Dynamic Coordination in Vehicle Routing

Catherine Rivers
Institute of Information Sciences and Technology
Massey University
New Zealand
crivers@clear.net.nz

Abstract

This study investigates vehicle routing where some of the node attributes of location, demand and time of availability are unknown at the start of the work cycle. The situation is dynamic with an element of stochasticity. Coordination in dynamic vehicle routing is implemented by estimating unknown attributes and developing contingency plans to cope with differences between estimated and actual values. Such implementation can have a beneficial influence on the costs of vehicle routing by reducing the total travel time or by rendering unnecessary an increase in fleet size that would have catered for unknown demand.

1 Introduction

In practice, it is not uncommon for customers in a vehicle routing scheme to change their demands, the times at which their demands become known or even their very existence.

This paper investigates coordination [2] where at least some of the node attributes of location, demand and time are unknown at the start of the work cycle. It aims to show that coordination exists in a dynamic environment and that implementation of coordination can have a beneficial influence on the costs of vehicle routing by reducing the total travel time or by obviating an increase in fleet size due to catering for unknown demand.

The scheme using dynamic coordination can be compared with the non-coordinated dynamic scheme and with the coordinated scheme that would have existed if all attributes had been pre-determined.

1.1 Nomenclature

Each situation is categorised by the prior knowledge (or lack of it) of the node attributes mentioned above. Nomenclature is shown in Table 1. Possible approaches for each situation are outlined below. Following this, three representative situations will be investigated in more detail in a full simulation environment.

location	demand	time	
√	√	√	Not dynamic
√	√	×	DC1
√	×	√	DC2
×	√	√	DC3
√	×	×	DC4
×	√	×	DC5
×	×	√	DC6
×	×	×	DC7

Table1. Dynamic coordination classification.

DC1

Where only the time is unknown, the node in question is included in routes that approach it at different times. As soon as the demand exists, the best of the tentative routes services it; the remainder of the tentative routes are amended, and dynamic checking no longer exists.

DC2

The demand of a particular node is not known in advance, only that the node will have a demand. In this case, tentative planning deals with a projected demand equivalent to

- the last demand,
- the average of the last few demands,
- the most common of the last few demands, or
- some other estimate

dependent on the number of previous demands by this node and the pattern which those demands make (if any). For example, if there are fewer than 4 previous occurrences, use the last occurrence; if there are more, search for a pattern. If no pattern emerges, use the average of the last four. If there have been no previous demands from this node, a similar projection could be made based on like nodes (if any) or on the average of all nodes.

DC3

Only the location is unknown. If possible, an estimate could be made on the likely location. Otherwise, routes could leave spare capacity to accommodate the known demand. As soon as the location is revealed, the best route to service it can be confirmed and the others cleared.

DC4

Only the location is known. A combination of the approaches used by DC1 and DC2 might be helpful. Knowing the location enables the node to be included in several routes. This has the advantage that, where the demand exceeds projections, other vehicles may be able to carry surplus.

DC5

Only the demand is known. Carriage can be guaranteed by reducing the capacity of all vehicles by the expected demand. This essentially includes the node in all routes, which may not be sensible since distant routes would be affected even though there was no chance of their involvement.

DC6

Only the time is known. If the demand can be projected, spare capacity can be reserved in vehicles for it, as in DC5. The quality of the projection is vital to the success of this approach. This is a time-limited method; after the time at which the demand is revealed, all routes can be confirmed.

Reduction of vehicle capacity in DC5 and DC6 may not be realistic where there are many routes that will not be involved in servicing the unknown node but which will be penalised unnecessarily in their carrying capacity.

DC7

Nothing is known in advance. An unexpected node is revealed with a demand after the routes have been scheduled. At the time of its disclosure, routes near it are investigated to determine their suitability for detouring to service all or part of its demand. A decision needs to be made on which of the existing load or the revealed demand has priority. Normally, the new demand would have to fit in with the existing load without displacing it.

2 Representative situations

Three situations are investigated in detail: in DC2 only one of the three attributes, the demand, is unknown; in DC4 two attributes are unknown and only the location is known; in DC7 none of the three attributes is known.

2.1 DC2 Unknown demand

2.1.2 Problem generation

A single depot/ fifteen node system is generated in a rectilinear grid [3]. The attributes of Nodes 1 - 14 are fully determined, Node 15 has a known location and a demand that will become apparent at time t=15. Node 15 has had only one previous demand of value 50. Although coordination is possible in the form of split deliveries [1], that technique will not be applied in the initial construction phase since it might serve to cloud the effect of the dynamic nature of the demand. There is no load/unload time at nodes.

2.1.3 Problem solution

An estimation is made of the value of the unknown demand. After initial routes are determined, a coordination plan is constructed and the preliminary estimation is adjusted to a greater value; the sizes of loads delivered to and collected from coordination points are modified. At a given time, the actual demand of all nodes is revealed and the necessity of applying the coordination is determined.

Solution is a four step process:

Step One	Initial routes are constructed and scheduled using an estimated value for the unknown demand.
Step Two	Routes are adjusted so that the node in question is not visited before the time at which its true demand is revealed.
Step Three	Coordination points are then established at likely sites so that they will be available for use when the mystery demand is revealed.
Step Four	At the time of revelation, routes are adjusted to accommodate, if possible, the known demand. Coordination sites are emptied before the end of the work cycle.

Table 2. Dynamic coordination

2.1.4 Demand estimation

The unknown demand is estimated by considering the number of past occurrences of demand from that particular node:

- >3: use the average of the last four demands
- <4: use the latest demand
- no previous demands: use the average of all demands for the known nodes.

Where there is no pattern, using the last four values accommodates extremes of demand.

2.1.5 Coordination site location

Site location consists first of determining the position of all vehicles and the serviced/unserviced state of nodes, and then preparing for large differences in the values of the estimated and actual demands of Node 15.

Since it is known that the true value of Demand(15) will be revealed at time t=15, the locations of all vehicles at that time can be determined in advance (see Table 3). The five vehicles have completed 15 steps of each route and some nodes have already been serviced.

Vehicle	Progress at t=15
Blue	(0) (7) (10) 13 0
Green	(0) (11) 9 8 5 0
Pink	(0) (12) 6 0
Black	(0) (3) 1 15 0
Red	(0) (2) (4) 0 14 0

Table 3. Progress at t=15. Serviced nodes are enclosed in parentheses.

The basis for coordination site location is contingency planning. If the actual demand is no greater than the estimated value, coordination is not necessary, but the possibility of an actual value far greater than the estimated value must be considered and planned for in advance of any scheduling. All routes must be considered, since the use of split deliveries can have a ripple effect on distant routes.

Accommodation of a much greater value of Demand(15) involves the creation of two coordination points, X and Y, one of which must be used before the blue vehicle services its original route. This vehicle delivers to X a load with value maxDemand-estDemand(15) before returning to the depot and resuming its scheduled route to nodes 7, 10 and 13. When the actual value of Demand(15) is revealed, another vehicle can collect the required amount from X. The best vehicle to collect the load from X is the pink vehicle which services route 4 in reverse order so that it can collect the balance of actual Demand(15)-estimated Demand(15) as it passes through X en route from node 6 to node 12. Pink delivers the extra load to coordination point Y, where the black vehicle collects it, combines it with estDemand(15) and services node 15. Demand(15) is carried split but delivered intact.

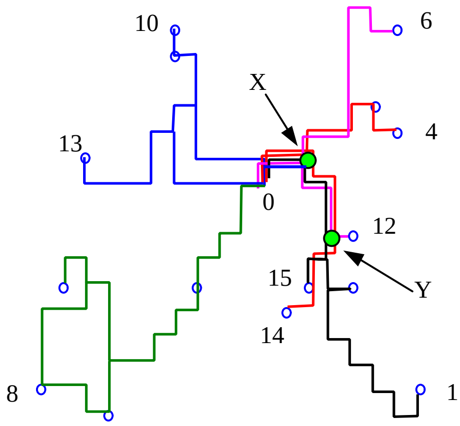

Figure 1. Initial routes showing potential coordination sites (filled discs).

2.1.6 Incorporation of coordination

At time t=15, the actual demand of Node 15 is known to be 150, necessitating the use of coordination. After Demand(15) is revealed as being much greater than the estimated value the routes are those shown in Table 4.

Vehicle	Distance	Route(s)	Nodes
blue	42	1	0 X 0 7 10 13 0
green	50	2	0 11 9 8 5 0
pink	41	4	0 6 X Y 12 0
black	42	5	0 3 1 Y 15 0
red	44	3 6	0 2 4 0 14 0

Table 4. Route structure at time t>= 15.

2.1.7 Limitations and accommodations

Of the three vehicles involved in coordination, only the pink vehicle poses a limitation on the viability of this scheme. It can deliver to Y only maxVehCap-Demand(12) without violating the vehicle capacity constraint. Because the blue vehicle carries nothing other than maxVehCap-estDemand(15) and the black vehicle carries a maximum of Demand(15) while engaged in coordination, neither can exceed the maximum vehicle capacity. This dynamic coordination scheme is feasible while

$$\text{Demand}(15) \le \text{estDemand}(15)+\text{maxVehCap}-\text{Demand}(12).$$

For values of Demand(15) greater than that, a new accommodation must be made by refining the estimate of the unknown demand. There are two ways of achieving this.

One way is to use the value of Demand(12) as an estimate for Demand(15). This is only possible where the vehicle servicing node 15 can accommodate such a load. If it can, the above coordination scheme will be viable for any value of Demand(15) since no demands in this example exceed maxVehCap.

The other method adjusts estDemand(15) with regard to Demand(1) and Demand(3). Once the initial routes have been created at t=0, the estimated value of Demand(15) is increased to the maximum that the allocated vehicle can carry. In this example, the new estimate is 105 (maxVehCap-Demand(1)-Demand(3)). Since the pink vehicle can carry up to 137 units, there is no limit beyond maxVehCap on Demand(15).

Since each method is effective and each depends on the spare capacity of the black vehicle, either could be used. Adjustment of the estimated value of Demand(15) takes place before t=0.

2.1.8 DC2 Conclusions

In deterministic coordination, all factors (demand, location of nodes, availability) are known before the start of a work cycle. All vehicles leave the depot with their routes firmly decided. In contrast, dynamic coordination exists in an uncertain atmosphere where at least one of the factors is unknown.

In this example, dynamic coordination has produced a most acceptable result. In accommodating the dynamic demand of a single node, the fleet size has not been increased, there is only a small (<2%) increase in distance and the requirements of all nodes have been met with very little disruption to the original routes. No nodes have been rerouted and only one route was reversed.

If Demand(15) is such that the coordination points are not required, any returning vehicle can collect the load from X and deliver it to the depot without penalty; point Y will not be used.

The distance of 226, which would have resulted from a non-coordinated scheme where Demand(15) was known in advance, has been reduced to 220 through the use of coordination.

The black vehicle delivers a total of 205 units to nodes 1, 3 and 15 without violating the vehicle capacity constraint because it has already delivered Demand(3) and Demand(1) before collecting the balance of Demand(15).

The difference between the estimated value and actual value of Demand(15) determines the routing scheme:

estDemand(15) ≤ Demand(15) ≤ 145	original scheme
145 ≤ Demand(15) ≤ maxVehCap	coordination scheme

Table 5. DC2 options.

There is no limitation on the use of coordination in the initial routing scheme. It has not been applied in this example in order to highlight the effect of the dynamic variant.

2.2 DC4 Only location is known

2.2.1 Problem generation

A single depot/fifteen node system similar to that above is generated with the exception that fourteen nodes have pre-determined demands and Node 15 has a demand that will become apparent at some stage. As above, Node 15 has had only one previous demand (of value 50), so the preliminary estimate of Demand(15) is 50.

2.2.2 Problem solution

Initial routes are constructed using a preliminary estimation for demand, calculated as in DC2 above. For each of several time values, a coordination plan is made and the estimated demand is modified to be as large as possible within the plan. When the actual demand is revealed at $t=t_{rev}$ the coordination plan for the value of t close to t_{rev} is considered for adoption. Generally, a coordination plan for smaller values of t can be modified for larger values. Each coordination plan is determined as for DC2 where the time attribute had value 15.

2.2.3 Choosing time values

Times vary from t=0 to t=50. Node 15 is 10 units away from the depot, so there is no point in considering any time greater than t=40 since no vehicle can service node 15 and still return to the depot without violating the constraint specifying that all vehicles must return to the depot within the maximum time of 50. The coordination plan for t=15 was constructed in the discussion of DC2, above. Times t=25 and t=35 will be considered since no revealed time will be more than 5 time units away from one of them.

2.2.4 Initial route construction

This is the same as that determined for DC2 above.

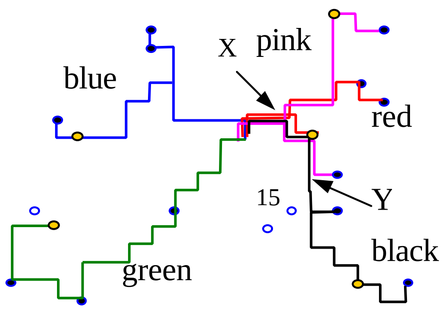

Figure 2. DC4 initial routes. Light discs indicate vehicle positions at t=25. Filled discs indicated servided nodes.

2.2.5 Time t=25 scheme and t=35 scheme

For each scheme, the location of all vehicles is determined as in Figure 2 (black discs indicate serviced nodes). In the t=25 scheme, most nodes have already been serviced, leaving only nodes 5, 14 and 15, and all vehicles are still en route. Because the black vehicle has not yet reached the point Y created for t=15, the earlier coordination scheme can still be used.

Since Black reaches Y at t=33, the t=15 scheme can be used for t=35 with the condition that the black vehicle either

- waits for 2 time units for Demand(15) to be revealed and then collects the appropriate balance from Y, or
- collects the maximum load from Y and continues to node 15 to wait there for the demand to be revealed.

The second option is preferred since node 15 is serviced as soon as possible and the vehicle can return to the depot.

2.2.6 Limitations and accommodations

As long as t<41, either the t=15 scheme or the t=35 scheme can be used to accommodate the dynamic requirements of node 15. If the demand is revealed after t=40, the black vehicle cannot return to the depot within the time limit of 50.

Time, t	Demand(15), D: D=0	0 < D ≤ est(Demand)	Est(Demand) < D ≤ 145	145 < D ≤ maxVehCap
t=0	Plan A	Plan C	Plan C	Plan C
1<t≤33	Plan B	Plan D	Plan E	Plan F
33<t≤40	Plan B	Plan D	Plan G	Plan G
t>40	Plan B	Plan B	Plan B	Plan B

Table 6. DC4 options (see nomenclature in Table 7).

Plan A.	Remove node 15 from consideration.
Plan B.	Do not visit node 15; return to depot with estDemand(15).
Plan C.	Not dynamic.
Plan D.	Deliver Demand(15) and return to depot with surplus.
Plan E.	Original routes with no dynamic coordination scheme. Can involve static coordination.
Plan F.	Scheme as for t=15 above.
Plan G.	As for Plan F except that the vehicle collects the maximum from Y and waits at node 15 until D is known.

Table 7. Nomenclature for Table 6.

2.2.7 DC4 Conclusions

Coordination is a valid technique for vehicle routing where the demand of a particular node is unknown until some unpredictable time after the vehicles have left the depot. It consists of estimating the demand, modifying the estimation, and developing coordination schemes at a selection of different times, in order to be able to judge the worth of coordination when the actual demand is revealed.

2.3 DC7 Nothing is known

2.3.1 Problem generation

A single depot/ fifteen node system is generated similar to that employed above with the exception that all fifteen nodes have known time, demand and location. At some unknown time, a sixteenth node might appear with an as yet unknown demand.

2.3.2 Possible approaches

When demand, time and location are all unknown, estimates can be made only if the appearance of a node is anticipated. Several methods suggest themselves, some of which require planning in advance of the vehicles' departure from the depot.

i) Demand is estimated and space reserved in vehicles for it. This penalises most vehicles since they would not be used in servicing the node when it appears anyway. The most likely vehicles cannot be guessed at since the node's location is unknown.

ii) Coordination sites are created at several locations. At the time of disclosure, nearby vehicles are investigated to determine their suitability for servicing this node.

iii) All vehicles carry maximum payloads. This means that many will have surplus, available for the new node, at no penalty cost.

iv) Nearby vehicles service the node from their own payloads which are replenished later. Pre-determined nodes affected by this technique are serviced later.

A combination of these methods will be employed. The problem is solved as for DC2 and DC4 above. The coordination plan for a specific small time value is developed and this can then be adjusted to accommodate higher time values.

2.3.4 Coordination site location

Several coordination points are established at which the vehicles deposit their surplus loads. These should be reasonably well spread out, preferably at the junctions of natural features, and placed early in the schedule to allow for early declaration of a new node. They must be at locations from which their loads can be retrieved if unused before the end of the work cycle. All materiel must be returned to the depot before the end of the work cycle.

2.3.5 Location of new node

When the new node is further from the depot, there is more chance of a satisfactory solution where there are more vehicles in the scenario since there is less distance for one or more of them to detour to the new node. Similarly, there is greater success where the new node lies close to established routes that have still to be travelled at the time of disclosure since the detour distance is reduced.

2.3.6 DC7 Conclusions

- It is important to have potential detour ability in order to service new nodes appearing away from established routes. Without the ability to detour, alternative routes (of equal length) need to be employed.
- Coordination sites need to be spread around in locations that are either close to established routes (so that they may be used by more than one vehicle) or at natural junctions (such as bridges between pockets) in the rectilinear grid.
- Coordination sites need to be established early so that they may be used over a wider time frame. In contrast, it can be helpful to have some surplus held on board in case new nodes appear towards the end of established routes.

3 Conclusions

Coordination exists in the dynamic environment as well as in the fully deterministic one and is implemented by estimating unknown attributes and developing contingency plans to cope with differences between estimated and actual values.

When demand is unknown, it can be estimated in the first instance by considering the demand history of nodes. This preliminary estimation is adjusted to increase it to its maximum size depending on the capacities of vehicles and their service histories. Initial routing and scheduling is modified to accommodate the estimated demand and a coordination plan is then developed. When the actual value is known, if it exceeds the estimated value, the most appropriate coordination plan can be adopted.

When more than one attribute is unknown, the simpler situation of the one-attribute-unknown system is considered at a variety of values in the second attribute and the most appropriate scheme adopted when all attributes are known.

Where time is unknown, coordination sites should be established early in the process to enable their use over as wide an interval as possible.Coordination sites need to be emptied before the end of the work cycle.

Coordination in the dynamic environment can reduce the total travelling time of a vehicle fleet. It does not depend on the method of construction of the initial routing scheme.

References

[1] M. Dror and P. Trudeau. *Savings by split delivery routing*, Transportation Science 23 (1989), pp141-145

[2] C.M. Rivers, *The En Route Re-distribution of Payload*, APORS ’97 Conference (1997)

[3] C.M. Rivers, *Coordination in the Rectilinear Grid*, ORSNZ Conference (1998)