An Overview of Models and Techniques for Integrating Vehicle and Crew Scheduling

Richard Freling[1], Albert P.M. Wagelmans[1], and José M. Pinto Paixão[2]
[1]Econometric Institute, Erasmus University Rotterdam, The Netherlands email: freling@few.eur.nl
[2]DEIO, Faculdade de Ciências da Uiversidade de Lisboa, Portugal

Abstract: In this paper, the problem of integrating vehicle and crew scheduling is considered. Traditionally, vehicle and crew scheduling have been dealt with in a sequential manner, where vehicle schedules are determined before the crew schedules. The few papers that have appeared in the literature have in common that no comparison is made between simultaneous and sequential scheduling, so there is no indication of the benefit of a simultaneous approach. In order to get such an indication before even solving the integrated problem, we propose a method to solve crew scheduling independently of vehicle scheduling. We introduce a mathematical formulation for the integrated problem, and briefly outline algorithms. The paper concludes with computational results for an application to bus scheduling at the public transport company RET in Rotterdam, The Netherlands. The results show that the proposed techniques are applicable in practice. Furthermore, we conclude that the effectiveness of integration as compared to a sequential approach is mainly dependent on the flexibility in changing buses during a duty.

1 Introduction

In this paper we consider an overview of models and techniques for the integration of bus and driver scheduling. Technical details are not considered here. For a more comprehensive description of the subject, including the technical details, we refer the readers to Freling (1997).

In Fig. 1.1 we show the relation between four operational planning problems which typically arise in public transport organisations. Decisions about which routes or *lines* to operate, and with what frequency, are based on the available infrastructure, service requirements for the passengers, and demand aspects. We assume that these are known for the operational planning phase. Also known are the travel times between various points on the route. These may differ for various parts of the planning period (e.g., the travel times may be longer during busy hours). Based on the lines and frequencies, timetables are determined resulting in *trips* with corresponding starting and ending times and locations. An example of a trip is departure from location 1 at 9:00 am and arrival at location 2 at 10:00 am.

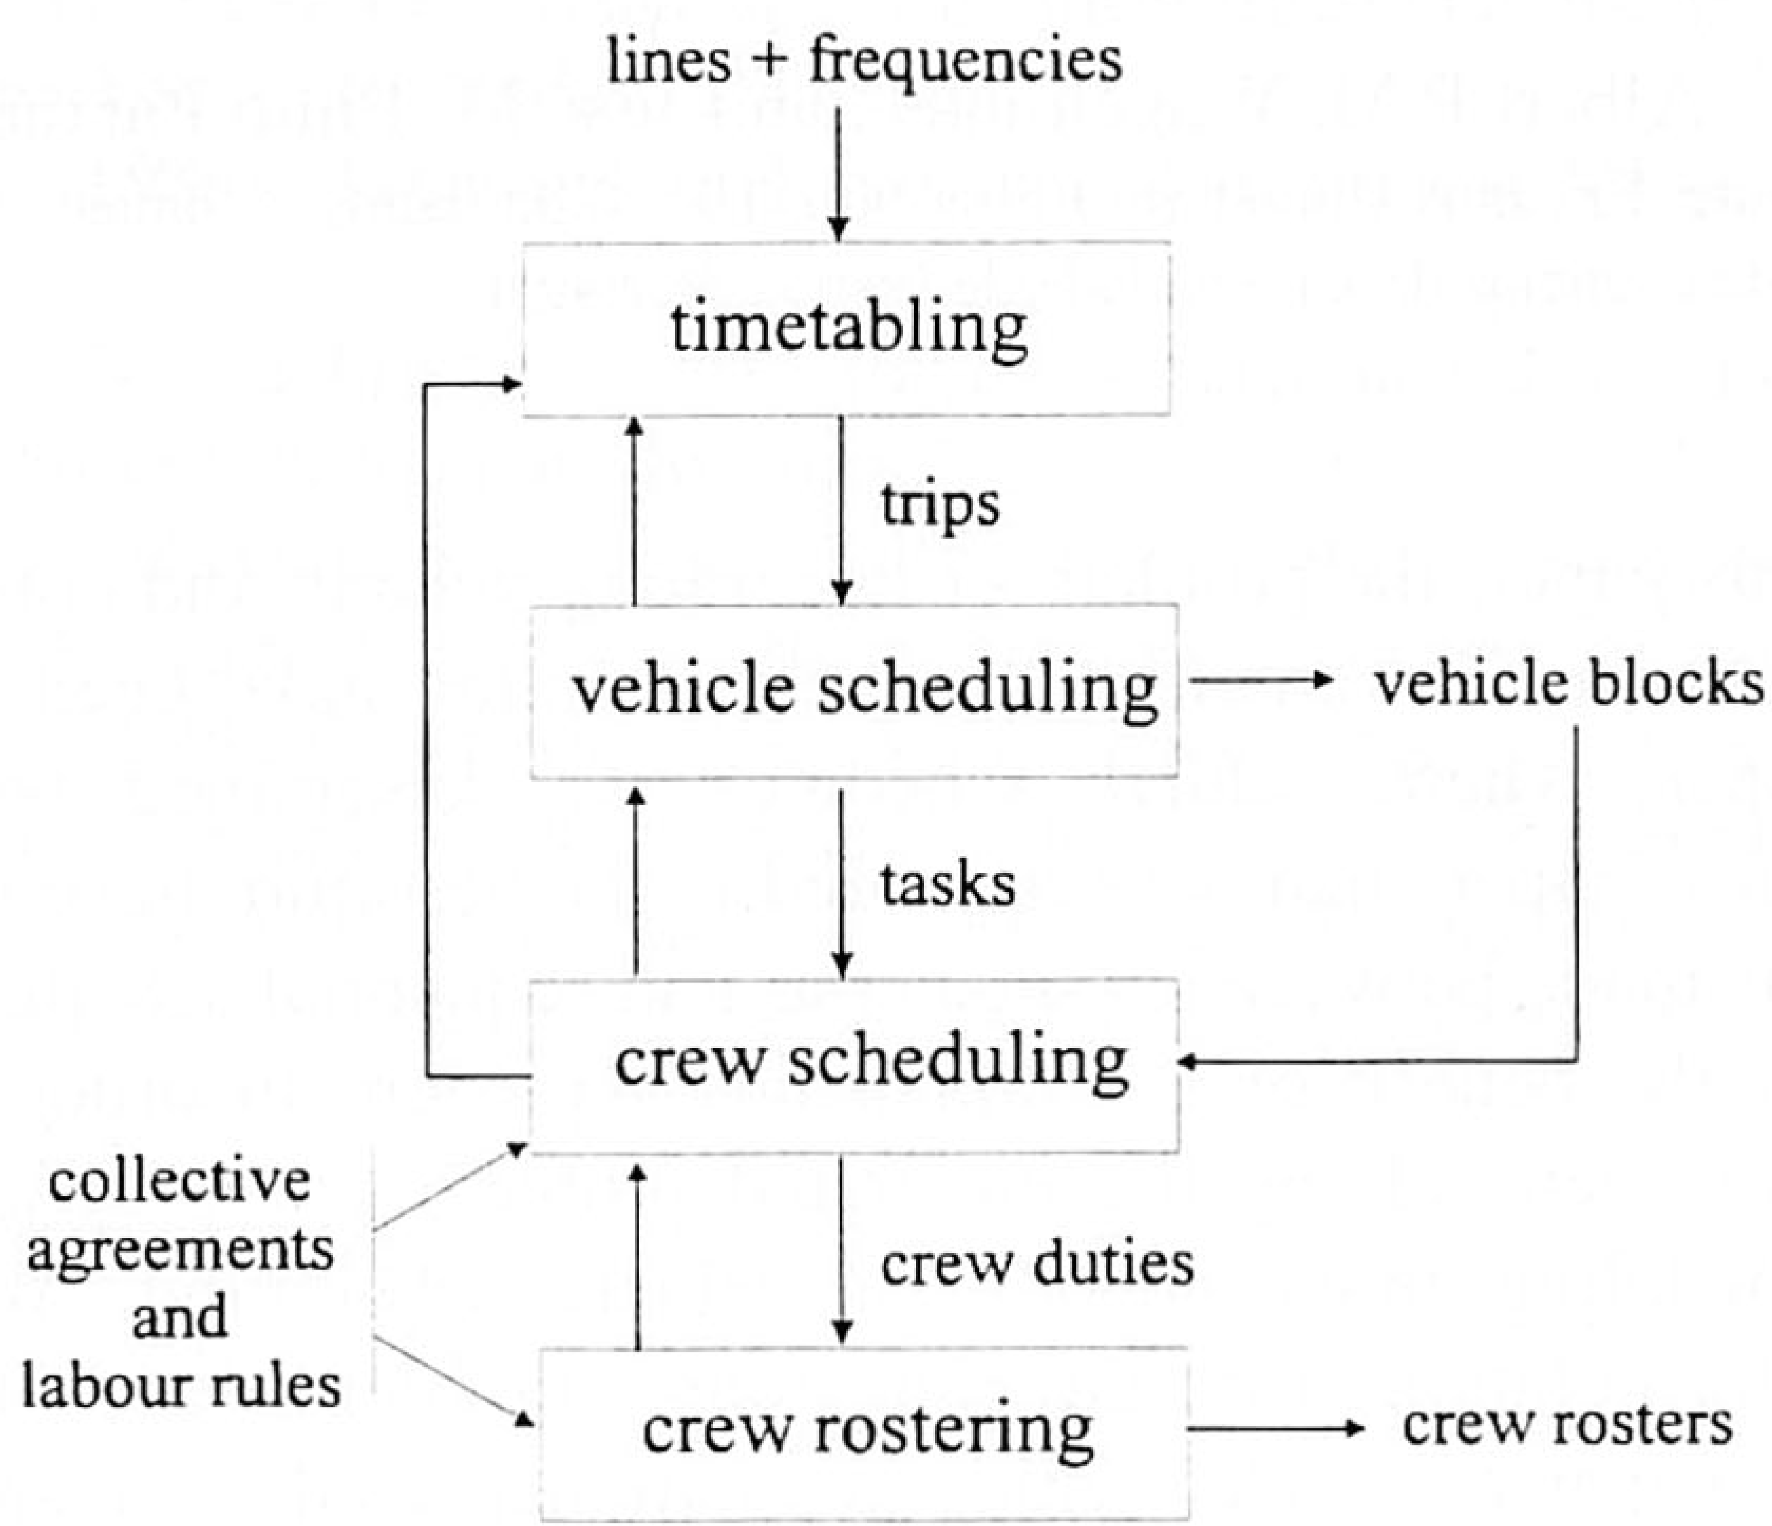

Fig. 1.1

The next two planning processes are vehicle and crew scheduling, which consist of assigning vehicles to the trips and crews to vehicles, respectively. Crew scheduling is short-term crew planning (e.g. one or several days), while the crew rostering process is long term crew planning (e.g. one month or half a year).

In most practical situations these processes interact with each other as shown in the figure. Presumably, the planning could be more efficient if an integrated approach were taken; however, because it is mathematically not feasible to consider the whole process at once, most theoretical studies deal with these processes in a sequential manner. Some approaches have been proposed in the operations research literature which deal with an integration of two of the planning problems. Several papers consider the vehicle scheduling problem with time windows, which is an example of the integration of timetabling and vehicle scheduling; that is, starting times of trips are not fixed but can vary within time windows (see Ferland/Fortin (1989)).

The integration of vehicle and crew scheduling is the topic of this paper, which is organized as follows. In Sect. 2, we discuss several approaches for vehicle and crew scheduling. Recent theoretical developments in the field of optimization, as well as rapid developments in computer technology in terms of power and speed at much lower cost, inspired us to tackle vehicle and crew scheduling in an integrated manner. One of our primary research objectives was to obtain insight into the effectiveness of integrating vehicle and crew scheduling compared with sequential approaches. In Sect. 3 we discuss the potential benefits of integrating vehicle and crew scheduling. The integration of vehicle and crew scheduling in general has received little attention in the literature. We provide an overview of the scarce literature in Sect. 4. In Sections 5 and 6, we outline models and algorithms. Another research objective was to test the practical applicability of the proposed techniques. In Sect. 7 we consider application of bus and driver scheduling at RET, the urban public transport company in Rotterdam, the Netherlands.

2 Problem Definition

Although in the early eighties several researchers recognized the advantage of integrating vehicle and crew scheduling, most of the algorithms published in the literature follow the sequential approach where vehicles are scheduled before and independent of crews. Algorithms incorporated in successful computer packages use this sequential approach as well, while sometimes integration is dealt with at the user level (see Darby-Dowman/Jachnik,/Lewis/Mitra (1988)). In the operations research literature, only a few publications address simultaneous approaches to vehicle and crew scheduling. None of those publications makes a comparison between simultaneous and sequential scheduling. Hence, they do not provide any indication of the benefit of a simultaneous approach.

In this paper, we consider vehicle and crew scheduling from a different angle. Besides a complete integration of vehicle and crew scheduling, we also consider the reverse sequential approach of scheduling crews before and independent of vehicles.

2.1 Traditional Sequential Approach

In the traditional sequential approach vehicles are scheduled before and independent of crews. The *single depot vehicle scheduling problem* (SDVSP) is defined as follows: given a depot and a set of trips with fixed starting and ending times, and given travel times between all pairs of locations, find a feasible minimum cost schedule such that (1) each trip is assigned to a vehicle, and (2) each vehicle performs a feasible sequence of trips. All the vehicles are assumed to be identical. A schedule for a vehicle is composed of *vehicle blocks*, where each block consists of a departure from the depot, the service of a feasible sequence of trips and the return to the depot. The cost function is usually a combination of vehicle capital (fixed) and operational (variable) cost. The SDVSP is known to be solvable in polynomial time.

For the sequential approach it is assumed that the vehicle scheduling problem has been solved when considering the scheduling of crews; that is, the set of vehicle blocks defining the vehicle schedule is known. The *Crew Scheduling Problem* (CSP) is defined as follows: given a set of *tasks*, find a minimum cost set of *duties*, such that (1) each task is assigned to one duty, and (2) each duty is a sequence of tasks that can be performed by a single crew. The vehicle schedule defines vehicle blocks which should be covered by duties at minimum cost. The blocks are subdivided at *relief points*, defined by location and time, at which a change of driver may occur. A task is defined by two consecutive relief points and represents the minimum portion of work that can be assigned to a crew. Each duty must satisfy several complicating constraints corresponding to work rules for crews. Typical examples of such constraints are maximum working time without a break, minimum break duration, maximum total working time, and maximum duration. The cost function is usually a combination of fixed costs such as wages, and variable costs such as overtime payment.

Crew scheduling has similarities to vehicle scheduling but is more complex due to several complicating constraints such as the aforementioned work rules for crews. Instead of assigning trips to vehicles, crew scheduling involves the assignment of tasks to crews or, better, to duties. A basic assumption is that every crew is the same, i.e. individual crew members are not considered. Since the type of constraints differ from application to application it is difficult to define a generic crew scheduling problem. Beasley/Cao (1996) propose the CSP with only spread time constraints as the generic CSP. In fact, this is the SDVSP with time constraints (see Freling/Paixão (1995)), which Fischetti/Martello/Toth (1987) have shown is NP-hard. The last group of authors have shown in Fischetti/Martello/Toth (1989) that the CSP with only working time constraints is NP-hard. Here we assume that the CSP has at least spread time or working time constraints, and is therefore NP-hard. Furthermore, we assume that at least one break must occur in a duty, which adds considerably to the complexity of the problem. We believe that, as long as one aims at developing techniques which are generally useful, a generic crew scheduling problem should at least have working or spread time and break constraints.

Since a duty is defined as a sequence of tasks and breaks, we can consider constraints which define the feasibility of a duty as resource constraints. Each task and break consumes a certain amount of a resource, and the total amount must be within the allowed interval. Below we present a summary of such *resource* or *local constraints* which define the feasibility of a single duty:

- *Time constraints*, placing limitations on the time in a duty. Examples are limited duration (i.e. *spread time*), limited working time, limited paid time, limited working/spread time without rest period or break.
- A minimum number of *breaks*. For example a minimum number of rest periods, coffee/tea breaks, meal breaks, overnight rests, etc.
- *Location constraints*, placing limitations on the locations visited during a duty. For example, a duty must originate and terminate at the same location.
- *Vehicle constraints*, defining links between crews and vehicles. For example, a limited number of changes of vehicle during a duty (*changeovers*), or *vehicle attendance* when a vehicle is stationed at a location other than the depot.
- *Crew dead-heading constraints*, placing limitations on crew transportation when the crew is not working on a vehicle.

Besides local constraints, *global constraints* also may exist which deal with groups of duties at once:

- *Time constraints*, e.g. limited average working time.
- *Location constraints*, e.g. limited number of crews available at a crew base.
- *Duty type constraints*, e.g. limited number of duties with overtime.

A particular application of the CSP is the Bus Driver Scheduling Problem (BDSP). For the BDSP the planning horizon is usually one day, that is, the BDSP consists of determining a set of duties or workdays which will then be assigned to individual bus drivers. An important notion here is the definition of a *piece of work*, which is a set of consecutive tasks in a vehicle block to be performed by a single driver. Sometimes a piece of work is defined more generally as a set of tasks without a break. A duty consists of one or more pieces of work. If a duty contains more than one piece of work, then these pieces are separated by breaks or free time periods. The duration of pieces of work and breaks in a duty is usually limited. Typical global constraints restrict the percentage of certain types of duties in the solution, such as split duties (i.e. with a large unworked period between pieces) and *trippers* (i.e. a single piece or short duty).

2.2 Integrated Approach

The *vehicle and crew scheduling problem* (VCSP) is the following: given a set of service requirements or trips within a fixed planning horizon, find a minimum cost schedule for the vehicles and the crews, such that both the vehicle and the crew schedule are feasible and mutually compatible. We make the following assumptions:

1. The vehicle scheduling characteristics correspond to the SDVSP as defined previously, that is, one depot, identical vehicles, fixed starting times of trips, and no time constraints.

2. The cost function for the VCSP is simply the summation of the vehicle and crew scheduling cost functions defined previously. The primary vehicle scheduling objective is to minimize the number of vehicles, while the primary crew scheduling objective is to minimize the number of crews.
3. A piece of work is defined as a sequence of tasks in a vehicle block which can be performed by one crew without interruption. This sequence of tasks is only restricted by its duration which must be within certain time limits.

The last assumption is not restrictive. For example, it can incorporate a restricted number of pieces, restricted spread time, and restricted working time. The three assumptions make our approach in principle applicable to bus and driver scheduling, although the crew scheduling characteristics are general.

We distinguish between two types of tasks, i.e., *trip tasks* corresponding to (parts of) trips, and *dh-trip tasks* corresponding to deadheading trips (*dh-trips*). All trip tasks need to be covered by a crew, while the covering of dh-trip tasks depends on the vehicle schedules and determines the compatibility between vehicle and crew schedules. In particular, each dh-trip task needs to be assigned to a crew if and only if its corresponding dh-trip is assigned to a vehicle. Note that one or more trip tasks may correspond to one trip, depending on the relief points along that trip. Similarly, one or more dh-trip tasks may correspond to one dh-trip. For example, a dh-trip between locations e_i and b_j which passes the depot corresponds to two dh-trip tasks, one from e_i to the depot and the other from the depot to location b_j. Here we assume that waiting at the depot is not a task because vehicle attendance at the depot is not necessary.

2.3 Scheduling Crews Independent of Vehicles

As an alternative to the crew scheduling problem, we consider crew scheduling independent of vehicle scheduling. The *independent crew scheduling problem* (ICSP) is the following: given a set of trip tasks corresponding to a set of trips, and given the travelling times between each pair of locations, find a minimum cost crew schedule such that all trip tasks are covered in exactly one duty and all duties satisfy crew feasibility constraints. When the ICSP is used as a method for determining crew schedules, vehicles need to be scheduled afterwards such that the vehicle and crew schedules are compatible. Therefore, the duties need to satisfy extra requirements in order to assure that a crew is available for each task induced by the vehicle schedule. In Sect. 3 we discuss the use of the ICSP for determining the potential benefits of integration.

3 Potential Benefits of the Integration of Vehicle and Crew Scheduling

Vehicle and crew scheduling problems often interact with each other: the specification of vehicle schedules will place certain constraints on the crew schedules and vice versa. Because vehicles are often much more flexible to schedule than crews, it may be inefficient to schedule vehicles without considering crew scheduling. Vehicle oriented characteristics of crew scheduling may affect the extent to which the integration of vehicle and crew scheduling is beneficial compared with the traditional sequential approach. Examples of such characteristics are the following:

- A restricted number of *changeovers*, that is, a crew is restricted in changing vehicles during a duty.
- Restricted crew dead-heading.
- Extra start-up time on a vehicle when a new crew is assigned to it.
- Compulsory continuous attendance when a vehicle is waiting.
- Minimum duration of a piece of work.
- Domination of crew costs over vehicle costs.
- Minimum break with the vehicle if a piece of work is longer than a certain duration.
- Crew relieves only occur at the depot, that is, changeovers are not allowed outside the depot.

If none of these characteristics exist, there is probably no need to integrate vehicle and crew scheduling because crews can move independently of vehicles. However, combinations of the mentioned characteristics appear frequently in practice. Integration may serve two purposes: feasibility and/or cost efficiency. We illustrate this with two examples.

Tosini/Vercellis (1988) consider the extra-urban bus driver scheduling problem, which is an example where integration is necessary for obtaining feasible solutions. In this situation, driver dead-heading is not allowed due to long distances, while driver relieves can occur only at crew bases. If buses are scheduled without attention to the driver scheduling, it may be that no feasible driver schedule exists, because a bus may be away from a crew base too long to be serviced by a driver. Therefore, the bus should pass by a crew base once in a while to change drivers. Figure 3.1 illustrates this case with six trips on one vehicle block, marked by their locations *L1*, *L2* and *L3*.

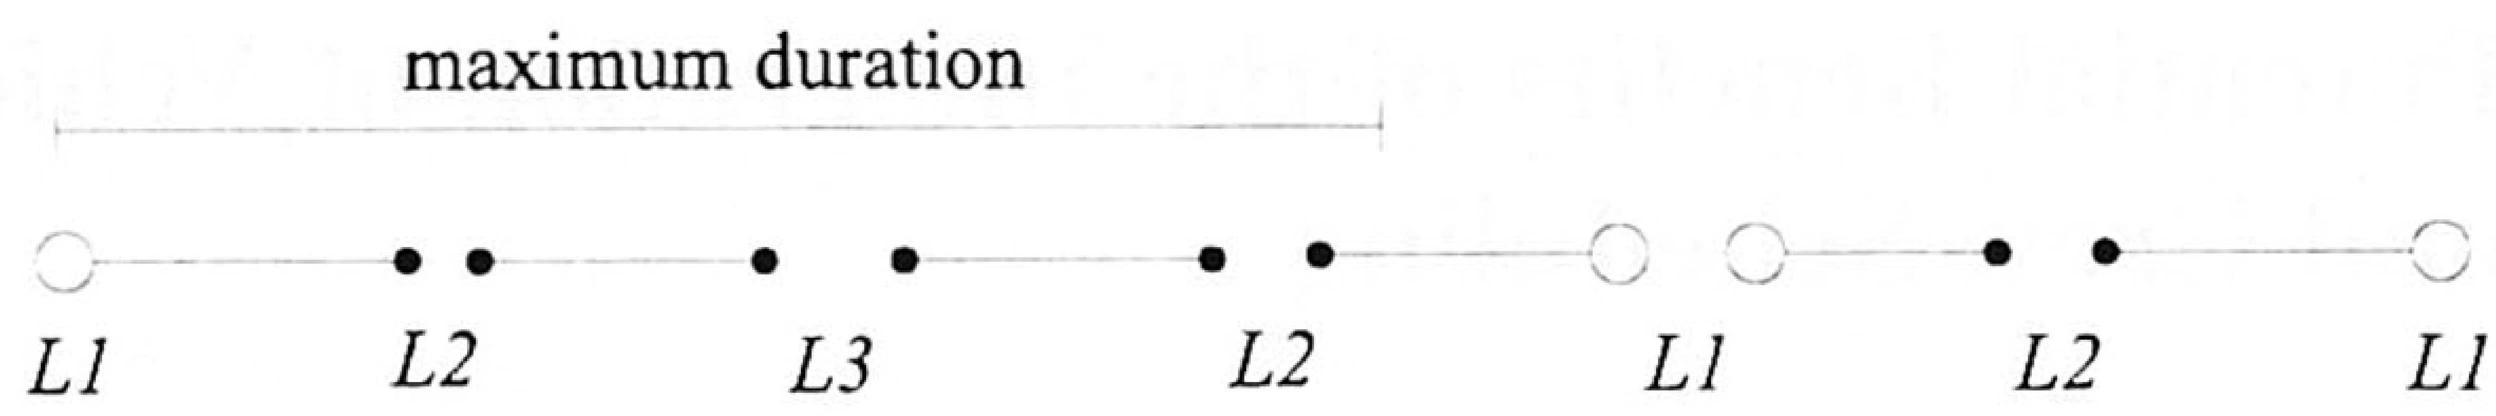

Fig. 3.1

Suppose that the only crew base is at location *L1*. Then, the problem is unfeasible if the duration of the first four trips away from the base exceeds the maximum duration allowed. Note that a possible approach for the VCSP is to consider crew rules for in the vehicle scheduling step. The example above can be tackled by determining vehicle schedules subject to a time constraint (see Freling (1995)).

An example of integration being more cost efficient is illustrated by two vehicle scheduling networks depicted in Figures 3.2 and 3.3, respectively. A directed vehicle scheduling network has nodes corresponding to a source s and a sink t and one node for each trip, where it is assumed that trips are numbered according to increasing starting time. An arc from the source to a trip corresponds to a vehicle leaving the depot to perform the trip, an arc from a trip to the depot corresponds to a vehicle entering the depot after performing the trip, while an arc between two trips corresponds to a vehicle performing both trips in sequence (see also Sect. 5). Only the arcs are drawn which correspond to two different vehicle schedules with three vehicles each. The idea is that the two different vehicle schedules may be served with crew schedules with different degrees of efficiency. Suppose that changeovers are not allowed and the maximum duty duration is such that at most two trips can be included in a duty. Then, at least five duties are necessary to cover the three vehicles in the vehicle schedule of Fig. 3.2. In Fig. 3.3 it is shown that it might be possible to save one duty by adjusting the vehicle schedule with the same number of vehicles. At least four duties are necessary to cover this schedule.

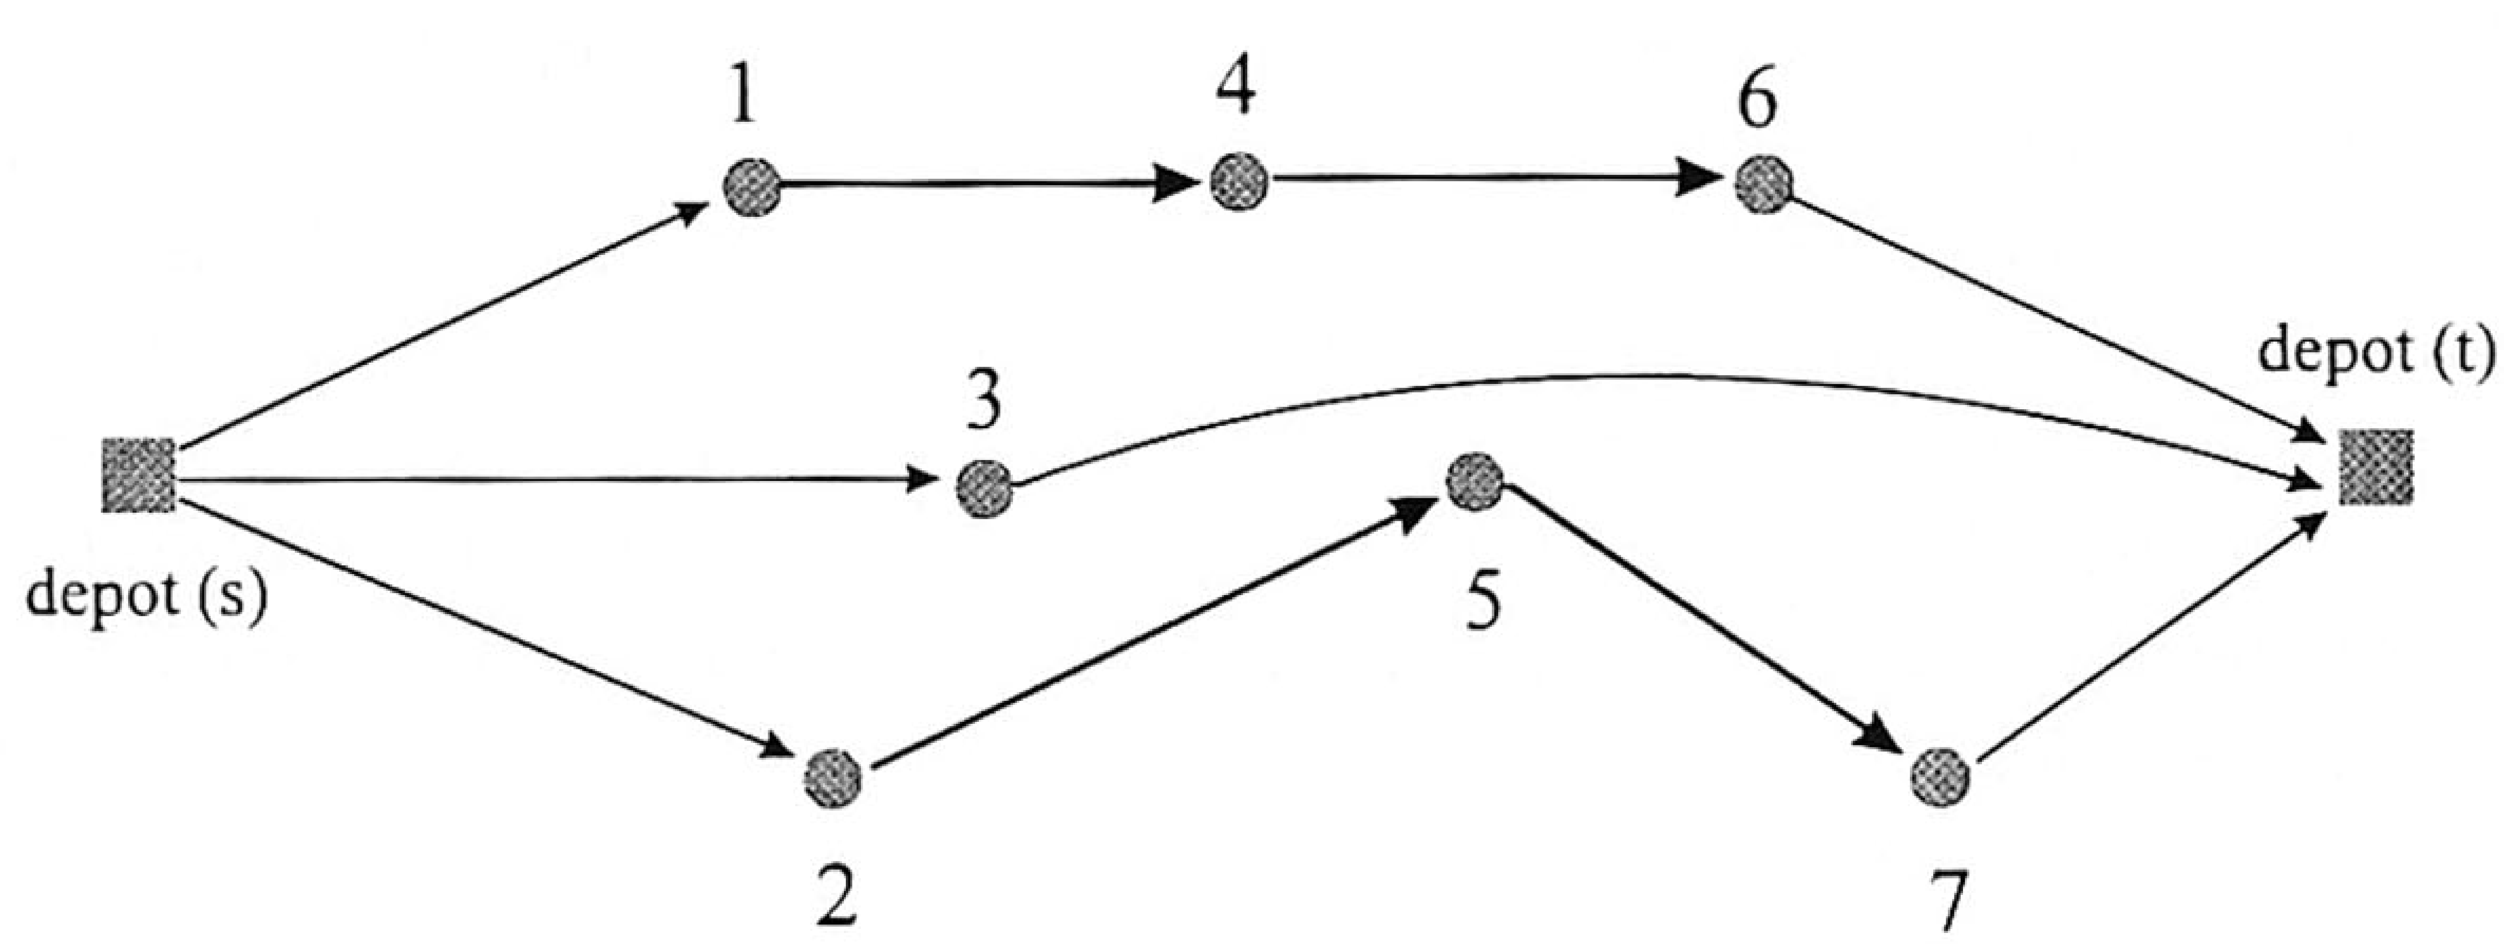

Fig. 3.2

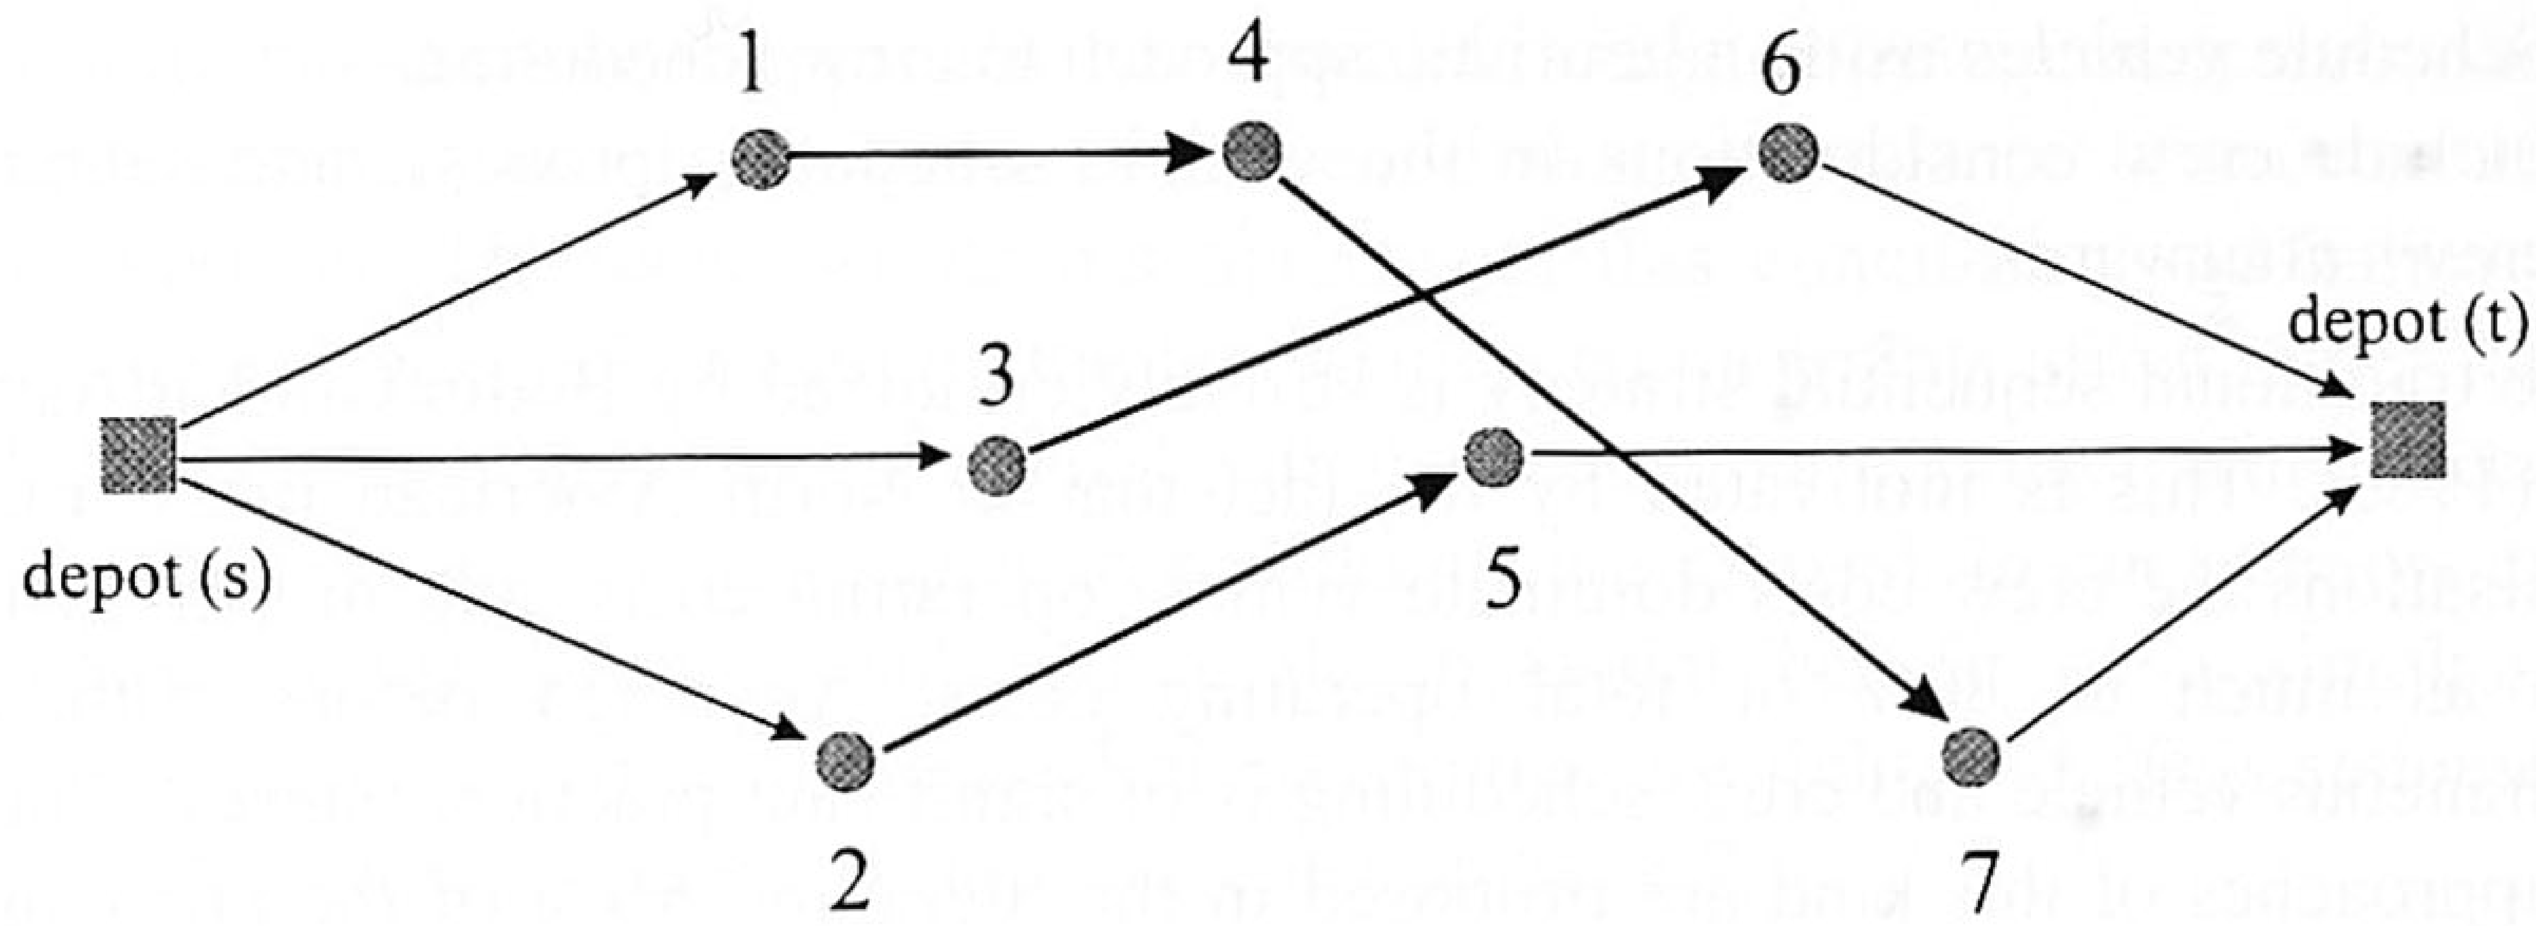

Fig. 3.3

It may be useful to have a measure of the potential benefit of integration with respect to cost efficiency, without needing to solve the VCSP. Consider the solution values of the following three approaches:

1. Traditional sequential approach (solution value v1): first solve the SDVSP and then the CSP.
2. Independent approach (solution value v2): independently solve the SDVSP and the ICSP.
3. Integrated approach (solution value v3): solve the VCSP.

The solutions of the independent approach are of no practical use since the resulting vehicle and crew schedules are usually not compatible. However, it is easier to obtain solutions for the first two approaches than for the third approach, and we know that $v2 \leq v3 \leq v1$. Thus, if v2 is significantly less than v1, it may be that the crew scheduling solution will improve significantly when considering integration compared with the sequential approach. On the other hand, we know that there is no need to integrate if v2 and v1 do not differ much.

4 Literature Review

Overviews of algorithms and applications for the SDVSP and some of its extensions can be found in Daduna/Paixão (1995) and in Desrosiers/Dumas/Solomon/Soumis (1995). Recent surveys on solution methods for the BDSP can be found in Odoni/Wilson/Rousseau (1994) and in Wren/Rousseau (1995). In this section, we discuss the literature on simultaneous scheduling of vehicles and crews. To our knowledge, this literature deals mainly with bus and driver scheduling, and all approaches proposed in the literature belong to one of the following two categories:

1. Schedule vehicles using a heuristic approach to crew scheduling.
2. Include crew considerations in the vehicle scheduling process, and schedule crews afterwards.

The traditional sequential strategy is strongly criticized by Bodin/Golden/Assad/Ball (1983). This is motivated by the fact that in North American mass transit organisations the crew costs dominate vehicle operating costs, and in some cases reach as much as 80% of total operating costs. As noted before, although simultaneous vehicle and crew scheduling is of significant practical interest, only a few approaches of this kind are proposed in the literature. Most of the procedures are of the first category and are based on a heuristic procedure proposed by Ball/Bodin/Dial (1983). This procedure involves the definition of a scheduling network, which consists of vertices characterised by parts of trips called *d-trips* that have to be executed by one vehicle and crew, and two vertices s and t representing the depot. Several types of arcs can be grouped into two categories, those which indicate that a crew and vehicle proceed from one d-trip to another and those which indicate that only the crew proceeds from one d-trip to another (*crew-only arcs*). The solution procedure is decomposed into three components, emphasising the crew scheduling problem: a piece construction component, a piece improvement component and a duty generation component. All three components are solved using matching algorithms. The piece construction routine generates a set of pieces whose time duration is less than some constant T. Note that this corresponds to vehicle scheduling with time constraints. In the second step pairs of short pieces are combined into partial duties, while in a third step pairs of these duties and longer pieces are combined into two- and three-piece duties. Vehicle schedules are generated simultaneously by deleting the crew-only arcs and fixing arcs used by pieces in the solution. This procedure is applied to large VCSP instances which correspond to the entire physical network, i.e., all lines are considered at once, while no restrictions are placed on interlining, i.e., a crew may work on an arbitrary number of lines.

Similar heuristic approaches of the first category are proposed by Tosini/Vercellis (1988), Falkner/Ryan (1992), and Patrikalakis/Xerocostas (1992). All these approaches use a similar crew scheduling network as in Ball/Bodin/Dial (1983). For the sake of illustration we briefly discuss a three phase procedure proposed in Patrikalakis/Xerocostas (1992). In the first phase, a set covering problem is solved to determine a set of crew duties which cover all timetabled trips. Because the vehicle movements are not known, the actual starting and ending times of crew duties and other parameters, such as idle time, are calculated approximately at this stage. In the second phase, a set of compatible vehicle schedules are built around the resulting duties by solving a minimum cost network flow problem. The compatibility of the crew and vehicle scheduling solutions is ensured by providing vehicles to all required crew movements. In the third phase, the crew duties are reconsidered using a restricted crew scheduling network to generate complete duties. Their conclusion is that the proposed approach can be more efficient than the traditional sequential approach when vehicle oriented crew constraints, such as a maximum number of changeovers and continuous attendance of crews, are not very important. However, we do not agree with this conclusion and believe that their conclusion is only based on the fact that such constraints affect the efficiency of the three-phase approach compared with the sequential approach which becomes more complex when vehicle oriented constraints are relaxed. In our opinion, and as discussed in the previous subsection, the potential benefit of a simultaneous approach increases when these kind of constraints are tightened. This statement is supported by computational results presented in Sect. 7.

Approaches of the second category are proposed by Darby-Dowman/Jachnik/ Lewis/Mitra (1988) as an interactive part of a decision support system, and by Scott (1985) who heuristically determines vehicle schedules which consider crew costs. An initial vehicle schedule is heuristically modified according to estimated marginal costs associated with a small change in the current vehicle schedule. The estimated marginal costs are obtained by solving the linear programming dual of the HASTUS crew scheduling model (see Rousseau/Blais/ HASTUS (1985)). Results obtained with public transport scheduling problems in Montréal, show a slight decrease in estimated crew costs which is mainly due to relatively high estimated marginal costs in the periods before the morning and evening peak hours.

5 Modelling

In Freling (1997) we propose several mathematical formulations for the VCSP and ICSP. We will briefly discuss a slightly simplified version the most important of these formulations. This formulation includes a formulation for the SDVSP. For the SDVSP, trips serviced by the same vehicle are linked by *deadheading trips* (*dh-trips*), that is, movements of vehicles without serving passengers. Dh-trips consist of travel time (or vehicle deadheading) and/or *idle time*. Idle time is defined as the time a vehicle is idle at a location other than the depot. Let b_i and e_i be the start and end locations, and let bt_i and et_i be the starting and ending times of a trip i, respectively. Two trips i and j are said to be *compatible* if the same vehicle can cover these trips in sequence, that is, if $et_i + trav(e_i, b_j) \leq bt_j$, where $trav(e_j, b_i)$ is the deadheading travel time from location e_j to location b_i. A sequence of trips is feasible if each consecutive pair of trips in the sequence is compatible.

Let $N=\{1,2,...,n\}$ be the set of trips, numbered according to increasing starting time, and let $E=\{(i,j) \mid i<j,\ i,j \text{ compatible},\ i\in N, j\in N\}$ be the set of arcs corresponding to dh-trips. The nodes s and t both represent the depot at location d. We define the vehicle scheduling network $G=(V,A)$, which is an acyclic directed network with nodes $V=N\cup\{s,t\}$, and arcs $A=E\cup(s\times N)\cup(N\times t)$. A path from s to t

in the network represents a feasible schedule for one vehicle, and a complete feasible vehicle schedule is a set of disjoint paths from s to t such that each node in N is covered. Let c_{ij} be the vehicle cost of arc $(i,j) \in$ A, which is usually some function of travel and idle time. Furthermore, a fixed cost K for using a vehicle can be added to the cost of arcs (s,i) or (j,t) for all $i,j \in$ N. For the remainder of this paper, we assume that the primary objective is to minimize the number of vehicles. This means that K is high enough to guarantee that this minimum number will be achieved. In Sect. 3, we have shown two examples of network G.

An important consideration when formulating the VCSP or ICSP is the way crew tasks are defined without knowing the vehicle schedules in advance. Recall that we consider two types of tasks, namely a set of trip tasks denoted by I_1 of which we can be sure that they have to be serviced by a crew, and a set of dh-trip tasks denoted by I_2 that need to be covered by a crew if and only if a vehicle traverses this dh-trip.

To our knowledge, only Patrikalakis/Xerocostas (1992) propose a formulation for the VCSP, but this model is only used for illustration and is computationally intractable. The mathematical formulation we propose for the VCSP contains a vehicle scheduling formulation based on network G=(V,A), which assures the feasibility of vehicle schedules. The remaining constraints in the formulation assure that each trip task is assigned to a duty and each dh-trip task is assigned to a duty if and only if its corresponding dh-trip is part of the vehicle schedule. Before providing the mathematical formulation, we need to introduce some notation. K denotes the set of all feasible duties, and K(p) is the set of duties covering trip task $p \in I_1$ or dh-trip tasks $p \in I_2$, and $I_2(i,j)$ denotes the set of dh-trip tasks corresponding to dh-trip $(i,j) \in A$. Decision variables y_{ij} and x_k are defined as follows: y_{ij} indicates whether a vehicle covers trip j directly after trip i or not, while x_k indicates whether duty k is selected in the solution or not. The VCSP can be formulated as follows (model VCSP1):

$$\min \sum_{(i,j)\in A} c_{ij} x_{ij} + \sum_{k\in K} d_k x_k$$

$$\sum_{\{j:(i,j)\in A\}} y_{ij} = 1 \qquad \forall i \in N, \tag{1}$$

$$\sum_{\{i:(i,j)\in A\}} y_{ij} = 1 \qquad \forall j \in N, \tag{2}$$

$$\sum_{k\in K(q)} x_k = 1 \qquad \forall p \in I_1, \tag{3}$$

$$\sum_{k\in K(q)} x_k - y_{ij} = 0 \qquad \forall (i,j) \in A, \forall q \in I_2(i,j), \tag{4}$$

$$x_k, y_{ij} \in \{0,1\} \qquad \forall k \in K, \forall (i,j) \in A.$$

The objective coefficients c_{ij} and d_k denote the vehicle cost of arc $(i,j) \in A$, and the crew cost of duty $k \in K$, respectively. The objective of this 0-1 linear programming problem is to minimize the sum of total vehicle and crew costs. The first two sets of constraints (1) and (2) correspond to a *quasi-assignment formulation* for the SDVSP, that is, each trip is assigned a predecessor and a successor in order to ensure that the nodes in network G are covered by a set of *s-t* paths which together include every trip node once. Constraints (3) ensures that each trip task p will be covered by one duty in the set $K(p)$. Furthermore, constraints (4) guarantees the link between dh-trip tasks and dh-trips in the solution; that is, the constraints guarantee that each dh-trip task q is covered by a duty in the set $K(q)$ only if the corresponding arc is in the vehicle solution. The model contains $\Theta(|A|+|K|)$ variables and $\Theta(|N|+|A|)$ constraints, which may already be quite large for instances with a small number of trips. This is probably the main reason why complete integration has not been considered previously. In Freling (1997) we show how to reduce the number of constraints considerably, that is, instead of $|A|$ being of order $|N|^2$ it is in practice of order $|N|$.

It is not obvious how to incorporate a restricted number of changeovers in the model. For the CSP, a restricted number of changeovers is often dealt with by restricting the number of pieces of work in a duty, that is, it is implicitly in the model due to the definition of the variables. This is valid since a piece of work corresponds to one vehicle, but also more restrictive because two pieces of work may correspond to one vehicle. However, this is not valid if the minimum number of breaks is larger than the maximum number of changeovers. For the VCSP, we can also restrict the number of pieces of work in a duty in order to deal with a restricted number of changeovers. A piece of work is here defined as a sequence of tasks in $I_1 \cup I_2$ so that a piece of work corresponds to one vehicle. In Freling (1997) we propose a mathematical formulation for the particular situation where no changeovers are allowed, while at least one crew break is required, that is, restricting the number of changeovers by restricting the number of pieces of work is not valid. We also propose a formulation for the ICSP. Both are set partitioning formulations.

6 Algorithms

In Freling et al. (1997) we propose three new algorithms for the SDVSP: an auction algorithm, a two-phase approach in case of a special cost structure, and a core oriented approach dealing with a large number of arcs in network G. For one algorithm for the VCSP we solve the SDVSP up to hundreds of times using the auction algorithm. Therefore, the computational speed of such algorithm is very important. For real life data from the RET in Rotterdam, we have solved problems with up to 1328 trips within 17 seconds on a Pentium 90 PC. For the computational tests in the next section we consider problems with up to 148 trips. The auction algorithm solves each of these problems within 0.11 seconds.

The algorithms we developed and applied to the CSP, VCSP and the ICSP are all Lagrangian heuristics with a quality guarantee. That is, we do not attempt to solve the problems to optimality, but instead use optimisation based techniques in a heuristic algorithm which produces a guarantee of the quality of the solution in terms of the difference with a lower bound. Since the late eighties there has been increasing interest among operations researchers in exact optimisation techniques for crew scheduling applications (see Desrochers/Soumis (1989) and Paixão (1990)). From an application point of view, the main reason is a rapid changing market with increased competition. On the other hand, from a technical point of view, the main reasons are the application of column generation techniques, and the ability to solve large scale integer linear programs using improved computer hardware and software. We refer to Desrosiers/Dumas/Solomon/Soumis (1995), Carraresi/Girardi/Nonato (1995), and Caprara/Fischetti/Toth (1996) for examples of successful applications of column generation to routing and scheduling problems.

We developed and implemented algorithms which consist of column generation applied to Lagrangian relaxations. The column generation is necessary since the models usually contain a huge number of variables. For each model we have at least a set of variables which corresponds to the set of all feasible duties. The motivation for considering a huge number of variables in the formulation, is that this allows for considering complex resource constraints in the pricing problem of a column generation approach. Such an approach starts with a small set of initial variables (columns) and iteratively updates the set of variables by solving a master and a subproblem while keeping the number of variables small.

The master problem corresponds to a Lagrangian relaxation, and the subproblem to one or more shortest path type of problems. The aim of column generation is to end with the set of variables that contain the optimal solution. This is achieved when no *duties price out*, including those outside the set of variables currently contained in the master problem (see Freling (1997) for more details).

In case of the CSP, the subproblem is a constrained shortest path problem (see Desrochers/Soumis (1989)). For the VCSP and ICSP we developed an approach for the subproblem which is depicted in Fig. 6.1.

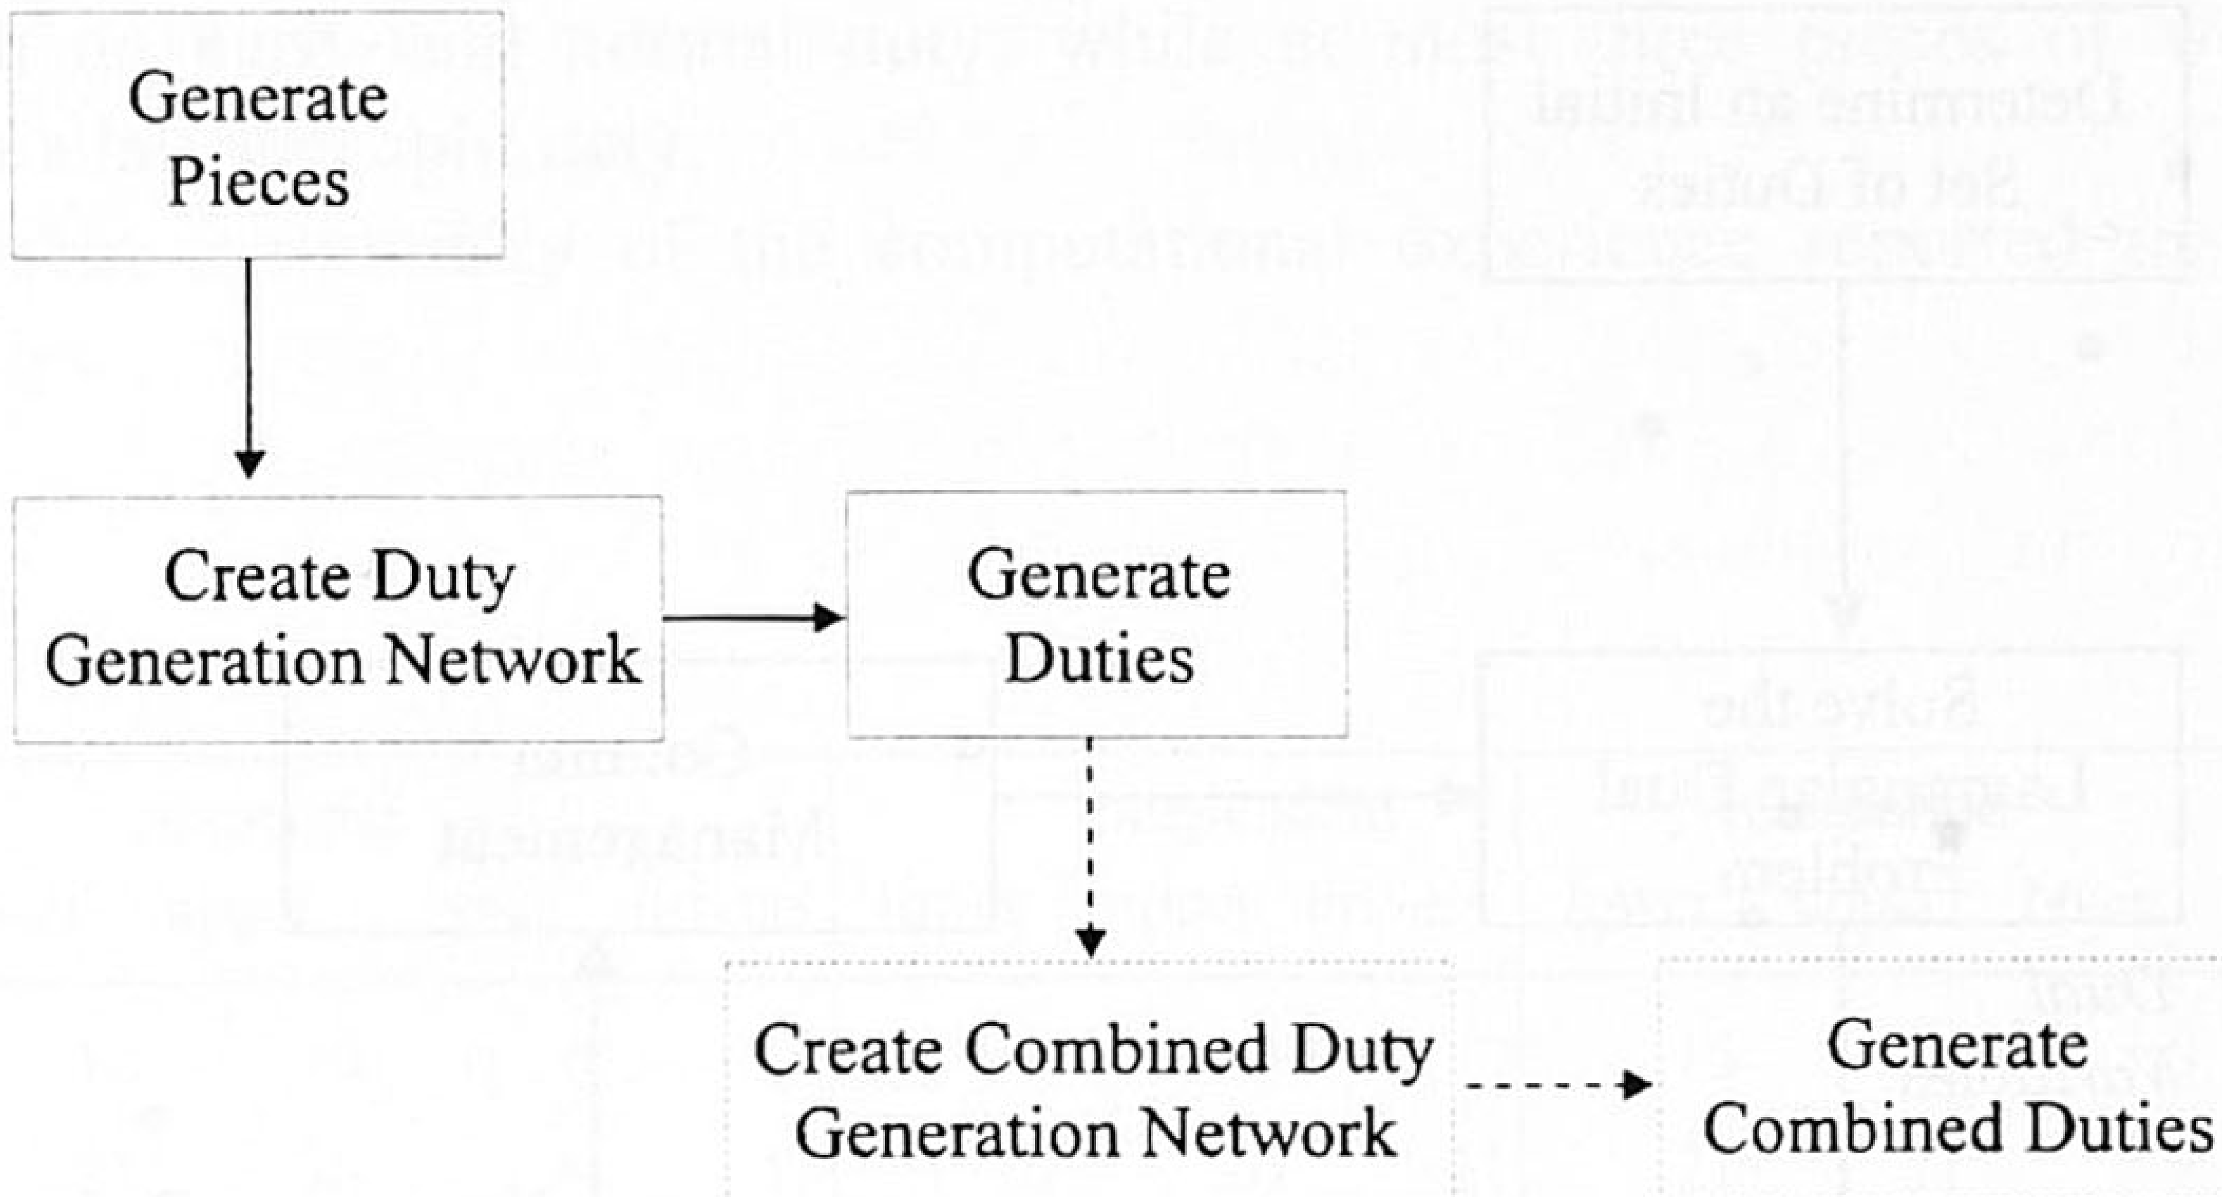

Fig. 6.1

The idea is that duties are generated in three or five steps:

1. Generate pieces: solve an all-pairs shortest path problem in an acyclic network.
2. Construct a network for duty generation.
3. Generate duties: solve a constrained shortest path problem in an acyclic network.
4. Construct a network for combined duty generation.
5. Generate combined duties: solve a sequence of shortest path problems.

Steps 4 and 5 are only necessary for the VCSP when changeovers are not allowed. In that case a duty is entirely assigned to one vehicle block, and a *combined duty* consists of one vehicle block and the crew duties assigned to it.

The structure of the algorithms for the CSP, VCSP and ICSP is depicted in Fig. 6.2. After convergence of the column generation algorithm, a heuristic algorithm is used to determine a feasible solution with the duties in the final master problem as input.

For the CSP, the ICSP and the VCSP without changeovers, we use a set covering heuristic. For the VCSP with changeovers, the heuristic consists of the following three steps:

1. determine a crew schedule using a set covering heuristic.
2. based on this crew schedule, determine a vehicle schedule which is possibly incompatible.
3. if the vehicle and crew schedules are incompatible, solve the CSP based on the vehicle schedule to get a compatible crew schedule.

Again, we refer to Freling (1997) for more details

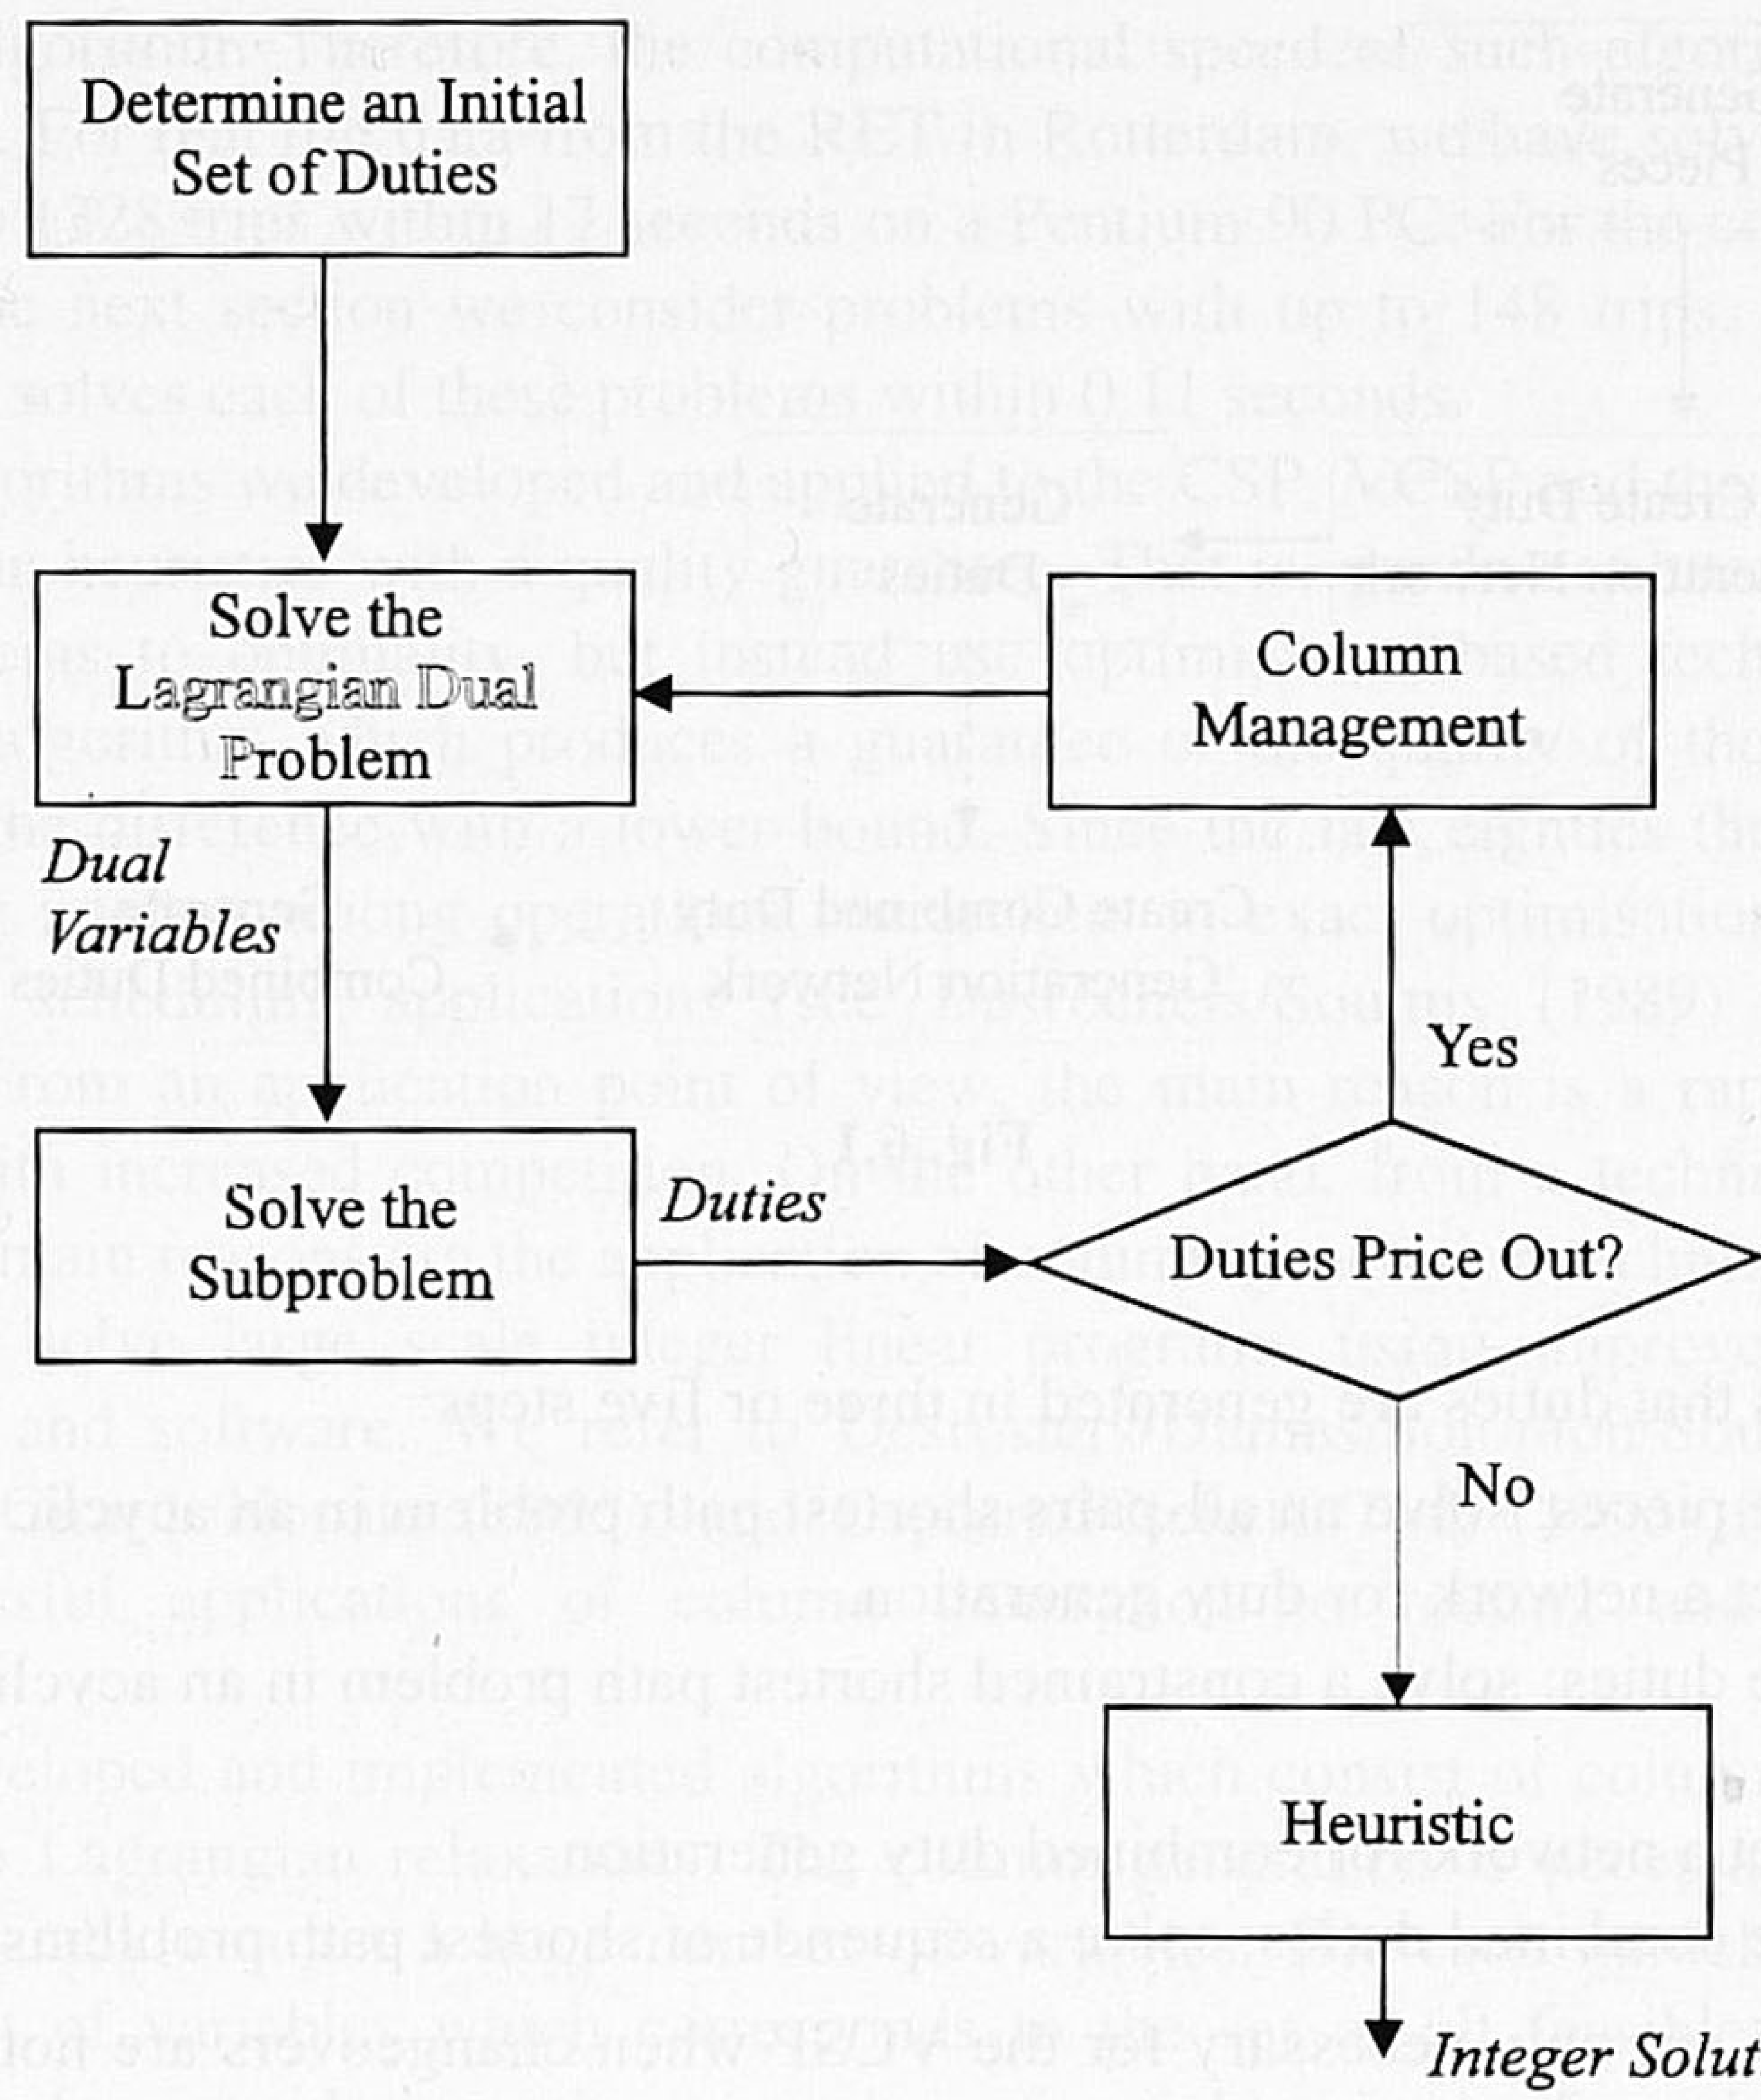

Fig. 6.2

7 Computational Experience at RET Rotterdam

We have tested the algorithms on a set of data provided by RET. This public transport company in Rotterdam provides passenger service with bus, metro and tram. The data corresponds to bus and driver scheduling for individual bus lines where the objective is to minimize the sum of the number of buses and drivers. Five duty types are allowed: a tripper, an early duty (start before 10:50 am), a normal duty (start before 3.15 pm), a late duty (start after 3.15 pm), and a split duty (break duration at least 3.5 hours). A tripper consists of one piece of work with a limited duration. The other duty types have limited duration of pieces of work, break, duration (spread time), and working time. At most two pieces of work are allowed in an early and normal duty, while at most three pieces of work are allowed in a late and split duty.

We present a summary of the computational experience reported in Freling (1997).

Table 7.1

	sequential				independent			integration			
trips	lower	upper	buses	drivers	lower	upper	drivers	lower	upper	buses	drivers
24	15	15	6	9	13	13	7	15	15	6	9
42	21	22	9	13	20	20	11	21	21	9	12
72	15	16	5	11	14	14	9	15	15	5	10
113	23	25	8	17	21	22	14	24	24	9	17
148	34	35	11	24	31	33	22	34	34	11	23

In Table 7.1 we show results obtained with five data sets, defined by the number of trips from 24 to 148. We compare solutions obtained with the sequential approach (columns 2-5), independent approach (columns 6-8) and the integrated approach (columns 9-11). Columns 2-3, 6-7 and 9-10 show the lower and upper bounds for the three approaches, respectively. In case of the sequential approach, these bounds are obtained by the summation of the minimum number of vehicles obtained by solving the SDVSP and the number of drivers corresponding to the solution obtained with the algorithm discussed in the previous section. Columns 3 and 4 show the number of buses and drivers in the solution. In case of the independent approach the bounds are obtained by solving the SDVSP and the ICSP independently. This solution is not useful in practice since bus and driver schedules are generally not compatible. However, it gives an indication of the potential benefit of integration as discussed in Sect. 3. The number of buses is the same as in column 4. The computation times on a Pentium 90 PC with 32Mb RAM are:

1. sequential approach: up to 54 seconds for the lower bounds, and up to 27 seconds for the upper bounds.
2. independent approach : up to 2814 seconds for the lower bounds, and up to 21 seconds for the upper bounds.
3. integrated approach : up to 3374 seconds for the lower bounds, and up to 63 seconds for the upper bounds.

The heuristic algorithms perform well, as can be seen from the table. All solutions for the integrated approach are optimal, and the gap between lower and upper bounds is at most two (8%) for the sequential and independent approach. Comparing the sequential with the independent approach, shows that the potential savings when integrating vehicle and crew scheduling are not very large. This is also clear from the results of the integrated approach. The upper bound is never below the lower bound of the sequential approach. In three out of five test cases, the integrated solution was better than the sequential solution due to the savings of one driver. However, it is very well possible that this difference is caused by the quality of the heuristics and not by the effectiveness of integrating vehicle and crew scheduling.

Table 7.2

	sequential				integration			
trips	lower	upper	buses	drivers	lower	upper	buses	drivers
24	22	22	6	16	16	16	7	9
42	28	30	9	21	22	22	11	11
72	18	22	5	17	15	15	5	10
113	30	33	8	25	30	34	10	24
148	46	47	11	36	39	46	14	32

In Table 7.2 we show results for the case where changeovers are not allowed. Computation times for the sequential approach are up to 10 seconds for the lower bounds and up to 27 seconds for the upper bounds, and for the integrated approach up to 1255 seconds for the lower bounds and up to 11 seconds for the upper bounds. The gap between lower and upper bounds for the integrated approach is up to seven (15%) which makes it more difficult to draw conclusions for the larger two instances. For the smaller three instances, the impact of integration is clearly more substantial compared with the results in Table 7.1. In four out five test cases, the integrated solution was better than the sequential solution due to the saving of up to 10 drivers at the cost of at most three extra buses. We can conclude that integration is better since heuristic solutions are below lower bounds of the sequential approach for three of the problems. This also confirms the expectation of the potential savings resulting from comparing the sequential with the independent approach. For example, in the case of 72 trips, the solution of the integrated approach saves seven drivers compared with the sequential approach, with the same number of buses. Interestingly, for this problem the effect of not allowing changeovers for the sequential approach is completely eliminated by considering the integrated approach.

Comparing Tables 7.1 and 7.2, it is clear that the impact of not allowing changeovers is huge.

8 Conclusion

Our first goal was to develop techniques which are applicable in practice. The results in the previous section show that we can get good solutions within reasonable computation times on a personal computer. Larger instances can be tackled by adding heuristic features to the column generation algorithm. Our second goal was to get an indication about when it is useful to integrate vehicle and crew scheduling. Although we can not draw general conclusions based upon one application, we can see that it is clear that the benefits of integration are greatest if changeovers are not allowed. In practice, it often occurs that either changeovers are not possible due to long distances or changeovers are not allowed for juridical or technical reasons In Freling (1997) we have performed tests for randomly generated data. The results support the conclusions obtained with the real life data.

For future research, we suggest to test and improve the algorithms for other, possibly larger sized, applications.

Acknowledgements: The authors would like to thank an anonymous referee for several helpful comments, and the RET company for providing the data used in the computational study.

References

Ball, M.O./Bodin, L.D./Dial, R. (1983): A Matching Based Heuristic for Scheduling Mass Transit Crews and Vehicles. in: Transportation Science 17, 4 - 31.

Beasley, J.E./Cao, B. (1996): A tree search algorithm for the crew scheduling problem. in: European Journal of Operational Research 94, 517 - 526.

Bodin, L.D./Golden, B./Assad A./Ball, M.O. (1983): Routing and Scheduling of Vehicles and Crews: The State of the Art. in: Computers and Operations Research 10, 63 - 211.

Caprara, A./Fischetti, M./Toth, P. (1996): A heuristic algorithm for the set covering problem, Working paper, DEIS, University of Bologna, Italy.

Carraresi, P./Girardi, L./Nonato, M (1995): Network Models, Lagrangean Relaxation and Subgradients Bundle Approach in Crew Scheduling Problems. in: J.R. Daduna/I. Branco/J.M. Pinto Paixão (eds.): Computer-Aided Transit Scheduling, Proceedings of the Sixth International Workshop. (Springer Verlag) 188 - 212

Daduna, J.R./Pinto Paixão, J.M. (1995): Vehicle scheduling for public mass transit - an overview. in: J.R. Daduna/I. Branco/J.M. Pinto Paixão (eds.): Computer-Aided Transit Scheduling, Proceedings of the Sixth International Workshop. (Springer Verlag) 76 - 90.

Darby-Dowman, K./Jachnik, J.K./Lewis, R.L./Mitra G. (1988): Integrated Decision Support Systems for Urban Transport Scheduling: Discussion of Implementation and Experience. in: J.R. Daduna/A. Wren (eds.): Computer-Aided Transit Scheduling: Proceedings of the Fourth International Workshop. (Springer Verlag) 226 - 239.

Desrochers, M./Soumis, F. (1989): A Column Generation Approach to the Urban Transit Crew Scheduling Problem. in: Transportation Science 23, 1 - 13

Desrosiers, J./Dumas, Y./Solomon, M.M./Soumis F. (1995): Time constrained routing and scheduling. in: M.O. Ball/T.L. Magnanti/C.L. Monma/G.L. Nemhauser (eds.): Handbooks in Operations Research and Management Science 8: Network Routing. (North-Holland) 35 - 139.

Falkner, J.C./Ryan, D.M. (1992): EXPRESS: Set partitioning for bus crew scheduling in Christchurch. in: M. Desrochers/J.M. Rousseau (eds.): Computer-Aided Transit Scheduling: Proceedings of the Fifth International Workshop. (Springer Verlag) 359- 378.

Ferland, J.A./Fortin, L. (1989): Vehicle scheduling with sliding time windows. in: European Journal of Operational Research 38, 213 - 226.

Fischetti, M./Martello, S./Toth, P. (1987): The fixed job schedule problem with spread-time constraints. in: Operations Research 35, 849 - 858.

Fischetti, M./Martello, S./Toth, P. (1989): The fixed job schedule problem with working-time constraints. in: Operations Research 37, 395 - 403.

Freling, R./Pinto Paixão, J.M. (1995): Vehicle scheduling with time constraint. in: J.R. Daduna/I. Branco/J.M. Pinto Paixão (eds.): Computer-Aided Transit Scheduling, Proceedings of the Sixth International Workshop. (Springer Verlag) 130 - 144.

Freling, R./Pinto Paixão, J.M./Wagelmans, A.P.M. (1995): Models and algorithms for single depot vehicle scheduling. Internal Report 9562/A, Econometric Institute, Erasmus University Rotterdam, Rotterdam, submitted for publication in Transportation Science.

Freling (1997): Models and techniques for integrating vehicle and crew scheduling. PhD Thesis, Tinbergen Institute, Erasmus University, Rotterdam.

Odoni, A.R./Rousseau, J.M. (1994): Models in urban and air transportation. in: S.M. Pollock/ M.H. Rothkopf/A. Barnett (eds.): Handbooks in Operations Research and Management Science 6: Operations Research and the Public Sector. (North-Holland) 129-150.

Paixão, J.M. (1990): Transit crew scheduling on a personal workstation. in: Operational Research '90, Selected Papers from the Twelfth IFORS International Conference on Operational Research, Athens, Greece.

Patrikalakis, I./Xerocostas, D. (1992): A New Decomposition Scheme of the Urban Public Transport Scheduling Problem. in: M. Desrochers/J.M. Rousseau (eds.): Computer-Aided Transit Scheduling: Proceedings of the Fifth International Workshop. (Springer Verlag) 407 - 425.

Rousseau, J.M./Blais, J.Y./HASTUS (1985): An interactive system for buses and crew scheduling. in: J.M. Rousseau (ed.): Computer Scheduling of Public Transport 2. (North-Holland) 473 - 491.

Scott, D. (1985): A Large Linear Programming Approach to the Public Transport Scheduling and Cost Model. in: J.M. Rousseau (ed.): Computer Scheduling of Public Transport 2. (North-Holland) 473 - 491.

Tosini, E./Vercellis, C. (1988): An Interactive System for Extra-Urban Vehicle and Crew Scheduling Problems. in: J.R. Daduna/A. Wren (eds.) Computer-Aided Transit Scheduling: Proceedings of the Fourth International Workshop. (Springer Verlag) 41 - 53.

Wren, A./Rousseau, J.M. (1995): Bus Driver Scheduling - An Overview. in: J.R. Daduna/I. Branco/J.M. Pinto Paixão (eds.): Computer-Aided Transit Scheduling, Proceedings of the Sixth International Workshop. (Springer Verlag) 173-187.