ANWB automates and improves repair men dispatching

Citation for published version (APA):

Document status and date:
Published: 01/01/2008

DOI:

Document Version:
Publisher's PDF, also known as Version of record

Please check the document version of this publication:

Peter van Huigenbosch, Joris van de Klundert,
Laurens Wormer

ANWB automates and improves repair men
dispatching

RM/08/055

JEL code: C61

ANWB automates and improves repair men dispatching

Peter van Huigenbosch
ANWB, Wolfheze, The Netherlands

Joris van de Klundert
Maastricht University, Faculty of Economics & Business, Maastricht, The Netherlands

Laurens Wormer
Mateum, Maastricht, The Netherlands

Abstract

ANWB, the Dutch automobile association, provides assistance, car repair and replacement services to its nearly 4 million members. ANWB services around 1.3 million requests per year in The Netherlands. Historically, the operational planning process of assigning requests to service men was regionally organized, and human planners solved the sometimes large and hectic planning situations in real time. At a national level, some 50 planners were required to do the job, and the quality of the planning and operations were largely unknown. In a large business process reengineering project, ANWB redesigned the planning processes, leveraging state of the art IT and operations research techniques. As a result, the 24/7 planning processes are smoothened, can be executed by as few as 14 planners who work at a national level, and the operational planning and performance have improved. As new competitors entered the market, ANWB has been able to sustain its extraordinary high customer ratings and market share, while adapting its proposition to the competitive prices dictated by the market.

Keywords: Repair men, dispatching, service responsiveness

Road Side Service Industry

Since the early days of automobile manufacturing, the automotive industry has been focusing on producing and selling cars. Maintenance and repair of cars has been the business of garages and road side service providers. In many European countries, road side services are provided through national associations which car owners join as members. In return for the membership contributions, the not for profit associations service their members whenever required. Despite the lack of a competitive environment in the road side repair service industry, the Royal Dutch Touring Club, ANWB, who provides road side services in the Netherlands, has a long track record of high customer satisfaction and loyalty, as do other associations in other countries. ANWB has almost 4 million members, and services nearly 1.3 million service request per year. On peak days, request arrival rates easily exceed a level of one per ten seconds.

Over the last decade, this quiet industry setting is being disturbed by a number of business trends. First of all, several automobile manufacturers aim to establish a long term relationship with their customer which includes maintenance and repair services. Second, car lease companies have emerged and have large fleets on the road. For various reasons lease companies usually choose not to deliver road side services via the traditional private membership associations. Thirdly, new purely commercial parties enter the market, challenging the associations which are historically not very competition oriented. Hence, traditional associations such as ANWB have more difficulty attracting and retaining good customers. Finally, the improved mechanical quality of cars and higher proportion of software related defects threaten to reduce the number of service requests that can be effectively handled by a road side repairman. All in all, ANWB faces a more competitive market with new entry by competitors and more substitutes becoming available.

Major competitive weapons in the road side service market are price, quality and response time. Hence, in order to defend market share, ANWB needs to effectively combine low cost with responsiveness. Operating costs are largely determined by personnel cost, service men as well as staff, and by equipment, such as various types of repair cars and devices. Hence, operating costs are by and large determined at strategic and tactical planning levels. Of course, the capacity decisions made at the strategic and tactical planning levels influence response time as well. Short term however, response time is highly dependent on operational decisions, in particular the dispatching: the assignment of service men to service requests.

In view of the industry developments and resulting need for improved effectiveness, ANWB started an ambitious program in 2001, consisting of several projects involving operational management and planning processes, including the supporting IT. One of the projects in the program was the project 'Dispatch' which aimed to improve the service men responsiveness by redesigning the operational planning processes, using innovative information, communication and operations research technology. This paper describes the operational planning problem, the original and the redesigned planning processes and

operational processes. Further we describe into some detail how the project Dispatch started ANWB to better steer its operations in the direction of the desired balance between customer satisfaction and cost effectiveness, by choosing its performance indicators and planning objectives accordingly.

Dispatching Road Side Service Men

The operational planning problem of ANWB is dynamic, and evolves in real time: around 1.300.000 times per year a new customer request arrives, expecting timely service, and therefore leading the ANWB to reconsider its current operational planning. At any given time instant, the real time problem instance the dispatchers face can be described by the following data:

1. A set of unattended service requests, described by the time of notification, location, and perhaps an indication of the requested service,
2. A set of repair men (with vehicles), described by current location, and a status indicating whether they are busy, pausing, driving to a request, or available.

Solving a real time problem instance consists in assigning service men to unattended service requests. Moreover, when several service requests are assigned to a same service man, the requests have to be sequenced, or at least a first request has to be selected.

Regardless of the solution approach, the problem evolves dynamically. The current problem instance changes since new service requests or repair men arrive, or the status of one of the service men changes. For instance it changes from busy to available when a service request is successfully completed. Dispatching is therefore to solve a never ending series of real time problem instances, so as to best service customer requests. In these somewhat abstract terms, an important objective is therefore to maximize the customer satisfaction, as it is measured on a routine basis. This operational planning problem falls into the category of dynamic vehicle dispatching problems. Many of the approaches proposed in the literature (see below) use models for real time instances which require to solve NP-hard problems, and hence solving a real time problem instance may be time consuming. An important aspect of this operational planning problem is however that customer satisfaction depends on the solutions provided to a series of successive real time instances, rather than on the solution to either one of them. This dynamic setting makes the problem even harder, since planning decisions have to be taken without knowledge of future data. On the other hand, the occurrence of unforeseen events weakens the concept of optimality at any specific time instant in the dynamic setting.

Literature

Much of the research on routing is viewed from traveling salesman and vehicle routing frameworks. The repairmen dispatching problem considered in this paper bears a certain

reminiscence with the stochastic dynamic vehicle routing problem as it is discussed for instance in [2], [3], and [5] and the references therein. At the same time it is quite different, since vehicle routing problems, or dial a ride problems, assume that loads or people have to be transported, and therefore the capacity of the vehicle and the transportation movement play an important role. In repairmen dispatching, transportation is absent. Instead a service of unknown duration must be provided. Krumke and various others, see e.g. [4],[8], [9], have done extensive theoretical and practical work on this problem for ADAC, the German sister organization of ANWB. They also report [8] that the integer linear programming based approach they propose, and which we will also address below, outperforms various heuristic approaches as they had been previously used, and are more commonly used in dynamic vehicle routing [3],[5].

The work of Krumke et al. is closely related to the work reported in this paper both from a theoretical and a practical viewpoint. However, since their work is customized towards a particular application, as is ours, there are many differences. In particular, the work reported in this document focuses on the successful implementation, rather than on the algorithmic details, however important they are.

In view of the present importance of (after sales) service processes in many industries, it is remarkable how little attention customer service models have received in operations research literature. We will address this issue in some more depth when discussing performance indicators and objective functions.

Description of the Original Process

The original operational process is best described from the customer perspective. A customer is driving or intends to start driving, but cannot continue because of a defect. The customer calls ANWB, and reaches a call centre agent. The customer provides the required information, such as name, car, license plate, and location. Since customers are sometimes in a hurry to be serviced, or in an unsafe situation, asking detailed information about the request and/or defect is often not desirable. Hence, at the moment of dispatch, such information is usually not available.

Before the call is ended and the request is forwarded to the planning, the call centre agent usually gives an indication of the waiting time. This indication is very important for customer satisfaction. Typically customers who are serviced within 30 minutes are very satisfied about the timeliness of the response. Waiting times smaller than 60 minutes are also considered to be acceptable, as this is the maximum waiting time ANWB promises to its members. If waiting time is more than one hour, customer satisfaction typically decreases, even to dissatisfaction. Waiting for more than 90 minutes is considered to be extremely disappointing.

Dissatisfaction ([1]) can sometimes be undone by a friendly and professional service man. Another way to handle dissatisfaction caused by long waiting is by correctly informing the customer on beforehand of long waiting times. Originally, estimates based

on historical data, such as the average waiting time over the passed hour, have been provided. The predictive value of this average waiting time at time of request is limited. First of all, the variance of the waiting time is considerable as can be expected in a dynamic context. Second, when many incoming requests arrive in a short period of time, which is not uncommon (heavy traffic, bad weather conditions), the average waiting at the time of request is much lower than the average waiting time at the time of service. This implies that the predictive value is especially poor when waiting times are long. We conclude that the average present waiting time is not at all a good predictor of waiting time for any specific new customer. It is a process-oriented, internal, historical performance indicator, rather than an external, customer oriented indicator, which enables the management of the present and future service performance.

When the call is completely registered, it will automatically be made available for planning. In the original process, a human planner responsible for the region in which the request emerged handles the operational planning of the request. The planner uses a geographical user interface displaying the status quo, as well as a list of unattended requests. The planner assigns unattended requests to service men. Service men would usually not be assigned more than one future task.

Since ANWB services on average 3500 service request per day, the complexity of the planning is considerable. To cope with this complexity, ANWB divided The Netherlands into regions, each with its own operations and planning. Each region was planned by 3-4 planners, yielding a total of around 50 planners. Every service man is based in one of the regions, and in principle service men don't service requests outside of the region where they are based. Servicing requests outside of the region is considered to be undesirable, e.g. because it may cause the service man to be far away from his end of duty location while approaching end of duty. By consequence, the last half hour of duty might be spent driving back to base, instead of servicing customers.

The project Dispatch

The direct objectives that ANWB had with respect to the project Dispatch were:

1. Improve timeliness of response to service requests
2. Lower operational costs
3. Speed up the dispatch process
4. Lower cost of the dispatch process

The latter two objectives were to be attained through developing and implementing an IT system in which timely availability of service request information was enhanced. Moreover the system was to make 90% of the dispatching decisions automatically. Assuming that the remaining 10 % would require as much human interaction as before, this would yield a reduction of more than 30 FTEs in dispatching., or in other words, a yearly reduction of 2 million euro's in planning costs.

Besides the labor intensity, the way in which the dispatch process was organized originally was perceived to lead to several operational inefficiencies. First of all, service men were restricted to stay in their regions, and it frequently occurred that one region had an excess of service men, while a neighboring one had an excess of requests and hence long waiting times. (Although practical solutions were implemented for such problems, they were not systematically and satisfactorily addressed.) Second, each planner had his or her own working methods, and perhaps their own subjective preferences in assigning service men to requests. This made it hard to implement corporate service priorities in dispatching and in operations. Obviously, this provided an opportunity for an automated dispatching system to increase responsiveness.

The usual real time instances of each region were quite manageable and capacity was such that response times were usually at an acceptable level under these circumstances. However, during rush hours, and in bad weather conditions (e.g. after the first frost) demand would become more hectic, customers were waiting in bad or unsafe conditions, and the complexity of the planning tasks increased. Of course, human dispatchers make mistakes under pressure, and hence an adequate automated system might do better. On the other hand, the human dispatchers were skillful and experienced, knew their regions and could anticipate future events. Hence, it was by no means clear that an automated dispatching system would perform comparable to (leave alone better than) human dispatchers. Thus, the quality of the solutions generated by an automated system was considered as an opportunity and a threat at the same time. It was considered a critical success factor throughout the various stages of the project, and we will explicitly describe how it was managed below.

Project plan and scope

A large part of the Project Dispatch was devoted to restructuring the IT and information flow, including the dispatch processes. We will not address these IT related issues in this paper, apart from the Dispatch control room that ANWB implemented for its nationwide dispatch processes. In this control room which was realized around 2004, two to four dispatchers guide the planning, the bulk of which is taken care of automatically (as explained below.) The dispatchers take care of exceptions and stay in touch with the repairmen, using personal computers with real time information, as well as a giant screen of three by four meters, depicting the present situation in the Netherlands (Figure 1.)

Instead of focusing on the IT, we consider the dispatch problem, its solution methods, and their implementation. In the project, we considered several of the methods proposed in the literature, as well as a method proposed by ANWB itself.

Initially, we considered two groups of evaluation criteria for the dispatching process and the resulting operational performance. A first group was the acceptance of the planning. Since the new software was a direct threat for the number of jobs in the dispatch process and for the number of repairmen needed, resistance in the organization was expected. Moreover, poor performance of the system might lead to low acceptance and

subsequently not using the system, thereby destroying a large investment and its future potential. To solve possible problems in advance, we compiled a list of possible 'sins' that we wanted an automated dispatch algorithm to avoid.

Figure 1: ANWB Dispatch Center

A first example of such a sin is unbalanced division of workload between service men. Another example is the event of two service men crossing one another coming from opposite directions while driving to their assigned requests. Obviously, when service men cross one another on their way to their respective assignments, exchanging the requests among the service men yields a better assignment, yielding the dispatch decisions made to be suboptimal. A good solution approach should provide cross-free solutions. However, since the problem is solved dynamically and ANWB maintained the policy of not reassigning service men which are already on their way to a request, crossing couldn't be ruled out.

We build a simulation environment in which we tested the candidate algorithms before releasing them. The simulations studies revealed that the intended dispatch algorithms had low probability of committing any of the sins on the list.

Operational Objectives

As pointed out above, dispatch decisions impact operational performance: they effect how effectively resources are utilizes to service the customer requests. As is common in many service organizations, ANWB considered and considers average waiting time to be an important performance indicator to measure this operational effectiveness. As the examples in Figures 2-5 point out, it may however be in conflict with customer

satisfaction, rather than a representative of it. Minimizing average waiting time can lead to long waiting for some customers, and thereby to unnecessary loss of customer satisfaction. Hence, before choosing solution methods, we spend time on selecting the right performance measures, and hence the dispatching objectives, which fit the service organization ANWB.

The search for appropriate objectives, and hence for performance indicators, has not been an easy and preliminary step. It is an ongoing debate which is influenced by the needs and expectations of the customer, the operational organization of ANWB, but also by the competitive landscape ANWB operates in, and is benchmarked in. While the project dispatch was in execution, new players entered the market of offering road side services to private car owners at a lower price. This aroused benchmarking studies by national media on cost versus performance. The costs of assistance is not necessarily free for members, since depending on the membership, repair, towing away, alternative transportation, et cetera may come at varying additional costs. Moreover, there are differences in the response times of various providers, and also in the equipment and training of the service men, leading to differences in the capability to actually repair the car.

These changes in the competitive environment required ANWB to reconsider its competitive position. While in the process of executing project Dispatch and other related projects, and thereby being able to provide better service and/or lower operating costs, it had the opportunity to align its operations steering with its strategic priorities. Knowing that simple and straightforward indicators such as average waiting time might fall short of capturing these strategic priorities we thoroughly analyzed a variety of objectives and their interrelationships.

After concluding that customer satisfaction is less dependent on average waiting time but rather on waiting longer than the 60 minutes mentioned in the membership agreement [1], we considered to measure the service level in a variety of ways, such as:

- Percentage of customers waiting longer than 60 minutes
- Total number of minutes per day waited in excess of 60 minutes

- Percentage of customers waiting longer than 60 minutes + b ´ percentage of customers waiting longer than 90 minutes
- Total number of minutes per day waited in excess of 60 minutes + b ´ Total number of minutes per day waited in excess of 60 minutes

In practice, this often meant that we considered histograms of end of day waiting time distributions and compared solution approaches, and parameter settings, comparing the histograms.

On the other hand, since price is such an important competitive weapon, we considered efficiency related operational performance indicators such as:

- Average number of requests serviced per day per road patrol,
- Average travel time to next request
- Average percentage of idle time per road patrol (that is time not spend servicing or travel to a request).
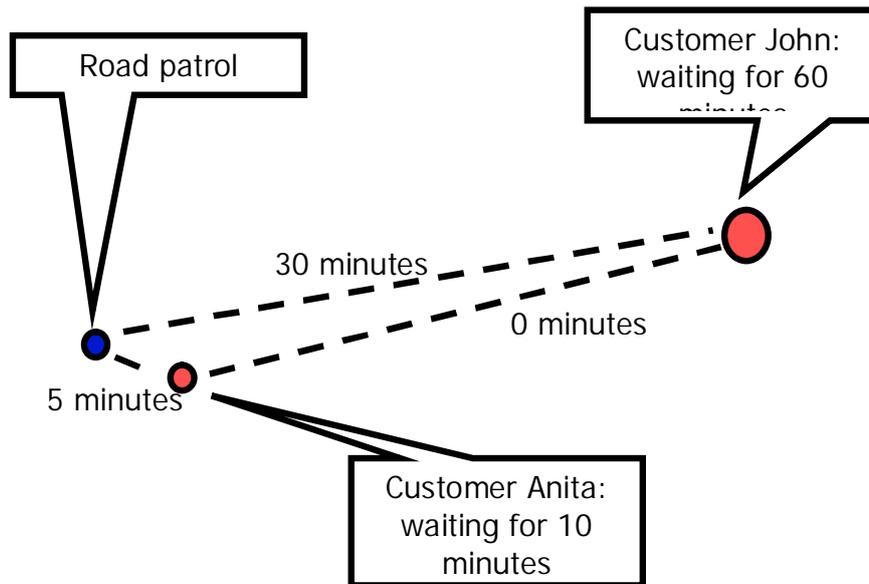

Figure 2: Instance at T=0

Solution

1. Go to Anita, repair (20 minutes), Go to John

Waiting time Anita:	10+5	=	15
Waiting time John:	60+5+20+30	=	**115**
Average waiting time:			65

2. Got to John, repair (20 minutes), go to Anita

Waiting Time John:	60+30	=	90
Waiting time Anita:	10+30+20+30	=	90
Average waiting time:			90

Figure3: Instance at T=0, minimizing waiting time

To the best of our knowledge, service level based criteria are quite rare in the operations research literature on vehicle routing or repair men dispatching. Vehicle routing problems usually involve around minimizing travel distance, perhaps with some time window constraints. Repair men dispatching problems often also consider minimization of the average waiting time, or the travel distance. Minimizing travel distance is attractive from

the viewpoint that a repair man is either traveling or repairing (assuming he is never idle.) Hence, the shorter the traveling, the sooner the services. Thus, still assuming repairmen are never idle, minimizing travel distance minimizes the sum of the flow times of the customers, and hence the average waiting times. From a direct cost perspective it turned out to be less interesting: only a very small percentage of operational costs are distance dependent (e.g. fuel). The cost perspective entered the discussions more indirectly. An improved dispatching might open the possibility to attain the aspired service level with less service men and equipment, thereby reducing costs by reducing capacity.
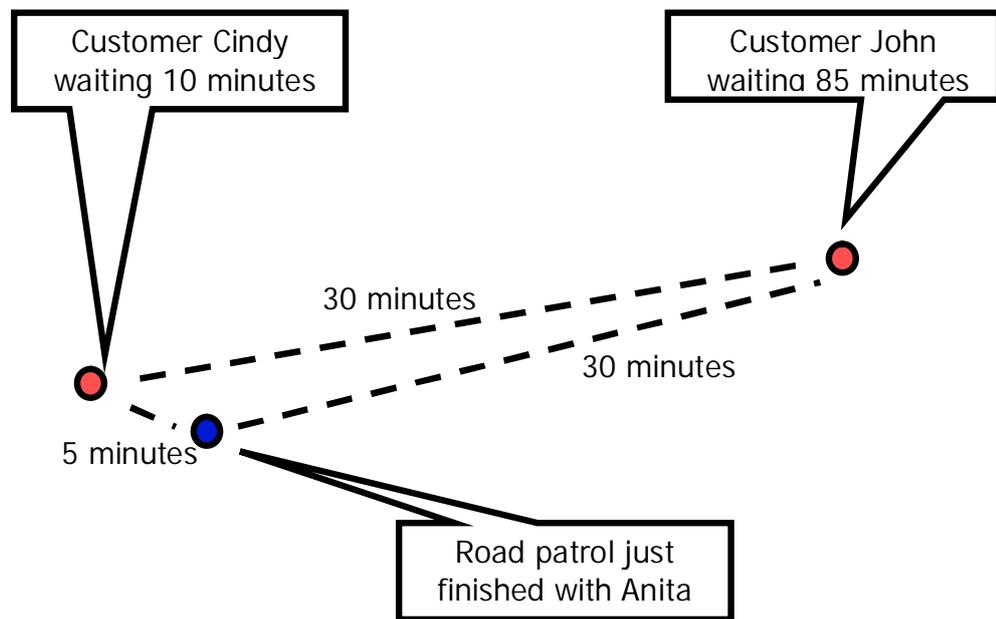

Figure 4: Instance at T=25

Solution continued

1.   Go tot John, repair (20 minutes), go to Cindy

Waiting time John: 85+30   =   115
Waiting time Cindy: 10+30+20+30   =   90
Average waiting time: 102.5

2.   Go to Cindy, repair (2 minutes), go to John

Waiting time Cindy: 10+5   =   15
Waiting time John: 85+5+20+30   =   **140**
Average waiting time: 77.5

Figure 5: Instance at T=25 , minimizing waiting time

Before diving further into the prioritizing of performance indicators and subsequently selecting dispatch objectives, it is important to understand the influence of the real time character of the problem. The operational performance is evaluated using end of day measurements. Since the problem reveals itself in real time, the direct effect of dispatch decisions on end of day performance is unclear when the dispatch decisions are made. By consequence, it is not necessarily the case that a dispatch decisions which optimizes performance with respect to a certain objective at some moment during the day, leads to optimal end of day results with respect to this objective. For example it is not clear whether optimizing the end of day performance with respect to the 60 and 90 minutes service levels is best achieved by taking the same objective functions for the real time dispatching problems which are solved during the day. In fact, from our simulation experiments it became clear that this was not the case. Regardless of the algorithmic approach, minimizing waiting in excess of 60 and 90 minutes turned out not to be a good objective for the real time instances. Intuitively this can be explained as follows. Since long waiting only starts after 60 minutes, this objective function doesn't distinguish between arriving after 15 or after 50 minutes. Hence it allows for long travel times, which lead to inefficiencies that might cause problems afterwards. Instead, in order to minimize end of day long waiting, it turned out to be favorable to place an emphasis on minimizing the average waiting times when solving the real life instances. Further discussion and computational experiments can be found in [7] and [10].

The simulation environment we developed allowed analyzing and learning these performance indicator relationships while developing algorithms and before putting the automatic dispatch software into practice. As it deepened the discussion on prioritizing operational performance indicators and linking them with the strategic objectives of ANWB, doing this phase thoroughly laid the foundation for subsequent achievements. We'll continue the discussion after introducing the models and algorithms we considered.

Models and algorithms

Initially we considered three different algorithmic approaches:

1. Multiserver approaches. In this approach the service men are modeled as servers who service requests that pop up in a two dimensional space, as considered in [2] and [6]. Such models are typically solved using 'server policies' which are simple dispatching rules, describing how to select requests which are served next, and by which server. Various authors ([2],[6]) report that a simple First Come First Serve (FCFS), where requests are serviced by the nearest available free service men, performs well from a theoretical as well as a practical viewpoint. Notice that such server policies are quite myopic and disregard some of the requests while making dispatching decisions for others.
2. Matching approaches. In this approach, the dispatching problem at a given time instant is modeled by a bipartite graph, in which the service men form one vertex class, and the request the other vertex class. The weight of an arc between a service man and a request equals the (expected) contribution to the objective

function when the service men is assigned to the request. For instance, the weight might be the expected waiting time exceeding 60 minutes if the corresponding service man travels to the request after finishing the request he is currently servicing (of which the duration is unknown). The solution approach is now to find a maximum cardinality matching of optimum weight. Notice that when the number of request exceeds the number of service men, the maximum cardinality matching leaves some requests unassigned. Therefore, the value doesn't take each of the customers into account, especially when the work load is high and long waiting is likely.

3. A set partitioning model, as proposed in [8][9]). In this model the set of unassigned requests is partitioned into subsets, exactly one service man is assigned to each subset, and a service men is assigned at most one subset. The value of a partition is obtained through adding the values of the subset-serviceman combinations. The value of a subset serviceman combination is determined by calculating the expected 'cost' when the serviceman services the requests in the subset. Since the cost typically involve waiting time or travel time, this requires to determine an order in which the requests in the subset are visited, and estimates on service time. The cost for a given combination can subsequently be optimized by solving a (small) TSP. In contrast to the matching approach and the multi server approach, the set partitioning approach always takes into account all requests when determining a next incident for each serviceman.

We didn't consider local search approaches, since the results of Krumke et al. [8],[9] indicated that the set partitioning approach clearly outperformed an advanced local search algorithm
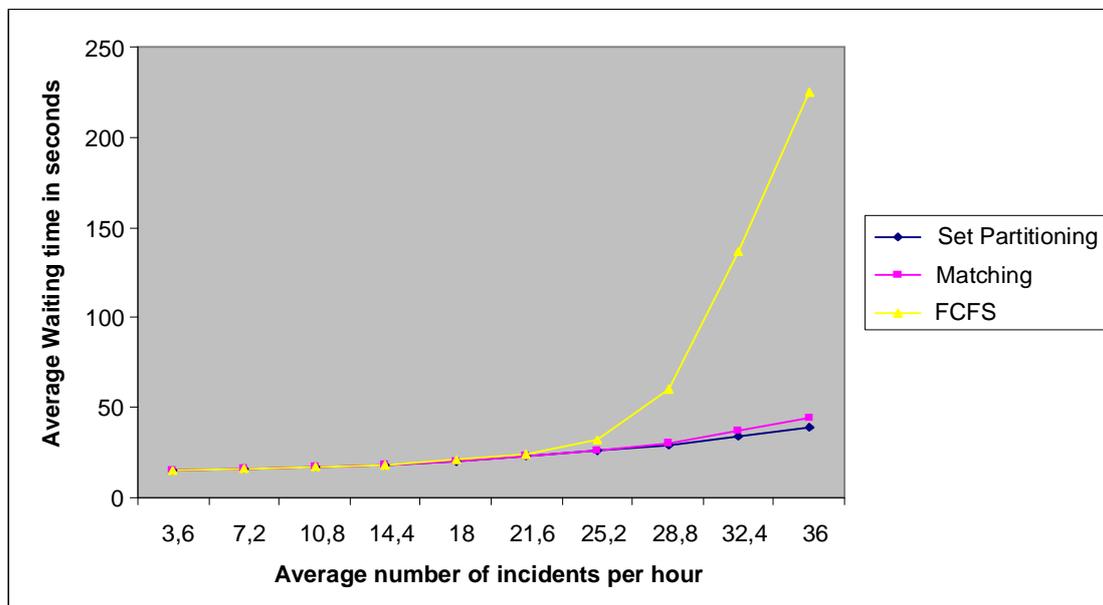

Figure 6: Comparing average waiting times for different solution approaches on test instances

When the work load is modest to average, each of the three considered algorithms performed equally well in our simulation experiments. However under heavy work load, the multi server policies performed poorly, leading to growing and unacceptably high waiting times even when the matching and set partitioning approaches were still providing satisfactory performance. Figure 6 illustrates how average waiting times of requests gets out of hand with the FCFS policy as the number of expected requests per hour increases in a stylizes setting which is derived from practical data. At the same time it illustrates that the matching and set partitioning approach still deliver acceptable average waiting times, the set partitioning approach performing slightly better in this respect.

Figure 7 displays the waiting time results for the densest of the instances in which 36 incidents arrive per hour. Despite the small difference in average waiting time, the percentage of customers waiting more than 60 minutes is around 15 % in the set partitioning algorithm, whereas it is more than 22 % under the matching approach.
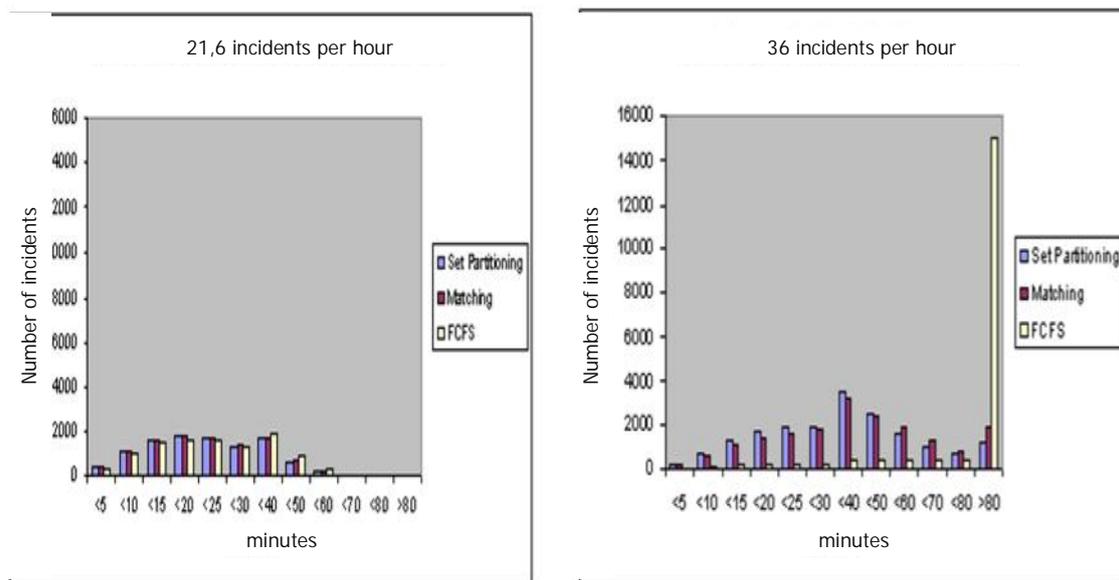

Figure 7: Comparing waiting times distributions for different solution approaches on two test instances

In the simulation experiments, the matching and set partitioning approaches gave comparable results with respect to the average waiting time, while the set partitioning approach performed slightly better with respect to reducing waiting in excess of 60, respectively 90 minutes.

Implementation

Because the matching approach was much simpler, and therefore appeared more flexible in the short run, we decided to use the matching approach in the first releases of the

automated dispatch software. It went life late 2003 and performed satisfactory. In various stages, the planning process was centralized and soon afterwards the whole of The Netherlands was dispatched from the central dispatch center. The percentage of automatic dispatches easily reached a level above 80 %. Thus from a technological viewpoint, and from the viewpoint of the dispatch process, the success was clear.

The impact on operational performance was harder to assess. During this initial phase of using the new dispatching technology the number of service requests dropped by more than 5 %, and the number of service men was lowered even more. At the same time, ANWB was replacing some repair cars by tow away vehicles and services. Nevertheless, ANWB managed to sustain its responsiveness according to agreed company targets, while reducing operating costs.

Set partitioning

While taking the Matching algorithm in operation, we worked on the implementation of the set partitioning algorithm. In order to have a reliable and fast solution method designed to solve exactly the huge problem instances every 30 seconds, we implemented an a column generation and ILP based branch & bound algorithm which was solved using Cplex [7,8,10]. This set partitioning algorithm was implemented and went life in early 2006. To ensure reoptimization within 30 seconds required advanced integration with the shortest route calculation engine. The performance improvement over the matching approach was not the only reason to use the set partitioning algorithm. We expected that the set partitioning formed a much better basis for implementing service and process innovations that ANWB had planned for the near future. For instance, the set partitioning algorithm enables to provide a personal expected waiting time for each newly arriving customer request. This is valuable because of the aforementioned notion that customer dissatisfaction from long waiting is less if they are informed of it in advance. Simulation experiments indicated that such customized waiting time estimates could be reliably met. Moreover, a next project for ANWB has been to introduce a separate operational organization of less extensively trained and equipped service operators, to service a large set of simple requests such as battery problems. Taking full advantage of the possibilities of this new organization required the dispatching of these two organizations to be integrated. The set partitioning algorithm has been adapted to deal with this new situation, whereas the matching algorithm couldn't.

Organizational issues

Project Dispatch was part of a much larger program which ANWB executed in order to adapt itself to the changing and more competitive industry circumstances. Although it was widely accepted that change was necessary, it obviously wasn't greeted with enthusiasm at all stages. Although none of the employees was fired, many positions disappeared, some people were relocated, and for others their jobs changed. The dispatch project had of course a huge impact on the dispatch organization, rendering 70% of the

dispatchers to become obsolete. Moreover, it implied that the service men which form the heart of the primary process received work instructions from a software system. Hence there was much skepticism, and the process was rolled out very carefully. Mistakes could easily lead to serious delay and/or not realizing the original objectives. The fact that the dispatching algorithms were quite advanced, even for the professionals involved, was and is considered as a risk. The results cannot be easily interpreted or explained. Moreover, the advanced technology sometimes makes it more difficult to communicate about operational and technical issues in clear and simple statements. All in all, we conclude that executing the implementation at a pace which was organizationally feasible, instead of at the pace which was technologically feasibly, has delivered ANWB the success of bringing its dispatching from a relatively outdated process to a state where they are years ahead of their competitors and counterparts in other European countries – not only in technology but also in aligning dispatch with operations strategy.

Business Results

As described in the introduction the road repair service industry landscape in The Netherlands is changing. ANWB is defending its former monopoly earned market share and customer base, while being under attack by new players and substitutes. The responsiveness of ANWB and its competitors has drawn considerable attention by national media ranging from leisure magazines to national television and radio. ANWB has performed well in the tests which media have reported on. Of course this service delivery performance is the result of the operational performance and the dispatching. In addition, responsiveness plays a dominant role in the marketing campaigns and has as such even lead to court cases regarding unjustified claims (which ANWB has won). The importance attached to responsiveness in public relations and marketing, combined with the pressure competitors put on it by competing heavily on price, makes responsiveness a very strategic and sensitive issue, on which none of the competitors publicly reports.
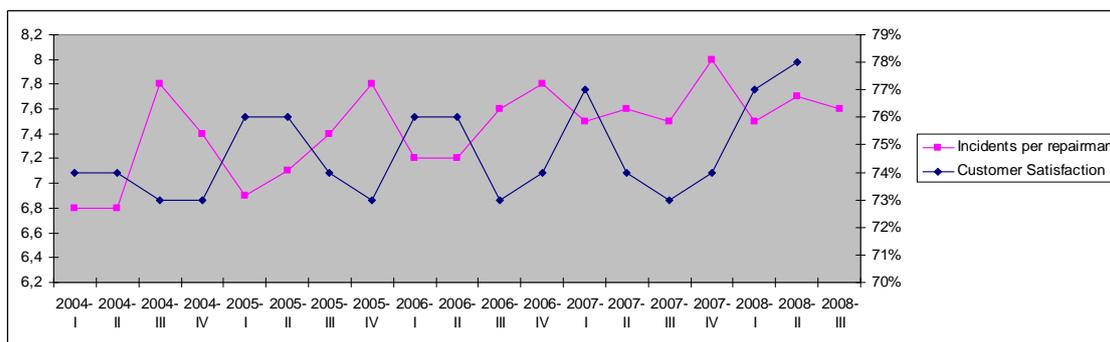

Figure 8: Quarterly results

In view of the above, and since we have argued that average waiting time is not a key driver of customer satisfaction, we report directly on customer satisfaction on the one hand, and on operational efficiency on the other hand in Figures 8 and 9 below, in which. Customer Satisfaction is measured by the percentage of customers which have requested

service and rate the service very good or excellent. Productivity is measured in average number of requests serviced per day per repairman. The figures clearly display that ANWB has succeeded in increasing Customer Satisfaction and efficiency simultaneously. Especially Figure 8 which provides a moving average over the last four quarters, thereby eliminating seasonal effects, shows how productivity has steadily improved, especially after the introduction of the set partitioning based method.
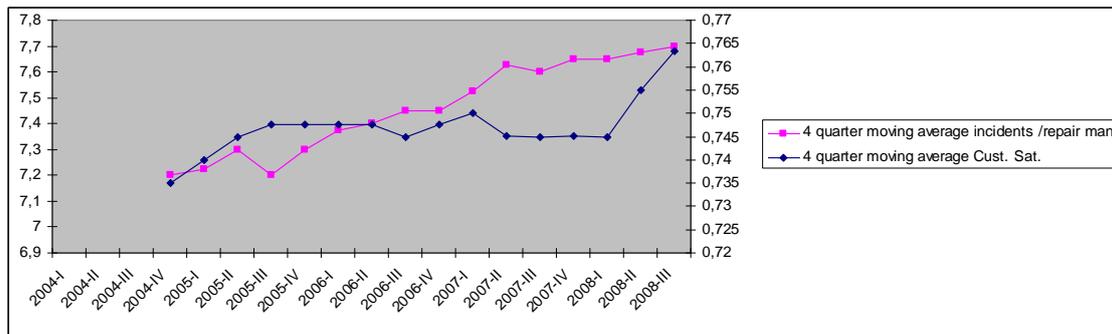

Figure 9: Yearly results (by 4 quarter moving averages)

Of course, implementing a new dispatch method in itself hasn't increased productivity of the repairmen. As demand is exogenous, productivity is simply determined by the number of incidents divided by the number of repairmen scheduled to be on duty. Hence the increased productivity results from reducing the workforce relative to the number of incidents, enjoying the capability of the set partitioning method to maintain the same level of responsiveness. Figure 10 displays how the number of incidents, the workforce, and the productivity have developed since in the introduction of the set partitioning algorithm. It clearly displays the slight reduction in incidents as it results from the developments discussed in the introduction. Moreover, we notice that ANWB has reduced its number of repairmen more than the decline in demand and indeed attained a productivity which is higher by around 14%. This capacity reduction resulted in a yearly reduction in operating costs of several million euro, which exceeds the total project cost.

Customer satisfaction is not only viewed from the perspective of the percentage of serviced customers who rate the service to be very good or excellent. As less than satisfactory service is a key driver of churn, we also considered the percentage of serviced customers who rated the service less than good. Since 2006 this percentage has been stable around 4%. Of course absolute waiting time is not the only cause of not meeting customer expectations. For instance the repair activity itself and its outcome play a role,, next to circumstances which are beyond the control of ANWB. However, additional customer satisfaction research has revealed that dissatisfaction is especially likely when waiting times are substantially longer than the average present waiting which is currently reported as an indication. Hence the possibility of the set partitioning based method to provide a reliable individual waiting time estimate allows ANWB to further improve customer satisfaction in the near future.

Although we judge it impossible to isolate the contribution of the new dispatch methods and software, it is clear that since its introduction ANWB has considerably improved its

combined level of customer satisfaction and operational efficiency. Moreover it has achieved this while reducing the required number of planners from over 50 to 14. Altogether we therefore conclude that it has played an important role in considerably reducing costs, while slightly improving customer satisfaction as required by the very challenging competitive environment. Thus, the use of advanced state of the art combinatorial optimization techniques and technology has played a crucial role in the realization an ambitious corporate strategy. The system has by now several years of reliably performance in dispatching repairman for almost all of 1300000 requests per year, 24/7, 365 days per year. To the best of our knowledge, this makes ANWB the first large service organization to implement such an advanced real time dispatching algorithm.
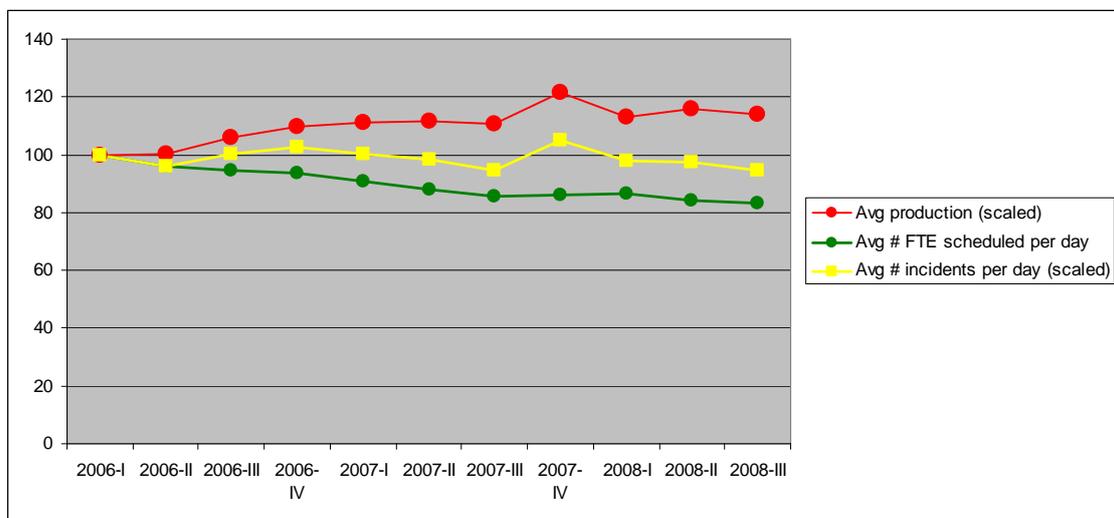

Figure 10: Quarterly results on operational efficiency

The achieved business results have been necessary for ANWB to deal with the increased competition. Despite the entrance of new substitutes and competitors, ANWB lost less than 5 % of the member population from former monopoly-like times, and helped it to sustain its dominant market position despite fierce and aggressive competition.

Acknowledgement

Literature

[1] ANWB, Klanttevredenheidsonderzoek, Kwartaalrapportage, 2003, (In Dutch), ANWB.

[2] Bertsimas, Dimitri, J., Van Ryzin, Garret., A stochastic and dynamic vehicle routing problem in the plane with multiple capacitated vehicles, Operations Research, Vol 41., Nr 1, pp 60-76.

[3] Gendreau, Michel, Potvin, Jean-Yves, 1998, Dynamic Vehicle Routing and Dispatching, in Fleet Management and Logistics, pp 115-226, Kluwer Academic Press.

[4] Hiller, Benjamin, Krumke, Sven, O., Rambau, Jorg, 2006, Reoptimization gaps versus error models in online-dispatching of service units for ADAC, Discrete Applied Mathematics, 154, pp 1897-1907.

[5] Ichoua, Soumia, Gendreau, Michel, Potvin, Jean-Yves, 2000, Diversion Issues in Real Time Vehicle Dispatching, Transportation Science, Volume 34 , Issue 4 , pp 426 – 438.

[6] Irani, Sandy, Lu, Xiangwen, Regan, Amelia, 2004, On-Line Algorithms for the Dynamic Traveling Repair Problem, Journal of Scheduling, Volume 7, Nr 3, pp 243-258.

[7] Klundert, Joris van de, Wormer, Laurens, The after salesman problem, 2008, Working Paper, Faculty of Economics and Business Administration, Universiteit Maastricht.

[8] Krumke, Sven, O., Rambau, Jorg, Torres, Luis, M., 2002, Real-Time Dispatching of Guided and Unguided Automobile Service Units with Soft Time Windows, in Lecture Notes in Computer Science, Algorithms - ESA 2002: 10th Annual European Symposium, Rome, Italy, September 17-21, 2002. Proceedings, pp 417-424, Springer Berlin/Heidelberg.

[9] Krumke, Sven O., Saliba, Sleman, Vredeveld, Tjark, Westphal, Stephan, 2006, Approximation Algorithms for a Vehicle Routing Problem, Working Paper, Faculty of Economics and Business Administration, Universiteit Maastricht.

[10] Wormer, Laurens, 2005, On line toewijzen van Autoreparatiediensten (in Dutch), Master's Thesis, Faculty of Economics and Business Administration, Universiteit Maastricht.