

Schirmer, Andreas; Potzhar, Kathrin

Working Paper — Digitized Version

Professional course scheduling in airline transport pilot training: A case from Lufthansa flight training

Manuskripte aus den Instituten für Betriebswirtschaftslehre der Universität Kiel, No. 539

Provided in Cooperation with:

Christian-Albrechts-University of Kiel, Institute of Business Administration

Suggested Citation: Schirmer, Andreas; Potzhar, Kathrin (2001) : Professional course scheduling in airline transport pilot training: A case from Lufthansa flight training, Manuskripte aus den Instituten für Betriebswirtschaftslehre der Universität Kiel, No. 539, Universität Kiel, Institut für Betriebswirtschaftslehre, Kiel

This Version is available at:

<https://hdl.handle.net/10419/147617>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Manuskripte
aus den
Instituten für Betriebswirtschaftslehre
der Universität Kiel

No. 539

**Professional Course Scheduling
in Airline Transport Pilot Training
- A Case from Lufthansa Flight Training**

Schirmer, Potzahr



**Manuskripte
aus den
Instituten für Betriebswirtschaftslehre
der Universität Kiel**

No. 539

**Professional Course Scheduling
in Airline Transport Pilot Training
- A Case from Lufthansa Flight Training**

Schirmer, Potzahr

June 2001

Please do not copy, publish or distribute without permission of the authors.

Dr. Andreas Schirmer ^{a,b,*}, Dipl.-Inform. Kathrin Potzahr ^c

^a Deutsche Lufthansa AG, Informatiksysteme Personal, HAM PI,
Weg beim Jäger 193, D-22313 Hamburg, andreas.schirmer@dlh.de

^b Institut für Betriebswirtschaftslehre, Christian-Albrechts-Universität zu Kiel,
Wilhelm-Seelig-Platz 1, D-24098 Kiel, Germany, schirmer@bwl.uni-kiel.de

^c sd&m software design & management AG,
Lübecker Straße 1, D-22087 Hamburg, kathrin.potzahr@sdm.de

www.bwl.uni-kiel.de/bwlinstitute/prod/mab/schirmer

Contents

1. Introduction.....	1
2. Educational Scheduling.....	3
2.1. University Scheduling	4
2.2. School Scheduling.....	5
2.3. Conference Scheduling.....	7
2.4. Professional Course Scheduling.....	8
3. Problem Setting.....	11
3.1. Course Scheduling Subproblem.....	11
3.2. Room Assignment Subproblem.....	13
3.3. Modelling.....	14
4. Algorithmic Approaches.....	15
4.1. Course Scheduling Methods.....	16
4.2. Room Assignment Methods.....	20
5. Computational Evaluation.....	21
5.1. Test Instances.....	21
5.2. Computational Results.....	22
6. Summary and Outlook.....	23

Abstract: Several well-known and well-researched problem fields, such as school scheduling or university scheduling, reside in the realm of educational scheduling. Recently, some new course scheduling problems have begun to draw attention, which turn out to sport substantially more complex requirements and objectives than the classical course scheduling problems. We refer to their problem field as professional course scheduling. In order to demonstrate its practical relevance, we describe a real-world application. Lufthansa Flight Training GmbH (LFT) offers license, recurrency, emergency, and human factors training for airline, navy, and air force pilots as well as service and emergency training for cabin attendants of more than 50 airlines worldwide. This charges LFT with the problem to develop a monthly schedule for courses, instructors, and rooms that optimizes an objective function measuring adherence to seven different soft constraints while meeting a number of complex precedence, temporal, and resource-related constraints. In the past, LFT did all its scheduling manually, but management was dissatisfied with this process due to the significant cost and time involved. LFT commissioned us to carry out a feasibility study in which the applicability of operations research methods was to be demonstrated. We developed a prototype decision support system which utilizes construction and neighborhood search methods based upon concepts from project scheduling and graph theory. It turned out that significant improvements over the manual process could be realized; in addition, the algorithmic ideas employed are general enough to be easily adapted to other problems in the field of professional course scheduling. The development of a full-fledged decision support system is currently in progress at LFT.

Keywords: DECISION SUPPORT SYSTEM; COURSE SCHEDULING; AVIATION; TRAINING; EDUCATION

*"A planning question, after all, is likely to have an answer."
(Solow 1963)*

1. Introduction

The Lufthansa Group has more about 70,000 personnel, making it one of the largest employers in Germany and one of the largest airlines worldwide. With flight and cabin crews of about 3,000 and 11,000, respectively, training its own crews has a long tradition with Lufthansa German Airlines and is considered key to safety and service quality. All related activities are combined in Lufthansa Flight Training GmbH (LFT). LFT mainly offers license, recurrency, emergency, and human factors training for airline pilots as well as service and emergency training for cabin attendants; other training areas are being added to the product portfolio. The customer base currently lists more than 50 airlines, along with more than a dozen other companies and organisations, not all of them aviation-related. Part of LFT is the Verkehrsfliegerschule (air transport pilot school) which provides the license training that prepares student pilots *ab initio*, i.e. without prior flying experience, for the airline transport pilot license (ATPL). In addition to this civilian side, since 1963 the German Air Force and Navy entrust the complete license training of their transport pilots to the LFT. Founded in May 1956, the Verkehrsfliegerschule has trained more than 5000 pilots to the ATPL level. The school enlists a staff of about 200, a third of which are instructors.

Pilot training is categorized primarily into ground and flight training, the former focussing on the theoretical, the latter on the practical aspects of commandeering an aircraft, so each course is divided into ground and flight training phases. While the Verkehrsfliegerschule conducts

ground training in its Bremen facilities, for economical and weather-related reasons most of the flight training takes part in two LFT subsidiaries in Arizona. The contents of each course are spelled out in a syllabus which details the subjects to be covered, the topics of each, as well as the total number of lessons per topic. Note that, due to the long duration of courses (22 months in most cases), each course is designed individually, allowing to accommodate upcoming legislative changes as well as customer wishes as the need arises. Currently, the ground training comprises a total of 1248 lessons of 45 minutes each, for a total of about 208 working days plus some additional buffer days allowing to compensate for unforeseen events.

Throughout this paper, we concentrate on the problem of scheduling the complete ground training activities of the LFT. Such a schedule defines for each course when and where which topic is taught by which instructor. Operational schedules have to meet a variety of constraints and are judged against seven objectives, some of these mutually conflicting. A schedule usually covers a planning interval of 20 workdays. Hence, due to the repetitive nature of the scheduling process, instruction already completed must be tracked to allow feasibility checking.

LFT management was dissatisfied with the existing manual procedure because of its significant cost. Currently, schedules are constructed manually which is an extraordinarily tedious and time-consuming task that monopolizes the five ground training department heads for almost a full day each time. As this process repeats itself every 20 working days, their total load amounts to about 55 man-days per year. In addition, one staff member is fully occupied by clerical tasks related to scheduling, half of which are directly attributable to the manual character of the scheduling process. Management felt that reducing the time invested by the department heads alone would be well worth the cost of this project. Also, due to cost and duration of the procedure, it is virtually impossible to apply and assess different planning strategies. This, however, becomes more and more desirable since - following the current upswing in air transport demand - the pilots school is working near its capacity limits w.r.t. several resources. In addition, a computer-supported scheduling process should allow for faster answers to questions like, Can we fit an additional course into our ongoing commitments?, which is important from the marketing perspective. In contrast, currently an aggregate summary of all courses is kept on a planning board which takes several square meters of wallspace. Some years ago, LFT had asked external consultants to look into ways to provide computer-based decision support for this process, a task which the consultants verbatim reported to be "impossible". Recently, LFT staff had scrutinized several timetabling software packages available on the marketplace - all of these had been found to be unsuitable for the problem faced. Now, LFT management commissioned us to carry out a feasibility study, in cooperation with LFT staff, on the applicability of operations research methods for the problem setting outlined. The following objectives were established for the study:

- Generate usable schedules, i.e. respect all planning constraints (of course, this objective holds for manual planning, too);
- Reduce the planning time required as compared to manual planning;
- Improve upon the quality of manual planning, i.e. achieve better results with respect to the planning objectives;
- Generate several alternative course schedules.

All these objectives were achieved. The findings of the study demonstrate that a decision support system employing appropriate heuristic algorithms can substantially reduce time and effort required to come up with solutions while improving the quality of the solutions produced as well.

Finally, we provide an outline of this contribution. While the field of educational scheduling is diverse, offering many interesting methodologies to address such problems, we found LFT's problem too different from the classical problems to warrant adopting their methodologies. To demonstrate this, in Section 2 we survey the literature on educational scheduling, exposing the substantial differences between professional course scheduling problems (such as the one of LFT) and more traditional course scheduling problems. In the remainder, we describe the major deliverables of the project. In Section 3, we introduce the problem setting of LFT which for convenience and efficiency we divide into a scheduling and a room assignment subproblem. In Section 4, we then adopt techniques from project scheduling and graph theory to develop construction and neighborhood search methods for both subproblems. Our purpose is not only to report from a successful scheduling implementation but also to provide sufficient detail so that others facing similar problems can develop a scheduling tool suited to their needs. Section 5 reports on the results of our computational experiments. Section 6 concludes the paper with some final remarks.

2. Educational Scheduling

Educational or course scheduling embraces a wide field of applications. Mooney et al. (1996) provide an abstract characterization of this field as "a familiar and difficult scheduling problem where the 'tasks' to be scheduled are meetings. Resources include instructors, classrooms, and groups of students, or classes. In the broadest sense, course scheduling includes many related problems such as examination, conference, university, and school scheduling." Indeed, a variety of articles have been published on problems of educational scheduling within the last 15 years. Virtually all of these gravitate around university and school scheduling. In the following, we review the relevant literature. For a more detailed introduction, we refer the reader to Schmidt, Ströhlein (1980), de Werra (1985), and Schaerf (1999). In order to clarify the similarities and

dissimilarities between the problem fields, we will rely on using some pivotal terms, viz. *course*, *teacher*, *period*, *subject*, and *topic*, universally, and we will add specific interpretations in the context of each problem field. Let us add that we consider the classification of educational subject-matter into topics to be a more fine-grained division than the one into subjects; in other words, we regard a subject to comprise one or more topics.

2.1. University Scheduling

Objects

University scheduling or academic course scheduling (Tripathy 1984; Aubin, Ferland 1989; Kang, White 1992; Sampson et al. 1995) addresses the planning of lectures and seminars at universities. In practical terms, this amounts to constructing an assignment of teachers and courses to rooms and periods.

Interpretations

Note that, in this context, the term *teachers* refers to professors and lecturers, while the duration of a *period* is defined by the length of one lesson. The planning horizon comprises one week, and the process of planning repeats every semester. A *course* means a group of students combined by the desire (or necessity) to be instructed on a specific topic; *subjects* are not considered explicitly since the division of educational matter into subjects alone would be too broad for planning purposes. Many *topics*, esp. preparatory or introductory ones, are covered in several lectures or classes and may even be given by several teachers in parallel. Nevertheless, most authors assume each lecture to be characterized by a unique topic and a unique circle of participants, an assumption that seems justifiable if and only if lectures of professors are considered. In that case, courses and topics may be identified with each other, so only one of the planning objects courses and topics must be taken into account explicitly (we will see below that this is not always the case); our choice rests with the courses.

Assigning students to courses

While the planning horizon is only one week, attendance at academic courses is for several months, viz. the duration of a semester or term. Hence, the assignment of students to courses is not made part of university scheduling problems as considered here. Rather, this organizational task is usually delegated to the individuals concerned which have to select a set of courses from the lecture timetable instead. Any conflicts possibly arising from having chosen some courses scheduled to take place simultaneously must be resolved by appropriately restricting the selection made.

Constraints

A summary of the constraints considered in the literature is given in Table 1. A straightforward requirement is to prevent double assignments of courses or rooms, i.e. none of these may be assigned more than once per period. Of course, in principle the ban on double assignments applies to teachers, too. However, all authors address only the planning of lectures given by professors. Therefore, for each course there is exactly one corresponding teacher, so for teachers this requirement need not be stated explicitly. All courses must be scheduled; the availability of rooms to be used has to be verified. In some papers, researchers also allow teachers to be unavailable in some periods, e.g. due to vacations or other obligations (Aubin, Ferland 1989; Kang, White 1992). Also, course capacities may be capped by placing limits on the number of students admissible to a course (Sampson et al. 1995).

Objectives

In addition to the feasibility of the solutions found, which is the paramount goal, only one objective is pursued. It addresses preferences established by the teachers for periods as well as rooms, the total sum of all preferences is to be maximized. Objectives and constraints apply to the objects bracketed in (); items bracketed in [] are not found in all applications cited.

Objects	Objectives	Constraints
Teachers (T)	Feasibility	Complete scheduling (C)
Courses (C)	Preference maximization	No double assignments (C, R)
Rooms (R)		Availabilities (R)
Periods (P)		[Availabilities (T)]
		[Capacities (C)]

Table 1: Characteristics of University Scheduling

2.2. School Scheduling

Objects

School scheduling or timetabling (de Gans 1981; Abramson 1991; Cangalovic, Schreuder 1991, 1992; Hertz 1992; Costa 1994; Alvarez-Valdes et al. 1996) addresses the planning of school lessons which involves the assignment of teachers, subjects, and courses to rooms and periods.

Interpretations

The meaning of *teacher* is pretty straightforward in this context. A *course* is understood to denote a class of students. A *period* represents a lesson. The planning horizon comprises a week. The results of the planning process are valid for the complete school term. Most *subjects*, such as mathematics or history, are taught in several courses and can be administered by several alternative teachers. Hence, in school scheduling courses, subjects, and teachers cannot be identified with each other, so all these planning objects must be dealt with explicitly. On the other hand, there is no need to consider *topics* in this context: the selection of specific topics within the subjects assigned is laid down in a syllabus by the respective school authority and thus is not up to decision.

Assigning students to courses

Again, throughout the school year classes usually remain constant in their composition, so course composition forms no part of the planning process. In contrast to university scheduling, where students follow their individual leanings and interests in selecting lectures, here the assignment of students to courses is performed by some central authority in each school, often some teacher commissioned with this task by the principal. For purposes of planning, partitioning the set of students into subsets has the benefit of substantially reducing the complexity of the problem as now the task of constructing non-conflicting assignments of students to subjects can be handled more easily on the aggregate level of whole courses.

Constraints

In schools, most courses could be handled by several teachers, so the situation here differs from that in university scheduling where teachers (professors) and courses (lectures) can be identified with each other. Therefore, in addition to the constraints that apply there, the no-double-assignments-constraint is extended to apply to teachers as well. A new requirement is best characterized as schedule compactness, i.e. the students' timetables, which are induced by a schedule, may contain no idle periods, with possible exceptions at the beginning or the end of a day. Again, some approaches allow teachers' absence in certain periods to be taken into account; the same may be possible for courses (Alvarez-Valdés et al. 1996). Some authors also take into account assignments prescribed in advance (Costa 1994) or precedence constraints defining a particular sequence between certain topics (Hertz 1992). Finally, some models seek an even distribution of assigned periods for teachers and students (de Gans 1981, Hertz 1992, Alvarez-Valdes et al. 1996).

Objectives

All these problems belong to the class of feasibility problems, so apart from finding a feasible course schedule no further objectives are pursued. A summary is given in Table 2.

Objects	Objective	Constraints
Teachers (T)	Feasibility	Complete scheduling (C)
Classes (C)		No double assignments (T, C, R)
Subjects (S)		Availabilities (R)
Rooms (R)		Compactness (P)
Periods (P)		[Availabilities (T, C)]
		[Even distribution (T, C)]

Table 2: Characteristics of School Scheduling

2.3. Conference Scheduling

Objects

The only publication in this field we are aware of is by Eglese, Rand (1987) who address the scheduling of conference seminars. While the parallels between this and the problem fields visited above are not obvious at first glance, they can be readily seen once appropriate interpretations in terms of educational scheduling are given for each planning object. For each participant (student) a list of seminars (lectures) he wishes to attend is known, also the speaker (teacher) of each seminar is known. In order to meet overwhelming demand, seminars can be held repeatedly within the same conference. The problem consists of finding an assignment of participants to teachers, rooms, and periods. In what follows, each seminar is characterized as a combination of teacher and period.

Interpretations

Because of the simple structure of the problem, most planning objects that have been of importance so far need not be considered explicitly here. The equivalent of the students are the attendees of the conference. A *course* would translate into the set of all attendees of a particular seminar, but courses are not regarded since the attendees are scheduled individually. Each seminar is devoted to a particular topic, so once the seminars are considered, *neither subjects nor topics* must be taken into account. Each seminar has a uniform duration of one *period* and thus is scheduled as a whole. As each seminar can be conducted by only one particular speaker (*teacher*) rather than a group of alternative speakers, selection of the speakers is not up for decision. The planning horizon covers the duration of the conference.

Assigning students to courses

In contrast to the more classical problem fields discussed above, planning here directly addresses the assignment of individuals rather than of groups of individual students.

Constraints

Some constraints prevent participants from being assigned more than once to the same seminar as well as rooms from being assigned to several seminars at the same time. In addition, another constraint prevents feasible schedules from scheduling several seminars to be given by the same speaker to the same period. Recall that each seminar is associated with exactly one speaker; for this reason, the same constraint also excludes any two same seminars from being scheduled to one period. As is true in university scheduling, all seminars are mutually independent with respect to their content, so precedence or unavailability constraints are no part of this problem.

Objectives

Many scheduling problems charge us with finding a complete solution where complete means that all objects to be assigned must be assigned, a prominent example is the *job completion constraint* commonly found in problems of job shop scheduling or project scheduling. However, rather than attempting to find a schedule that completely covers all attendance wishes of all participants, the authors follow a relaxed version of this requirement. Each participant is asked to state a preference value for each of the seminars, the total preference sum of all wishes complied with is to be maximized.

Objects	Objective	Constraints
Teachers (T)	Feasibility	No double assignments (T, A, R)
Students (A)		Availabilities (R)
Rooms (R)		
Periods (P)		

Table 3: Characteristics of Conference Scheduling

2.4. Professional Course Scheduling

Objects

In broad terms, this field may be characterized as dealing with the question, How to schedule courses that are conducted by professional training providers? We chose the term *professional course scheduling* as it neatly implies several things: The courses are held for monetary compensation, to be obtained from either the students themselves or from the organizations to which they are affiliated. The courses serve to confer knowledge or skills and thus qualify as training rather than education; they are focussed and goal-oriented, often preparing students for some kind of examination or certification. Thus, training contents are largely structured; often the

course components follow a standard sequence. Another interesting consequence of the students' - or their organizations' - payments is that these will often have a much larger say in which schedules will be acceptable than other students. This influence implies that usually a larger number of constraint types and objectives will have to be dealt with in the scheduling process.

Several other differences are worth pointing out: The relation between teachers and subjects is more flexible in professional course scheduling. Different topics of the same subject may and will be instructed by different teachers, albeit teachers will usually be qualified to cover only topics from one or two subjects. Also, both timetabling and academic course scheduling intend to find periodic schedules; these have to be repeated at regular intervals, often as small as one week. In contrast, professional course schedules are non-repetitive (although the scheduling process itself may be repeated, viz. if the planning interval is smaller than the course duration), usually covering a longer planning horizon such as one month or even one year. Subjects are therefore broken down into separate topics, which also allows to account for the fact that different topics will make different demands on certain resources (e.g. installations, rooms, simulators). Still completing a topic may take longer than just one period, hence topics are knocked down into component lessons which then can be assigned to periods on the time axis. Consequently, professional course scheduling intends to find assignments between teachers, courses, lessons, rooms, and periods.

Interpretations

A *course* refers to a number of students which receive the same training and are instructed as a group. Note that, different from university scheduling, most *subjects*, *topics* and lessons will be given in several courses, so identifying topics (or lessons) with students is not viable here. As *teachers* we refer to instructors or tutors charged with conducting the training. Again, the duration of a *period* is defined by the length of one lesson.

Assigning students to courses

As in school scheduling, the composition of courses as a group remains constant over time (usually throughout the whole training). Therefore, the assignment of students to courses is no part of the planning considered here; rather it is done before a new course commences, usually by the planning department of the training provider.

Constraints

Some constraints are similar to those of the classical problem fields (see the overview in Table 4): teachers must be available in order to be scheduled, rooms can only be assigned once per period. In addition to these classical constraints, several others come into play: the availabilities not only of rooms but also of teachers and courses must be taken care of.

Professional course programs usually follow a prescribed syllabus - which may even be dictated by governmental regulations, as happens to be the case in many aviation applications – so precedence relations exist between certain topics or even lessons. Some of these precedences may reflect logical or didactic considerations, others may simply be of a practical nature (theoretical before practical training; training before exams). Other possible restrictions involve time windows, defined by earliest start and latest finish time, within which certain courses must be scheduled. Additional restrictions may apply.

Objectives

For many traditional course scheduling problems even the associated feasibility problem is (strongly) NP-complete, owing to the combination of scarce resources and a fixed planning horizon within which all courses must be scheduled. In other words, even finding out whether a feasible solution exists at all (to say nothing of actually finding such a solution) is as hard as the most difficult planning problems known. Yet for professional course scheduling problems, there is a trivial feasible solution for each instance, which amounts to scheduling no courses at all; this is possible because requests for courses need not be honored.

Therefore, the main thrust of research is not merely towards finding feasible solutions, rather researchers strive for schedules which optimize a variety of other objectives along the way. Among these are maximization of profit, service level, or course priorities, an even distribution of teacher workload, or minimization of course duration. Indeed, in a recent project on a professional course scheduling problem in the aviation industry, ten objectives were found to be of relevance for scheduling purposes (Haase et al. 1999).

Objects	Objectives	Constraints
Teachers (T)	Profit	No double assignments (T, C, R)
Courses (C)	Service level	Availabilities (T, C, R)
Lessons (L)	Standard sequence (L)	Precedence relations (L)
Rooms (R)	Even distribution (T)	Temporal restrictions (C)
Periods (P)	Redundancy (T)	[Capacities (C)]
	Compactness (L)	
	Teacher continuity (T)	

Table 4: Characteristics of Professional Course Scheduling

3. Problem Setting

In the sequel, we describe the problem setting faced by LFT. In doing so, we separate the problem into two subproblems, viz. one of course scheduling proper and one of room assignment. This is motivated by the fact that room demand is essentially independent of the particular schedule implemented since the number of courses and the number of students to be scheduled are exogenous parameters not subject to decision. The only room-related aspect impinging upon the course scheduling process is that certain rooms with specific equipment are scarce such that any feasible schedule must respect their available number. Yet also for these, the actual assignment of rooms may take place after the scheduling.

3.1. Course Scheduling Subproblem

The course scheduling subproblem can be characterized as follows:

- *Planning interval, slots, lessons, and periods:* Scheduling is done on a repetitive basis, each schedule covers an interval of four weeks. Each working day within this interval is divided into several slots, each of which comprises two lessons of 45 minutes. An exception holds for military aircrew training where the last slot of each day comprises three lessons. The number of periods in a planning interval is the product of working days; slots per day, and the number of lessons per slot. Planning is done on the level of slots; however, contractual regulations on instructor availability are formulated on the level of periods, so the temporal granularity is defined by the length of one lesson, i.e. a period.
- *Courses:* Within a planning interval, usually several courses are in a ground training phase and thus need to be considered. Recall that each course is designed individually, so the concept of different course types, which groups courses of identical structure and thus facilitates some aspects of scheduling in other training firms (Haase et al. 1999), does not apply. Albeit their design will often resemble that of others, all courses are considered unique.
- *Subjects and topics:* Instruction is logically divided into subjects (e.g. navigation, meteorology), each of which is divided further into a number of constituent topics. For scheduling purposes, each topic is characterized by specific requirements i.t.o. number and qualification of instructors as well type (and sometimes number) of rooms. For each course, there are precedences between specific topics, reflecting didactic insight or legal requirements (e.g. theoretical instruction *before* practical training *before* examination). To facilitate matters for the scheduling, there is a *standard sequence* of certain topics that meets all precedence requirements and has been found to perform well under practical considerations.

- *Instructors:* Training is carried out exclusively by LFT staff instructors; notable exceptions are some seminars conducted by external psychologists and, of course, official exams administered by federal examiners. Although most instructors are associated with one subject only, there are cases where an instructor either cannot cover certain topics of his subject, or where he may also instruct certain topics of another subject. For planning purposes, we therefore consider each instructor to be associated with a specific set of topics.
- *Course availability:* Courses may be blocked in certain periods, when they are already entered for exams or certain seminars, whose dates cannot be freely set by LFT. Except for these cases, instruction always commences in the first slot of each day and proceeds throughout the subsequent slots, implying that there are no idle slots.
- *Room availability:* Rooms are grouped into types according to characteristics such as seating capacity and specific equipment (required e.g. for radiotelephony or computer-based training). The number of rooms available per type varies over time since some rooms are temporarily used for other purposes.
- *Instructor availability:* No instructor can be assigned to more than one course per lesson. In addition, several contractual restrictions apply: Instructors are available for a number of periods per planning interval. Yet, they may be blocked in specific periods (e.g. for approved vacation). Instructors may teach no more than a certain maximum number of lessons per day. Further, the number of days on which they actually give the maximum number of lessons is limited per week as well as for the whole planning interval. Finally, instructors may ask for leave in particular periods (e.g. for private reasons); such wishes are prioritized according to their importance as perceived by LFT.

Note that these requirements charge us with tracking all lessons given per instructor, as to verify that plans are in accordance with contractual regulations.

- *Modes:* For each combination of course and topic, a number of modes are defined reflecting possible instructor and room assignments. Note that once an instructor has been assigned to a specific topic within in a course, this assignment remains fixed throughout the duration of the course.
- *Prescribed assignments:* Some combinations of course, topic, instructor, and room may need to be assigned to some period before the actual scheduling process begins.

The merits of a schedule are judged against several objectives. In total, LFT management identified and prioritized seven objectives, both objectives and priorities represent partly didactic experiences, partly management or instructor interests. We should emphasize that LFT management felt all these objectives should be included as shown below, even if some of them are not mutually exclusive (we list such objectives under the same name). Note that each of

these objectives is actually a soft constraint (Dige et al. 1993) such that the objective function - which is to be minimized - consists of the weighted sum of soft constraints' violations.

- *Standard sequence:* Schedules should respect the standard sequence of topics (priority 10).
- *Subject distribution:* For each course, instruction on the five subjects should be spread evenly over the planning interval; this objective is operationalized by requiring that no subject should be addressed more than once per day (priority 2). In particular, consecutive slots should not be devoted to the same topic (priority 4).
- *Instructor continuity:* In each course, as few instructors as possible should be involved to avoid changes unduly for instructors and students alike (priority 5).
- *Workload distribution:* The total workload should be evenly distributed among the staff, taking into account periods where some are unavailable (priority 2). Also, for each instructor the individual workload should be evenly spread over the planning interval, by keeping the number of days as small as possible on which he instructs for the maximum number of daily periods (priority 5).
- *Instructor requests:* The requests of instructors for exemption from being assigned to certain periods should be met, taking into account respective priorities (priority 5).

3.2. Room Assignment Subproblem

The room assignment subproblem can be characterized as follows:

- *Rooms:* Several rooms exist in which instruction or training may take place.
- *Room types:* Rooms are grouped into types according to their characteristics, as outlined above.
- *Room assignments:* In each period, a course requires one or several rooms of a particular type. Each room to be assigned must be available in that period and may not have been assigned to another course already.

The room assignment pursues just one objective, measured by a function counting the number of - undesired - room changes.

- *Room changes:* Over all courses, the need to change rooms for different periods should be as small as possible.

The final schedule represents an assignment of topics, rooms, and instructors to combinations of courses and periods.

3.3. *Modelling*

As a foundation for algorithmic development, the objectives and constraints of the problem were translated into a mixed integer model. For the sake of brevity, however, we refrain from presenting the model which comprises a total of 17 constraints and an objective function composed of seven terms. Rather, we confine ourselves to demonstrating that many of the constraints can be formulated using appropriate resource concepts.

Whenever lessons are scheduled, several objects are involved: a course, a topic, one or more teachers, and one room (only rarely will several rooms be required). Constraints verify the availability of teachers, courses, and rooms; they also ensure that the number of lessons to be scheduled for each topic, as laid down in the syllabus, is taken into account. All these constraints can be interpreted either as capacity limitations or as logical relations, all of which can be formalized in terms of appropriate resource concepts (Schirmer, Drexel 2001).

Resource concepts are categorized according to the way in which their availability is restricted: The available amount of nonrenewable resources is limited over the complete planning horizon, by a total capacity. The available amount of renewable resources is limited in every period, by a period capacity (Slowinski 1981). The recent concept of partially renewable resources, however, allows to formulate limited capacities over arbitrary sets of periods within the planning horizon (Böttcher et al. 1999; Schirmer, Drexel 2001).

- The availability of teachers, courses, and rooms is expressed in terms of renewable resources. Since rooms of the same type are *mutually exchangeable*, we define one such resource per room type. Recall that rooms may be blocked in some periods, hence we allow time-variant capacity profiles by setting the capacity per period for each type to the number of rooms available in that period. Of course, teachers and courses are not mutually exchangeable, so we define a separate resource for each of these, with a capacity of one if it is available, and a capacity of zero if it is blocked (recall that teachers may be blocked for vacations, courses for external seminars, the dates of which are beyond the authority of LFT). Also, if the end of a ground training phase of a course falls within the planning horizon, that course is blocked on all unused buffer days, such that the ensuing idle days arise at the end of the phase, thus resulting in a more compact schedule. Finally, for each course the slots are fixed in which the training is to take place (usually in the first three or four slots of each day), in all other slots the course is blocked as well.
- Another restriction, namely that in no course more than the prescribed number of lessons may be accorded per topic, is easily handled if we introduce one nonrenewable resource for each combination of course and topic; its capacity is determined by the number of lessons prescribed. Also, the maximum number of lessons that each teacher may give within the planning horizon can be

couched in terms of one nonrenewable resource per teacher, with a capacity determined by the maximum number of lessons.

- Additional, more complex contractual regulations are formalized in terms of partially renewable resources. The reason is that some of these limitations pertain to certain sets of slots, others to certain sets of working days; no such restrictions can be expressed in terms of the classical resource concepts (Schirmer, Drex1 2001). One of these is the requirement that each teacher may only give a certain number of lessons per working day. This can be expressed by defining one partially renewable resource per teacher, where for each working day the capacity over the set of all blocks of the day equals the number of daily lessons allowed. Other regulations limit the number of days on which a teacher may give the above number of lessons, as well over the whole planning horizon as per week. To meet this requirement we define sets of periods comprising all the blocks of the planning horizon as well as sets covering the blocks of each single week, their capacities being set to the respective values set down in the contractual regulations.

4. Algorithmic Approaches

Formal analysis of the problem's algorithmic complexity shows it to contain several strongly NP-equivalent problems as special cases, so developing exact optimization algorithms is hardly a promising avenue (Schirmer 1999, S. 27-31, 159-162). Neither does applying solution methods from other problems of educational scheduling constitute a viable approach since these algorithms would fail to accommodate the more complex objectives and constraints of the problem considered here. We therefore developed tailored heuristic construction as well as improvement algorithms for both subproblems, which we discuss in the following. Note that the algorithmic ideas employed are general enough to be easily modified to adapt them to other problem settings in professional course scheduling.

Following the characterization of Mooney et al. (1996), a course scheduling problem may contain time-assignment type subproblems (examination scheduling, timetabling) as well as resource-assignment type subproblems (instructor, classroom, or student group assignment scheduling). The authors report that due to the algorithmic intractability of such problems large-scale instances are usually solved by iterating between the timetabling and single-resource assignment subproblems. Although the problem at hand clearly belongs to this category, we will demonstrate in the sequel that both time- and resource assignments can be performed simultaneously by appropriately adapting well-known concepts from project scheduling, allowing to dispense with the need to alternate between different phases. Again, for ease of presentation we distinguish between the course scheduling problem proper and the room assignment problem associated with it. Both problems are only loosely connected, so solving them simultaneously would hardly pose higher obstacles for algorithmic tractability but would

make the presentation less obvious. The overall procedure is straightforward: First, we construct a solution for the course scheduling subproblem which we then improve. Second, we construct a feasible room assignment for the resulting course schedule which we then improve as well.

Let us point out one particularly noteworthy aspect of our solution process: While the search proceeds, we do not store only the best solution found so far but a specified number of best schedules found, the intention being to provide LFT management with several good alternatives to choose from. In order to facilitate the tedious and error-prone task of analyzing several different schedules, for each schedule some statistics summarize the most important information for management purposes, such as utilization of instructor capacities, distribution of instructor workload, and number of leave requests respected. In this way, the final selection of a schedule can be based largely on analyzing and comparing these statistics. Although to the best of our knowledge this approach is rarely used in conjunction with decision support systems, it does offer significant benefits concerning acceptance and actual use of such systems which stem from several motivational and behavioral effects (Schirmer 2000a). It also draws on two insights into the representation of complex real-world problems in terms of formal models; in our context, one of them pertains to the formulation of constraints, the other to that of objective functions consisting of multiple objectives. The first is a fundamental insight that originated from artificial intelligence research, in particular from the design of expert systems: For all but the most simple settings, it is hardly possible to extract all potentially relevant requirements of a problem and integrate them into a formal model (Kim, Courtney 1988; Boose 1989), in other words: some constraints may be so rarely of relevance in the real world that they are overlooked when building a formal model. The second insight concerns the necessity, common to all decision problems under multiple objectives, to determine the individual preference order of the objectives as seen by the deciders. While methods for this task exist, they are rather complex, so managers' willingness to undergo such procedures can reasonably be expected to be less than enthusiastic. Therefore, a formal objective function may or may not be an exact representation of management preferences. Both phenomena bear the question whether a formally optimal solution is also the best solution in practical terms. It is thus desirable to proffer several good alternative solutions; doing so offers the freedom to sometimes select a schedule that does not coincide with the formal optimum but that also meets objectives or constraints not properly included in the formal model.

4.1. Course Scheduling Methods

4.1.1. Construction Method

We developed a priority rule-based construction algorithm with a backtracking capability based upon the serial scheduling scheme (cf. e.g. Kolisch 1996). Our use of a serial

construction method was motivated by several reasons. First of all, on several capacitated planning problems such algorithms have been found to belong to the most effective algorithms currently known (Kolisch, Hartmann 1999; Schirmer 2000b). Also, augmenting partial schedules in a serial manner is quite similar to the manual procedure employed by LFT, thus facilitating understanding and acceptance of our approach. Also, a serial algorithm easily allows to concentrate early on potential bottlenecks: E.g., if only few feasible periods remain for a topic, then that topic can be scheduled immediately; under a parallel regime, other topics might get scheduled earlier, effectively blocking required instructors.

As is common, partial schedules, starting from the empty one, are augmented in a stage-wise fashion until all topics of all courses have been considered. Since the planning interval covers only a small portion of the duration of each course, even the final schedule for a planning interval will be a partial one (except in December, see below); topics left unscheduled will then be reconsidered in the following planning interval. We divide the set of all topics into four disjoint subsets or states, viz. *scheduled*, *resource-infeasible*, *eligible*, and *remaining*. The first set comprises all topics which have already been scheduled. Resource-infeasible topics cannot be scheduled within the current planning interval due to resource restrictions, and must be relegated to a later interval. The decision set essentially comprises all precedence-feasible topics which are not resource-infeasible. Now, in each stage priority rules are used to select one of the eligible topics from the decision set, a mode for that topic, and finally a feasible period for both. The priority rules are employed deterministically, only ties are broken randomly.

In more detail, the algorithm begins by scheduling all prescribed assignments. As these cannot be changed by later algorithmic steps, they are straightforward to 'get them out of the way' soon. Note, however, that due to precedence relations some prescribed assignments may be infeasible until all their predecessors have been scheduled. Such assignments will then be scheduled as soon as possible by restricting the decision set to predecessors of such assignments. If no such topics exist, another restriction applies which admits only topics to the decision set where some lessons were already scheduled, with the intention to conclude topics before commencing new ones. If no such topics exist either, the decision set comprises all feasible course-topic combinations. In order to restrict its cardinality, however, it is built in a two-step process, allowing to select a course in the first step, and then one of the topics of only this course in the second. Selecting a mode is then straightforward: if the topic has already been taught in that course, the same mode is assigned to avoid instructor changes, otherwise a mode is chosen by a priority rule. Finally, we choose a slot in which the topic can be feasibly scheduled, given the mode assignment selected before. If such a slot is found, the topic is scheduled in that mode and slot and the resources are updated. Otherwise we try a limited backtracking approach.

The idea central to this approach is to free a feasible slot by moving an already scheduled topic to a free slot, and then to schedule the topic at hand to the newly freed slot. To do so, we first of all check to which slots the topic at hand could be feasibly scheduled if the topics currently scheduled would be descheduled. For all these topics we then check in which free slots these could be scheduled as well. If several such operations exist, we select the one resulting in the best objective function value, ties to be broken arbitrarily. If no such operations exists, the topic is deleted from the decision set and marked as resource-infeasible.

To illustrate this, consider the (partial) schedule depicted in Figure 1, where course names are listed along with the initials of the corresponding instructor; shaded fields represent blocked courses. Let us also assume that instructor AB is unavailable in slot 1 of day 1. Now, topic MET 2.0 is to be scheduled for Course 1. As this topic has already been taught by AB, he should continue to do so, but this is resource-infeasible. Now we try to backtrack earlier assignments. We begin by finding those slots to which the topic could be assigned, disregarding all other assignments which are (Day 1, Slot 2), (Day 2, Slot 1), and (Day 2, Slot 2). Of these, the topic from (Day 2, Slot 1) cannot be moved to (Day 1, Slot 1) as the corresponding instructor CD is already assigned to Course 2. Moving the topic from (Day 2, Slot 2) to the free slot and assigning MET 2.0 to the so-freed slot would result in two cases where the same topic is taught in two slots of the same day, which is regarded as undesirable i.t.o. the objective function. The other possibility, i.e. moving the topic assigned to (Day 1, Slot 2) to (Day 1, Slot 1) and using the freed slot, would not change the objective function value, so this operation would be implemented, as shown in Figure 2.

	Day 1				Day 2			
	1	2	3	4	1	2	3	4
Course 1		ERI 3.3 XY	ATC 3.1 EF		NAV 4.1 CD	ERI 3.3 XY	MET 2.0 AB	
Course 2	NAV 1.0 CD	NAV 1.0 CD	MET 1.0 AB		ERI 1.0 XY	ATC 2.1 JK	TEC 2.0 GH	

Figure 1: Backtracking - Initial Schedule

	Day 1				Day 2			
	1	2	3	4	1	2	3	4
Course 1	ERI 3.3 XY	MET 2.0 AB	ATC 3.1 EF		NAV 4.1 CD	ERI 3.3 XY	MET 2.0 AB	
Course 2	NAV 1.0 CD	NAV 1.0 CD	MET 1.0 AB		ERI 1.0 XY	ATC 2.1 JK	TEC 2.0 GH	

Figure 2: Backtracking - Final Schedule

4.1.2. Neighborhood Search Method

To improve the so-constructed initial course schedules, we employ a tabu search algorithm. Its constituent components can be described as follows.

- *Evaluation function:* We evaluate the schedules by their objective function value, to which a penalty term is added that multiplies the number of resource-infeasible topics by a sufficiently large weight; consequentially scheduling an additional topic is preferred over any other improvement of the objective function value.
- *Neighborhood:* The neighborhood of a schedule comprises all schedules which can be constructed from it by one of several swap and shift moves. Searching the complete neighborhood, though, would entail substantial computational effort as the relatively complex objective function must be computed for each candidate move in order to evaluate the corresponding schedule. For the sake of efficiency, we therefore use a *neighborhood decomposition strategy* as proposed by Glover et al. (1993) by maintaining one partial neighborhood for each course (cf. below).

It is well to point out that the space searched usually does not coincide with the space of all feasible schedules, for two reasons. On one hand, for efficiency we reduce neighborhood cardinality by considering only topics already scheduled in the initial schedule, plus all standard sequence-feasible ones; the latter criterion is motivated by the observation that, if one considers the priority assigned to standard sequence adherence within the objective function, one could hardly expect to improve upon the best known schedule by scheduling topics that would violate the standard sequence. On the other hand, the construction method may fail to generate a feasible initial schedule; in that case at least one course would comprise a topic where less than the required number of lessons are scheduled. To accommodate these cases, our definition of neighborhood includes all schedules where for each topic the number of scheduled lessons is at least as high as in the initial schedule but no higher than the required number of lessons (rather than only those schedules where the number is equal to the required number of lessons).

- *Tabu status:* The tabu status is accorded to all slots involved in a move; for each course, we maintain a separate tabu list of fixed length.
- *Aspiration criterion:* We use the standard definition of an aspiration criterion, disregarding the tabu status of a schedule if it has a better evaluation than the best one found so far. Of all moves which either are non-tabu or - if tabu - meet the aspiration criterion, the one with the lowest evaluation is selected; ties are broken randomly.
- *Iteration:* Rather than searching the complete neighborhood and selecting the best schedule from it, our following a neighborhood decomposition strategy implies that in each iteration

we search all partial neighborhoods (recall that each of these is associated with exactly one course, and vice versa). Within each partial neighborhood, the best schedule is selected; ties are broken arbitrarily. A complete iteration may thus be seen as a compound move which applies one move to each course. A *neighborhood sequence* then prescribes in which order these neighborhoods are searched. The most straightforward sequence, which is the one we use, is to examine the neighborhoods in ascending index order of their defining courses.

- *Termination criterion:* The search terminates after a specified number of iterations.

4.2. Room Assignment Methods

Given a feasible course schedule, the room assignment can be done separately for each slot, since for algorithmic purposes rooms are regarded as renewable resources.

4.2.1. Construction Method

The construction method applied is straightforward: to each course it assigns the required number of available and suited rooms in each slot. If for some course at hand this is impossible, we recursively try to find another feasible room for an already scheduled course and assign the room thus freed to the course at hand. This process is repeated until either a feasible solution is found or all possibilities to assign alternative rooms have been exhausted.

4.2.2. Neighborhood Search Method

We then try to improve the room assignments by means of neighborhood search, again considering the slots separately (neighborhood decomposition). Given a feasible room assignment for one slot, its neighborhood consists of all assignments which can be obtained by a cyclic shifting of rooms among several courses (the most simple of the compound moves induced would be an exchange of rooms between two courses). Using a weighted directed graph with nodes representing the rooms and an edge from a room r to a room r' if either the course assigned to r could be shifted to r' or r is "unoccupied", such cyclic shiftings can be found by a simple graph-theoretic algorithm detecting the presence of negative-length cycles.

A cycle in that graph can be interpreted as a feasible move where iteratively either a course or the status "unoccupied" is shifted to the next room on the cycle. If we define the weight of an edge as the change in the objective function value incurred by that shift, then each negative-length cycle represents a move resulting in a solution with a lower, i.e. better objective function value. We follow a first-improvement regime by implementing the first move improving the room assignment objective function. The search terminates if no further improvements are possible. By construction of the graph a negative-length cycle exists if and only if the

corresponding room assignment can be improved, hence our search algorithm always terminates with an optimal assignment.

5. Computational Evaluation

5.1. Test Instances

As is common with large-scale practice projects, the acquisition and aggregation of operational planning data proved to be a massively time-consuming and costly task. A considerable portion of the data was never gathered before. The majority of data that was available already was so in written form only and had to be brought into a computer-readable form.

A comprehensive characterization of a typical operational instance would be beyond the scope of this contribution, so we restrict our presentation to some informative figures listed in Table 5.

Periods	Courses	Instructors	Subjects	Lessons To be Scheduled	Blocked Days Per Instructor
216	15	30	5	1686	4.6

Table 5: Typical LFT Operational Instance - General Characteristics

The number of periods T reflects 27 working days, each comprising four slots of two lessons; note that the December planning interval always comprises all remaining working days of the year, and thus is longer than the usual 20 working days.

An experimental assessment of algorithms should be based on a sample of representative test instances. However, providing a comprehensive sample of such instances turned out to be a rather challenging task: Efforts to reproduce problem instances of past years and compare our algorithmic approach to the manual scheduling process of LFT failed due to the prohibitive cost involved to collect the planning data as well as to reconstruct and evaluate the corresponding manually constructed schedules. Using a general-purpose instance generator such as ProGen (Kolisch et al. 1995; Schirmer 1999) was prevented by the special structure of the problem. Due to the complexity of the problem, the development and implementation of a dedicated instance generator was considered beyond the scope of this project. Hence, we choose to generate additional test instances by carefully modifying characteristics of operational instances, as proposed by Haase et al. (1998).

5.2. Computational Results

All algorithms were implemented in C. The course schedule construction method was run twice, the better schedule was fed to the tabu search method. The tabu search algorithm was then run for 20 iterations. Control parameter settings for both course scheduling algorithms were determined in separate experiments. Subsequently, the construction and the improvement method were applied to obtain the room assignments. Note that these are deterministic algorithms without any control parameters, so no previous experiments were required to fine-tune these.

The benefits to be realized from our approach are best demonstrated by the following characteristics of the solutions found (cp. the objectives in Section 1):

- For all operational instances, all constraints were met.
- The total computation time required was about 150 seconds, measured on a Pentium 133 personal computer with 16 MB RAM under Windows 95, which is sufficiently short for all practical purposes.
- The objective function values of the best algorithmic schedules are markedly better than those of the corresponding manual schedules. We will demonstrate this below on the planning results for the planning interval of December 1997.
- For each operational instance, five alternative course schedules were determined.

In the remainder of this section, we compare the best algorithmic schedule for December 1997 to the corresponding manual schedule. For both schedules, Table 6 summarizes the violations of the soft constraints which define the objective function. Of these, we show those violations separately which are directly caused by prescribed assignments. We might emphasize that the ability to measure and quantify these violations constitutes another benefit of computer-based decision support systems as described here, since the violations induced by irreconcilable prescriptions would have gone undetected in the past.

The violations of both the manual and the algorithmic schedules must be discounted by these numbers to obtain those violations induced by decisions up to the respective scheduling procedure. Therefore, in the column marked "Relative 1" we exhibit the relative improvement of the algorithmic schedule against the manual schedule, so $D = (A-B)/A$. In column "Relative 2" essentially the same relative improvement is shown although it is discounted for the prescribed violations, i.e. $E = (A-B)/(A-C)$.

Objective	Priority	Violations			Improvement by Algorithmic Schedule	
		Manual Schedule	Algorithmic Schedule	Prescribed	Relative 1	Relative 2
		(A)	(B)	(C)	(D)	(E)
Standard sequence	10	unknown	102	71	N/A	N/A
Instructor requests	5	2	0	0	100%	100%
Workload distribution 1	2	358	192	0	46%	46%
Workload distribution 2	5	45	15	9	67%	83%
Instructor continuity	5	17	16	15	6%	N/A
Subject distribution 1	2	154	113	38	27%	35%
Subject distribution 2	4	72	40	38	44%	94%

Table 6: Best Algorithmic vs. Manually Constructed Schedule

Some remarks seem in order: No information is available on the number of standard sequence violations incurred by the manual schedule, due to the excessive effort it would have taken to extract this information. The rather high number of 102 violations for the algorithmic schedule, however, fares not as bad as it may seem: further analysis revealed 71 of them to be induced by standard sequence violations committed in the preceeding planning interval. To understand this, recall that for didactic reasons the algorithm attempts to continue topics that were already begun in the previous schedule. As that schedule had still been manually constructed by LFT, the algorithm was bound to continue violations made in the past.

For all other objectives, the algorithmic schedule yields substantially better results. Excluding the standard sequence violations, the corresponding objective function value of the manual schedule totals 1632 while that of the algorithmic schedule equals 925, for a reduction of 43%.

6. Summary and Outlook

In this contribution, we covered the problem of scheduling the complete ground training of LFT. This problem shares a common core of planning objects, objectives, and constraints with other scheduling problems (Haase et al. 1998, 1999) that are only loosely related to classical problems of educational scheduling. This insight led us to consider these problems as belonging to a separate problem field, which we refer to as professional course scheduling.

A comparative analysis of the relevant literature on educational scheduling revealed this field to be substantially more complex than the traditional problem fields. This fact also explains why the timetabling software considered by LFT was found to be inappropriate for the course scheduling problem at hand: Despite apparent similarities, both problem fields differ too much in terms of their structure. In order to be feasible, i.e. usable in the real world, course schedules must meet several complex constraints that are part of no school scheduling problem.

Therefore, course schedules constructed by some school scheduling software will in most cases violate one or several constraints that are essential for their practical use. Also, school scheduling problems make no provisions for objectives other than feasibility of the schedules derived.

Consequently, the task is set to devise algorithms appropriate for this field. In this contribution, we have used priority rule-based construction methods and tabu search to build and improve course schedules, as well as a straightforward assignment method and a graph-theoretic algorithm to find and improve corresponding room assignments. Together, these algorithms produced schedules that rank considerably better than the manually assembled ones w.r.t. all measured objectives. Also, they took only minutes to arrive at the final five alternative schedules (of which the one shown in Table 6 is the best one), compared to several man-days (full-time equivalents) required for the manual procedure. In addition, the algorithmic ideas employed are general enough to be easily adapted to other problems in the field of professional course scheduling.

In total, our approach has demonstrated that a computer-based decision support system for generating course schedules at LFT is both feasible and desirable. Benefits to be realized on a short-term scale include less time and resources to be invested in the planning process, better instructor utilization, a more balanced workload, and more clarity of planning results for planners, staff, and customers; on the long run, our approach can contribute towards more flexibility to react to last-minute changes and the ability to conduct what-if analyses of the effect of different planning parameters, e.g. modifications of labor restrictions and work rules.

The development of a full-fledged decision support system that folds in these and other features is currently in progress at LFT.

Acknowledgements

As anyone who has worked in industrial applications would know, many people were involved with this project. From LFT, we wish to mention Robert Nyenhuis, former chief ground instructor, who managed to wing this project; Detlev Carius, incumbent chief ground instructor, who kept it flying; and Reiner Brauwer whose attitude was instrumental. We also recognize the support of Lothar Martin, managing director. Finally we thank Peter Kandzia for his support and Jürgen Hauschildt for his advice on some details of this project.

Note

The project described in this article took place while the first author held a teaching position at the Institut für Betriebswirtschaftslehre and the Institut für Informatik und Praktische

Mathematik, both Christian-Albrechts-Universität zu Kiel. Thus, the opinions expressed herein are those of the authors alone, and do not necessarily reflect those of the Lufthansa Group or any of its companies.

References

- ABRAMSON, D. (1991), "Constructing school timetables using simulated annealing: sequential and parallel algorithms", *Management Science* 37, pp. 98-113.
- ALVAREZ-VALDES, R., G. MARTIN, AND J.M. TAMARIT (1996), "Constructing good solutions for a school timetabling problem", *Journal of the Operational Research Society* 47, pp. 1203-1215.
- AUBIN, J. AND J.A. FERLAND (1989), "A large scale timetabling problem", *Computers and Operations Research* 16, pp. 67-77.
- CANGALOVIC, M. AND J.A.M. SCHREUDER (1991), "Exact colouring algorithm for weighted graphs applied to timetabling problems of different length", *European Journal of Operational Research* 51, pp. 248-258.
- CANGALOVIC, M. AND J.A.M. SCHREUDER (1992), "Modelling and solving an acyclic multi-period timetabling problem", *Discrete Applied Mathematics* 35, pp. 177-195.
- COSTA, D. (1994), "A tabu search algorithm for computing an operational timetable", *European Journal of Operational Research* 76, pp. 98-110.
- DE GANS, O.B. (1981), "A computer timetabling system for secondary schools in the Netherlands", *European Journal of Operational Research* 7, pp. 175-182.
- DE WERRA, D. (1985), "An introduction to timetabling", *European Journal of Operational Research* 19, pp. 151-162.
- DIGE, P., C. LUND, AND H.C. RAVN (1993): "Timetabling by simulated annealing", in: *Applied simulated annealing*, Lecture Notes in Economics and Mathematical Systems 396, Springer, Berlin, pp. 151-174.
- EGLESE, R.W. AND G.K. RAND (1987), "Conference seminar timetabling", *Journal of the Operational Research Society* 38, pp. 591-598.
- GLOVER, F., E. TAILLARD, AND D. DE WERRA (1993), "A user's guide to tabu search", *Annals of Operations Research* 41, pp. 3-28.
- HAASE, K., J. LATTEIER, AND A. SCHIRMER (1998), "The course scheduling problem at Lufthansa Technical Training", *European Journal of Operational Research* 110, pp. 441-456.
- HAASE, K., J. LATTEIER, AND A. SCHIRMER (1999), "Course planning at Lufthansa Technical Training - Constructing more profitable schedules", *Interfaces* 29:5, pp. 95-109.
- HERTZ, A. (1992), "Finding a feasible course schedule using tabu search", *Discrete Applied Mathematics* 35, pp. 255-270.
- KANG, L. AND G.M. WHITE (1992), "A logic approach to the resolution of constraints in timetabling", *European Journal of Operational Research* 61, pp. 306-317.
- KOLISCH, R. (1996), "Serial and parallel resource-constrained project scheduling methods revisited: Theory and computation", *European Journal of Operational Research* 90, pp. 320-333.
- KOLISCH, R. AND S. HARTMANN (1999), "Heuristic algorithms for solving the resource-constrained project scheduling problem: Classification and computational analysis", in: *Project scheduling - Recent models, algorithms and applications*, J. Weglarz (ed.), Kluwer, Boston, pp. 147-178.
- KOLISCH, R., A. SPRECHER, AND A. DREXL (1995), "Characterization and generation of a general class of resource-constrained project scheduling problems", *Management Science* 41, pp. 1693-1703.
- MOONEY, E.L., R.L. RARDIN, AND W.J. PARMENTER (1996), "Large-scale classroom scheduling", *IIE Transactions* 28, pp. 369-378.
- SAMPSON, S.E., J.R. FREELAND, AND E.N. WEISS (1995), "Class scheduling to maximize participant satisfaction", *Interfaces* 25, pp. 30-41.

- SCHAERF, A. (1999), "A survey of automated timetabling", *Artificial Intelligence Review* 13, pp. 87-127.
- SCHIRMER, A. (1999), *Project scheduling with scarce resources - Models, methods, and applications*, Kovac, Hamburg.
- SCHIRMER, A. (2000a), "Widerstände gegen Innovationen - Erfahrungen aus der Einführung von Entscheidungsunterstützungssystemen für die Personaleinsatzplanung", *Zeitschrift Führung + Organisation* 69, pp. 340-348 (in German).
- SCHIRMER, A. (2000b), "Case-based reasoning and improved adaptive search for project scheduling", *Naval Research Logistics* 47, pp. 201-222.
- SCHIRMER, A. AND A. DREXL (2001), "Allocation of partially renewable resources - Concept, capabilities, and applications", *Networks* 37, pp. 21-34.
- SCHMIDT, G. AND T. STRÖHLEIN (1980), "Timetable construction - An annotated bibliography", *The Computer Journal* 23, pp. 307-316.
- TRIPATHY, A. (1984), "School timetabling - A case in large binary integer linear programming", *Management Science* 30, pp. 1473-1489.