

The Livestock Collection Problem

Johan Oppen, Arne Løkketangen
Molde University College
6402 Molde, Norway
Email: Johan.Oppen@hiMolde.no
Arne.Lokketangen@hiMolde.no

Abstract

We present a problem dealing with transportation of live animals to slaughterhouses. The problem is taken from the Norwegian meat industry, and may be viewed as a vehicle routing problem extended with constraints to ensure a smooth production flow at the abattoir. In addition, several constraints to ensure animal welfare have to be met. These include duration limits for how long animals may stay on the vehicle as well as rules for mixing different animal categories. A tabu search based solution method for the problem is presented together with some preliminary test results.

1 Introduction

The project “Transportation of living animals - reduced transportation costs, good animal welfare and first-class meat quality” is a co-operation between Norwegian Meat Research Centre, Gilde Norsk Kjøtt, Fatland and Molde University College. Norwegian Meat Research Centre is a centre for competence and emergency control in the Norwegian meat industry. Gilde Norsk Kjøtt and Fatland are Norwegian meat companies. The project was launched in 2003 and stipulated to last for five years. The main goal of the project is to develop a decision support system to reduce transportation costs and secure good animal welfare and meat quality, as these are three main factors for the profitability of both farmers and the meat industry. The Livestock Collection Problem (LCP) may informally be described as the problem of constructing the best possible set of routes for transportation of living animals from farmers to one or more slaughterhouses. Each route must be feasible according to constraints on duration, vehicle capacity, mix of animal types, and visiting order of the customers. The whole solution as a set of routes must adapt to both the production plan and the lairage capacity at the slaughterhouse. A solver for the LCP will be an important part of the decision support system referred to above, and the main focus in this paper will be to describe the problem and suggest possible solution methods. The rest of the paper is divided as follows. In section 2 follows a closer description of the Livestock Collection Problem. An overview of solution methods is given in section 3, followed by a description of our solution approach in section 4. In section 5 we describe test instances and our computational results, and conclusions and suggestions for future work are given in section 6.

2 Problem description

The LCP may be viewed as a vehicle routing problem extended with constraints regarding production and inventory. In this section we first present the basic Vehicle Routing Problem (VRP). We then give a more detailed presentation of the Livestock Collection Problem, and discuss how the basic VRP model should be extended in this setting. In the end we add constraints to model the connection between the routing part and the inventory and production part of the problem.

The Vehicle Routing Problem

The Vehicle Routing Problem (VRP) deals with the allocation of transportation tasks to a fleet of vehicles, and the simultaneous routing for each vehicle. The VRP was first described by Dantzig and Ramser [2], and is a computationally hard optimization problem with high industrial relevance. The classical VRP is defined on a graph $G = (N, A)$ where $N = \{0, \dots, n\}$ is a vertex set and $A = \{(i, j) : i, j \in N\}$ is an arc set. Vertex 0 is the depot; the other vertices are the customers. The travel cost between customer i and j is defined by $c_{ij} > 0$ and d_i is the demand for customer i . The vehicles are usually identical, each with a capacity q . The goal is then to design a least cost set of routes, all starting and ending at the depot, where each customer is visited exactly once. The total demand of all customers on a route must be within the capacity q . This classical formulation is often referred to as the capacitated VRP or CVRP. If $c_{ij} \neq c_{ji}$ for at least one pair of customers, e.g. due to a one-way road, we have an *asymmetric VRP*.

The Livestock Collection Problem

The Livestock Collection Problem deals with transportation of livestock from farmers to slaughterhouses. This problem is an example of a *rich problem*, meaning that we have to add extensions to the basic model in order to describe the problem we are trying to solve. See Toth and Vigo [3] for an overview of extensions to the VRP.

The objective function

In a standard VRP, the objective is to minimize costs, subject to constraints. It is not evident what should be the objective in this problem. Cost minimization should obviously be part of it, but maximization of animal welfare may also be included, even if most animal welfare aspects are treated by constraints. The cost structure is also somewhat unclear for the time being, and as a rather simple starting point we choose to minimize the total distance travelled by all vehicles.

Registration and planning

Farmers register animals for slaughter via telephone or web, usually during the week before the animals are collected. The slaughterhouse constructs a slaughter plan for the following week based on the available animals and the demand for meat products. A transportation plan is then set up to provide the slaughterhouse with animals according to this plan. In practice, the slaughter plan may also be adjusted to fit to the transport plan. In this phase of our work, we find it reasonable to view the slaughter plan as input to the model, together with information about what animals to collect during the week. This means that decisions concerning which animals to collect during a given week are presumed to be taken before any route construction starts. At a later point in time we may

consider also to include decisions about the slaughter plan and what animals to collect in our model.

Animal types and categories

When living animals are transported, they have to be divided into *types* and *categories* for animal welfare reasons. The types are by and large the different animal species; these are divided further into categories by age/size, gender, whether sheep are sheared or not, and whether the animals have horns or not. Mixing of different animal types, and to a certain degree categories, in the same compartment of a vehicle is not allowed. It is not allowed to mix animals with and without horns, even if they belong to the same category. Groups of different gender cannot be mixed, neither can animals of substantially different size, e.g. calves and bulls. Pigs from different herds must be in different compartments to avoid fighting. The categories require different amounts of space in the vehicle during transport. The slaughterhouse needs to know the number of each animal type and category to be collected from each farmer in order to plan both the collection of animals and the production process at the abattoir. For the time being, we consider the three animal types bovine, pigs and ovine. When these types are further divided into categories, we get 16 categories all together.

The farmers are supposed to give information about animals with and without wool and horns when animals are registered for slaughter. Today, this information is only partly available for the slaughterhouse before the animals are collected.

The production process

Each slaughterhouse processes one or more of the animal types. If more than one type is processed, there is a separate production line for each animal type. For each production line, its capacity measured in animals per hour is known. One or more of the production lines may operate simultaneously. The slaughter plan tells what production lines to operate at what times and at what speed during the time horizon, usually one week. It is very important for the abattoir to have a steady supply of animals for slaughter to avoid production breaks, which is very costly. The supply of animals must be according to the slaughter plan to avoid situations where e.g. a large number of pigs are available for slaughter, but the production line for pigs is idle.

This coupling between what is transported to the slaughterhouse and what is slaughtered by each production line has consequences for the route planning. Routes have to be planned so as both to avoid stock out situations and to avoid situations where animals are lairaged too long before slaughter, or where the lairage capacity is exceeded. These constraints may be viewed as a time window type of constraints for the depot (abattoir), but they must be considered as a whole, meaning that it makes no sense to point at a specific route and state that “this route must be back at the depot between time e and time l in order to be feasible”. We will not have a situation where a single route is feasible or infeasible because of these constraints, it will rather be a question if the whole solution, viewed as a set of routes, is feasible or not.

As there is one production line for each animal type, we have to keep track of the inventory level for each type. This means that we have to keep track of three animal types in the lairage, as opposed to during transport, where we must divide the animals into 16 categories. For each animal type, the inventory level at the end of a day must be less than or equal to the amount that is going to be slaughtered the following day. The end

inventory of one day must of course be equal to the start inventory on the following day, and the end inventory on Friday must be zero for all animal types.

The vehicle fleet

The vehicle fleet is heterogeneous. The vehicles have a fixed number of sections (usually three) with a fixed area in each section. Most sections may be split in two tiers; some may be split in three. Bovine will in most cases need the full height of a compartment, while pigs and ovine may be stacked in two or three tiers. In order to utilize the vehicle's capacity to the full, it has to be "packed" in a clever way. This is in itself an optimization problem.

Some of the vehicles have trailers. It is not allowed to leave the trailer loaded with animals behind and collect it later, so the trailers may be viewed merely as an extension of the vehicle capacity. It should be noted, however, that a vehicle with and without a trailer may need different amounts of time to load and unload the same amount of animals due to different manoeuvring characteristics.

The vehicles may be used for up to three tours per day, and the first tour each day starts from the driver's home. The rest of the tours start from the abattoir, and all tours end at the abattoir. The drivers bring the vehicle with them when the working day ends, but this dead-heading is not part of the problem as no collection of animals is done.

Unloading and cleaning of vehicles

When the vehicles arrive at the abattoir, they are unloaded as soon as possible. The animals are brought into the lairage where they stay until they are slaughtered.

The time needed to unload a vehicle is uncertain; it depends both on the number and category of animals as well as how willing the animals are to leave the vehicle. For small animals it is usually easy to get them into the lairage, for large animals this may sometimes be quite hard.

The vehicles also need to be cleaned and disinfected between tours. The time needed to clean and disinfect a vehicle and prepare it for the next tour depends on the load it was carrying on the last tour, but may be considered fairly constant.

Both the capacity at the unloading ramps and the cleaning capacity are limited resources, and may cause queuing problems if many vehicles arrive at the abattoir for unloading and cleaning in a short time period. It is thus desirable to avoid such problems by having a route plan that spreads the arrival times evenly throughout the day.

Visiting and loading order

Due to differences in health status for the different herds, there may be restrictions on the visiting order. If livestock from an infected herd is to be collected, this farm has to be the last one on the route. Farms with breeding herds must be visited by an empty and clean vehicle, and thus need to be first on the route.

Time limits

There is a general rule stating that no animal is allowed to stay on the vehicle for more than eight hours. In addition, it is beneficial to keep the time on the vehicle as low as possible for certain categories of animals. This is because the time on the vehicle may affect the animal welfare which in its turn may affect the meat quality. The importance of how long the animal stays on the vehicle varies over the categories of animals. If the

animal welfare is affected without having consequences for the meat quality, it may still be a goal to minimize the time on the vehicle.

When the animals arrive at the abattoir, they are unloaded and lairaged until the time of slaughter. Lairaging serves as a temporary store between farm and slaughter, and the abattoir has a certain lairage capacity for each animal type. The animals may be lairaged overnight. For some animal categories it is beneficial to keep the lairage time as short as possible, or at least avoid staying overnight, in order to keep the animal welfare, and thereby the meat quality, at a high level. In order to have enough animals to start the slaughter process on Monday morning, some livestock collection is done on Sunday afternoon. No animals can be lairaged over the weekend.

For the time being we choose not to include time windows for the farmers, as in most cases the slaughterhouse decides when the different farmers should be visited, and it seems like this is accepted by the farmers. What we do have is a time window for the slaughterhouse stating that no vehicle should arrive later than some point in time in the evening.

Problem size

The size of a Livestock Collection Problem, measured in number of customers and number of tours, may vary a lot for Norwegian slaughterhouses. The size also depends on the length of the planning horizon.

For Gilde Hed-Opp at Rudshøgda and Fatland Jæren at Hommersåk, which are the two abattoirs used as examples in this work, the number of customers visited per day will typically be between 40 and 200. The number of tours per day will then be between 8 and 40. This means that problem instances including data for one week may have up to 1000 customers, and thus need up to about 200 tours.

3 Solution methods

The VRP has become one of the most widely studied problems in combinatorial optimization, and much effort has been put into developing both exact and heuristic methods to solve the VRP. The classical VRP is known to be strongly NP-hard.

Solution methods for the VRP

Exact methods

An exact solution method is able to find the optimal solution of a problem and to prove the optimality of the obtained solution. Exact solution methods for the VRP include Branch & Bound and Branch & Cut (Toth and Vigo [3]). These techniques are known to consistently solve instances with up to 50 customers in reasonable time.

Heuristic methods

Heuristic solution methods cannot guarantee that an optimal, or even feasible, solution is found; neither can optimality be proved when an optimal solution is actually found. The reason for using heuristics is the assumption that a suitable heuristic can be able to find reasonably good solutions relatively quickly. This can be very useful in cases where exact methods perform poorly due to the size or structure of the problem.

Bräysy, Gendreau, Hasle and Løkketangen ([7, 8]) give an overview of heuristics for the VRP. This includes heuristics for rich models where many of the possible extensions are included, which is important in our setting.

Solution methods for the Livestock Collection Problem

Not much seem to have been done to solve VRPs extended with constraints regarding production and inventory. Gullberg and Hovden [1] solved a simplified version of the Livestock Collection Problem with seven customers using CPLEX, and found that adding the eighth customer made the problem too hard for CPLEX to find a feasible solution in reasonable time. They formulated the problem as a Mixed Integer Problem and simplified it by dealing only with one animal category, a homogeneous vehicle fleet and no precedence constraints.

Sigurd, Pisinger and Sig [6] describe and propose a solution method for a pickup-and-delivery problem with time windows and precedence constraints. The application is transportation of live pigs between farms in Denmark according to a number of veterinary restrictions. To avoid the spread of diseases, the vehicles must visit the farms in a non-decreasing order of health levels. A low health level corresponds to a healthy livestock.

4 Our approach

We propose a tabu search based heuristic for the Livestock Collection Problem. There are two main reasons for choosing a heuristic in the first place, namely the richness of the model and the problem size. Both aspects strongly suggests that exact methods are going to fail in finding solutions to our problem, at least if we use standard software packages like CPLEX. The reasons for choosing Tabu Search as our first method to apply on this problem, is both that we have some experience from using Tabu Search for another variant of the VRP, and that a substantial amount of research shows that Tabu Search in general is well suited for solving rich problems. This means that we hope to find it relatively easy to apply Tabu Search to the Livestock Collection Problem, and that we have fairly good reasons to believe that we will get acceptable results.

In this section, we first give a general description of tabu search. We then describe in more detail how our algorithm for the LCP works.

Tabu Search

Heuristics for the VRP are usually divided into *classical* and *modern* heuristics. Classical heuristics have been used since the 1960s, modern heuristics from 1990 on. Toth and Vigo [3] provide an overview of both classical and modern heuristics for the VRP.

Modern heuristics can be divided into two groups, *local search* based methods and population based methods. We need the following in a local search.

- An *initial solution* as a starting point for the search.
- A defined search *neighborhood*, a set of neighboring solutions.
- A *move*, how to change a solution to get to a neighboring solution.
- A *move evaluation function* to tell how good the different moves are, or how good they are believed to be.
- A *stopping criterion* to tell when to stop the search.

An *initial solution* can be obtained in a number of ways, the most common being a randomly generated solution or a solution generated by a construction heuristic. If a random solution is chosen as a starting solution, one must be aware of the possibility of

getting an infeasible solution. If a feasible starting solution is needed, steps must be taken to make the solution feasible. For a VRP, one of the classical heuristics could be used to make a starting solution to be improved by another heuristic.

A search *neighborhood* is defined by *moves*. A move is a way to change a solution to get a different solution, and the neighborhood is the set of all possible solutions that can be reached from the current solution by performing such a move. The *size* of the neighborhood is the number of different solutions that can be reached by performing a move. Clearly, different moves define different neighborhoods, often of significantly different size. A *move evaluation function* is used to give a measure of how good a move is, and need not necessarily be based on the objective function value only. A *search strategy* has to be applied to guide the move selection, and a *stopping criterion* is needed to determine when to stop the search. Typically, the search stops when a predefined number of iterations (moves) have been performed without any improvement of the solution, when a time limit is reached, or when a quality measure is fulfilled. During the search, *diversification techniques* are often applied to lead the search into new and unexplored regions of the search space.

The simplest version of local search is *steepest descent*, where all the neighbors are evaluated in each iteration. The best neighbor leading to a better solution is chosen, and the search stops if no improving neighbor can be found. This makes the search stop in a local optimum, and this feature is a major problem with local search.

A *metaheuristic* is a strategy that is used to guide other heuristics and keep the search from getting stuck in local optima. Several such strategies have been developed since the early 80's, and many of them use local search as the inner heuristic.

Tabu Search is a local search based metaheuristic, and was introduced by Fred Glover in 1986. The main ideas are to avoid recently visited parts of the solution space and to guide the search towards new and promising areas. Non-improving moves are allowed in order to escape from local optima, and attributes of recently performed moves are declared *tabu* or forbidden for a number of iterations to avoid cycling. Usually all neighbors of the current solution are evaluated, and the best non-tabu solution is chosen. It is common to disregard the tabu status of a move if the move leads to a better solution than the best known solution so far, this is an example of an *aspiration criterion*. Tabu Search does not rely very much on randomness, but tries to be "intelligent" in its way to perform the search. Tabu Search has been successfully adapted to numerous discrete optimization problems, VRP included. See Glover and Laguna [4] for more information about Tabu Search.

A tabu search based heuristic for the LCP

Our algorithm for the Livestock Collection Problem is based on a tabu search algorithm for the CVRP described by Oppen and Løkketangen [5]. Their method has been extended to deal with the constraints described in section 2.

Time and inventory

The time horizon for a problem instance is one or more days, usually six (one week, starting on Sunday). This is due to the fact that some collection is done on Sunday, in order to have animals to start the production on Monday morning. Because no animals can be lairaged during week ends, less collection is done on Friday. For each tour, we keep track of the time the vehicle leaves the driver's home or the slaughterhouse, when the vehicle

arrives and leaves each of the customers, when the vehicle arrives at the slaughterhouse and when the vehicle is ready for the next tour after unloading and cleaning.

At the slaughterhouse, each day is split into four time periods. The first is from 6.00 to 9.30, the second from 9.30 to 12.00, the third from 12.00 to 14.30 and the fourth is from 14.30 to the working day ends. The slaughtering usually stops at 15.00 or 16.00 in the afternoon, but vehicles with animals may arrive later. The slaughter plan tells how many animals of each type to slaughter in each time period. When a vehicle is unloaded, the animals are added to the inventory in the corresponding time period. The inventory level is calculated at the end of each time period. The inventory level from the previous period plus the animals received during the period must be sufficient to meet the slaughter plan for the period plus a safety stock. The inventory level must not exceed the lairage capacity in any time period.

Moves

A tour has a legal starting and ending location if it satisfies the following. The first tour each day starts at the driver's home, the following tours start at the slaughterhouse. All tours end at the slaughterhouse. Let S be the set of all solutions that satisfy the following constraints: every tour has a legal starting and ending location, every customer belongs to exactly one tour, and two tours that use the same vehicle are separated in time. In addition, all farms with breeding herds are visited before the first farm with a non-breeding herd, and all farms with infected herds are visited after the last farm with a non-infected herd. This means the tours need not be feasible according to the precedence constraints, but all farms with breeding herds and infected herds are visited in the start and the end of the tours, respectively.

In each move, a customer is moved from its current tour to a different tour. More formally, let $A(s) = \{(i, j, k, l) : \text{customer } i \text{ is visited by vehicle } j \text{ on day } k \text{ and tour number } l\}$ be an attribute set associated with each solution $s \in S$. The neighbourhood $N(s)$ of a solution s is defined by applying an operator that removes an attribute (i, j, k, l) from $A(s)$ and replaces it with a different attribute (i, j', k', l') , where $j \neq j' \vee k \neq k' \vee l \neq l'$. This gives a neighbourhood size of $|N| = n(m-1)$, where n is equal to the number of customers and m is equal to the total number of tours for all vehicles. When a customer is removed from a tour, the tour is reconnected by linking the predecessor and successor of the removed customer. The insertion of a customer into a tour is done so as to minimize the increase in the length of the tour, but without changing the order of the customers already in the tour. If the customer to be inserted represents a breeding herd or an infected herd, it is inserted as the first or last customer, respectively.

Move evaluation and diversification

For each possible move from s to $s' \in N(s)$, let $\Delta c(s')$ be the change in total length of all tours in the solution, and let $\Delta q(s')$ be the change in violation of capacity constraints, measured as the total number of animals exceeding the vehicles' capacities. Let $\Delta r(s')$ be the change in total number of visits that violate precedence constraints, let $\Delta s(s')$ be the change in total violation of duration constraints, and let $\Delta t(s')$ be the change in total violation of the time window at the abattoir. All these five terms are normalized to have terms of approximately equal magnitude. Moves are evaluated using a function $f(s') = \Delta c(s) + \alpha \Delta q(s) + \beta \Delta r(s) + \gamma \Delta s(s) + \chi \Delta t(s)$, where α, β, γ and χ are positive parameters that are dynamically adjusted during the search.

In addition to the four criteria for infeasibility described above, a move may result in a solution that violates inventory constraints by bringing too many or too few animals of one or more types to the slaughterhouse in one or more time periods. This is taken care of by a separate procedure that shift tours backwards or forwards in time, and this procedure is run after every iteration.

To diversify the search, any solution $s' \in N(s)$ such that $f(s') > 0$, is given a penalty $p(s') = \lambda \sum_{(i,j,k,l) \in A(s')} \rho_{ijkl}$ that is added to $f(s')$. Here, ρ_{ijkl} is the number of times the attribute (i, j, k, l) has been part of a *good solution*, that is, a solution that is feasible and has a total length less than η times the length of the best solution found so far. The parameter λ is used to control the intensity of the diversification. These penalties are used to lead the search into less explored parts of the solution space whenever a local optimum is found. If $f(s') \leq 0$, the penalty term is not added to $f(s')$.

Initial solution

The initial solution is generated by a greedy approach, where all customers first are assigned to the closest vehicle, as long as this vehicle does not get too many customers. The closest vehicle is the vehicle whose driver lives closest to the customer. Each vehicle then has a list of unassigned customers, and tours are constructed by taking the last customer in the list and inserting this customer into the current tour as long as the vehicle has enough capacity. The customer is inserted into the tour in the best possible way, that is, the increase in distance is minimized. The order of the customers that have already been inserted is not changed. If the customer represents a breeding herd or an infected herd, it is inserted as the first or last customer, respectively. If the vehicle is full, the customer is inserted into a new tour.

If the problem instance has a time horizon of one week, the vehicles with most customers are first given one tour on Sunday. Then all vehicles construct tours from Monday to Friday, starting with one tour each day, then a second, and so on, until all customers are inserted into a tour.

The resulting initial solution may be infeasible for several reasons. Some tours may last too long, meaning that animals stay on the vehicle for more than eight hours, and some tours may end too late in the evening. Some tours may visit more than one breeding herd or more than one infected herd, which is not allowed in a feasible solution. In addition, the whole solution, viewed as a set of tours, may violate inventory and production constraints either by not providing enough animals according to the slaughter plan for one or more time periods, or by overfilling the lairage during one or more time periods.

Tabu search

The tabu search starts from the initial solution and moves, at each iteration, to the first improving non-tabu neighbor such that $f(s') < f(s)$. One iteration starts from where the previous ended in order to evaluate a larger part of the solution space. The attribute (i, j, k, l) that was removed from $A(s)$ is now declared tabu for *tabu tenure* iterations, where tabu tenure is the tabu length or duration. During these iterations, it is not allowed to move customer i back to vehicle j 's tour l on day k . By the use of a simple aspiration criterion, a tabu move can still be chosen if this leads to a solution that is the best found so far in the search. After each move, the values of the parameters α, β, γ and χ are adjusted. If the current solution is feasible with respect to vehicle capacities, the value of α is decreased to make it less costly to violate vehicle capacity constraints. If the current solution does violate such constraints, α is increased to make it more costly. In

the same way, β is decreased if the solution violates inventory and production constraints and increased otherwise, γ is adjusted according to whether precedence constraints are violated or not, and χ is decreased or increased according to the violation of the duration constraints.

If the current solution is feasible, has a total length less than η times the length of the best feasible solution found so far during the search, and the number of iterations performed has reached κ , the solution is considered *good*. Whenever a good solution is found, the ρ values for the attributes of the solution are incremented. In addition, a 2-opt procedure is applied to the tours of the solution. It would be a waste of time to apply the 2-opt procedure for all feasible solutions, but it is also important to apply it often enough to capture the solutions that become a new best solution after a 2-opt procedure is applied. This means that the parameter η is used to identify good solutions for two different purposes, and it might well be that different values should be used for these purposes. The reason to wait until at least κ iterations are performed is the belief that during the first phase of the search, a lot of improvement is found. It is thus likely that any feasible solution found this early is a new best, but that an even better one will be found soon.

The search continues until a preset time limit is reached, or until a preset number of moves are performed.

5 Computational testing

Test instances

To test the performance of our algorithm, we will use data from the abattoirs Gilde Hed-Opp at Rudshøgda and Fatland Jæren at Hommersåk. For both abattoirs, we have information for week 19 and 38 in 2004 about what farmers were visited during the week and how many of each animal category was collected from each farmer. The number of visits during a week varies from about 300 to about 750; the number of visits per tour usually is from one to six. This means the sizes of the problems are quite large.

In addition, we have information about available vehicles (home location and capacity), slaughter plans and lairage capacities. We also have estimates on parameters regarding vehicle speed and loading, unloading and cleaning times. For this first phase of testing we use rather rough estimates for some parameters, and plan to refine these later.

Results

Our test results so far clearly show that the problem is solvable in reasonable time, also for instances of realistic size. This is by itself an important result. The current version of our solver for the LCP performs about 40 iterations per minute on a problem instance with 600 customers. However, to be able to discuss the quality of our solutions, we need more than just the ability to create what appears to be feasible solutions.

We need more accurate data to be able to generate plans/solutions that are implementable in practice. This means e.g. better information about the vehicle fleet (speed, capacities, etc.) and time needed for loading, unloading, cleaning and so on. We should probably use travel times instead of distances in the objective, and we need to compare our plans to plans created by the current manual system. We also need the ability to visualize our plans in order get them evaluated by people from the industry. In addition, there is still a lot of work to do on the solver to improve its performance.

6 Conclusions and future work

We have given a description of the Livestock Collection Problem, which can be characterized as a rich VRP extended with inventory constraints. The problem is taken from the Norwegian meat industry, which demands better solutions to reduce transportation costs. We propose a tabu search based solution approach for the LCP, and preliminary test results show that the problem is solvable in reasonable time.

There is still much work to do, both to improve our solution methods and to improve the quality of data. The latter is necessary to get solutions that are useful in practice, which is important in this setting.

References

- [1] Gullberg BO, Hovden KJ. Transportation of Livestock to Slaughterhouses - a Mathematical Model (In Norwegian). Molde University College, 2004.
- [2] Dantzig GB, Ramser JH. The Truck Dispatching Problem. *Management Science* 1959; 6:80.
- [3] Toth P, Vigo D (Eds.). *The Vehicle Routing Problem*. Philadelphia: Society for Industrial and Applied Mathematics, 2002.
- [4] Glover F, Laguna M. *Tabu Search*. Kluwer Academic Publishers, 1997.
- [5] Oppen J, Løkketangen A. Arc routing in a node routing environment. Forthcoming in *Computers & Operations Research*, 2005.
- [6] Sigurd M, Pisinger D, Sig M. Sceduling Transportation of Live Animals to Avoid the Spread of Diseases. *Transportation Science* 2004; 38: 197-209.
- [7] Bräysy O, Gendreau M, Hasle G, Løkketangen A. A Survey of Heuristics for the Vehicle Routing Problem, Part I: Basic Problems and Supply Side Extensions. Technical Report. Molde University College, 2003.
- [8] Bräysy O, Gendreau M, Hasle G, Løkketangen A. A Survey of Heuristics for the Vehicle Routing Problem, Part II: Demand Side Extensions. Technical Report. Molde University College, 2003.