Logist. Res. (2012) 5:47–63
DOI 10.1007/s12159-012-0080-2

ORIGINAL PAPER

Rich vehicle routing in theory and practice

Michael Drexl

Received: 6 September 2011 / Accepted: 14 June 2012 / Published online: 3 July 2012

Abstract The contribution of this paper is a comparison of the state of the art of scientific research on and commercial software for modelling and solving vehicle routing problems. To this end, the paper presents a compact review of vehicle routing literature and an overview of the results of a recent study of commercial vehicle routing software systems with respect to the problem features these systems are able to handle and the solution methods the systems use for automatic generation of vehicle routes. In this way, existing application and research gaps are identified.

Keywords Rich vehicle routing · Commercial vehicle routing software · Heuristics

1 Introduction

Vehicle routing is a central task in a large number of private and public corporations. Routes have to be planned in very diverse sectors of the economy, not only in the logistics and transport business, but in virtually all industrial sectors producing physical goods. In addition to transport on public roads, applications of vehicle routing can also be found in intra-plant logistics, that is, local transport within a factory or warehouse building or on company premises.

Beside the considerable importance of effective and efficient vehicle routing for the enterprises themselves, the macroeconomic relevance of vehicle routing must not be overlooked: the avoidance of unnecessary or unnecessarily long routes with low capacity utilization removes pressure from road infrastructure, improves traffic flow for freight as well as passenger transport, and, by reducing emissions, makes a sustained contribution to decrease the harmful effects of transportation.

For operational research (OR), vehicle routing constitutes one of its great success stories. Vehicle routing problems (VRPs) in their many variants have been the subject of intensive study for more than half a century now. This has led to the publication of thousands of scientific papers and to the foundation of numerous software companies worldwide selling commercial vehicle routing software (CVRS). This development is certainly due to the intellectual challenge VRPs pose as well as to their practical relevance in logistics and transport. Research on VRPs is incessantly ongoing, stimulated by unsolved theoretical problems and continuous input from logistics practice.

The contribution of this paper is a comparison of scientific research on VRPs and commercial software for modelling and solving VRPs. To this end, the paper presents an overview of the results of a recent study of CVRS with respect to the problem features these systems are able to handle and the solution methods the systems use for automatic generation of vehicle routes (Drexl [31]). These findings are contrasted with the state of the art of scientific VRP research. In this way, existing application and research gaps are identified. This should be of interest for VRP researchers, and also logistics practitioners using or

M. Drexl (✉)
Chair of Logistics Management, Gutenberg School of Management and Economics, Johannes Gutenberg University, Mainz, Germany
e-mail: michaeldrexl@web.de

M. Drexl
Fraunhofer Centre for Applied Research on Supply Chain Services SCS, Nuremberg, Germany

planning to use CVRS should benefit from this paper, by learning about the potential of modern CVRS.

Throughout the paper, the following definitions apply. The fundamental activity to be planned in vehicle routing is called a *request*. A request may be a transport order, such as the delivery of a shipment from a central depot to a recipient, the pickup of a shipment from a consignor and the transfer to a central depot, the pickup of a shipment at some location and the transport to some other location, or a visit at a location to perform a service there, without picking up or delivering a physical good. *Vehicle routing* means to group requests into clusters performed by one vehicle each, and to determine, for each cluster, a complete sequence of the resulting locations to be visited. This process can be performed manually by a human planner, automatically by a computer program executing an algorithm, or by a combination of both. The goal of vehicle routing is the optimization of an objective function. This will regularly be the minimization of a cost function, of the number of used vehicles, of the total distance travelled, etc.

The rest of the paper is structured as follows. The next section describes the decisive aspects by which the numerous variants of VRPs can be distinguished. Section 3 then gives a brief overview of the state of the art of scientific VRP research. Section 4 presents the results of a comprehensive study of the German CVRS market, focussing on modelling and algorithmic aspects for the automatic solution of VRPs. Section 5 discusses the gaps between theory and practice, and Sect. 6 gives a conclusion and an outlook.

2 An overview of VRPs: dimensions of richness

The archetypal, fundamental VRP, the *capacitated vehicle routing problem* (CVRP), is as follows. Given are a set of identical vehicles stationed at one depot and equipped with a limited loading capacity, and a set of geographically dispersed customers with a certain demand for a homogeneous good. The task is to determine an optimal (with respect to an objective function) route plan, that is, a set of vehicle routes, specifying which customers are visited by which vehicle in which sequence, such that each customer is visited exactly once, the complete demand of each customer is satisfied, and the loading capacity of the vehicles is maintained on each route. The objective is to minimize overall cost or travelled distance.

As mentioned, there are a huge number of variants, extensions, and generalizations of the CVRP. VRPs can be categorized according to their properties with respect to the requests to be fulfilled, the fleet available for doing so, the desired route structure, the objectives pursued, and the considered planning horizon. An overview of these dimensions of richness in real-world VRPs is given in Fig. 1, and a discussion of these characteristics follows. (An orthogonal characterization is by application area. Industry sectors where scientific VRP research is particularly widespread are discussed in Sect. 3.2; industry sectors and fields of application where the use of CVRS is common are presented in Sect. 4.4.2.) Characteristics of rich VRPs are also discussed in Hasle and Kloster [51], Sect. 3, and Sörensen et al. [79], who state (p. 241): ‘Although there is an increasing scientific focus on so-called “rich” VRPs (that incorporate more complex constraints and objectives), they have not in any way caught up with the whole complexity of real-life routing problems’. To a large extent, this point is also supported by the results described in the present paper.

2.1 Requests

There are a large number of different aspects of requests. First of all, *time windows* are central properties of requests. Time windows can be caused by the request itself (e.g. earliest ready-time of a manufactured good to be picked up or latest delivery time of a component needed at the destination) or by the location where the request is to be performed (opening hours). There can be one or more disjoint time window(s) for a request (opening hours in the morning and in the afternoon). Moreover, time windows can be vehicle-dependent (e.g. large delivery vehicles having more restrictive access to customers in inner-city zones than small ones).

Another important aspect is *pairing and precedence*: if a request consists in the transport of a good from a pickup location to a delivery location, then, if no transshipments are allowed, one and the same vehicle must visit both locations. Moreover, it is obvious that the pickup must be performed before the delivery. There are also *complex requests*, which consist of more than one pickup and one delivery location. Often, a nested execution (so-called *LIFO loading*: pickup request A, pickup request B, deliver B, deliver A) is required (e.g. if vehicles can be loaded and unloaded only from behind).

Vehicle-driver-request compatibilities are the third fundamental aspect. Depending on the vehicle characteristics and the driver qualifications, not all requests can or may be performed by all vehicles and drivers, even if the request locations are accessible to the vehicle and the driver.

Further important types of requests are *optional requests* (requests that need not be assigned to a route, but whose execution brings a bonus), *periodic requests* (requests that have to be executed several times within a planning horizon, mostly according to visitation patterns,

Characteristics of VRPs

Requests
- Time windows
  - Single
  - Multiple
- Pairing and precedence constraints
  - Pickup before delivery
  - LIFO loading
- Vehicle-driver-request-compatibilities
- Special types of requests
  - Optional
  - Periodic
  - Expected
  - Incompatible
- Different ways of performing a request
  - Request splitting
  - Alternative delivery locations
- Request-dependent vehicle itineraries

Fleet
- Costs
  - Fixed
  - Variable
  - Penalty
  - Tariffs
- Capacity constraints
  - Weight
  - Volume
  - Loading metres
  - Pallet places
- Driving speed
  - Same for all vehicles
  - Different for different vehicles
  - Constant
  - Time-dependent
  - Load-dependent
- Temporal availability
- Initial and final vehicle positions
- Type
  - Lorry
  - Trailer
  - Train
  - Ship
  - Aircraft etc.
- Technical equipment
- Number of vehicles of each type
  - Limited
  - (Virtually) unlimited
- Drivers
  - Qualifications
  - Social legislation

Route structure
- Closed/open
- Multiple routes per vehicle
- Fixed route zones
- Visually attractive routes
- Interdependent routes with synchronization
  - Between autonomous and non-autonomous vehicles
  - Between elementary and composite vehicles
  - To perform transshipments
- Inter-route resource constraints
- Balanced routes, i.e., all routes should be similar with respect to
  - Capacity utilization
  - Number of stops
  - Duration
  - Costs etc.

Objectives
- One-dimensional
  - Min. number vehicles
  - Min. cost
  - Min. distance
  - Max. profit
- Multi-dimensional
  - Weighted sum of single objectives
  - Hierarchical ordering of objectives
  - Multi-criteria objective searching for Pareto-optimal solutions
- Hardness of constraints
  - Hard
  - Soft (with penalty in objective function)

Scope of planning
- Time horizon
  - Tactical
  - Operational
  - Real-time
- Data availability and accuracy
  - Deterministic
  - Stochastic
  - Uncertain or unknown

Fig. 1 Dimensions of richness in VRPs

for example, twice a week, but not on consecutive days), *expected requests* (requests that have not yet been issued by the customer, but will probably be), *incompatible requests* (parallel incompatibility: do not transport the requests at the same time with one vehicle; sequential incompatibility: do not transport request *B* on a route that has transported request *A* before), and indirect requests (e.g. automatic generation of empty container balancing requests).

A further very difficult aspect is when there are *different ways of performing a request*. This refers to the possibility to split up the fulfilment of a request between several vehicles, or to the possibility to perform a request by different operations altogether. An example for the first case is multi-modal transport, where a request to bring a consignment from *A* to *B* can be performed by a direct transport from *A* to *B*, via a meet-and-turn operation, or via one or several hubs. Another is that the request to deliver *x* units of a good to a customer can be fulfilled by one delivery of *x* units by one vehicle or by several deliveries by several vehicles. An example for the second case is parcel delivery, where a package must be delivered to the recipient's office address from nine to five o'clock, and to his home address after six o'clock or on Saturdays. This raises the additional question of how to choose a way of performing the request (where to perform a meet-and-turn operation, how to split a request into sub-requests, when to deliver a package).

Finally, there is the aspect of *request-dependent vehicle itineraries*. This refers to situations where the transport links a vehicle is able or allowed to use depend on the requests it is carrying, which means that requests determine the distance and the travel time between locations. For example, when a vehicle for transporting bulky goods is empty, it may be able to pass through a low undergrade crossing, but when the vehicle is loaded, it may be too high to pass and may thus have to make a detour. Similarly, if a tank vehicle for oil or fuel delivery is empty, it is allowed to travel through a water protection area; if such a vehicle is loaded, it must take an alternative, longer way.

2.2 Fleet

The term 'fleet' refers to the resources available for fulfilling the requests. These resources comprise vehicles of different types as well as drivers operating them.

2.2.1 Vehicles

Vehicles may differ with respect to several criteria, the most important ones being costs, capacities, driving speed, temporal availability, actual and desired position at the beginning and the end of a planning process, type and technical equipment, and the (in)ability to visit certain locations and use certain transport links.

Relevant *cost categories* are fixed costs for using a vehicle, and distance-, time-, and stop-dependent variable costs. Distance-dependent costs may include road tolls. Time-dependent costs may be linear or nonlinear and may include overtime pay or daily allowances for drivers. Moreover, costs may be calculated based on tariffs. Tariffs used to be mandatory in Germany until the end of the twentieth century and were dependent on goods types, weight, distance, time, etc.; although the numeric values have decreased sharply, the calculation formulas are still common in practice. For planning purposes, penalty costs are often used to consider soft constraints or undesired properties of solutions, or to allow infeasible solutions during the solution process.

The most common *capacity constraints* in goods transport are weight or payload, volume, loading metres, and number of pallet places. Several of these may be relevant at the same time.

The *driving speed* may be the same for all vehicles or differ between vehicles or vehicle types. It may also depend on the load a vehicle is carrying (the fuller, the slower). Moreover, the driving speed may be constant for all vehicles throughout the complete planning horizon, or be time-dependent. This is particularly important for short-distance and inner-city transport, where travel times are significantly higher during rush-hours compared to off-peak times.

Also, the *temporal availability* of a vehicle may be limited, for example, due to scheduled maintenance or Sunday driving bans for heavy lorries, but not for smaller vans.

In operational, short-term planning, initial *vehicle positions* (depots) are given, whereas in tactical, mid-term planning, it is often part of the planning task to determine appropriate locations for the vehicles. In operational as well as tactical planning, locations at the end of the planning horizon are arbitrary if open routes are allowed.

Concerning *type and technical equipment*, there are different criteria that determine whether or not a vehicle is in principle able to perform a request, disregarding the current point in time, location, or capacity utilization. Among these criteria are the vehicle type (lorry, train, ship, etc.), the vehicle class (swap-body vehicle, tank vehicle, etc.), vehicle dimensions and weight, and technical equipment such as a fixed installed tail-lift, a fork-lift on board, dangerous-goods equipment, etc. The dimensions of a vehicle also influence which transport links can be used (large lorries cannot use small inner-city roads, super-tankers cannot use the Panama canal). Moreover, depending on its weight or emission level, a vehicle may be unable or not allowed to use certain roads.

In tank vehicles, there are often several *compartments* that can be filled separately to allow the simultaneous transport of different goods or products or requests. When there are n compartments, n different products can be transported. When all compartments are used, no request can be executed that requires the transport of another good, even when none of the capacity constraints listed above would be violated.

The *number of vehicles* of each type and class is important, too. In reality, the number of vehicles is of course always limited. For tactical planning of the fleet size and mix, it may, however, be interesting to allow an unlimited number of vehicles of each type and class.

2.2.2 Drivers

As far as *drivers* are concerned, restrictions regarding *qualifications* limit the compatibility between drivers and vehicles as well as between drivers and requests. Such qualifications may be the type of driving licence a driver possesses, whether or not a training for the transport of dangerous goods was completed, or knowledge of customer or region specifics.

Another matter of utmost importance in real-world lorry road transport are *driver rules*. In the European Union and in other parts of the world, there is extensive social legislation on driving, working, break and rest times for drivers; see Humphreys [56] for an overview. The automatic tachograph introduced in the European Union nowadays allows for much tighter supervision of compliance with social legislation for drivers, and the road transport industry in Europe is acknowledging that today, it has to comply with the regulations very exactly. It is important to note that an algorithm cannot determine a ‘legal schedule’ for a route, because the term ‘legal schedule’ has no legally binding mathematical definition; it is a purely juristic concept. In an unlucky attempt to provide flexibility in practice, the European Union has introduced an intractably complex set of optional rules along with the mandatory ones. These rules leave a lot of room for interpretation, and a dispute about the legality of a schedule will eventually have to be settled in court. For practitioners, this means that trying to exploit the optional rules is dangerous. For algorithm developers, the optional rules mean a lot of tedious work: on the one hand, for considering them, on the other hand, for ensuring that the overall algorithm is not slowed down too much.

2.3 Route structure

Some aspects of the route structure concern each route individually, others lead to interdependencies between routes.

2.3.1 Individual routes

The standard case is a *closed route* (loop), starting and ending at the same location (depot). Nevertheless, also *open routes*, where a vehicle may be at any location at the end of its route, are relevant in many situations. For example, in long-distance road transport, vehicles are en route for a complete week, but routes are planned only for the next day, so the routes for Monday to Thursday end at the last customer scheduled for the respective day. The converse, that is, the planning of *multiple closed routes for one vehicle*, is also possible, for example, in local delivery applications, where vehicles return to the depot more than once during a day to reload.

Further possible types are routes with special geographic properties, such as the consideration of *fixed route zones* in tactical planning, and routes with limits on total duration or waiting time.

2.3.2 Interdependent routes

The usual assumption in almost all VRPs is that the only coupling or linking or joint constraints between the routes of different vehicles are related to request covering, to ensure that each request is performed exactly once. The preceding aspects leave routes *independent* of one another in this sense. The feasibility of one route does not affect the feasibility of another. However, there are also requirements that lead to route *interdependencies*, to routes that must be *synchronized*. In such cases, the feasibility of one route may depend on the feasibility of one or more other routes. Multiple synchronization of vehicles or routes may be relevant with respect to space, time, load, or common scarce resources.

One example of such a requirement is that of a 'visually attractive' *route plan*, which often means intersection-free routes. Synchronization requirements also occur when there are *different types of elementary vehicles* that may or must join and form a composite vehicle to be able to move in space or to perform a request. An example is the planning of separate routes for lorries and trailers (or tractors and semi-trailers). Each lorry and each trailer is an elementary vehicle and can be used to perform requests, and a route is computed for each lorry and each trailer that is actually used. Naturally, the route of a trailer must be synchronized with the routes of one or more lorries that must pull the trailer on the whole or on a part of its itinerary. A very similar case is the planning of separate routes or rotations for vehicles and drivers, where, during the planning horizon, a vehicle may be operated by different drivers, and a driver may drive several vehicles. This improves the temporal capacity utilization of vehicles, since these can essentially be used 24 h a day, whereas drivers need regular breaks and rests and have to obey the above-mentioned driver rules.

Furthermore, allowing *transshipments* of load between vehicles leads to interdependent routes. Transshipments occur in the form of meet-and-turn routes with exchange of complete swap-body platforms or as partial exchanges of single consignments or one-way transfer of load from one vehicle to another. Multi-modal transport, by definition, requires transshipments of load.

Also *inter-route resource constraints* such as processing capacities at depots, a maximum number of vehicles arriving at a depot per time unit due to limited number of ramps or conveyor belt capacities, etc., make synchronization between routes necessary.

Finally, there is the requirement of *balanced routes*. This refers to the stipulation that all routes of a plan be similar with respect to covered distance, duration, number of requests, costs, etc.

2.4 Objectives

Objective functions may be *one- or multidimensional*. Potential one-dimensional objectives pursued are the minimization of the number of vehicles used, of the overall distance covered by all vehicles, and of the total cost of all vehicles. If not all requests are mandatory, the objective may be the maximization of the difference between the profit obtained from the fulfilled requests and the costs incurred for fulfilling them.

When the objective comprises several dimensions, it is possible to consider a *weighted sum* of one-dimensional objectives, to have a *hierarchical* (lexicographic) ordering of the dimensions (e.g. the minimization of the number of vehicles used as the most important criterion and the minimization of cost as the second one, as the tie-breaker in the case of two route plans using the same number of vehicles), or a *multi-criteria* objective searching for Pareto-optimal solutions.

An important aspect that must be mentioned here is that many constraints and requirements discussed above can be considered as *hard* or *soft constraints*. Any violation of a hard constraint is strictly forbidden and invalidates a route plan. Constraints such as technical or logical restrictions (e.g. pickup before delivery) or legal obligations (e.g. working and rest times of drivers) are always hard. A violation of a soft constraint does not directly invalidate a route plan, but is undesired and thus considered with a penalty in the objective function. The penalty usually increases with increasing degree of violation, and if the degree of violation exceeds a certain threshold, the constraint becomes hard and invalidates the route plan. Time windows, for example, are sometimes considered as soft constraints.

2.5 Planning horizon and data availability

Depending on the *frequency of planning* and the duration during which plans remain valid, or, put differently, the range and scope of the decisions taken, there are *tactical* (medium-term), *operational* (short-term), and *real-time* (dynamic) VRPs. Tactical decisions encompass the size and composition of the fleet (number of vehicles of each type, size and technical configuration, assignment to depots) and the preparation of 'framework routes' based on average data for application areas with periodic supply or demand variation (e.g. seasonal route plans for raw milk collection at farms with high volume in spring and low volume in winter or different routes for postal distribution on different weekdays). The resulting plans may cover multiple periods. Operational vehicle routing is concerned with the planning of routes for the next day(s), based on concrete data. Real-time routing takes into account new or changing data (such as, for example, incoming requests, vehicle breakdowns, traffic congestions) and adapts plans while these are being executed.

A related aspect is the *availability and accuracy of the data* on requests, vehicles, relevant locations, and traffic links. There are three cases: the *deterministic* case, where all data are known in advance, the *stochastic* case, where some data are known in advance only in the form of probability distributions, and the case of *uncertainty*, where some data are unknown and become known only during planning or during execution of a route plan.

3 Scientific VRP research

As stated in the introduction, over the last half century, there have been thousands of scientific publications on vehicle routing, starting with the famous paper by Dantzig and Ramser [25]. Therefore, the following elaborations only give a very rough overview, a 'survey of surveys', and necessarily refer the reader to the literature for details. The existing VRP literature can be divided into theoretical papers studying models or methods for idealized or standardized problems and problem-oriented case studies dealing with concrete real-world applications. The former class considers exact as well as heuristic solution approaches and uses theoretical benchmark instances to measure the effectiveness of the devised algorithms. (A large number of benchmark instances for different types of VRP can be found at http://people.brunel.ac.uk/~mastjjb/jeb/info.html.) For the latter class, the term 'rich vehicle routing' has been coined rather recently to denote models and solution approaches for problems that feature several or all aspects of a real-world application. Most papers belonging to the latter class focus on one new or particularly interesting or difficult aspect. A number of important such aspects were queried in the CVRS study mentioned above and was discussed in detail in the previous section.

3.1 Problem variants

Researchers have devoted a lot of effort to the study of a rather small number of abstract, generic, and well-defined extensions of the CVRP, and rightly so: although these scientific variants of the CVRP hardly ever appear in practice in their pure form, their study is worthwhile, because the results and insights obtained can serve as a basis for tackling the numerous detailed and specific real-world routing problems. The most important such theoretical CVRP extensions are briefly described in the following, by pointing out in which respect these problems extend the CVRP. The cited references are surveys of the respective problem.

In the *VRP with time windows* (VRPTW, Bräysy and Gendreau [10, 11], Cordeau et al. [20]), the service at each customer must start within a given single hard time window. In the *split-delivery VRP* (SDVRP, (Archetti and Speranza [4]), customers may be visited more than once by more than one vehicle. Each vehicle may deliver a fraction of a customer's demand.

In *Pickup-and-delivery problems* (PDP, Parragh et al. [68, 69]), the tasks consist in the transport of shipments from one location to another, that is, not only the delivery locations are all different, but also the pickup locations. *Dial-a-ride problems* (DARP, cf. ib.) consider the transport of persons and, in contrast to PDPs, usually feature constraints restricting passenger inconvenience, for example, by limiting the maximum ride duration. It must be noted that there are different sub-types of PDPs, such as the VRP with backhauls or the VRP with simultaneous delivery and pickup. The reader is referred to the above surveys for a complete taxonomy of PDPs and DARPs.

In *Periodic VRPs* (PVRP, Francis et al. [37]), several visits are required to serve a customer during a multi-period planning horizon. These visits must take place in different periods. An interesting variant is the *consistent VRP* (Groër et al. [47]). This is a periodic VRP where each customer must always be visited by the same vehicle in the different periods and each customer must be visited at 'roughly the same time' on each visit.

Heterogeneous fleet VRPs (HVRP, Hoff et al. [54]), as the name implies, consider the case that not all vehicles are identical. The *fleet size and mix VRP* (FSMVRP, cf. ib.) is the tactical variant of the HVRP and considers (different) fixed costs for using (different types of) vehicles.

The *capacitated arc routing problem* (CARP, Corberán and Prins [18]) is a variant of the CVRP where the tasks are not to visit customers to perform a service, but where the service is performed while travelling along the links of a (road) network.

Location-routing problems (LRP, Nagy and Salhi [66]) combine routing and locational decisions. The task is to determine a set of vehicle routes and, for each route, the location where it starts and ends. Using a location by stationing a vehicle there incurs fixed costs.

In *stochastic VRPs* (Flatberg et al. [35], Cordeau et al. [20]), information on occurrence and volume of customer demand or travel times between customers is given by probability distributions. In *Dynamic VRPs* (Powell et al. [72]), as already outlined in Sect. 2.5, the planner is forced to make decisions before all relevant information becomes available; decisions must then be modified as new information is received. Essentially, planning is performed parallel to plan execution.

The *inventory routing problem* (IRP, Moin and Salhi [65], Andersson et al. [2]) is a very special type of VRP. In IRPs, there are no customer demands. Instead, each customer has a given consumption rate of a good, a given initial stock and a given storage capacity. The depot has to perform zero or more deliveries to each customer during a multi-period planning horizon to ensure that no customer runs out of stock. The objective is to plan routes of minimal cost for the deliveries.

There is also an increasing number of publications on algorithms for considering *VRPs with driver rules* (see, e.g. Archetti and Savelsbergh [3], Goel [43], Drexl and Prescott-Gagnon [33], Goel [44], Kok et al. [61], Prescott-Gagnon et al. [73]).

In addition to the above-mentioned surveys, Toth and Vigo [82] and Golden et al. [46] are recent monographs on VRPs and their variants. All of these references contain results on exact as well as heuristic methods.

3.2 Application-oriented research

There are some application areas where OR methods have a long-standing tradition (not only in the context of vehicle routing), and where there is a particularly large number of application-oriented papers. Such niches of applied VRP research can be found in the airline industry (Klabjan [60], Ball et al. [7]), public transport (Desaulniers and Hickman [29], Hickman et al. [53]), ship routing (Christiansen et al. [17], Hennig [52]), rail transport (Cordeau et al. [21], Caprara et al. [14]), and letter mail or parcel delivery (Bodin and Levy [9], Wong [84]).

Seminal case studies describing the successful solution of rich real-world VRPs are listed in Table 1, sorted chronologically. Note that this list is necessarily incomplete.

3.3 Solution methods

VRPs are usually modelled using graphs or networks and formulated as mixed-integer programs (MIPs). As regards solution methods, there are two fundamental approaches: *Mathematical-programming-based algorithms* on the one hand, and *heuristics and meta-heuristics* on the other.

Table 1 Selected case studies on rich VRPs

Paper	Application
Savelsbergh and Sol [77]	Dynamic, multi-period pickup-and-delivery with complex requests
Xu et al. [85]	Pickup-and-delivery with complex cost functions and LIFO loading
Hollis et al. [55]	Simultaneous and interdependent vehicle and crew routing and scheduling
Cheung et al. [16]	Synchronized routing of lorries and trailers
Irnich [58]	Arc routing with turn and street crossing restrictions, cluster constraints, and alternative service modes
Zäpfel and Bögl [86]	Simultaneous and interdependent vehicle and crew routing and scheduling with outsourcing options and working time regulations
Ceselli et al. [15]	Heterogeneous fleet, multi-depot, split-delivery VRPTW with open routes and request incompatibilities
Bock [8]	PDP with time windows and transshipment options
Oppen et al. [67]	IRP with route duration and precedence constraints using heterogeneous vehicles with compartments
Rieck and Zimmermann [74]	Simultaneous delivery and pickup with synchronization constraints at loading docks
Schmid et al. [78]	VRP with splitting of loads and synchronization of different vehicles at customer sites
Derigs et al. [26]	VRP with multiple use of tractors and trailers

3.3.1 Exact approaches

Mathematical programming algorithms are based on MIP models and, in theory, guarantee to find an optimal solution if one exists. The most successful exact algorithms for VRPs are branch-and-cut-and-price methods, which combine cut and column generation with branch-and-bound. However, mathematical programming algorithms typically require too much time and memory for large instances. Moreover, the computation times for instances of the same size and structure often vary to a large degree. At the time of this writing, CVRP and VRPTW instances of more than 200 customers cannot be consistently solved to optimality. Rich real-world instances with many complicating constraints and a realistic number of requests are still untractable with exact methods.

A milestone in the field of exact methods for VRPs is the paper by Desaulniers et al. [27], which discusses issues arising in the modelling and solution of time-constrained vehicle routing and scheduling problems using mixed-integer programming and column generation. Desaulniers et al. [28] present a monograph on column generation and branch-and-price, Røpke [75] describes several exact algorithms for VRPs and PDPs, Spoorendonk [80] treats issues related to cut and column generation, and Baldacci et al. [5] describe an exact solution framework for different types of VRP that outperforms all other exact methods published so far and solves several previously unsolved benchmark instances. Finally, Baldacci et al. [6] provide an up-to-date review of the state-of-the-art exact algorithms for the CVRP and the VRPTW.

3.3.2 Heuristics and metaheuristics

Heuristics and metaheuristics do not offer an optimality guarantee, but they overcome the limitations of exact algorithms and are able to find close-to-optimal solutions in short time, even for very large instances. Heuristic methods can be divided into constructive procedures, which are used to compute an initial feasible solution, and improvement procedures, which iteratively try to improve a given solution by systematically modifying it. Metaheuristics are superordinate procedures that control the search processes performed by constructive and improvement heuristics. Section 4.6 contains an extensive list of constructive and improvement heuristics as well as metaheuristics. Best-known heuristic solutions to benchmark instances for the problems described in Sect. 3.1 have been computed with many different methods, so there is definitely no silver bullet. However, it must be noted that most successful heuristic approaches are so-called *hybrid* procedures combining several 'classical' ones.

Method-oriented surveys or tutorials on heuristics and meta-heuristics are given by Funke et al. [38] (local search), Røpke [75] (large neighbourhood search), Ahuja et al. [1] (very large-scale neighbourhood search), Powell [71] (adaptive dynamic programming), Cotta et al. [22] (metaheuristics), and Gendreau and Potvin [41] (metaheuristics). Gendreau and Potvin [39] develop an integrating and unifying overview of metaheuristics, and Gendreau et al. [40] present a categorized bibliography of metaheuristics for several types of VRP.

Hasle and Kloster [51], Sect. 4, give a description of a commercial software for solving rich VRPs. In particular, the conceptual approach for modelling and representing rich VRPs in a software tool (as opposed to in a mathematical model) is described in detail. Moreover, the implementation of the solution algorithms used is explained. Groër et al. [48] describe a publicly available programming framework for solving VRPs and give detailed explanations of the framework's design. Both codes were not part of the survey presented in the next section. Both papers treat aspects of vehicle routing software that could not be queried in the survey.

3.4 Trends in VRP research

3.4.1 Richness and robustness

With respect to models, there is a clear trend towards considering ever 'richer' problems (Hartl et al. [50]), and towards developing generic, unified modelling frameworks (Irnich [59]) for representing these rich problems. With respect to methods, considerable progress has been made concerning the development of exact as well as heuristic solution algorithms that are *robust*, that is, work well for a broad range of problems both in terms of running time and solution quality (Pisinger and Røpke [70], Baldacci et al. [5]).

3.4.2 Self-adaptation and hyperheuristics

A related aspect is the trend towards *self-calibrating and/or self-adapting algorithms*: Many of the metaheuristics developed over the years are highly sophisticated and contain a large number of parameters for which sensible values must be set to obtain good solutions. Different problems or instances with different data characteristics require different parameter settings. Similarly, the recently proposed so-called *hyperheuristics* adapt the search space continuously in the course of the computations, based on the previous solution progress. The monograph edited by Cotta et al. [22] describes several approaches in the fields of self-calibrating and self-adapting meta- and hyperheuristic methods (see also Burke et al. [13]).

3.4.3 Matheuristics

As mentioned, the most successful heuristic algorithms are nearly always hybrid methods combining different constructive and improvement procedures, and sometimes also different meta-heuristics. A new class of hybrid meta-heuristics has emerged only recently: it was stated in the preceding section that there are two fundamental types of solution method, mathematical programming and (meta-)heuristics. Due to advances in mathematical programming theory as well as computer hardware, these two worlds have begun to merge. Hybrid methods that use mathematical programming models and algorithms as subroutines are now commonly subsumed under the term *matheuristics*. A pertinent monograph is Maniezzo et al. [64]. The paper by Doerner and Schmid [30] gives a survey on matheuristics for rich VRPs.

3.4.4 Parallel algorithms

On the technical side, one of the more recent advances in computer hardware is multi-core processors, allowing real multi-threaded processing on single, standard personal computers. This has lead to a renewed and increased interest in *parallel algorithms*. Both exact and heuristic methods can benefit from multi-threaded implementations. Crainic [23] presents a survey of parallel solution methods for VRPs.

4 CVRS: a comprehensive study of the German market

A CVRS is a computer program that allows to (1) read in and display data on vehicle depots, customers, distances, and travel times between locations, on requests, vehicles, and drivers, (2) construct, save, and display vehicle routes, and (3) determine a complete route plan for a given data set (a *problem instance*) by executing construction and improvement algorithms, possibly after entering a set of parameters, without further user interaction.

4.1 Components of CVRS

A CVRS typically consists of the following five components: an *interface to a database* or enterprise resource planning (ERP) system allows reading in the relevant data and writing back the solution. A *geographical information system* (GIS) is necessary for geocoding address data, computing distance and travel time matrices, and visualizing data and solutions in digital maps. A *planning module*, the heart of the system, supports automatic, manual, and interactive planning. A *telematics module* allows a data exchange between vehicles and the dispatching office as well as the tracking and tracing of vehicles (see Goel [42] for details). Lastly, a *statistics module* serves to compute key performance indicators and to create reports.

Usually, but not exclusively, CVRS is used for planning routes of motor vehicles on public roads. In this case, a CVRS is often embedded into a *transport management system* (TMS). A TMS contains components for data entry, planning, administration, execution, control, and billing of transport services. TMSs and other software systems for logistics and transport are thoroughly discussed in Crainic et al. [24].

4.2 Reasons for using CVRS

VRPs are, in essence, highly complex mathematical optimization problems. Because of this complexity, software for supporting human planners and decision-makers has been widely used for years. There is a considerable number of manufacturers of CVRS, and many of these manufacturers have been in the business for decades. Reasons for the use of CVRS, as specified by users, follow. The most important reason is surely that CVRS helps to reduce the costs of executing and planning routes, and to increase efficiency. Moreover, through automatic planning, dispatchers are relieved from routine jobs. Telematics and statistics functionality of CVRS improve the possibilities for transport monitoring and surveillance as well as for statistics and controlling, so that the quality and the transparency of the overall planning process is improved. An important point often raised by senior executives is that the dependency of the company on single persons (experienced dispatchers) and their knowledge is reduced. Also, the work of the sales department (of freight forwarders) is simplified, because faster and more precise pricing of ad hoc customer requests becomes possible. Finally, work processes in general are unified and streamlined. The benefit of the practical use of CVRS is substantiated by several scientific studies. See, for example, the literature survey in Eibl [34], p. 45 ff.

4.3 Structure of the study

For the compilation of the study, a thorough search for CVRS manufacturers active on the German market was performed, and no less than 50 firms could be found. All of them were asked to fill in a detailed questionnaire containing more than 500 pieces of information on relevant aspects of the company and the CVRS in nine categories:

1. Company
2. Product
3. Information technology and software engineering
4. User interface
5. Geographical information system

6. Telematics
7. Models and algorithms for automatic vehicle routing
8. Reporting, key performance indicators, and statistics
9. Prices

The questions posed were, to a large extent, either of the multiple-choice or the yes–no type. The obtained information was evaluated on an aggregate level, by summing or averaging over all questionnaires. No information on single manufacturers or systems is given. The answers in the returned questionnaires were checked for plausibility; nevertheless, correctness could of course not be verified. However, the published results being aggregated, no vendor had anything to gain from exaggerating the capabilities of his product. With respect to content, aspects that were considered relevant by the study author, based on his own professional experience, were queried.

OR/MS Today, the journal edited by the Institute for Operations Research and the Management Sciences (INFORMS), features, in a 2-year cycle, a survey on CVRS for the North American market. The latest one, from February 2012, comprises 12 vendors (see http://www.orms-today.org/surveys/Vehicle_Routing/vrss.html). The results of both studies are hard to compare, because different information was gathered and organized differently. Only one company appeared in both surveys.

4.4 General results

Twenty-eight companies sent back a filled-in questionnaire. This is a return rate of 56 %, which is acceptable.

4.4.1 Company structure and size

Most CVRS companies have their headquarters in Germany. The number of employees is 36 on average. The first manufacturer of 'software for logistics' was founded in 1961. The first CVRS, that is, vehicle routing software featuring an automatic, algorithm-based planning component, was offered in 1979.

All firms offer launching and roll-out support as well as user-specific adaptation and customization of their software. (This shows that CVRS is (still) not a standard, off-the-shelf product.) In addition, most companies use their own software for project work and consulting services.

Most firms, but not all, consider the algorithms used in their systems a core competence. Only four companies do not possess the source code of the algorithms and do not hold exclusive rights on the code. These companies specialize in transport management systems (TMS) and use third-party components for automatic planning.

Cooperation with academia is common. More than three quarters of all firms stated that they cooperate with at least one university. This is mostly done in the form of master's and Ph.D. theses. Twelve firms offer a free test or demo version, and several firms provide low-priced or free licences for use in teaching.

A basic, single-user licence for commercial use costs 15,000 Euros on average. This does not include customizing and preparatory training of users.

4.4.2 Industry sectors using CVRS

As for the industry sectors using CVRS and the respective number of sold licences, no reliable data could be gathered. Some firms provided detailed data, others did not or gave only aggregated information. However, a quintessential finding is that numerous CVRS firms have customers in the following sectors: *Industrial firms producing physical goods* use CVRS to plan the sourcing of raw materials and the transport of (semi-)finished goods between plants and warehouses or to wholesalers. Concrete applications are as diverse as milk collection at farmyards or wood transport from forests to mills on the supply side, and delivery of finished cars or ready-made concrete garages on the distribution side. In the *wholesale and retail trade*, the distribution of consumer goods such as drinks, frozen food, furniture, or heating oil, to name but a few, is planned with CVRS. *Freight forwarders* in the less-than-truckload as well as the full-truckload business determine routes for local feeding and distribution as well as linehaul and long-distance tramp transports with CVRS. The same holds for *parcel delivery companies* and *letter mail services*. Also many *reverse logistics and waste collection firms* plan routes with the help of CVRS. An important field across sector boundaries is the solution of service technician, salesman, and other *staff dispatching problems*. Finally, in *intra-plant logistics* of industrial firms, routes of automated guided vehicles used to fulfil transport orders between warehouses and production sites are planned with CVRS.

4.5 Models and algorithms for automatic vehicle routing in CVRS

Obviously, a central part of a CVRS and the most interesting one from an OR perspective is the automatic planning component. Therefore, detailed information on this aspect was queried. This part of the questionnaire was answered by 27 participants.

4.5.1 General features of the automatic planning component

Most systems possess the following functionality: they are capable of solving general PDP as described above with up to 10,000 requests, including problems with a heterogeneous fleet, a multi-period planning horizon, and multiple use of vehicles, providing a feasible solution within five minutes when distance and travel time matrices are given. Most systems can be used to determine only one part of the solution of a VRP, namely the clustering of the requests into groups to be performed by one vehicle (whereas the other part, the sequencing of requests, is left to the dispatcher, or, even more often, to the drivers), and allow the automatic assignment of vehicles and drivers to such groups. Moreover, a re-optimization component capable of computing a new feasible solution from an existing one after small changes to the instance, such as the arrival of some new requests, is a standard feature. Another is the possibility to limit the duration of an optimization run by specifying a maximal number of iterations and a maximal running time. Lastly, a batch mode for automatic execution of optimization runs with specific parameter settings is also usual. Interestingly, only eight firms claim to be able to solve arc routing and postman problems. Moreover, although there is a considerable amount of scientific literature on the topic (see the survey by Wäscher et al. [83]), only six systems contain a module for optimizing storage space utilization (using 3D packing algorithms).

4.5.2 Modelling features

4.5.2.1 Request-related features Most systems can handle single as well as multiple time windows for requests and locations, whereas vehicle-dependent time windows are only seldom considered.

The obvious ‘pickup-before-delivery’ precedence constraint can be handled by all systems. The requirement of request precedence within a route is also commonly covered. A nested execution is supported by half of all systems, but only few systems can deal with precedence constraints of requests on different routes.

The large majority of the systems allows the consideration of vehicle-driver-request compatibilities, and more than two-thirds of all systems are able to handle optional requests and parallel incompatibility between requests. About half of all systems can deal with periodic requests and with complex requests consisting of more than one pickup and one delivery location. This indicates that these request types are commonly encountered in practice. On the other hand, only few systems can consider expected or indirect requests, sequential request incompatibility, and different ways of performing a request. No system supports request-dependent vehicle itineraries.

4.5.2.2 Fleet-related features A homogeneous fleet is rarely found in the real world. Therefore, the ability to consider heterogeneous vehicles, at least with respect to costs and capacity, is an absolute must for a CVRS. The surveyed systems generally support different fixed, distance- and time-dependent costs. More than half of all systems also offer different stop-dependent costs and the use of tariffs; penalty costs for considering soft constraints are prevalent, too. The simultaneous consideration of different capacity constraints is standard as well. Additionally, almost two-thirds of the systems offer support for multi-compartment vehicles.

An issue pointed out by several participants of the study is that the available geographical data for lorry routing still do not cover all relevant attributes (passage heights, barred roads, etc.) in a truly reliable fashion even in Western Europe. This is because the commercial providers of geographical data have concentrated on the much larger car navigation system market. This situation is about to change, though, and many systems already consider these data in the preprocessing phase, when distance and travel time matrices are computed.

Different driving speeds per vehicle and time- or load-dependent travel times are rarely supported. This is probably because the necessary data are still difficult to obtain in a sufficiently high quality. Tactical fleet planning by directly specifying an unlimited number of vehicles of each type is offered by half of all systems; in the other systems, a sufficiently large number of vehicles must be specified by hand if the fleet size and mix is to be determined.

Although the concrete type and class of a vehicle are irrelevant for a solution algorithm, the surrounding software must be able to manage these attributes in order to provide sensible feedback to the user. In this respect, most systems are capable of considering lorries/tractors, trailers/semi-trailers, and cars, but only few systems can also manage pedestrians, bicycles (both of which are relevant, for example, in mail delivery), trains, ships, or aircraft. Similarly, the technical equipment of a vehicle is only relevant to determine vehicle-request compatibility, and the dimensions and weight of a vehicle are only relevant to determine vehicle-location compatibility. The consideration of vehicle class, type, and technical equipment attributes for specifying vehicle-request-location compatibilities is possible with more than four-fifths of the systems.

All CVRS manufacturers claim that their systems possess the feature of considering the European Union social legislation on driving, break, and rest times for drivers. During the study, however, the author has gained the impression that many systems consider these regulations only incompletely. In particular, only few systems, according to the study, contain the rules for double-manned vehicles.

4.5.2.3 Route-related features Most systems allow the planning of closed and open routes as well as multiple routes for one vehicle. Further types of route supported by most systems are routes for lorries and trailers, where the trailer can be uncoupled and left behind at parking places, but with a fixed assignment of lorry and trailer (so that only a route for the lorry is computed, and the route of the trailer is a part of its lorry's route), the consideration of fixed route zones from tactical planning, and routes with a maximum waiting time. Almost half of all systems support the computation of balanced routes with similar capacity utilization, number of stops, duration, and cost.

The necessity to compute interdependent routes with synchronization requirements between autonomous and non-autonomous vehicles such as lorries and trailers, between elementary and composite 'vehicles' such as lorries and drivers, or between arbitrary vehicles to perform transshipments is frequent in practice. However, these features are supported by few systems only. Mostly, these requirements are left to manual planning. The same holds for inter-route resource constraints.

4.5.3 Objective functions

With respect to objective functions, almost everything that is reported in the literature is also available in all or most systems: it is possible to minimize the number of vehicles used, the overall distance covered by all vehicles, and the total cost of all vehicles. In addition, about half of all systems support a weighted sum of one-dimensional objective functions, hierarchical, or multi-criteria objective functions.

4.5.4 Planning modes

The classical CVRP corresponds to an operational, single-period, static, deterministic planning situation. Given that reality is neither static nor deterministic, it is not surprising that CVRS supports different planning modes.

With respect to the frequency of planning, tactical planning of standard or base routes using aggregate, average data as well as operational, day-by-day planning is supported by all systems. Some systems use different algorithms for these two modes, taking into account that for tactical planning, running time is not critical. Multi-period or rolling horizon planning is also supported by most systems.

In addition, dynamic or real-time planning is possible with most systems (changing assignments of requests to vehicles while the latter are already en route, caused by events such as new requests or breakdowns of vehicles).

Interactive planning is also supported in most cases. This means that the user can make small changes to an existing plan proposed by the algorithm, such as fix assignments of requests to routes, fix the sequence of partial routes, manually assign a certain vehicle to a route, etc. Both features, dynamic and interactive planning, require the capability to re-optimize an existing plan after small changes, without changing the fixed parts.

Only three companies state that their algorithms are capable of handling stochastic customers (where the necessity to visit a location is stochastic) or stochastic demand/supply (where the amount of demand/supply is stochastic).

4.6 Algorithmic features

Whereas participating in the study was seen as a marketing measure by most CVRS manufacturers, several firms were reluctant to specify details about the algorithms used in their software. However, the questions in this part of the questionnaire were still answered by 21 firms, so that also these results may be considered representative.

Of course, any algorithm used for vehicle routing in practice is necessarily a heuristic. Therefore, the questions on algorithms asked which construction and improvement procedures and which metaheuristics were used.

4.6.1 Constructive procedures

The ranking of constructive procedures is as follows (number of mentions in parentheses):

1. Parallel savings (16)
2. Insertion (11)
3. Cluster first, route second (10)
   Nearest Neighbour (10)
5. Proprietary (7)
   Sequential savings (7)
7. Dynamic programming (6)
8. GENI intra-route (5)
9. Regret (4)
   Route first, cluster second (4)

4.6.2 Improvement procedures

The ranking of improvement procedures is as follows:

1. Relocate (move a request to another route) (16)
2. *k*-opt (15)
3. Swap/exchange (exchange two requests between two different routes) (13)
4. String-relocate (move a route segment to another route) (12)
5. Cross/string-exchange (exchange two route segments between two different routes) (11)

6. Or-opt (8)
7. k-opt* (generalization of k-opt to capacitated problems) (7)
8. Lin-Kernighan (6)
9. GENI inter-route (5)
   λ-interchange (exchange at most λ requests between two routes) (5)
   (Very) Large-scale neighbourhood search ((V)LSNS, exponential-size neighbourhoods) (5)
12. Double-bridge move (3)
    Ejection chains/cyclic transfers (move a fixed number of requests from route 1 to route 2, then the same number from route 2 to route 3 etc.) (3)
14. Proprietary (2)

4.6.3 *Metaheuristics*

The ranking of metaheuristics is as follows:

1. Tabu search (10)
2. Genetic algorithms (8)
3. Threshold accepting (7)
4. Proprietary (6)
   Ruin-and-recreate/fix-and-optimize/ripup-and-reroute (6)
   Simulated annealing (6)
7. Adaptive large neighbourhood search (5)
8. Ant colony systems (4)
   Guided local search (4)
   Variable neighbourhood descent (4)
   Variable neighbourhood search (4)
12. Greedy randomized adaptive search procedure (3)
    Memetic algorithms (3)
14. Adaptive guided evolution strategies (2)
    Attribute-based hill climber (2)
    Backbone search (2)
    Great deluge (2)
    Indirect search (decoder) (2)
    Neural networks (2)
    Scatter search (2)
21. Adaptive/approximate dynamic programming (1)
    Adaptive memory programming (1)
    Artificial immune systems (1)
    Particle swarm optimization (1)
    Record-to-record travel (1)

4.6.4 *Mathematical-programming-based approaches*

The ranking of mathematical-programming-based approaches is as follows:

1. Branch-and-cut (5)
   Constraint programming (5)
3. Column generation (4)
4. Branch-and-price (3)
5. Benders decomposition (1)
   Lagrangian relaxation (1)

4.6.5 *Components, libraries, and benchmarks used*

The ranking of solvers, programming frameworks, and algorithm libraries is as follows:

1. CPLEX (3)
2. Boost (2)
   COIN (2)
   LEDA (2)
   XPRESS (2)
6. BCP (1)
   CBC (1)
   Gurobi (1)
   lp_solve (1)
   SCIP (1)
   SoPlex (1)

Eight firms stated that they tested their algorithms with the Solomon VRPTW benchmarks: six have used the Gehring/Homberger VRPTW problems and two the Li/Lim PDPTW instances. No firm was willing to tell anything about the results.

5 The gaps between theory and practice

It goes without saying that any CVRS is a commercial product for end-users without programming and OR skills. Therefore, it must offer an up-to-date graphical user interface with adaptable look-and-feel as well as a help system. Moreover, a comfortable interface to common TMS or ERP systems is also an essential feature. (Together with the usual GIS, telematics, and statistics modules, this means that, typically, less than 10 % of the code of a CVRS is for the VRP solution algorithms.) Moreover, aspects such as versatility, genericity, and maintainability are fundamental for any commercial software and must be ensured through adequate design, thorough testing, and extensive documentation. This slows down the development process considerably. A scientific code, on the other hand, is often written as a prototype, a proof-of-concept, to be used only by the developers themselves.

CVRS must be able to handle many different problems. It is not an option to develop and implement a specialized algorithm for each new customer. Therefore, algorithms used in CVRS must necessarily be generic and easily extendable to new problem features. Summing up the previous sections, an algorithm for use in a state-of-the-art CVRS supports the following features:

- Pickup-and-delivery requests
- Compatibility between locations, requests, vehicles and drivers
- Multiple time windows for locations and requests
- Consideration of service times
- Heterogeneous fleet with respect to cost, capacity, start and end depots
- Fixed, distance-, time-, stop-dependent, penalty costs, tariffs
- Multiple capacity constraints
- Multiple use of vehicles
- Driver rules
- Weighted and hierarchical cost functions
- Dynamic planning over a one-week planning horizon with event- or time-based rolling horizon planning
- Re-optimization options
- Interactive planning

Few algorithms described in the literature, if any, are able to deal with all these features. Research algorithms usually work for special, mostly idealized, types of VRP only.

Practitioners need robust, fast, extensible, and simple (parameter-free) algorithms capable of solving instances with thousands of requests. The last 0.1 % in solution quality to be gained from an additional complex algorithmic device are insignificant, since the data available in practice are never 100 % accurate. On the other hand, when using exact methods, the scientific world strives to find 'provably optimal' solutions for small, idealized problems; using heuristics, researchers are, to a large extent, focussed on improving best-known results for benchmark instances or solving concrete real-world problems with prototypical implementations of specialized algorithms. This point is further elaborated in Cordeau et al. [19] and Pisinger and Røpke [70].

A central argument in the above-mentioned paper by Sörensen et al. [79] is that, according to these authors' experience, commercial CVRS uses quite a large number of improvement heuristics to improve initial solutions determined by constructive procedures. This is in contrast to scientific codes, which tend to use few, but rather complex and sophisticated techniques. The reasons for this are, according to Sörensen et al., that (1) an approach using many improvement procedures can overcome the greedy behaviour of an approach that uses only a single one and (2) supplying a large arsenal of diverse search strategies allows a flexible adaptation of the software to the specific requirements of each customer. The results of the study described in the present paper support this observation: as can be seen in Sect. 4.6.2, 21 routing tools use 111 different improvement procedures, which corresponds to an average of more than five.

5.1 Application gaps

Aspects that are rather well studied in theory but have not yet found widespread use in CVRS and where, consequently, there is an application gap are *stochastic vehicle routing* (Flatberg et al. [35], Cordeau et al. [20]), *time-dependent travel times* (Fleischmann et al. [36], Taniguchi and Shimamoto [81], Haghani and Jung [49]), and *mathematical-programming-based approaches* (Maniezzo et al. [64]). As far as stochastic VRPs are concerned, this is probably because in most cases, it is extremely difficult, if not impossible, to provide sufficient data in sufficient quality to derive useful probability distributions for customer demands/supplies, and it is doubtful whether this is going to change soon. The consideration of variable travel times (driving speeds) depending on the time of day also requires reliable data. Here, the outlook is more optimistic. Detailed information will be available in the foreseeable future, at least for large urban regions, which is where peak and off-peak times are most pronounced anyway. Mathematical-programming-based approaches (matheuristics) are still a rather new field of research, but more and more pertinent publications appear. CVRS manufacturers will not ignore this trend.

5.2 Research gaps

On the other hand, a research gap is apparent with respect to the following aspects: there are only very few papers on problems with *optional, expected, indirect, or complex requests*. The existing literature is mostly concerned with deterministic problems where all requests must be fulfilled and consist of single visits or one pickup and one delivery. Exceptions to this rule are Savelsbergh and Sol [76], Røpke [75] and Goel and Gruhn [45]. What is more, practitioners often find it difficult to give a clear-cut definition of their problem's objectives and constraints. Therefore, *soft constraints* such as visual attractiveness of routes (Lu and Dessouky [62]) or a preferred assignment of certain drivers or vehicles to customers (Groër et al. [47]) are quite important in real-world applications. For the same reason, a 'fair' and *balanced sharing of the workload* between different routes is important in practice, but seldom considered in the literature (an exception is the paper by Bredström and Rönnqvist [12]). *Tariffs and complex cost functions* are often used in practice, but rarely considered in the literature. See Ceselli et al. [15] for a striking exception.

Most importantly, models and algorithms for *integrated and synchronized vehicle routing* are still scarce: in almost all vehicle routing models and algorithms, the routes of the different vehicles are assumed to be independent of one another, so that modifying one route does not have any effects on other routes. However, in a surprisingly high number of cases, this assumption does not hold. Examples for practical applications requiring a spacial, temporal, and in some cases also load-related synchronization or coordination of routes are the planning of inter- and multi-modal transports, the planning of meet-and-turn routes, transports over hubs or cross-docking locations, simultaneous planning of routes for lorries and trailers, if trailers may be pulled by different lorries, simultaneous vehicle and driver routing, if drivers may change vehicles, and automatic planning of multiple types of resources (drivers, lorries/tractors, trailers/semi-trailers, swap-bodies/containers).

The following quote taken from Irnich [57], p. 9, still holds: ‘While research on integrated models and solution methods for combined vehicle and crew *scheduling* has made some remarkable advances ..., the literature on integrated vehicle *routing* still mainly focuses on *location routing problems* and *inventory routing problems*. Literature on other forms of integration is scarce. There is a need for new and improved techniques to attack integrated planning problems. As far as we can see, there is no convincing concept for dealing with VRPs with load transfer at hubs or consolidation points, especially in the context of bimodal or multimodal traffic. The same is true for long-haul goods traffic, which requires the coordination between feeder processes, linehaul, and distribution’.

More specifically, Macharis and Bontekoning [63], p. 400, state in their survey of inter-modal freight transport that ‘intermodal freight transportation research is emerging as a new transportation research application field, that it still is in a preparadigmatic phase, and that it needs a different type of models than those applicated to uni-modal transport’.

In short, a general, unifying modelling and solution concept for integrated and synchronized vehicle routing is still missing; science has to catch up in this respect. This statement is further supported by the fact that the survey article by Gendreau et al. [40], which presents an extensive literature list on VRPs, does not mention a reference on VRPs with multiple synchronization constraints. Moreover, the recent monograph by Golden et al. [46] does not contain a paper on this topic. This does not mean that there are no such papers at all. Rather, it shows that no systematic study of this problem class has yet been performed. Such research is now beginning to emerge (see the survey by Drexl [32]). On the other hand, the results of the study show that some CVRS systems contain solutions to concrete problems in this area, and judging from this author’s professional experience, there is considerable demand for powerful decision support tools for integrated and synchronized vehicle routing in practice.

6 Conclusions and outlook

The exact solution of even the basic variants of VRPs is still impossible for instances of realistic size. An exact solution of real-world problems with many additional side constraints will remain impossible in the short and medium term. However, close-to-optimal solutions of more and more complex and integrated problems, increasingly based on incomplete optimization approaches and mathematical-programming-based heuristics, are possible, and this is sufficient to provide useful decision support in practice. Nevertheless, as has already been alluded, in some areas there are gaps between industrial needs and the state-of-the-art CVRS of today. A detailed discussion of these issues is, however, beyond the scope of the present paper, but constitutes an interesting topic for further research.

For the foreseeable future, CVRS will remain a decision *support* system in almost all application areas. Essentially, fully automatic planning is possible only in some special cases, most notably in intra-plant logistics. In road and inter-modal transport, interactive planning with a human dispatcher having the final say is and will remain the rule. A modern CVRS, however, can considerably facilitate the daily routine work for human decision-makers. The systems have become so mature and user-friendly that, nowadays, after introducing a CVRS, nobody wants to return to purely manual planning any more. The concerns often voiced by many dispatchers, CVRS would invalidate their knowledge and experience or even make them lose their jobs, are unfounded. This has become evident in many discussions with manufacturers as well as users of CVRS.

Summing up, CVRS constitutes a fixed and indispensable component of logistics planning in practice. Just like any other product, CVRS has to adapt to ever-changing customer needs and expectations. This requires constant further development, both with respect to information technology and to OR models and algorithms. Consequently, even more than half a century after the first OR paper in this field, the practice of vehicle routing will continue to provide interesting and challenging problems for OR researchers.

References

1. Ahuja R, Ergun Ö, Orlin J, Punnen A (2002) A survey of very large-scale neighborhood search techniques. Discret Appl Math 123:75–102

2. Andersson H, Hoff A, Christiansen M, Hasle G, Løkketangen A (2010) Industrial aspects and literature survey: combined inventory management and routing. Comput Oper Res 37:1515–1536
3. Archetti C, Savelsbergh M (2009) The trip scheduling problem. Transp Sci 43:417–431
4. Archetti C, Speranza M (2008) The split delivery vehicle routing problem: a survey. In: Golden B, Raghavan S, Wasil E (eds) The vehicle routing problem: latest advances and new challenges. Springer, New York, pp 103–122
5. Baldacci R, Bartolini E, Mingozzi A, Roberti R (2010) An exact solution framework for a broad class of vehicle routing problems. Comput Manag Sci 7:229–268
6. Baldacci R, Mingozzi A, Roberti R (2012) Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. Eur J Oper Res 218:1–6
7. Ball M, Barnhart C, Nemhauser G, Odoni A (2007) Air transportation: irregular operations and control. In: Barnhart C, Laporte G (eds) Transportation, volume 14 of handbooks in operations research and management science. Elsevier, Amsterdam, pp 1–67
8. Bock S (2010) Real-time control of freight forwarder transportation networks by integrating multimodal transport chains. Eur J Oper Res 200:733–746
9. Bodin L, Levy L (2000) Scheduling of local delivery carrier routes for the united states postal service. In: Dror M (ed) Arc routing: theory, solutions, and applications. Kluwer, Boston, pp 419–442
10. Bräysy O, Gendreau M (2005) Vehicle routing problem with time windows, part i: route construction and local search algorithms. Transp Sci 39:104–118
11. Bräysy O, Gendreau M (2005) Vehicle routing problem with time windows, part ii: metaheuristics. Transp Sci 39:119–139
12. Bredström D, Rönnqvist M (2008) Combined vehicle routing and scheduling with temporal precedence and synchronization constraints. Eur J Oper Res 191:19–29
13. Burke E, Hyde M, Kendall G, Ochoa G, Özcan E, Woodward J (2010) A classification of hyper-heuristic approaches. In: Gendreau M, Potvin J-Y (eds) Handbook of metaheuristics, 2nd edn. Springer, New York, pp 449–468
14. Caprara A, Kroon L, Monaci M, Peeters M, Toth P (2007) Passenger railway optimization. In: Barnhart C, Laporte G (eds) Transportation, volume 14 of handbooks in operations research and management science. Elsevier, Amsterdam, pp 129–187
15. Ceselli A, Righini G, Salani M (2009) A column generation algorithm for a rich vehicle-routing problem. Transp Sci 43:56–69
16. Cheung R, Shi N, Powell W, Simão H (2008) An attribute-decision model for cross-border drayage problem. Transp Res Part E 44:217–234
17. Christiansen M, Fagerholt K, Nygreen B, Ronen D (2007) Maritime transportation. In: Barnhart C, Laporte G (eds) Transportation, volume 14 of handbooks in operations research and management science. Elsevier, Amsterdam, pp 189–284
18. Corberán A, Prins C (2010) Recent results on arc routing problems: an annotated bibliography. Networks 56:50–69
19. Cordeau J, Gendreau M, Laporte G, Potvin J, Semet F (2002) A guide to vehicle routing heuristics. J Oper Res Soc 53:512–522
20. Cordeau J, Laporte G, Savelsbergh M, Vigo D (2007) Vehicle routing. In: Barnhart C, Laporte G (eds) Transportation, volume 14 of handbooks in operations research and management science. Elsevier, Amsterdam, pp 367–428
21. Cordeau J, Toth P, Vigo D (1998) A survey of optimization models for train routing and scheduling. Transp Sci 32:380–404
22. Cotta C, Sevaux M, Sörensen K (eds) (2008) Adaptive and multilevel metaheuristics, volume 136 of studies in computational intelligence. Springer, Berlin
23. Crainic T (2008) Parallel solution methods for vehicle routing problems. In: Golden B, Raghavan S, Wasil E (eds) The vehicle routing problem: latest advances and new challenges. Springer, New York, pp 171–198
24. Crainic T, Gendreau M, Potvin J (2009) Intelligent freight-transportation systems: assessment and the contribution of operations research. Transp Res Part C 17:541–557
25. Dantzig G, Ramser J (1959) The truck dispatching problem. Manag Sci 6:80–91
26. Derigs U, Kurowsky U, Vogel U (2011) Solving a real-world vehicle routing problem with multiple use of tractors and trailers and eu-regulations for drivers arising in air cargo road feeder services. Eur J Oper Res 213:309–319
27. Desaulniers G, Desrosiers J, Ioachim I, Solomon M, Soumis F, Villeneuve D (1998) A unified framework for deterministic time constrained vehicle routing and crew scheduling problems. In: Crainic T, Laporte G (eds) Fleet management and logistics. Kluwer, Boston, pp 57–93
28. Desaulniers, G, Desrosiers, J, Solomon, M (eds) (2005) Column generation. Springer, New York
29. Desaulniers G, Hickman M (2007) Public transit. In: Barnhart C, Laporte G (eds) Transportation, volume 14 of handbooks in operations research and management science. Elsevier, Amsterdam, pp 69–127
30. Doerner K, Schmid V (2010) Survey: matheuristics for rich vehicle routing problems volume 6373 of lecture notes in computer science. Springer, Berlin, pp 206–221
31. Drexl M (2010) Software zur Tourenplanung – Marktstudie 2010. Fraunhofer Verlag, Stuttgart
32. Drexl M (2012) Synchronization in vehicle routing—a survey of VRPs with multiple synchronization constraints. Transp Sci. doi: 10.1287/trsc.1110.0400 (in press)
33. Drexl M, Prescott-Gagnon E (2010) Labelling algorithms for the elementary shortest path problem with resource constraints considering EU drivers' rules. Logist Res 2:79–96
34. Eibl P (1996) Computerised vehicle routing and scheduling in road transport. Avebury, Aldershot
35. Flatberg T, Hasle G, Kloster O, Nilssen E, Riise A (2005) Dynamic and stochastic aspects in vehicle routing—a literature survey. Technical Report STF90A05413, SINTEF
36. Fleischmann B, Gietz M, Gnutzmann S (2004) Time-varying travel times in vehicle routing. Transp Sci 38:160–173
37. Francis P, Smilowitz K, Tzur M (2008) The period vehicle routing problem and its extensions. In: Golden B, Raghavan S, Wasil E (eds) The vehicle routing problem: latest advances and new challenges. Springer, New York, pp 73–102
38. Funke B, Grünert T, Irnich S (2005) Local search for vehicle routing and scheduling problems: review and conceptual integration. J Heuristics 11:267–306
39. Gendreau M, Potvin J (2005) Metaheuristics in combinatorial optimization. Ann Oper Res 140:189–213
40. Gendreau M, Potvin J, Bräysy O, Hasle G, Løkketangen A (2008) Metaheuristics for the vehicle routing problem and its extensions: a categorized bibliography. In: Golden B, Raghavan S, Wasil E (eds) The vehicle routing problem: latest advances and new challenges. Springer, New York, pp 143–169
41. Gendreau, M, Potvin, J-Y (eds) (2010) Handbook of metaheuristics, 2nd edn. Springer, New York
42. Goel A (2007) Fleet telematics: real-time management and planning of commercial vehicle operations, volume 40 of operations research/computer science interfaces series. Springer, Berlin
43. Goel A (2009) Vehicle scheduling and routing with drivers' working hours. Transp Sci 43:17–26
44. Goel A (2010) Truck driver scheduling in the European Union. Transp Sci 44:429–441

45. Goel A, Gruhn V (2008) A general vehicle routing problem. Eur J Oper Res 191:650–660
46. Golden B, Raghavan S, Wasil E (eds) (2008) The vehicle routing problem: latest advances and new challenges, volume 43 of operations research/computer science interfaces series. Springer, Berlin
47. Groër B, Golden B, Wasil E (2009) The consistent vehicle routing problem. Manuf Serv Oper Manag 11:630–643
48. Groër B, Golden B, Wasil E (2010) A library of local search heuristics for the vehicle routing problem. Math Program Comput 2:79–101
49. Haghani A, Jung S (2005) A dynamic vehicle routing problem with time-dependent travel times. Comput Oper Res 32:2959–2986
50. Hartl R, Hasle G, Janssens G (2006) Special issue on rich vehicle routing problems. Central Eur J Oper Res 14:103–104
51. Hasle G, Kloster O (2007) Industrial vehicle routing. In: Hasle G, Lie K, Quak E (eds) Geometric modelling, numerical simulation, and optimization: applied mathematics at SINTEF. Springer, Berlin, pp 397–435
52. Hennig F (2010) Optimization in maritime transportation: crude oil tanker routing and scheduling. PhD thesis, Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology
53. Hickman M, Mirchandani P, Voß S (eds) (2008) Computer-aided systems in public transport, volume 600 of lecture notes in economics and mathematical systems. Springer, Berlin
54. Hoff A, Andersson H, Christiansen M, Hasle G, Løkketangen A (2010) Industrial aspects and literature survey: fleet composition and routing. Comput Oper Res 37:2041–2061
55. Hollis B, Forbes M, Douglas B (2006) Vehicle routing and crew scheduling for metropolitan mail distribution at australia post. Eur J Oper Res 173:133–150
56. Humphreys G (2007) Tachobook. Foster Tachographs, Preston
57. Irnich S (2007) Optimization models and methods for node-routing and arc-routing problems: contributions to enhancements of efficiency and generality. Habilitation Thesis Faculty of Business and Economics, RWTH Aachen University
58. Irnich S (2008) Solution of real-world postman problems. Eur J Oper Res 190:52–67
59. Irnich S (2008) A unified modeling and solution framework for vehicle routing and local search-based metaheuristics. INFORMS J Comput 20:270–287
60. Klabjan D (2005) Large-scale models in the airline industry. In: Desaulniers G, Desrosiers J, Solomon M (eds) Column generation. Springer, New York, pp 163–195
61. Kok A, Meyer C, Kopfer H, Schutten J (2010) A dynamic programming heuristic for the vehicle routing problem with time windows and European Community social legislation. Transp Sci 44:442–454
62. Lu Q, Dessouky M (2006) A new insertion-based construction heuristic for solving the pickup and delivery problem with time windows. Eur J Oper Res 175:672–687
63. Macharis C, Bontekoning Y (2004) Opportunities for or in intermodal freight transport research: a review. Eur J Oper Res 153:400–416
64. Maniezzo V, Stützle T, Voß S (eds) (2010) Matheuristics, volume 10 of annals of information systems. Springer, New York
65. Moin N, Salhi S (2007) Inventory routing problems: a logistical overview. J Oper Res Soc 58:1185–1194
66. Nagy G, Salhi S (2007) Location-routing: issues, models and methods. Eur J Oper Res 177:649–672
67. Oppen J, Løkketangen A, Desrosiers J (2010) Solving a rich vehicle routing and inventory problem using column generation. Comput Oper Res 37:1308–1317
68. Parragh S, Doerner K, Hartl R (2008) A survey on pickup and delivery models part I: transportation between customers and depot. J für Betriebswirtschaft 58:21–51
69. Parragh S, Doerner K, Hartl R (2008) A survey on pickup and delivery models part II: transportation between pickup and delivery locations. J für Betriebswirtschaft 58:81–117
70. Pisinger D, Røpke S (2007) A general heuristic for vehicle routing problems. Comput Oper Res 34:2403–2435
71. Powell W (2007) Approximate dynamic programming. Wiley, Hoboken
72. Powell W, Bouzaïene-Ayari B, Simão H (2007) Dynamic models for freight transportation. In: Barnhart C, Laporte G (eds) Transportation, volume 14 of handbooks in operations research and management science. Elsevier, Amsterdam, pp 285–365
73. Prescott-Gagnon E, Desaulniers G, Drexl M, Rousseau L-M (2010) European driver rules in vehicle routing with time windows. Transp Sci 44:455–473
74. Rieck J, Zimmermann J (2010) A new mixed integer linear model for a rich vehicle routing problem with docking constraints. Ann Oper Res 181:337–358
75. Røpke S (2005) Heuristic and exact algorithms for vehicle routing problems. PhD thesis, Department of Computer Science, University of Copenhagen
76. Savelsbergh M, Sol M (1995) The general pickup and delivery problem. Transp Sci 29:17–29
77. Savelsbergh M, Sol M (1998) Drive: dynamic routing of independent vehicles. Oper Res 46:474–490
78. Schmid V, Doerner K, Hartl R, Salazar-González J (2010) Hybridization of very large neighborhood search for ready-mixed concrete delivery problems. Comput Oper Res 37:559–574
79. Sörensen K, Sevaux M, Schittekat P (2008) "Multiple neighbourhood" search in commercial VRP packages: evolving towards self-adaptive methods. In: Cotta C, Sevaux M, Sörensen K (eds) Adaptive and multilevel metaheuristics, volume 136 of studies in computational intelligence. Springer, Berlin, pp 239–253
80. Spoorendonk S (2008) Cut and column generation. PhD thesis, Faculty of Science, University of Copenhagen
81. Taniguchi E, Shimamoto H (2004) Intelligent transportation systems based dynamic vehicle routing and scheduling with variable travel times. Transp Res Part C 12:235–250
82. Toth P, Vigo D (eds) (2002) The vehicle routing problem. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia
83. Wäscher G, Haußner H, Schumann H (2007) An improved typology of cutting and packing problems. Eur J Oper Res 183:1109–1130
84. Wong R (2008) Vehicle routing for small package delivery and pickup services. In: Golden B, Raghavan S, Wasil E (eds) The vehicle routing problem: latest advances and new challenges. Springer, New York, pp 475–485
85. Xu H, Chen Z, Rajagopal S, Arunapuram S (2003) Solving a practical pickup and delivery problem. Transp Sci 37:347–364
86. Zäpfel G, Bögl M (2008) Multi-period vehicle routing and crew scheduling with outsourcing options. Int J Prod Econ 113:980–996