International Journal of Production Economics, 23 (1991) 69–78
Elsevier

A DSS for capacity planning of aircraft maintenance personnel

Matthijs C. Dijkstra, Leo G. Kroon, Jo A.E.E. van Nunen and Marc Salomon
Rotterdam School of Management, Erasmus University, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

Abstract

In this paper we describe a Decision Support System (DSS) that has been developed for the aircraft maintenance department of the Dutch national airline company at the main airport in the Netherlands. The aircraft maintenance department is responsible for carrying out the regular short inspections of aircraft between their arrival at and their consecutive departure from the airport. The main resource of the aircraft maintenance department is its workforce. The DSS that has been developed can be used to support the management of the maintenance department in solving several capacity planning problems related to the size and the composition of the workforce. In this paper we give a description of the capabilities of the DSS. Furthermore, we describe the solution technique that is applied within the DSS for determining the required size and composition of the workforce.

1. Introduction

In this paper we describe a Decision Support System[1] (DSS) that has been developed for the aircraft maintenance department of the Royal Dutch Airline Company at Schiphol Airport. The Royal Dutch Airline Company is the major airline company in the Netherlands. In Dutch the name of this company is "Koninklijke Luchtvaart Maatschappij" (KLM). Schiphol Airport is the principal Dutch airport, situated near the Dutch capital Amsterdam.

The aircraft maintenance department is responsible for carrying out the regular short inspections of aircraft between their arrival at and their consecutive departure from the airport. The aircraft are owned by KLM and several other airline companies, as far as these companies have a contract with the aircraft maintenance department. An important management problem to be solved in this context is to guarantee that always sufficient engineers with appropriate qualifications are available at the airport for carrying out the required inspections. The DSS is used for supporting the management of the maintenance department to solve this kind of capacity planning problems. The quality of the maintenance department as a whole is measured in terms of the service level of the department and in terms of the efficiency of the workforce.

The DSS that was developed for the maintenance department of KLM belongs to a class of DSSs that focus on situations, where the workload of a (service) department can be projected some time in advance, based on a (preferably cyclic) time table and a set of norms specifying in which time intervals and how much service has to be delivered by the department. Such situations are encountered often within railway and airline companies (maintenance, catering, cleaning). The mentioned DSSs can be used to determine the capacity of the main resource of the involved department in such a way that a global matching can be expected between the determined capacity and the projected workload.

---

[1] Keen and Scott Morton [13] define a DSS as a coherent system of computer based technology used by managers as an aid to their decision making in semi- or unstructured decision tasks with the objective to support rather than to replace managerial judgement and to improve the effectiveness of decision making.

For example, Jacquet-Lagrèze and Meziani [1] built a DSS with such objectives for the aircraft cleaning department of Air France. Another application of such a DSS can be found on Schiphol Airport. There a DSS is being developed which can be used to determine the required number of gates. These gates are used for transferring the passengers between the aircraft on the platform and the Arrival and Departure Halls of the airport.

The remainder of this paper is organized as follows. In Section 2 a more detailed description of the situation at the airport is given. This description highlights both the workload and the workforce of the maintenance department. In Section 3 an overview is presented of the functions, the appearance, and the implementation of the DSS. The mathematical model that is used within the analysis module of the DSS is considered in Section 4. In addition, an approximation technique based on Lagrangean Relaxation is proposed for solving this model. The paper is concluded with some evaluating comments in Section 5.

2. Case description

In this section a global description of the practical situation within the maintenance department of KLM at Schiphol Airport is given. First, we describe the principal components determining the workload and the workforce of the department. Thereafter, the main management problem of finding a global matching between workload and workforce is considered.

2.1 The workload of the maintenance department

Generally, it is required that between the time of arrival and the time of departure of any aircraft at Schiphol Airport the aircraft is inspected and, if necessary, repaired before it is allowed to take off. The inspections of the aircraft of KLM and several other airline companies are carried out by the aircraft maintenance department of KLM. The major *technical tasks* of the maintenance department are the *arrival services, platform inspections,* and *departure services* of aircraft.

Figure 1 shows the workload of the maintenance department on a typical saturday as it is represented by the DSS. The workload is based on (i) *time tables,* (ii) *contracts with foreign companies,* and (iii) *maintenance norms.*

Since the time tables of airline companies usually have a cyclic character with a cycle length of one week, the workload of the maintenance department has a similar character. Furthermore, the workload on an average day shows some clearly distinguishable peaks, which are caused by the desire of KLM to have short transfer times between intercontinental and continental flights.

The maintenance norms specify in which time interval each task must be scheduled and how much time must be spent for each task. They are dependent on the aircraft type and on the airline company owning the aircraft. The maintenance norms are provided by the aircraft manufacturers, by the government, and by the airline companies.

2.2 The workforce of the maintenance department

The technical workforce consists of *ground engineers.* These engineers are highly educated employees, as their job is a very responsible one. After the inspection of an aircraft the engineer that carried out the inspection is responsible for the technical condition of the aircraft. Therefore, a governmental rule specifies that an engineer is allowed to carry out the inspection of an aircraft completely individually only if he has a *license* for the corresponding aircraft type. An engineer can obtain a license for a specific aircraft type by attending the required courses, passing the examinations, and getting the required amount of experience in practice. This process takes several months to several years, depending on the previous experiences of the engineer. From an operational point of view, it would be optimal if all engineers would have licenses for all aircraft types. Then the flexibility of the engineers would be maximum. However, a governmental rule specifies that engineers are allowed to have *two* licenses at most.

The organization of the workforce is determined by a clustering of the engineers into a number of *teams,* which constitute the smallest

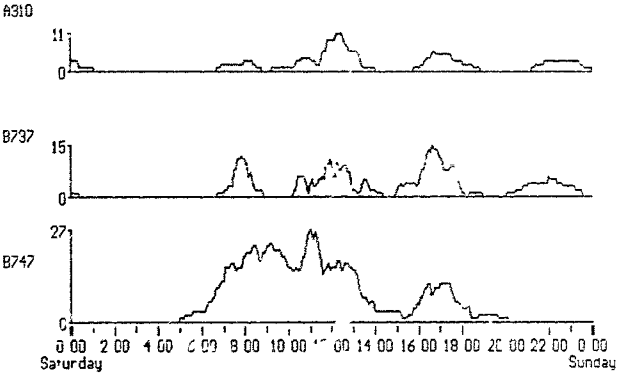

Fig 1 Workload on a typical saturday

organizational units At this moment there are 12 teams of engineers. The assignment of engineers to teams is such that the compositions of the teams are (almost) identical Each team consists of about 18 engineers The teams have to be available at the airport in specific time intervals of about eight consecutive hours (*shifts*). Figure 2 gives a graphical representation of the shifts that are carried out each day.

The assignment of teams to shifts is limited by several restrictions. For instance, for each team the average number of shifts per week should be equal to 5, and there should be at least one day off between a night shift and a consecutive day shift It is clear that it is a difficult management task to satisfy all these restrictions

2 3 Management problems

As was pointed out in Section 1, an important problem the management is confronted with is to guarantee that always sufficient engineers with appropriate qualifications are available at the airport, and thus to realize a global matching between workload and available workforce The *quality* of the matching can be expressed along several dimensions, such as the expected efficiency of the workforce and the expected service

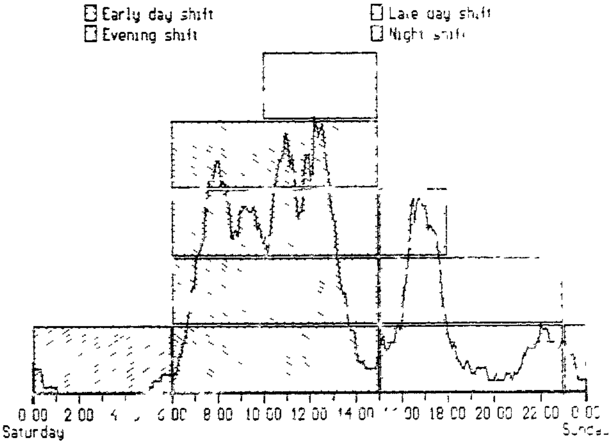

Fig 2 Daily shift scheme

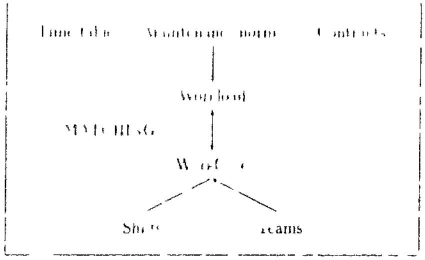

Fig 3 Scheme of problem components

vel of the maintenance department The major aspects concerning the capacity planning problem of the workforce of the maintenance department are summarized in Fig 3

Unfortunately not all aspects in Fig. 3 can be controlled within the maintenance department. With respect to the workload, decisions can be taken about the number of contracts with other airline companies and about the introduction of new technologies which may influence the maintenance norms However, there is little influence on the composition of time tables, both of KLM and contract companies On the other hand, most of the aspects that determine the size and the organization of the workforce can be controlled within the maintenance department to a large extent

3. Decision support system

In this section we give a description of the DSS that has been developed for the maintenance department of KLM We consider the functions, the appearance, and the implementation of the DSS

3.1 Functions of the DSS

The DSS can be used to make an *estimate* of the future workload, based on the time tables of KLM and other companies, on the maintenance norms, and on the contracts with other companies Moreover, the DSS can be used to *evaluate* the quality of the matching between a given workload and workforce In this way one can assess the [illegible] of the matching [illegible] to variations [illegible] the [illegible] the composition of the teams, the number of shifts per day, the begin and end times of the shifts and the number of teams per shift Finally, the DSS can be used to support the *determination* of the size and the composition of the teams, such that a global matching between workload and workforce can be expected Criteria used are the efficiency of the workforce and the service level.

3.2 Appearance of the DSS

The appearance of a DSS is determined by the User System Interface (USI). The USI provides the opportunity to specify commands in a *command language* and it responds to these commands in a *presentation language* (Be[illegible] [2]). We have chosen a menu driven command language which was accepted immediately by the users, as it is easy to learn and easy to use As we believe that a clear picture can tell more than a thousand words, we have chosen a presentation language that presents most of the results of the analyses in a graphical format The Figs 4 and 5 give an impression of the graphical capabilities of the DSS

The dark areas in Fig 4 show that part of the workload on a typical saturday that can be carried out by the available engineers. The small grey areas show that part of the workload that can not be carried out due to insufficient capacity of the workforce

Figure 5 shows the activities of the engineers on a typical saturday, split by license combination. Furthermore per license combination the activities are split by aircraft type. The various patterns represent the various aircraft types per license combination

However, a graphical format is valuable for providing a global overview of the quality of these results, but may be insufficient for providing detailed overview Therefore, some of the results of the analyses can be represented both in a graphical format and in a tabular format

3.3 Implementation of the DSS

From a technical point of view the DSS consists of a *database* and two modules a *database*

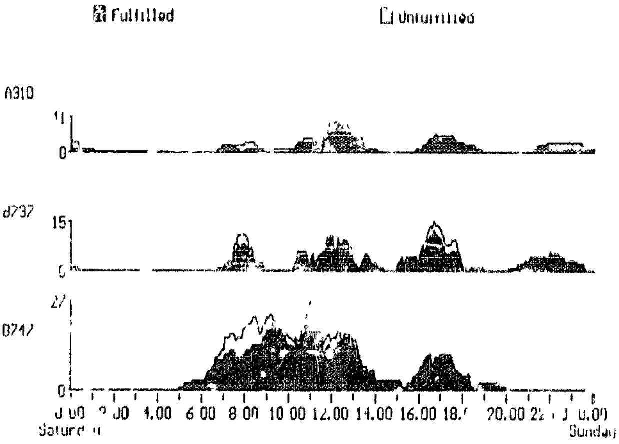

Fig. 4. Service level on a typical saturday

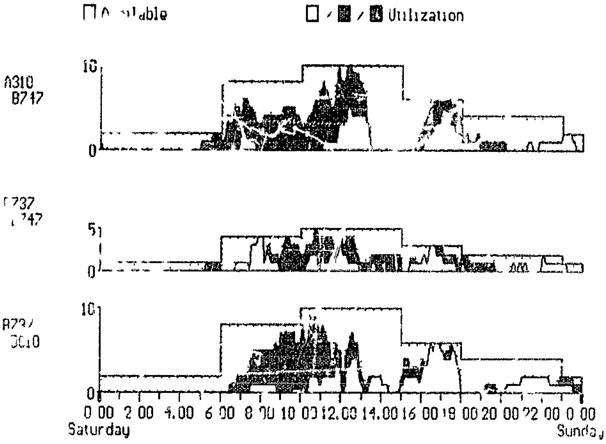

Fig. 5. Efficiency on a typical saturday

management module and an *analysis module*[2]. The database has a number of subdatabases. Each subdatabase contains a full set of data describing a *scenario*. The data per scenario are split up into three categories: basic data, data concerning the workload and data with respect to the workforce. The main categories of the data in the latter categories have been represented in Fig. 3.

The database management module provides functions for operations on complete scenarios and functions for operations on data of a selected scenario. A user system interface, which is part of the database management module, gives access to these functions.

The analysis module enables the analysis of a selected scenario through functions for *estimating* the workload and the workforce, for *evaluating* the matching between the workload and the workforce, and for *generating* specifications for

---

[2] The Clipper™ Compiler is used for the implementation of the database management module. Borland's Turbo™ Pascal 5.0 is used for the implementation of the analysis module. The DSS runs on an MS DOS™ personal computer.

the organization of the workforce. Part of the analysis module is a user system interface that is consistent in appearance with the user system interface of the database management module. The evaluation function simulates the maintenance process and results in a detailed assignment of jobs to engineers. The obtained schedule is assessed in terms of efficiency and service level, which are the main calculated performance indicators. The generation function results in a desirable size and composition of the teams based on the workload and the shift pattern. In the following section we give a more detailed description of the problem that is considered by the generation function. Furthermore, we give a description of an approximation algorithm that is used for solving this kind of problem.

4 A mathematical description

The aim of this section is to present a mathematical description of the optimization problem that is solved by the generation function of the analysis module. In order to keep the presentation clear, we make the simplifying assumption that there is only one shift and that, as a consequence, all engineers are continuously available at the airport. However, it is easy to modify the model in such a way that the shifts are taken into account correctly. We suppose that a set of jobs to be carried out has been determined. In addition, we suppose that all stochastic elements can be neglected. In practice this assumption is also made by the planners of the maintenance department.

More formally, we have a set J of jobs to be carried out. Job $j \in J$ requires continuous processing [illegible] in the interval (s_j, f_j) and is related to an aircraft of type a_j. Hence, each job j can be represented by a triple (s_j, f_j, a_j). The jobs must be carried out in a *non-preemptive* way by a number of engineers. Each engineer is assumed to have a combination of licenses which specifies the aircraft types the engineer is allowed to work on. The set of different license combinations is denoted by C. For $c \in C$ we associate a cost k_c with each engineer with license combination c.

[illegible] the problem is to determine a minimum cost composition of the workforce, such that all jobs can be carried out. [illegible] and Kroon [illegible] investigate the computational complexity of this problem. They show that it belongs to the class of NP-Hard problems. The problem is a generalization of the well-known Fixed Job Scheduling Problem (FSP), in which it is assumed that each engineer has a license for each aircraft type. FSP has been studied extensively by Dantzig and Fulkerson [4], by Gertsbakh and Stern [5], and by Gupta et al. [6]. The result of Lemma 1 describes the optimal solution of an instance of FSP. In this lemma the maximum job overlap is defined as the maximum number of jobs that must be carried out simultaneously.

Lemma 1. If each engineer has a license for each aircraft type, then the minimum number of engineers required for scheduling all jobs is equal to the maximum job overlap.

In [6] an $O(|J| \log |J|)$ time algorithm for determining the maximum job overlap is presented. Hence, FSP can be solved in $O(|J| \log |J|)$ time. Clearly, FSP provides a lower bound to the generalized optimization problem, where one has to take into account the aircraft types and the licenses of the engineers.

4.1 Formulation as Integer Program

The generalized optimization problem can be formulated as an Integer Program. In the description of this Integer Program we use the notation C_j for the set of license combinations that can be used for carrying out job $j \in J$. Conversely, for $c \in C$ the set J_c denotes the set of jobs that can be carried out by engineers with license combination c. Furthermore, we use the notation $\{t_p | p \in P\}$ for the set of start and finish times of the jobs. That is $\{t_p | p \in P\} = \{s_j, f_j | j \in J\}$. Now the decision variables of the Integer Program can be defined as

X_{jc} = binary variable indicating whether job $j \in J$ has to be carried out by an engineer with license combination $c \in C$.

Y_c = an integer variable indicating the required number of engineers with license combination $c \in C$.

In terms of these decision variables, the objective and the constraints can be stated as follows

$$\min Q = \sum_{c \in C} k_c Y_c \qquad (1)$$

Subject to

all variables are integer (2)

$$\sum_{c \in C_j} X_{jc} = 1 \quad \text{for } j \in J \qquad (3)$$

$$\sum_{\text{[illegible]}} X_{jc} \leqslant Y_c \quad \text{for } c \in C \text{ and } t \in T \qquad (4)$$

The objective function (1) expresses that we are interested in minimizing the total costs associated with the engineers. The integrality constraints (2) specify the integer character of the decision variables. The constraints (3) guarantee that each job is carried out exactly once. Finally, the constraints (4) specify that the maximum job overlap of the jobs that are assigned to the engineers with license combination c should not exceed the number of available engineers with license combination c. Hence, according to Lemma 1, any feasible solution to the Integer Program can be transformed into a feasible assignment of jobs to engineers and vice versa.

4.2 An approximation algorithm

Since the generalized optimization problem is NP-Hard, calculating optimal solutions can be time consuming for instances of the problem containing large numbers of jobs. Hence, for finding satisfactory solutions to real life problem instances within a reasonable amount of time, the use of an approximation algorithm is inevitable. The approximation algorithm that we use is based on Lagrangean Relaxation. Important components are (i) *a lower bounding procedure*, (ii) *an upper bounding procedure*, and (iii) *a procedure for updating the parameters*. These procedures are discussed in the following sections.

4.2.1 Lower bounding procedure

If Lagrangean Relaxation is applied to the constraints (3) specifying that each job must be carried out exactly once, then the following problem is obtained for a given set of Lagrange multipliers λ_j

$$\min Q(\lambda) = \sum_{c \in C} \text{[illegible]} - \sum_{j \in J_c} \lambda_j X_{jc}) + \sum_{j \in J} \lambda_j \qquad (5)$$

Subject to (2) and (4)

Now the problem decomposes into $|C|$ independent minimization subproblems. It is not difficult to see that each of these subproblems is a Minimum Cost Flow Problem in a directed graph with $O(|J|)$ nodes and $O(|J|)$ arcs. The details of the construction of the underlying network are given by Kroon [7]. It is well known, that Minimum Cost Flow Problems can be solved efficiently (Orlin [8]). Note that, due to the integrality property of these subproblems, the value of the Lagrangean dual equals the value of the Linear Programming [illegible] relaxation of Q (Fisher [9]).

4.2.2 Upper bounding procedure

As Q is a minimization problem, any feasible solution provides an upper bound to the optimal solution. In our case a feasible solution can be constructed [illegible] applying a "greedy" procedure. The procedure [illegible] be summarized as follows: initially the number of engineers is zero and all jobs are unassigned. As long as there are unassigned jobs, [illegible] workforce is enlarged by one engineer with a [illegible] *best* license combination. The jobs that can [illegible] carried out by this engineer are assigned to [illegible] These steps are repeated until all jobs ha[illegible] assigned. More formally, this procedure [illegible] as follows.

```
$k = 1$
$S = \{$all [illegible]
Repeat
    Cho[illegible] [illegible] best license combination $c^*$
    Add e[illegible] $E_k$ to the workforce
    Engin[illegible] obtains license combination $c^*$
    [illegible] that can be carried out by $E_k\}$
    [illegible]
Until [illegible]
```

In [illegible] description the set S represents the set of una[illegible] jobs. Now the only steps that are still uns[illegible]ed are the choice of a locally best license combination c^*, and the determination of a maximum weighted set of jobs assignable to c^*. These two steps are carried out simultaneously in the following way. Suppose that a partial workforce has been determined, and that the set S contains the jobs that are still unassigned. In this situation the priority of license combination c is measured as the value of a maximum weighted set of jobs $j \in S$, assignable to one engineer with license combination c minus the cost of one engineer with license combination c. The Lagrange multipliers λ_j, which are obtained from the lower bounding procedure, are interpreted as weights of the jobs. The priority of license combination c, given the set of unassigned jobs S and the Lagrange multipliers λ_j, is called $Q^{Sc}(\lambda)$. The given description implies that $Q^{Sc}(\lambda)$ is calculated as follows

$$\max Q^{Sc}(\lambda) = \sum_{j \in J_c \cap S} \lambda_j X_{jc} - k_c \tag{6}$$

Subject to

all variables are integer (7)

$$X_{jc} \leqslant 1 \quad \text{for } j \in J_c \cap S \tag{8}$$

$$\sum_{j \in J_c \cap S \mid s_j \leq t_p < f_j} X_{jc} \leqslant 1 \quad \text{for } c \in C \text{ and } p \in P \tag{9}$$

It is not difficult to see that $Q^{Sc}(\lambda)$ can be calculated as a shortest path in a directed graph with $O(|J|)$ nodes and $O(|J|)$ arcs (Kroon [7]). Now a locally best license combination c^* is defined as follows

$$c^* = \operatorname{argmax}\{Q^{Sc}(\lambda) \mid c \in C\}$$

4.2.3 Update procedure for the parameters

Initially, the Lagrange multipliers are set equal to 0. Thereafter they are updated iteratively by a subgradient procedure as described by Fisher [9]. The new values are based on the old values and on the solution obtained by the lower bounding procedure. This solution may be infeasible in the sense that some jobs are carried out not at all, whereas others are carried out several times. Now the Lagrange multipliers are updated in such a way, that jobs that are carried out not at all are stimulated in the next iteration and that jobs that are carried out several times are demotivated in the next iteration. More precisely, when going from iteration t to iteration $t+1$ we use the following update scheme

$$\lambda_j^{t+1} = \lambda_j^t + \omega^t \cdot \left(1 - \sum_{c \in C_j} X_{jc}\right) \quad \text{for } j \in J$$

In this formula the values of the variables X_{jc} were obtained by the lower bounding procedure. Furthermore, the multiplier ω^t is calculated as follows

$$\omega^t = \frac{\mu \cdot (Q_U - Q(\lambda^t))}{\sum_{j \in J}(1 - \sum_{c \in C_j} X_{jc})^2}$$

where Q_U is the best upper bound found so far and μ is a scaling parameter. Several experiments have shown that μ should have a value somewhere between 1 and 2.

4.2.4 Evaluation of the algorithm

The algorithm is stopped either if the gap between lower and upper bound is sufficiently small, or after a pre-specified number of iterations, whichever comes first. An extensive set of numerical results can be found in Kroon [7]. Based on these results we conclude that the described procedures are promising, both in terms of their efficiency and in terms of their effectiveness. Optimal solutions have been obtained several times within an acceptable amount of time. However, the feasible solutions are constructed by applying a greedy method and hence, one is not guaranteed to obtain optimal solutions always. Furthermore, a duality gap might exist and therefore an optimal solution might be not recognized as such (Shapiro [10]).

5. Discussion

Evaluation of a DSS is generally more difficult than evaluation of more traditional information systems, because for DSSs simple criteria like costs and revenues are hardly useful (Keen [11]). First, a DSS is never completely finished and thus the costs are difficult to specify. Second, the revenues of a DSS are found in qualitative aspects such as the impact the DSS has on the organization and on the quality of decision making and decisions. Evaluation of these aspects in quantitative terms is difficult and has not yet ob tained much attention in the literature (Elam et al [12])

As we put much effort into the user-friendliness of the DSS for KLM, the system was soon accepted by its proposed users The system is mainly used by staff employees to answer questions posed by the management of the maintenance department The DSS has been used to determine the impact of the contracts with other airline companies on the size and the organization of the workforce. Furthermore, the DSS is used currently to determine the required number of engineers and their qualifications during Summer 1992

In general the system provides the management with information that was not available before In this way it contributes to an increased insight into the various problems that have to be solved within the maintenance department As a consequence, the DSS is considered as a useful tool for analyzing such problems. The users of the system even advocated its use to other departments of KLM, whose workloads are mainly determined by the time tables of the involved companies, such as the helicopter department

One point of criticism concerning the DSS is, that it focuses too much on long term capacity planning. This is useful for the general management of the department, but the operational managers would like to use the system to simulate the day-to-day maintenance process, which is currently impossible. Furthermore, the DSS could be enhanced at several technical points, such as data management and scenario management A preliminary conclusion is, that the system provides the management of the maintenance department with appropriate support, although its impact should not be overestimated A more detailed evaluation of the system is a subject for further research

When considering the results of a DSS it is important to keep in mind that these results are based on mathematical models which are abstractions of reality As a result, optimality in mathematical terms need not match optimality in practical terms Furthermore, most of the calculations within a DSS are based on approximation algorithms Therefore the results of a DSS must be 'handled with care" It is always the user of the DSS who is in charge of judging the practical value of a solution by confronting it with qualitative or quantitative aspects that were not taken into account by the mathematical models within the DSS. Hence, a DSS must be used in an interactive way, where the intelligence of the user is combined with the capability of the DSS to organize and process enormous amounts of data, and to solve complex mathematical decision problems using sophisticated Operational Research techniques

Acknowledgement

The authors want to express their gratitude to Rene Kalmann, Paul Chün, Thom Grobben, and Jan Smit of the maintenance department of KLM for their many helpful comments during all development phases of the DSS.

References

1 Jacquet-Lagreze, E and Meziani, R , 1988 Linear programming and interactivity A manpower scheduling DSS Working Paper 8 , of Lamsade, Universite de Paris Dauphine

2 Bennett J L , 1983 Analysis and design of the user interface for Decision Support Systems In J L Bennett (Ed), Building Decision Support Systems Addison-Wesley, Reading MA

3 Kolen, A W J and Kroon, L G , 1989 License class design Complexity and algorithms Report 63 of the Rotterdam School of Management European J Oper Res , to appear

4 Dantzig, G B and Fulkerson, D R , 1954 Minimizing the number of tankers to meet a fixed schedule Naval Res Logist Quart , 1 217–222

5 Gertsbakh, I and Stern H I , 1978 Minimal resources for fixed and variable job schedules Oper Res , 18 68–85

6 Gupta, U I , Lee, D T and Leung, J Y -T 1979 An optimal solution to the channel assignment problem IEEE Trans Comput C-28 807–810

7 Kroon, L G 1990 Job scheduling and capacity planning in aircraft maintenance Ph D Thesis, Rotterdam School of Management

8 Orlin, J B , 1988 A faster strongly polynomial minimum cost flow algorithm Proceedings of the 20th ACM Symposium on the Theory of Computing, pp 377–387

9 Fisher, M L , 1981 The Lagrangean Relaxation method for solving integer programming problems Manage Sci 27 1–18

10 Shapiro, J , 1979 A survey of Lagrangean techniques for discrete optimization Ann Discrete Math 5 113–138

11 Keen, P.W.G. 1981. Value analysis: Justifying Decision Support Systems. MIS Quarterly 5, 1–15.

12 Elam, J.J., Huber, G.P. and Hurt, M.E. 1986. An examination of DSS literature (1975–1985). In: E.R. McLean and H.G. Sol (Eds.), Decision Support Systems: A decade in perspective. North Holland, Amsterdam.

13 Keen, P.W.G. and Scott Morton, M.S., 1978. Decision Support Systems: An organizational perspective. Addison-Wesley, Reading MA.