Science* 43(4):408–416.

Repoussis, P. P.; Tarantilis, C. D.; and Ioannou, G. 2006. A hybrid metaheuristic for a real life vehicle routing problem. In Boyanov, T.; Dimova, S.; Georgiev, K.; and Nikolov, G., eds., *Numerical Methods and Applications*, volume 4310 of *Lecture Notes in Computer Science*, 247–254. Springer.

Shaw, P. 1998. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In *CP*.

Solomon, M. M. 1995. Algorithms for the Vehicle Routing Problem with Time Windows. *Transportation Science* 29(2):156–166.