

**Dynamic Planning
of Pickup and Delivery Operations
by means of Genetic Algorithms**

Giselher Pankratz

Diskussionsbeitrag Nr. 361

Mai 2004

Diskussionsbeiträge des Fachbereichs Wirtschaftswissenschaft
der FernUniversität in Hagen

Herausgegeben vom Dekan des Fachbereichs

Dynamic Planning of Pickup and Delivery Operations by means of Genetic Algorithms

Giselher Pankratz

Abstract:

Customers are making increasing demands on the flexibility and the promptness of transportation services. In order to meet these requirements, planning methods are needed which are able to intelligently react on dynamically revealed information. Algorithms for dispatching transportation requests which arrive dynamically during the planning period can contribute to this goal. In this paper a Genetic Algorithm based approach for the Dynamic Pickup and Delivery Problem with Time Windows (DPDPTW) is developed. First, a Genetic Algorithm for the (static) PDPTW is described which features a special group-oriented genetic encoding. In order to be able to solve the dynamic problem, the Genetic Algorithm is embedded in a rolling horizon framework. Special synchronization and updating mechanisms are provided which assure that the reusable solution knowledge in the population is preserved over the plan revisions. The approach is evaluated using a large number of problem instances with varying degrees of dynamism.

Key words:

dynamic vehicle routing, pickup and delivery problem, genetic algorithms

Fachbereich Wirtschaftswissenschaft, FernUniversität in Hagen
Profilstraße 8, D – 58084 Hagen, BRD

Tel.: 02331/987-4432

Fax: 02331/987-4447

E-Mail: giselher.pankratz@fernuni-hagen.de

Dynamic Planning of Pickup and Delivery Operations by means of Genetic Algorithms

Giselher Pankratz

1 Introduction

While in the past research mainly concentrated on static variants of the vehicle routing problem (VRP) in which all requests are known in advance and no uncertainty exists, increasing interest in dynamic vehicle routing problems can be observed in recent years. The reason for this development is twofold (cf. e.g. Fleischmann et al. 2003): On the one hand, recent advances in communication and information technologies allow real time processing of information which is subject to permanent change. On the other hand, customers increasingly demand for fast and flexible fulfillment of their transportation requests. To meet these requirements, planning methods are needed which are able to intelligently react on dynamically revealed information.

This contribution reports on a Genetic Algorithm based approach for solving the dynamic Pickup and Delivery Problem with Time Windows (DPDPTW) which is a generalization of the Dynamic Vehicle Routing Problem with Time Windows (DVRPTW, cf. Larsen, 2000; Lund et al., 1996). Based on the definition of the (static) Pickup and Delivery Problem with Time Windows (PDPTW, Savelsbergh and Sol, 1995), the DPDPTW can be described as follows:

All transportation requests of a given planning period, e.g. a single day, have to be satisfied by a given fleet of vehicles. The requests are not completely known in advance, but become available during the planning period. Each request is characterized by its pickup location (origin), its delivery location (destination) and the size of the load that has to be transported from the origin to the destination. For each pickup and delivery location, a time window and loading and unloading times are specified. The load capacity, the maximum length of its operating interval, a start location and an end location are given for each vehicle. In order to fulfill the requests, a set of routes has to be planned such that each request is transported from its origin to its destination by exactly one vehicle.

The characteristic feature of the DPDPTW in which it differs from its static counterpart is the fact that a portion of the requests arrive successively during the planning period. That is, relevant information is revealed gradually during execution of the routes. As a consequence, planning and execution are overlapping processes, and planning decisions which possibly are irreversible have to be taken before all problem data becomes known.

In accordance with definitions of the static PDPTW usually found in the literature and in order to make it easier to evaluate the proposed approach when referring to standard benchmark problems, it is assumed that the number of vehicles is unlimited and all vehicles have the same capacity and the same home location where all routes start and end. In addition, it is assumed that no stochastic information is given or gathered which could be exploited in order to predict the arrival times or any other properties of future requests. Finally, for the purposes of this investigation, the dynamic arrival of requests is considered to be the only source of dynamics.

Applications of the DPDPTW range from local area courier services to less than truckload (LTL) transportation and long-distance haulage. A special case of the DPDPTW,

where a vehicle can only transport a single request at a time, is the dynamic Single Load Pickup and Delivery Problem with Time Windows (DSLDPDPTW), also referred to as dynamic Truckload PDPTW (DTPDPTW). In addition, the DPDPTW also models typical situations in public transit. For example, the dynamic Dial-a-Ride Problem with Time Windows (DDARPTW) is a DPDPTW in which people instead of goods are to be transported.

The rest of the paper is organized as follows. Section 2 gives a brief review of previous work on the DPDPTW and similar problems. In Section 3 the Genetic Algorithm is outlined which is used as a basic heuristic for solving the (static) PDPTW. Section 4 supplies a detailed description of the dynamic planning environment. Section 5 reports computational results obtained by the proposed Genetic Algorithm based approach for 5,600 DPDPTW instances which have been generated on the basis of 56 publicly available instances for the PDPTW. Finally, the main findings of the study are summarized in section 6.

2 Literature Review

In this section, a brief overview of important previous approaches for solving dynamic vehicle routing problems is given. With respect to the main topic of this investigation, the survey focuses on approaches to the DPDPTW and its variants.

Since the early 1970s, dynamic variants of the vehicle routing problem (VRP) are discussed in the literature. Frequently cited examples are the papers of Wilson et al. (1971), Psaraftis (1980), and Bell et al. (1983). An early survey of dynamic vehicle routing is provided by Psaraftis (1988). A further survey paper by the same author was published seven years later, taking into account more recent developments (Psaraftis, 1995). An overview of dynamic vehicle routing approaches focussing on local area routing is given by Gendreau and Potvin (1998). In their comprehensive survey of Pickup and Delivery Problems, Savelsbergh and Sol (1995) also consider dynamic variants of this problem class. A recent review of methods for dynamic vehicle routing problems which puts particular emphasis on parallel computing strategies is provided by Ghiani et al. (2003).

Throughout the literature, dynamic vehicle routing problems are solved in a sequential fashion by repeatedly updating the current plan, either at regular intervals or triggered by some external event. Basically, the dynamic problem is decomposed into a sequence of static subproblems which in turn are solved by a static algorithm. Often the static algorithm is embedded in some kind of a rolling horizon framework. Depending on which kind of basic procedure is applied in order to solve the respective static subproblems, the methods for solving dynamic vehicle routing problems can be divided into four main categories: exact algorithms, simple policy based techniques, problem specific heuristics and metaheuristics.

2.1 Exact algorithms

Due to the complexity of vehicle routing problems, exact algorithms are very rarely used as basic procedure for solving the static subproblems of a dynamic vehicle routing problem. One example is the approach proposed by Psaraftis (1980) for solving a dynamic dial-a-ride problem without time windows (DDARP), in which a dynamic programming algorithm is repeatedly applied each time a new request arrives. By means of this algorithm, only small problems with up to 10 requests could be treated.

2.2 Simple policy based techniques

In the context of dynamic vehicle routing, a policy is a simple decision rule which in each plan revision just determines the very next action of each vehicle without maintaining a comprehensive plan. Swihart and Papastavrou (1999) investigate several policies for the single-vehicle pickup and delivery problem without time windows (DPDP). The results prove the so-called nearest neighbor policy to be successful in many cases. Under this policy, after a request has been delivered, the vehicle proceeds to the nearest pickup location. Beside more elaborated techniques, Fleischmann et al. (2003) also make use of two simple rules for the assignment of requests when solving a special extension of the DSLPDPTW: In addition to the dynamic arrival of requests, they also consider fluctuating travel times as a further source of dynamics.

2.3 Problem specific heuristics

Among the problem specific heuristics for dynamic vehicle routing problems, insertion procedures are clearly the most popular. Early applications of insertion algorithms to the DDARP are reported by Wilson et al. (1971), Wilson and Weissberg (1976) and Wilson and Colvin (1977). For solving a DDARPTW, Madsen et al. (1995) adapt a static insertion procedure proposed by Jaw et al. (1986). Psaraftis (1988) describes a specialized assignment heuristic which is embedded in a rolling horizon framework in order to solve a military sealift emergency problem which can be interpreted as a DDARPTW. Savelsbergh and Sol (1998) address a DPDPTW which models a typical planning situation in less than truckload transportation. In this approach, an insertion procedure is embedded in rolling horizon framework which triggers plan revisions at regular time intervals. Yang et al. (2000) investigate several real-time planning policies for a DSLPDPTW. In contrast to the simple policies treated above these policies perform more sophisticated plan revisions, partially by the help of an insertion procedure. Powell et al. (2000) present two heuristics for solving a DSLPDPTW which is modeled as a generalized traveling salesman problem (TSP). Beside the two assignment rules already mentioned in section 2.2, Fleischmann et al. (2003) also develop an assignment algorithm and an insertion algorithm which can be integrated alternatively into a rolling horizon procedure for solving a DSLPDPTW with dynamic travel times.

2.4 Metaheuristics

Problem specific heuristics hold the danger of the search being trapped in a local optimum. Metaheuristics like Genetic Algorithms, Simulated Annealing or Tabu Search have become widely accepted as an appropriate means of avoiding this undesirable behavior. Despite this advantage, until now only a few metaheuristics have been reported for the DPDPTW.

Gendreau et al. (1998) propose a tabu search heuristic for the DPDPTW featuring a neighborhood definition based on the concept of ejection chains. As a diversification strategy, the authors make use of a so-called adaptive memory, which contains a number of the best solutions found thus far. After the search has terminated, a new solution is constructed from the routes in the memory, serving as a new starting solution. The algorithm is embedded into a rolling horizon framework and tested under different operating scenarios. In addition, the authors provide a parallel implementation of their algorithm.

Although the application of Genetic Algorithms and other evolutionary algorithms in dynamic environments attracts growing attention (Bierwirth, 1999; Branke, 2002; Branke, 2003), to the best of our knowledge, up to now no such algorithm has been published for the DPDPTW. However, Jih and Hsu (1999) have developed a hybrid Genetic Algorithm (GA) for solving the static single-vehicle PDPTW, which they pronounce to be appropriate also for dynamic scenarios. The initial population contains a set of routes generated by a dynamic programming algorithm which is aborted prior to completion. Unfortunately, the authors do not provide a description of the dynamic planning environment nor do they present any results for the dynamic case.

3 A Grouping Genetic Algorithm (GGA) for solving the PDPTW

As shown by the literature review, until now no Genetic Algorithm has been applied to the DPDPTW. Thus, it is an interesting task to verify the appropriateness of Genetic Algorithms for this problem class. The concept of Genetic Algorithms was first introduced by Holland (1975). In particular, Genetic Algorithms appear attractive with respect to dynamic problems like the DPDPTW because of their population-based approach, which allows the parallel exploration of different paths through the solution space of a problem. The fact that the population contains a large number of different solution fragments for the Genetic Algorithm to choose from can significantly contribute to the planning flexibility of the Genetic Algorithm with respect to unforeseen events. Apart from this, Genetic Algorithms are considered to be robust and easy to implement (Rayward-Smith, 1994).

This section gives a rough description of the Genetic Algorithm for solving the static subproblems of the DPDPTW. Further details of this algorithm can be found in Pankratz (2003). Throughout the following description it is assumed that the reader is familiar with the terminological and conceptual basics of Genetic Algorithms. Introductory treatments of Genetic Algorithms as a class of optimization algorithms can be found in Goldberg (1989) and Davis (1991). Because the proposed Genetic Algorithm employs an adaptation of the group-oriented genetic encoding introduced by Falkenauer (1998), it is denoted as a Grouping Genetic Algorithm (GGA).

In the following, the group-oriented genetic encoding for PDPTW solutions is explained first. After this, some configuration details of the GGA are outlined.

3.1 Group-oriented genetic encoding

The way in which solutions to the problem at hand are encoded is crucial for the success of any Genetic Algorithm. One difficulty of finding a good genetic encoding for solutions to the PDPTW is the fact that the PDPTW – like the VRPTW – is a combination of two interdependent subproblems: On the one hand, the subproblem of clustering requests and assigning them to a vehicle has to be solved (clustering or grouping problem). On the other hand, for each cluster of requests, a feasible route has to be specified in which the corresponding pickup and delivery nodes are visited (ordering or routing problem). It is not obvious how these two aspects of a solution can be represented simultaneously by a homogeneous encoding such as the standard binary representation of the classic Genetic Algorithm.

Therefore, a different way to encode solutions of the PDPTW is chosen which emphasizes the grouping aspect of the PDPTW. Obviously, the grouping part of the PDPTW is of great importance for the overall quality of a solution to the PDPTW (Savelsbergh and Sol, 1998). In particular, because both the time windows and the precedence constraints considerably restrict the number of routing alternatives for a given allocation of requests, the grouping aspect turns out to be dominant over the routing aspect in many cases. Because of this dominance of the grouping part of the PDPTW, a representation of solutions to the PDPTW is proposed which refers to the encoding principles of the so-called Grouping Genetic Algorithm (GGA) introduced by Falkenauer (1998). In this representation, each gene represents a group of objects instead of a single object. The following describes how the original encoding was adapted to the PDPTW.

Each gene in a chromosome represents a cluster of all requests that are assigned to a single vehicle. The length of a chromosome, i.e. the number of genes is variable and depends on the number of vehicles required by a given solution. Figure 1 illustrates the encoding of a solution by means of a simple example with seven requests that are carried out by three vehicles.

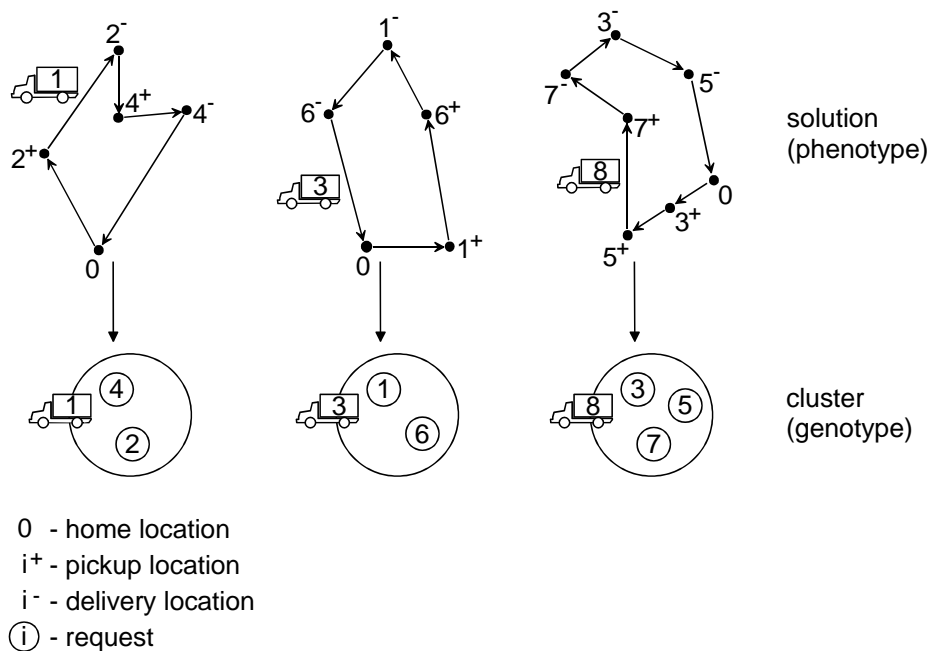


Figure 1: Group-oriented encoding (example)

By representing a partitioning of all transportation requests, a chromosome under this encoding only covers the grouping aspect of a PDPTW solution. As a consequence, the remaining routing aspect of a solution has to be added while decoding the chromosome. This is done by associating each chromosome with separate data structures containing the routing information for each gene. Since this information is hidden from the GGA, it cannot be manipulated directly by the genetic operators. Instead, the routes are constructed and maintained using a separate heuristic according to the grouping information encoded in the chromosome.

The general procedure of this heuristic can be described as follows: For each request to be inserted, all feasible insertions in all existing routes of the current (partial) solution are examined. To this end, all possible insertion positions for the pickup and the delivery node in

a route are tested taking into consideration precedence, capacity and time constraints. Additionally, a new vehicle is allocated and a new route for this vehicle is tentatively initialized with the current request. Among all feasible insertions, the one that causes minimal additional cost is selected and implemented.

3.2 Configuration of the GGA

Figure 2 describes the general search scheme of the proposed GGA.

Input: parameters n^{pop} , p^{cross} , p^{mut}

begin

initialize a population P of size n^{pop} ;

while not (termination criterion is met)

select a pair of individuals x, y from P as parents with regard to their fitness value;

generate two children x', y' applying the crossover operator to x and y with probability p^{cross} ;

generate two modified children x'', y'' applying the mutation operator to x' and y' , respectively, with probability p^{mut} ;

insert x'' and y'' into P and in turn remove the two worst individuals from P ;

return best individual from P as solution.

end

Figure 2: General search scheme of the proposed GGA.

An initial population is generated using the embedded insertion heuristic described above, considering all transportation requests in random order, starting with a single, empty vehicle. After all transportation requests have been inserted, a feasible initial solution is available. The procedure is repeated until an initial population of the desired size is obtained.


The population management is done following the steady-state approach without duplicates (Syswerda, 1989; Whitley, 1989). According to this approach, each newly generated pair of offspring is inserted immediately into the current population where it replaces the two worst individuals. This incremental approach ensures the survival of the best solution over the whole search and prevents the occurrence of duplicate individuals. Thus, it often reveals higher performance than generational replacement (Davis, 1991). In order to select the parent individuals, the binary tournament selection mechanism is applied (cf. Goldberg et al., 1990). In this scheme, two individuals are chosen randomly and the one which represents the better solution is selected as the first parent. To obtain the second parent, the procedure is repeated. The search terminates after a given total number, \tilde{n}^{max} , of individuals has been generated without improvement but no later than after a given maximum number, n^{max} , of generated individuals has been reached.

The crossover operator for the proposed GGA is an adaptation of the general group-oriented crossover scheme presented in Falkenauer (1998) and proceeds in five steps, as illustrated in Figure 3.


1. Specify a crossing section, i.e. a coherent segment of clusters, by randomly selecting two crossing points in each of the two parent chromosomes.

2. Insert the clusters in the crossing section of the second parent at the first crossing point in the chromosome of the first parent. The respective routes are directly adopted from the second parent without reconstruction. As a result of this operation a number of requests may occur twice (requests 1, 2, 6 and 7 in Figure 3). Moreover, two clusters may be assigned to the same vehicle (vehicle 1 in Figure 3).
3. To resolve the conflicts mentioned above, remove all clusters originally belonging to the first parent if they refer to a vehicle that is already allocated by any of the newly inserted clusters. If afterwards some requests still occur twice, eliminate them from those clusters that originally belong to the first parent. At the phenotype level, all eliminations are reproduced immediately in the routes appendant to the changed clusters. Note that the clusters and the respective routes imported from the second parent are left unchanged. After this step, some requests may remain unassigned (requests 5 and 8 in Figure 3).
4. Reinsert the unassigned requests in random order into the individual applying the insertion heuristic. This may require the allocation of an additional vehicle in order to assure feasibility. As a result of this step, a complete offspring is obtained.
5. Generate the second offspring by repeating steps 2 through 4 with reversed roles of the parents.


1. Specify crossing sections



2. Insert groups



3. Clean up chromosome



4. Re-insert unassigned requests




Figure 3: The group-oriented crossover operator

The group-oriented mutation operator for the PDPTW works as follows:

1. Select a cluster at random in the individual.
2. Eliminate this cluster from the chromosome and remove the associated route from the phenotype.
3. Re-insert all removed requests into the individual by means of the insertion heuristic, allocating a new vehicle if necessary to maintain feasibility.

In order to re-adjust the routing part of a solution due to grouping modifications caused by the genetic operators, the insertion procedure is also used. Instead of constructing a new solution from scratch every time the grouping is modified by crossover or mutation, each request that is subject to regrouping is deleted from its original route and reinserted without disturbing the relative order in which the remaining nodes are visited. Thus, the original routing decisions in the unchanged parts of a solution are preserved.

4 Embedding the GGA in a dynamic planning environment

In order to be able to solve the (dynamic) DPDPTW, the proposed GGA for solving the (static) PDPTW is embedded in a rolling horizon framework. This section provides a description of the dynamic planning environment. First, an overview of the dynamic planning procedure is given. After that, the mechanisms for both snapshot generation and synchronization are described in detail.

4.1 Overview of the rolling horizon framework

The course of the rolling horizon procedure can be outlined as follows (cf. Figure 4): In the beginning, the individuals for the initial population are created taking into account all requests which are known beforehand. Subsequently, the GGA is applied to the initial population. After the search has terminated, the best solution found so far constitutes the basis for the physical realization until the occurrence of new requests. Every time a new request arrives, a snapshot of the physical transportation process is taken in order to determine the current state of execution. Note that for the simulations carried out in this study, we assume the routes to be executed exactly as planned, so the snapshot can be done referring to the best solution found so far. Then, a new population is constructed by carefully synchronizing and updating the individuals in the population in order to reflect the new situation. The main idea behind this is to preserve as much as possible of the solution fragments which the individuals in the population have developed so far. Finally, the GGA is restarted using the adapted population as the new initial population.

Input: set of advance requests R_0

begin

initialize time index $i := 0$

initialize termination flag $stop := false$

initialize start population $P(t_0)$ on the basis of R_0

while $stop = false$

apply GGA to $P(t_i)$

pick best solution $\Phi^*(t_i)$ from $P(t_i)$ as the basis for physical execution

wait until next event e occurs

if $e =$ „arrival of new requests R^{new} “

$i := i + 1$

set t_i equal to current time

take a snapshot of the current state of physical execution $M(t_i)$

referring to $\Phi^*(t_i)$

initialize $P(t_i)$ by synchronizing the individuals in $P(t_{i-1})$ with $M(t_i)$

update $P(t_i)$ by inserting R^{new} in all individuals in $P(t_i)$

else

$stop := true$

return last found solution $\Phi^*(t_i)$

end

Figure 4: The rolling horizon procedure

Following a suggestion of Bierwirth (1999), Figure 5 illustrates the interaction between the planning system (i.e., the GGA) and the execution system (i.e., the transportation process) by two interconnected cycles which alternately affect the population: In the upper cycle, the GGA manipulates the individuals in the population by the application of the genetic operators selection, crossover and mutation. In the lower cycle, the current best solution in the population is used as the basis for execution, and the individuals in the population are adapted to the current state of the execution system every time the situation has changed. Due to this process of permanently adapting and re-using the individuals in the population, the population can be understood as an adaptive memory which is shared by both the planning and the execution system (Bierwirth, 1999).

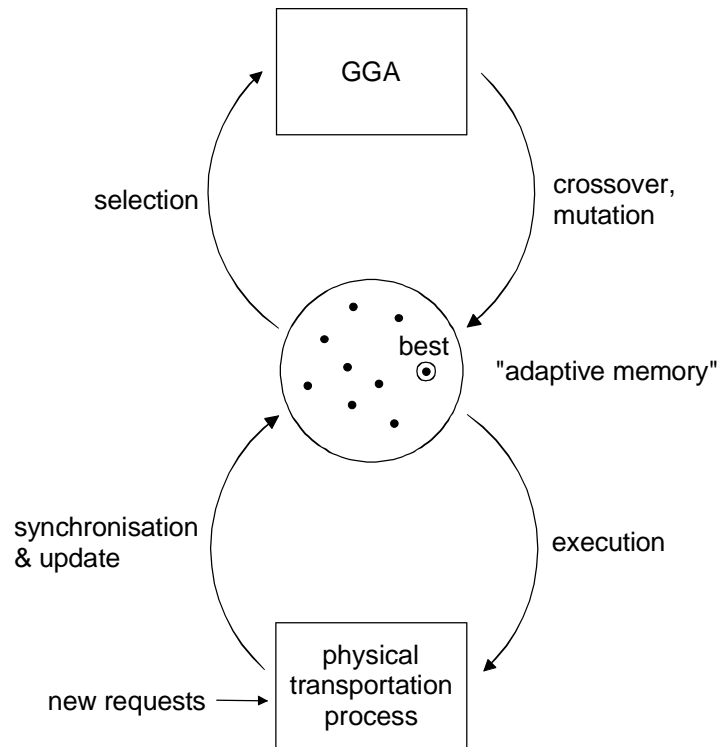


Figure 5: The population of the GGA coupling planning with execution

Obviously, the process of snapshot generation and synchronization plays an important role in the solution methodology proposed here. The following sections will look more closely at these steps.

4.2 Snapshot generation

Let t be the time of the occurrence of new requests. Then a snapshot of the execution system is generated according to the following four-stage procedure:

1. Create a copy of the current best individual in the population found so far.
2. Fix the route position of all nodes which have already been served at time t . A node in a route is considered to have been served at time t if the vehicle which serves the route has already visited this node or has already left the predecessor node for this node at time t .
3. Fix the vehicle assignment of all active and finished requests in the solution copy. A request is considered to be finished at time t if both the position of its pickup location and the position of its delivery location are fixed at time t whereas it is called active at t if its pickup node is fixed but its delivery node is not yet fixed at that time.
4. Delete all requests from the copied solution which are still free for replanning at time t , i.e. which are neither finished nor active at that time. These requests are called fully disposable at time t . If a route does no longer contain any requests after this step, delete the corresponding group in the copied individual and release the associated vehicle.

As a result, a truncated solution is obtained which merely contains the irreversible parts of the original best solution. This solution trunk represents the solution parts which must be implemented by all future solutions in order to maintain feasibility. Therefore, it serves as a pattern for the subsequent synchronization of the individuals in the population. Figure 6

demonstrates the generation of the truncated solution with a simple example. In this example, all active and finished requests are highlighted as gray circles.




Figure 6: Generation of the truncated solution

4.3 Synchronization and update

When synchronizing the individuals, the goal is to adapt all individuals in the population to the current state of execution. As stated above, this should be done in a conservative manner in order to utilize the inherent solution knowledge which each individual in the population has gathered in previous GGA runs. As far as this knowledge concerns those parts of the problem which still can be revised, this knowledge is still valid after the restart of the GGA and can be of great value when creating new solutions. Therefore, we propose a synchronization mechanism which tries to preserve as much as possible of the grouping information originally encoded in the individuals. This mechanism, which is applied to each individual in the population at time t , proceeds in six steps as follows:

1. Create a copy T of the solution trunk generated in the snapshot procedure.
2. Pick the first route of the solution represented by the individual I to be synchronized.
3. Select all routes from T which have at least one fixed request in common with the current route of I .
4. Insert as many fully disposable requests as possible from the current route of I into the routes of the T which have been selected in step 3, applying the insertion heuristic described in section 3.1. If there is no such route in T , i.e., the selection performed in step 3 is empty, allocate an additional vehicle and assign all requests under consideration to this extra vehicle.
5. Repeat steps 3 and 4 until all routes of I have been handled. After this, a number of fully disposable requests may remain unassigned because they could not be feasibly inserted in step 4.
6. Try to insert the unassigned requests into the routes of T regardless of the occurrence of common requests. This again may require the allocation of an additional vehicle in order to assure feasibility. Nevertheless, even this insertion may fail because after synchronization, the new schedule no longer gives enough leeway to integrate a highly urgent request. In this case, T is discarded. Otherwise, synchronization was successful, and I is replaced by T in the new population.

As a result, an individual is obtained which on the one hand is fully adapted to the current state of execution and, on the other hand, still exhibits similarities with the former individual. The synchronization scheme is illustrated by Figure 7, which continues the example of Figure

6. In this example, vehicle no. 3 has no request in common with any vehicle of the trunk, so a new vehicle is allocated incorporating all the requests of vehicle no. 5.




Figure 7: Synchronization of individuals

Note that at this stage, the population may contain fewer individuals than n^{pop} due to the possibility of synchronization failure. In order to compensate for this loss in population size, a number of additional individuals is generated by randomly inserting all disposable requests into a copy of the truncated solution.

Subsequent to the synchronization each individual in the new population is updated by inserting all requests which have newly arrived at time t .

5 Computational Study

This section reports the results of the proposed approach for a number of benchmark problems when compared to two competing heuristics. First, the benchmark data sets are described, then the numerical results are discussed.

5.1 Benchmark Data Sets

Since there are no benchmark data sets available in the literature for the DPDPTW, a total number of 5,600 different DPDPTW instances was generated. These problems were derived from 56 static PDPTW instances provided by Li and Lim (2001) which in turn are based on Solomon's 56 100-customer VRPTW problems (Solomon, 1987). Each of the generated DPDPTW instances comprises between 100 and 106 nodes, which is equivalent to 50 up to 53 requests. Following the classification of Solomon, Li and Lim organized their instances into six classes, denoted as LC1, LC2, LR1, LR2, LRC1 and LRC2, indicating the respective spatial distribution of the nodes and the length of the scheduling horizon. In LC1 and LC2 problems, nodes are clustered, whereas in LR1 and LR2 problems, they are randomly distributed. The nodes in LRC1 and LRC2 instances are partially clustered and partially randomly distributed. While the instances in LC1, LR1 and LRC1 have a short scheduling horizon, the horizon is longer in instances of classes LC2, LR2 and LRC2.

In order to obtain dynamic instances from the static PDPTW instances, an arrival time stamp was assigned to each request, leaving the rest of the problem data unchanged. During

the simulations, each request was announced to the system according to its specific arrival time stamp. The generated DPDPTW instances divide up in two problem sets with different dynamic properties.

The first problem set (P1) contains problem instances with varying degrees of urgency, i.e., when generating these instances, the time interval left to react to a new request after its announcement was varied. Before generating these problems, for each request r the latest possible arrival time t_r^{latest} was determined such that feasibility is provided. Let be $t^{close}(v_r^+)$ and $t^{close}(v_r^-)$ the due times of the pickup and the delivery time window of r , respectively, $t^{serve}(v)$ the service duration at a given location v , v^{home} the home location of the vehicles, and $t(u, v)$ the time required to travel between locations u and v . Then, t_r^{latest} can be defined as

$$t_r^{latest} := \min \{t^{close}(v_r^+), t^{close}(v_r^-) - t(v_r^+, v_r^-) - t^{serve}(v_r^+)\} - t(v^{home}, v_r^+),$$

which is the latest time at which a vehicle must depart from its home location in order to be able to reach both the pickup and the delivery location of the request just in time. Afterwards, each request was given a time stamp $t_r = a \cdot t_r^{latest}$, with a varying from 0.1 to 1.0 in steps of 0.1. According to these principles, from each of the 56 static PDPTW instances ten DPDPTW instances have been generated. Note that in all the problems of this problem set, no requests are known in advance.

In contrast to P1, the second problem set (P2) consists of problems with varying degrees of ex-ante knowledge. They were generated by stepwise increasing the portion q of requests known in advance by 0.1, with q ranging from 0.1 to 0.9. In the remainder, these requests are referred to as static requests whereas all other requests are called dynamic. All static requests, which were selected randomly from the total set of requests of the respective static instance, received a time stamp $t_r = 0$. The dynamic requests were given the highest possible degree of urgency by choosing a time stamp which is equal to the latest possible arrival time, $t_r = t_r^{latest}$. In order to reduce the risk of stochastic bias when picking the static requests, from each static instance, ten dynamic instances were generated for a given value of q . Under this generation scheme, a total number of 5040 instances have been generated for problem set P2.

In order to illustrate the differences between the two problem sets, Figure 8 shows the characteristics of two dynamic instances that were derived from the static problem lc201. The problem to the left is taken from problem set P1, whereas the instance to the right comes from problem set P2. Both problems represent a medium degree of urgency or ex-ante knowledge, respectively. For all the problems shown, the length of the time horizon is 3,390. For each point in time, the accumulated percentage of requests which have already become known to the system are plotted. Additionally, upper bounds for both the portion of known requests which are not yet completed and the portion of requests which still are fully disposable at a given time are provided. The former is calculated by subtracting all requests the delivery time windows of which have already closed at that time, whereas the latter results from deducting all requests the pickup time windows of which have already closed.




Figure 8: Characteristics of two medium-dynamic sample instances

Figure 8 reveals that under the conditions of problem set P1, the number of fully disposable requests, starting from a low level, increases and decreases in time, whereas in the case of problem set P2, the number of disposable requests starts at its maximum and decreases throughout the planning period.

5.2 Numerical results

The proposed GA based approach was evaluated using the 5,600 problem instances of problem set P1 and problem set P2. The parameter settings for the GGA, which were determined in a small number of preliminary experiments, are shown in Table 1. The values were chosen mainly with the idea of performing a thorough search but also to keep computation times reasonable. As a guide, it should take the system less than 60 seconds in the average to carry out a plan revision after the arrival of new requests.

Table 1: GGA parameter values

Parameter	value	description
n^{pop}	150	max. population size
p^{cross}	1.0	crossover probability
p^{mut}	0.5	mutation probability
n^{max}	1,500	max. number of individuals
\tilde{n}^{max}	500	max. number of individuals without improvement

The solution quality of the GA based approach is measured in terms of cumulative total travel distance required to serve all requests of a given instance. In order to judge the solution quality of the approach, the results obtained by two variations of the insertion heuristic described in section 3.2 are used as a standard of comparison.

The first heuristic, in the following denoted as H1, inserts each new request immediately after its arrival without fundamentally changing the routes of the previous plan, i.e., the assignment of requests to vehicles and the relative order of the locations in the routes are left unchanged. Thus, the plan revisions performed by this heuristic can be characterized as incremental.

By contrast, the second heuristic, referred to as H2, performs a total plan revision every time new requests arrive: all decisions of the previous plan, except for those which have been fixed due to the current state of execution, are discarded, and a new plan is generated from scratch.

All algorithms were implemented in JAVA using JDK 1.4. The tests were carried out on a PC (Pentium 4 processor, 2 GHz) under the objective of minimizing total travel distance. The results of the GA based approach were obtained by a single run on each instance.

In the average, the GA based approach requires computing times of less than 6 seconds for a single plan revision. Only in isolated cases computing times of more than a minute have been observed. This in particular concerns some problems of problem set P2 with a large portion of advance requests. The computing times of both H1 and H2 consistently stay below a second.

In the remainder, the results obtained by the GA based approach will be compared to the results of the two heuristics H1 and H2, respectively.

The diagrams in Figure 9 show aggregated results for problem sets P1 and P2, respectively. Note that in all curves, each point represents the average solution quality obtained by the respective method over all instances of identical degree of dynamism. For aggregation and comparison purposes, all results have been normalized using the best objective value obtained by the (static) GGA for the corresponding (static) PDPTW instance as a reference (cf. the results presented in Pankratz, 2004).

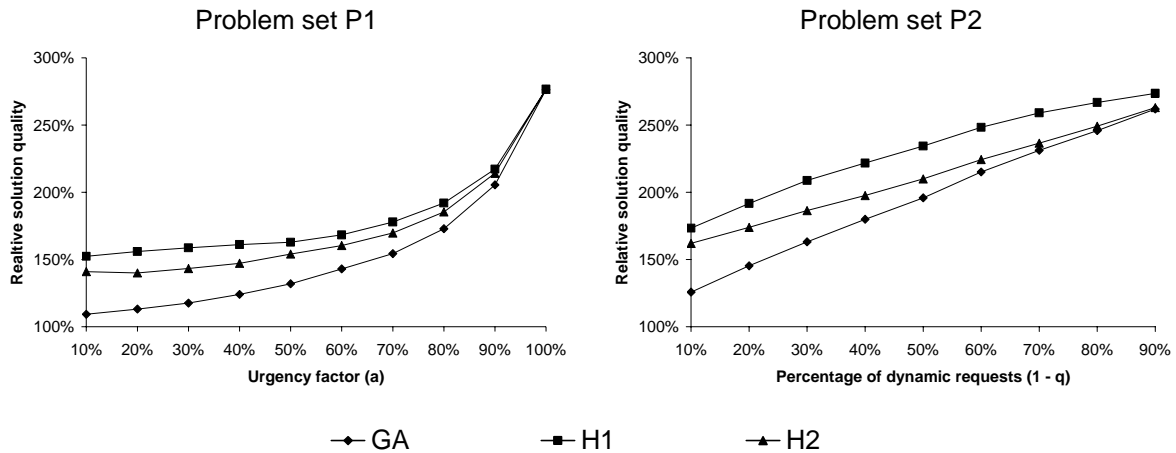


Figure 9: Aggregated results for problem sets P1 and P2, respectively

From Figure 9 the following conclusions can be drawn:

- Although the relative solution quality of all three algorithms gets worse when the problems get more dynamic, the GA based approach in the average performs significantly better than the two competing heuristics for both problem sets.
- Quite naturally, this advantage of the GA based approach diminishes with an increasing degree of dynamism. This is due to the fact that there is less and less room left to the GA based approach to optimize because more and more requests have to be served shortly after they have arrived, so they cannot be replanned even if it turns out later that they have been scheduled awkwardly with respect to subsequent requests.
- Nevertheless the GA based approach produces good results even under highly dynamic conditions. In the case of problem set P1, even for very high degrees of urgency ($a = 0.9$),

the GA based approach in the average produces results which are 5 percent better than those of H1 and 4 percent better than those of H2. Under the conditions of problem set P2, the GA based approach still performs 4 percent better than H1 when the portion of dynamic requests is 90 percent. However, when compared to H2, the advantage of the GA based approach drops below 2 percent as soon as the portion of dynamic requests reaches 80 percent.

- When comparing H1 to H2, H2 appears predominant throughout all runs. Note that this is in line with similar observations made by other authors who state that more flexible plan revisions consistently yield better results than simple, rather incremental plan revisions (cf., for example, Yang et al., 2000). It is particularly striking that H2 performs just slightly better than H1 in the case of problem set P1, whereas for problem set P2, it leaves H1 far behind. Obviously, H2 can profit by its higher degree of freedom in particular when a sediment of advance requests offers optimization opportunities.
- Interestingly, in the case of P1, the curves are comparatively flat in the beginning and ascend rapidly when the degree of urgency exceeds 70 percent, whereas in the case of problem set P2, total travel distance increases almost linearly with an increasing portion of dynamic requests. This can be explained by the way the arrival times are generated for the instances in problem set P1 (cf. Larsen, 2000): because the arrival time of a request r is set relative to the value of t_r^{latest} , requests with a greater value of t_r^{latest} have a longer reaction time than those with a smaller value of t_r^{latest} . Thus, if the values of t_r^{latest} are reasonably scattered over the planning period, enough long-term requests are left which assure considerable optimization opportunities even with values of a of about 0.5 to 0.7.

In order to further analyze the behavior of the GA based approach with respect to the spatial distribution of the nodes and the length of the scheduling horizon, Figure 10 and Figure 11 show averaged results obtained for problem sets P1 and P2, respectively, categorized by problem class. The observations made by the help of these charts can be summarized as follows:

- For problems with a longer scheduling horizon (problem classes lc2, lr2, and lrc2), high degrees of dynamism lead to higher losses in solution quality than for problems with a shorter scheduling horizon (problem classes lc1, lr1, and lrc1). This holds for both P1 and P2. The reason is that in these problems, the requests are spread over a longer period of time. As a consequence, requests which arrive late encounter a situation in which a considerable portion of requests already has been fixed. Obviously, this limits the optimization opportunities.
- Interestingly, the GA based approach behaves quite robust under these conditions: As the curves for problem classes lr2 and lrc2 demonstrate, the GA based approach has a considerable lead over the two comparative heuristics as long as the degree of dynamism is rather moderate.
- When analyzing the influence of the spatial distribution of the requests, it appears that if the scheduling horizon is short, problems exclusively containing clustered requests are slightly more difficult than problems where the requests are at least partially randomly distributed. This in particular becomes obvious when comparing the results of all three algorithms for problem class lc1 to those for problem class lr1 and lrc1, respectively. When the horizon is longer, one has to distinguish between P1 and P2. In the case of P2,

all three algorithms perform better if the problems include randomly distributed requests, i.e. for problem classes lr2 and lrc2. In the case of P1, the GA based approach yields better solutions for problem classes lr2 and lrc2 than for problem class lc2 if the degree of urgency is above 50 percent. Below this value the GA based approach performs almost identical for all three problems lc2, lr2, and lrc2. In contrast to the GA based approach, for lr2 and lrc2, both H1 and H2 obtain results of less quality than for lc2 if the urgency factor is below 60 percent. Interestingly, up to this value, the solution quality of both H1 and H2 is nearly independent of the degree of urgency. However, above this threshold, the respective curves ascend slower than for lc2.

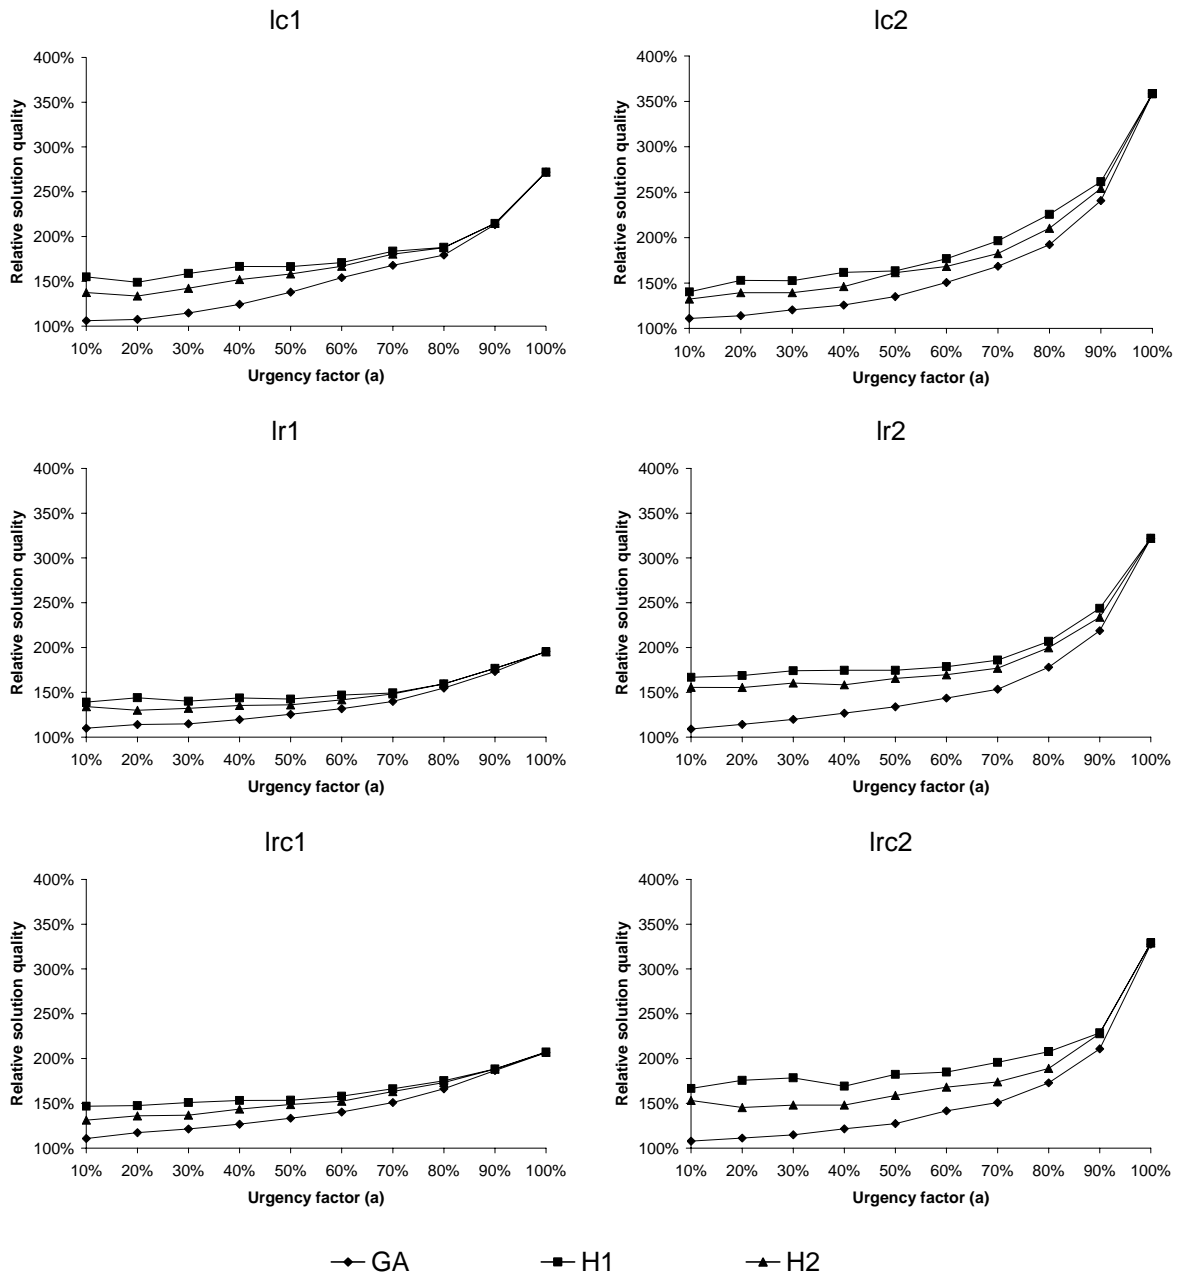


Figure 10: Results for problem set P1, categorized by problem class

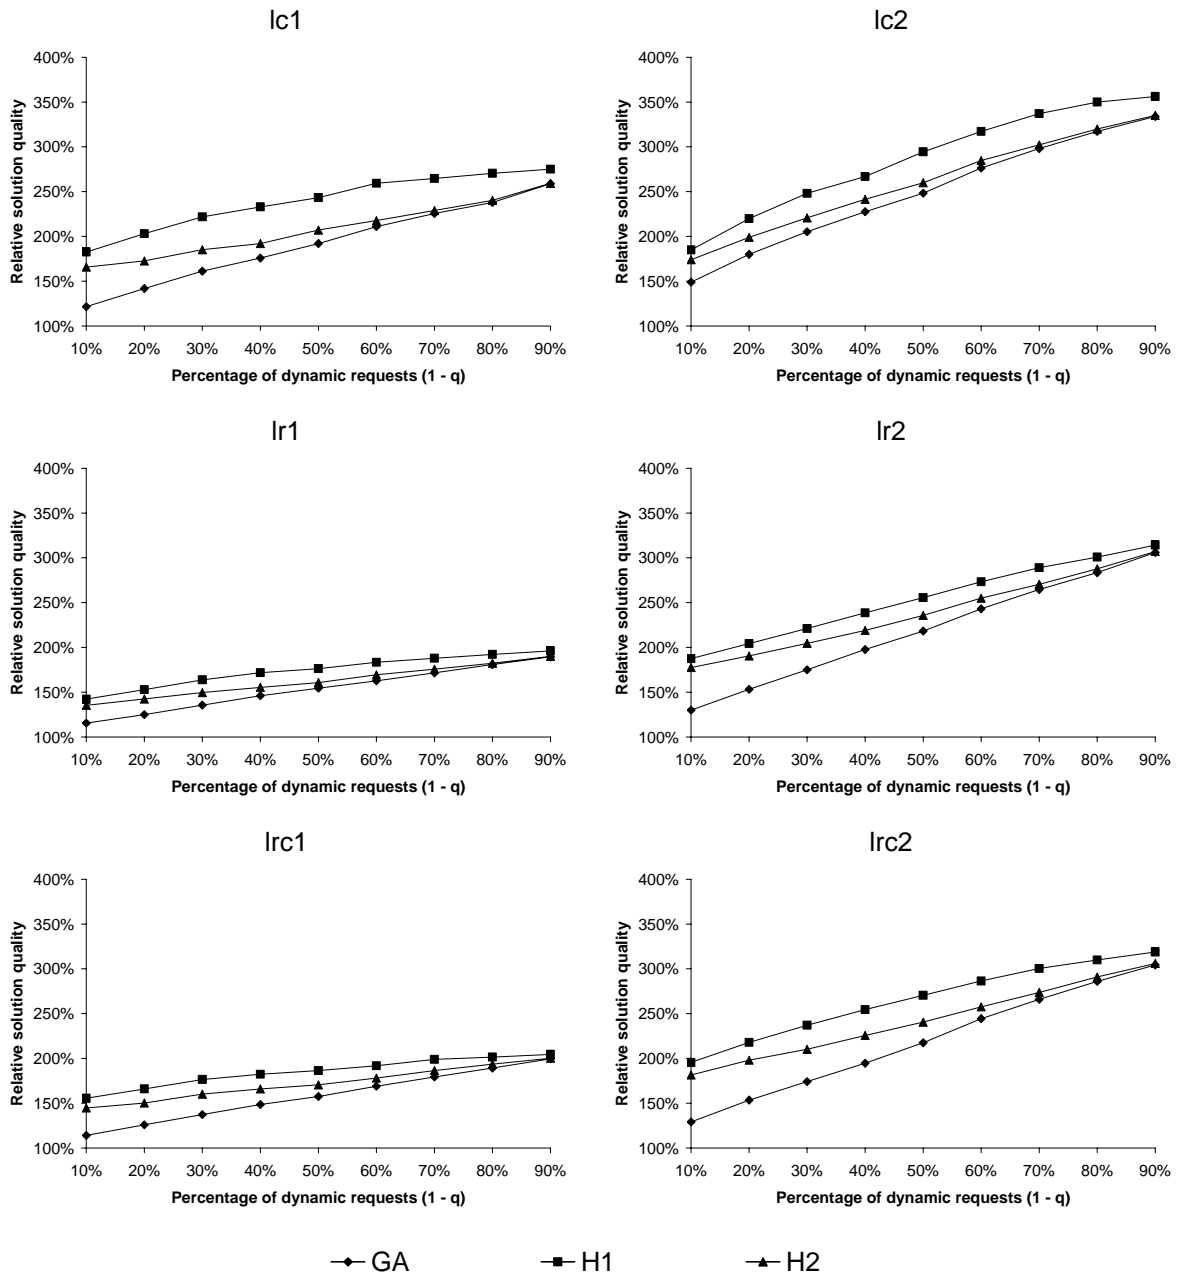


Figure 11: Results for problem set P2, categorized by problem class

6 Conclusions

In this paper, a Genetic Algorithm based approach for the Dynamic Pickup and Delivery Problem with Time Windows (DPDPTW) has been presented. To our knowledge this is the first time a Genetic Algorithm has been applied to the DPDPTW. As a specific feature, the Genetic Algorithm uses a group-oriented genetic encoding according to Falkenauer (1998). The approach was tested using 5,600 dynamic benchmark instances which were derived from 56 publicly available benchmark problems for the PDPTW. The experimental results have demonstrated that the proposed approach is able to find high quality solutions when compared to two comparative heuristics. The overall findings seem to justify the employment of Genetic

Algorithms for dynamic vehicle routing problems. Future research will be dedicated to the following issues:

- Generation of larger problem instances. Although the tests conducted in this contribution have provided valuable insights into the behaviour of the proposed method, it would be interesting to see the behavior with problem instances comprising several hundreds of requests.
- Development and evaluation of different objective criteria. Using specialized objective functions could help making the behaviour of the GA based approach less myopic.
- Optimizing the response times of the method. Since the current prototypical JAVA implementation has not been optimized with regard to time requirements, the algorithms should be re-implemented in C or C++ using optimized data structures.
- Further improvement of the embedded Grouping Genetic Algorithm. In particular, the integration of fast local search techniques into the genetic search appears promising.

At the moment, the proposed GA based approach is prepared for being integrated as local planning component into a multi agent system for cooperative transportation planning.

References

- Bell, WJ, Dalberto LM, Fisher, ML, Greenfield AJ, Jaikumar R, Kedia P, Mack RG, Prutzman PJ (1983) Improving the Distribution of Industrial Gases with an On-line Computerized Routing and Scheduling Optimizer. *Interfaces* 13:4–23
- Bierwirth, C (1999): Adaptive search and the management of logistics systems – base models for learning agents. Kluwer, Boston et al.
- Branke J (2002) *Evolutionary Optimization in Dynamic Environments*. Kluwer, Boston et al.
- Branke J (ed.) (2003) *Proceedings of the Workshop on Evolutionary Algorithms for Dynamic Optimization Problems (EvoDOP-2003)*, held in conjunction with the Genetic and Evolutionary Computation Conference (GECCO-2003). July 12, 2003, Chicago
- Davis L (ed) (1991) *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York
- Falkenauer E (1998) *Genetic Algorithms and Grouping Problems*. Wiley, Chichester
- Fleischmann B, Gnutzmann S, Sandvoß E (2003) Dynamic vehicle routing based on on-line traffic information. To appear in *Transportation science*.
- Gendreau M, Potvin JY (1998) *Dynamic Vehicle Routing and Dispatching*. Crainic TG, Laporte G (eds.) *Fleet Management and Logistics*. Kluwer, Boston et al.: 115–126
- Gendreau M, Guertin F, Potvin JY, Séguin, R (1998) Neighborhood Search Heuristics for a Dynamic Vehicle Dispatching Problem with Pick-ups and Deliveries. Technical Report CRT-98-10, Centre de recherche sur les transports, Université de Montréal, Montréal
- Ghiani G, Guerriero F, Laporte G, Musmanno R (2003) *Real-Time Vehicle Routing: Solution Concepts, Algorithms and Parallel Computing Strategies*. Research Report, Center of Excellence for High Performance Computing, University of Calabria, Italy
- Goldberg DE (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading/Massachusetts
- Goldberg DE, Korb B, Deb K (1990) Messy Genetic Algorithms: Motivation, Analysis, and First Results. *Complex Systems* 3: 493–530
- Holland JH (1975) *Adaptation in Natural and Artificial Systems – An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. The University of Michigan Press, Ann Arbor
- Jaw JJ, Odoni AR, Psaraftis HN, Wilson, NHM (1986) A Heuristic Algorithm for the Multi-Vehicle Advance Request Dial-A-Ride Problem with Time Windows. *Transportation Research Part B* 20: 243–257
- Jih, W.-R, Hsu, Y.-J (1999) Dynamic Vehicle Routing Using Hybrid Genetic Algorithms. In: IEEE Computer Society (ed) *Proceedings of the 1999 IEEE International Conference on Robotics & Automation*. IEEE Computer Society, Los Alamitos/California: 453–458
- Larsen A (2000) *The Dynamic Vehicle Routing Problem*. Ph.D. Thesis, Institute of Mathematical Modelling (IMM), Technical University of Denmark, Lyngby, Denmark
- Li H, Lim A (2001) A Metaheuristic for Solving the Pickup and Delivery Problem with Time Windows. In: IEEE Computer Society (ed) *Proceedings of the 13th IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE Computer Society, Los Alamitos, California: 160–167
- Lund K; Madsen OBG, Rygaard, JM (1996) *Vehicle Routing Problems with Varying Degrees of Dynamism*. Technical Report IMM-REP-1996-1, Institute of Mathematical Modelling (IMM), Technical University of Denmark, Lyngby, Denmark

- Madsen OBG, Ravn HF, Rygaard JM (1995) A Heuristic Algorithm for a Dial-a-Ride Problem with Time Windows, Multiple Capacities and Multiple Objectives. *Annals of Operations Research* 60: 193–208
- Pankratz, G (2003) A Grouping Genetic Algorithm for the Pickup and Delivery Problem with Time Windows. Submitted for publication (available from ftp://ftp.fernuni-hagen.de/pub/fachb/wiwi/win/forschng/publi/gp_p5_neu.pdf)
- Powell W, Snow W, Cheung R (2000) Adaptive Labeling Algorithms for the Dynamic Assignment Problem. *Transportation Science* 34: 50–66
- Psaraftis HN (1980) A Dynamic Programming Solution to the Single Vehicle Many-to-Many Immediate Request Dial-a-Ride Problem. *Transportation Science* 14: 130–154
- Psaraftis HN (1988) Dynamic vehicle routing problems. Golden BL, Assad AA (eds.) *Vehicle Routing: Methods and Studies*. North-Holland, Amsterdam et al.: 223–248
- Psaraftis HN (1995) Dynamic vehicle routing: Status and prospects. *Annals of Operations Research* 61: 143–164
- Rayward-Smith VJ (1994) A Unified Approach to Tabu Search, Simulated Annealing and Genetic Algorithms. In: Unicom Seminars Ltd (ed) *Adaptive Computing and Information Processing*. Unicom Seminars Ltd., London, Vol. I: 55–78
- Savelsbergh MWP, Sol M (1995) The General Pickup and Delivery Problem. *Transportation Science* 29: 17–29
- Savelsbergh MWP, Sol M (1998) DRIVE: Dynamic Routing of Independent Vehicles. *Operations Research* 46: 474–490
- Solomon M (1987) Algorithms for the Vehicle Routing and Scheduling Problem with Time Window Constraints. *Operations Research* 35(2): 254–265
- Swihart MR, Papastavrou JD (1999) A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *European Journal of Operational Research* 114: 447–464
- Syswerda G (1989) Uniform Crossover in Genetic Algorithms. In: Shaffer, JD (ed) *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo/California: 2–9
- Whitley LD (1989) The GENITOR Algorithm and Selection Pressure: Why Rank-Based Allocation of Reproductive Trials is Best. In: Shaffer JD (ed) *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo/California: 116–121
- Wilson NHM, Colvin NH (1977) Computer control of the Rochester dial-a-ride system. Report R77-31, Dept. of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts
- Wilson NHM, Weissberg H (1976) Advanced dial-a-ride algorithms research project: Final report. Report R76-20, Dept. of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts
- Wilson NHM, Sussman JM, Wong HK, Higonnet T (1971) Scheduling algorithms for dial-a-ride systems. Urban Systems Laboratory Report USL TR-70-13, Massachusetts Institute of Technology, Cambridge, Massachusetts
- Yang J, Jaillet P, Mahmassani HS (2000) Study of a Real-time Multi-vehicle Truckload Pickup-and-Delivery Problem, to appear in *Transportation Science*