Fast and exact audit scheduling optimization

Jan Motl¹ · Pavel Kordík¹

Received: 13 February 2021 / Accepted: 23 August 2021

Published online: 07 September 2021

© The Author(s) 2021, corrected publication 2021 **OPEN**

Abstract

This article is concerned with the cost and time-effective scheduling of financial auditors with integer linear programming. The schedule optimization considers 13 different constraints, staff scarcity, frequent alterations of the input data with the need to minimize the changes in the generated schedule, and scaling issues. The delivered implementation reduced the time to the first schedule from 3 man-days to 1 h and the schedule update time from 1 man-day to 4 min.

Keywords Audit-staff scheduling · Assignment problem · Integer linear programming · Multi-commodity network flow

Scheduling is an old optimization problem [5, 11, 18]. However, the difficulty with the optimization is that we have to carefully select what to optimize and which constraints to consider to get a solution for the optimization problem in an acceptable timeframe. Hence, the formulation of the schedule optimization necessarily differs from firm to firm as their needs differ. In this article, we list constraints encountered when optimizing a schedule for an auditing firm, and we illustrate how to formulate them as a multi-commodity network flow problem [9].

An unnamed auditing firm suffers from staff scarcity as the number of clients continues to increase. This has led to severe difficulties in planning activities. The last handmade schedule took 3 man-days to create, and each schedule update took another 1 man-day. It is predicted that the firm will continue to grow, and the situation will become even worse. Hence, the management decided that the firm needs a computer-aided planner, as presented in this article.

Our audit-staff scheduling problem may succinctly be characterized by the following statements:

1. The firm employs financial *auditors* for audit *engagements* (an audit engagement is an arrangement that the firm has with a client to conduct an audit of the client) during the following year.
2. Each engagement includes an availability calendar indicating when the audit can be performed.
3. Each engagement consists of *tasks*. Each task takes a specified amount of time. Each task is completed by a single auditor.
4. Each auditor has an availability calendar with the number of hours that the auditor can work on that day.
5. We want to find an assignment of the auditors to the engagements that respects the business rules further listed in Sect. 2.

A basic illustration of the desired schedule is provided in Fig. 1.

When we formulated the scheduling optimization problem as a multi-commodity network flow problem with 13 constraints, we reduced the time to the first schedule from 3 man-days with a handmade approach to 1 h and the schedule update time from 1 man-day to 4 min.

✉ Jan Motl, jan.motl@fit.cvut.cz; Pavel Kordík, pavel.kordik@fit.cvut.cz | ¹Czech Technical University in Prague, Thákurova 9, Prague 160 00, Czechia.

Fig. 1 An example of a schedule for two auditors and three engagements. Engagements can be divided into phases, for example, interim and final in engagements 1 and 3. The engagement availability window is depicted within black ranges

Paper structure First, we review the audit scheduling literature (Sect. 1). We then justify our choice of the optimization solver because it affects which constraints can be efficiently modeled and how they should be formulated (Sect. 2). In the requirements Sect. 2.2, we list the modeled constraints and provide examples to illustrate their relevance to audit scheduling. In Sect. 3, we describe the data model. The result Sect. 4 then provides an empirical comparison of the optimized schedule versus a handmade schedule, which was created by an experienced scheduler over a 1-month period. Finally, we end the article with a discussion Sect. 5 and concluding remarks in Sect. 6.

1 Review

This section reviews different methods that can be used to solve the audit scheduling methods.

Summers [20] used linear programming (LP) to minimize the cost of an auditor-project assignment based on the level and hourly cost of each auditor and the project's hourly requirements at each level.

Balachandran and Zoltners [1] extended Summer's LP problem formulation to an integer linear programming (ILP) problem formulation to avoid the expressivity limitations of LP.

Romeijn and Morales [14] introduced a class of greedy approximation algorithms for the assignment problem. As an advantage of greedy algorithms, they scale to large instances. However, they do not guarantee the return of a feasible, let alone optimal, schedule.

Wang and Kong [22] used a genetic algorithm (GA) to maximize the match between auditors' specialization and the project characteristics. They chose to use a GA to obtain an approximately optimal solution because they were concerned that obtaining an optimal solution would be prohibitively time-consuming [22][section 2.2].

An extensive review of staff scheduling is provided by Ernst et al. [10], Van Den Bergh et al. [21], and Kalra and Singh [13].

2 Method

We tested multiple optimization methods, including greedy and genetic algorithms. However, we were unable to obtain a feasible solution from either of these algorithms in real-world instances. We identified two reasons for this.

- Staff scarcity. To cover all the engagements with the limited number of the auditors, the generated schedule must be nearly perfect. For this reason, a pure greedy algorithm is unlikely to provide a valid schedule. This is in agreement with the observation by Drexler and Gruenewald [8].
- There is a high count of diverse hard constraints that must be satisfied to obtain a solution acceptable by the auditors. This makes it difficult to design mutation and crossover operators in a genetic algorithm that would maintain the validity of the solution.

For these reasons, we opted to use an ILP solver because it is guaranteed to (eventually) find a valid solution if such a solution exists. Once a valid solution is found, all improvements are guaranteed to be valid as well.

2.1 Formulation

We denote the sets by capital letter, for example, A , and their corresponding variables by lowercase characters, for example, a for $a \in A$.

We assume that the firm has A auditors, and we want to find a schedule for the auditors for the next D days. The firm has E engagements, where each engagement is divided into consecutive P phases (e.g., interim and final). Each engagement requires auditors at a specific level (e.g., analyst, senior manager). We assume that each auditor is exactly at one level out of all L levels. Because an engagement may require multiple auditors at the same level, we differentiate between them using I indexes.

By introducing aliases, task $T = (E, P, L, I)$ and sink $S = (E, P, L, I, D)$, where day d indicates when task t starts (Fig. 2), we can formulate the auditor schedule optimization as a multi-commodity network flow problem (a multi-commodity network is a network with one or more sources and one or more sinks [9, 17]).

The flow problem can be represented as a graph with auditors A on the left and sinks S on the right (see Fig. 3).

Fig. 2 Search space

Fig. 3 Scheduling problem formulated as a flow problem

If auditor a can be potentially assigned to a sink s , there is an edge x_{as} between the nodes. The task of the optimizer is then to decide which of the candidate assignments are selected and which are not:

$$x_{as} = \begin{cases} 1 & \text{if auditor } a \text{ is assigned to sink } s, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where the sink s tells us which task the auditor a should do and on which day d auditor a should start working.

The objective of the optimizer is to minimize the sum of the costs c_{as} associated with assigning auditor a to sink s :

$$\begin{aligned} & \underset{x}{\text{minimize}} \quad \sum_{a \in A} \sum_{s \in S} c_{as} x_{as} \\ & \text{subject to} \quad \sum_{a \in A} \sum_{d \in D} x_{as} = 1 \text{ for each task } t, \end{aligned} \quad (2)$$

where the constraint states that each task must be assigned. The following further describes the basic assignment model.

2.2 Requirements

Considerable effort was made to extract a complete set of optimization criteria and constraints from the auditors. We provide an exhaustive list of the auditor requirements in the hope that interested readers may take inspiration from them for their own needs. To facilitate

this, we do not only list the requirements, but we also justify the requirements, provide illustrative examples and implementation details.

Multitasking Based on the poor previous experience with auditors working on multiple tasks in a single day, it was required that a single auditor be assigned at most a single task per day.

Implementation. We enforced this with the following constraint:

$$\sum_{t \in O_{ad}} x_{as} \leq 1 \text{ for each auditor } a \text{ and day } d, \quad (3)$$

where O_{ad} is a set of overlapping candidate tasks for auditor a and day d .

Continuity It was also required that an auditor finishes one task before moving to another to minimize the context switch penalty. **Implementation.** We solve this by generating only edges, where an auditor utilizes all its available time to the potential task until the task is finished (see Fig. 4). The task durations are given by the engagement managers, and the durations are assumed to be the exact estimates, which are independent of the assigned auditor.

Auditor availability Each auditor has a calendar with the number of hours they can work on a specified day. This allowed us to model the following constraints:

Example. The auditor will be hired at the beginning of the next month.

Example. It is known that the auditor will leave the firm in 2 months.

Example. The auditor works part-time.

Example. Training, vacations, and national holidays.

Implementation. When the auditor does not work on day d , the auditor cannot start a new task on that day. This is ensured by not generating edges in the graph.

Fig. 4 Availability of each engagement in days is depicted by the black ranges. Possible sinks, assuming that auditor a can work 8 h every day, are represented by gray rectangles. Sinks 3 and 4 cannot be assigned to auditor a because of the overlap. However, sinks 2 and 4 can be assigned to auditor a because no overlap occurs. An example of an illegal sink is sink 6, because once an auditor starts working on an engagement, the auditor must finish it as soon as possible

Fig. 5 Either of the auditors working 8 or 4 h per day can be assigned to the depicted task. However, an auditor working only 2 h per day would not be able to finish the task in time

We also skip the generation of edges for tasks that auditor a may not finish in time (see Fig. 5).

Engagement availability Each engagement has a calendar of days when auditors can work on the engagement. This allowed us to model the following constraints:

Example. Availability of the client data.

Example. Due dates.

Example. Client's company-wide holidays.

Implementation. An edge can be created only on a day when engagement e is available. The edge can be created only if it is hypothetically possible for auditor a to finish the task in time.

Familiarity If possible, we prefer to assign auditors who have past experience with the engagement.

Implementation. We solve this by adjusting the cost matrix c .

Travel cost Auditor travel costs can have an important impact on the audit costs [4], particularly for audit companies with multiple offices. We approximate the travel cost by the flying distance between the auditor's office and the client in Euclidean projection and add it to the cost of the edge. Furthermore, not all auditors are willing to travel long distances.

Example. The auditor does not have a driving license.

Implementation. When an auditor is unwilling to travel, let us say, more than 10 km, we remove all the edges from the auditor to the clients that are more than 10 km away.

The travel cost between the office and the client is modeled by adjusting the cost matrix c .

Level substitutions If necessary, an auditor at one level can be substituted by an auditor at a different level.

Example. If no junior analyst is available, assign a senior analyst on the junior analyst task

Implementation. We solve it by adjusting the cost matrix c .

Updates The list of engagements and available auditors were not fixed. Hence, one of the requirements was the ability to change the engagement or the auditor lists without many changes to the already generated schedule.

Implementation. We approximate it by adjusting the cost matrix c to include a penalty if an auditor will suddenly work on a new task. Conversely, we include a reward in the cost matrix if the auditor will work on the same task as before.

Note that this formulation gives the optimizer a leeway to shift the tasks over time.

Hard task preference The firm may require certain auditors to participate in certain tasks. While this is a trivial constraint, it is a frequent one.

Example. The engagement manager must be assigned to his/her engagement.

Example. Conflict of interests.

Implementation. To enforce the assignment of auditor a to task t , we set the following:

$$\sum_{d \in D} x_{as} \geq 1 \text{ for auditor } a \text{ and all sinks } s \text{ with task } t \quad (4)$$

The forbidden combinations are simply excluded from the set of candidate edges.

Soft task preference The auditors may express preferences to certain tasks [15].

Example. The auditor wants to specialize in bank auditing.

Implementation. We solve this by adjusting the cost matrix c .

Staff scarcity As the customer base is growing, there is an insufficient number of auditors to work all of the required hours.

Implementation. We solve this through the introduction of "mock auditors" in the set of auditors, who if assigned will have to be hired. Of course, the recruitment and training of new auditors is a costly endeavor. Hence, we penalize the employment of mock auditors by extending the objective function to:

$$\text{minimize } \sum_{xy} \sum_{a \in A} \sum_{s \in S} c_{as} x_{as} + c_{\text{mock}} \sum_{a \in M} y_a \quad (5)$$

where M with $M \subset A$ is the set of mock auditors, c_{mock} is the cost per mock auditor, and y_a is a binary slack variable for each mock auditor. We use following constraints to force slack variables y_a to take value 1 when mock auditor a is assigned to at least one sink s :

$$y_a \geq x_{as} \text{ for each auditor } a \in M \text{ and sink } s. \quad (6)$$

Uncertainty As we plan further into the future, the chance of unforeseen disturbances will increase. To minimize the expected count of schedule changes caused by unforeseen disturbances, we favor scheduling tasks earlier rather than later, given the choice.

Example. Some auditor leaves the firm.

Implementation. We implement this by adjusting the cost matrix c with a hyperbolic discounting of the reward [19]:

$$c_{\text{reward}}(d) = c_{\text{reward}} \frac{1}{1 + kd} \quad (7)$$

where c_{reward} is a reward for starting a task on the first day in the schedule, d is the index of the day, and k is a

parameter governing the degree of discounting (we used $k = 0.01$ and constant c_{reward}).

Warm up It takes time to get ready to work on an engagement [12, 16]. This overhead is minimized by minimizing the number of distinct auditors at the engagement by assigning the same auditor to as many tasks as possible. To ensure that this criterion does not go against level specialization, this criterion uses a substantially smaller penalty than the level substitution penalty.

Implementation. We extend the objective function to the following:

$$\text{minimize}_{xyz} \sum_{a \in A} \sum_{s \in S} c_{as} x_{as} + c_{\text{mock}} \sum_{a \in M} y_a + c_{\text{warm up}} \sum_{a \in A} \sum_{e \in E} z_{ae}, \quad (8)$$

where $c_{\text{warm up}}$ is the cost of the “warming up” of one auditor to one engagement, and z_{ae} is a binary slack variable for each auditor and engagement. We use following constraints to force slack variables z_{ae} to take value 1 when auditor a is assigned to engagement e at least once:

$$z_{ae} \geq x_{as} \text{ for auditor } a \text{ and all sinks } s \text{ with engagement } e. \quad (9)$$

2.3 Rejected constraints

We also present a list of constraints, which were rejected by our auditors, but which, nevertheless, might be of interest to the reader.

Precedence Task j cannot start before task i is completed [2].

Example. The audit must be concluded before the documents get translated.

Justification. Our auditors rejected this constraint because it is already considered in the 2.2 engagement availability windows.

Lag Minimum and maximum time-lags between two tasks are given [15].

Example. By law, the period between two tasks cannot exceed 3 months.

Justification. Our auditors rejected this constraint because it is already considered in the 2.2 engagement availability windows.

Parallelity Tasks may be forced to be processed in parallel [2].

Example. An analyst and the analyst’s manager must work on the engagement on the same days.

Justification. Our auditors did not consider this to be an actual problem to solve.

Workload The total workload of an auditor is lower bounded [2].

Example. An auditor must have at least 30% utilization.

Justification. The firm does not suffer from the underutilization of the auditors but rather from the opposite problem. Consequently, 2.2 auditor availability constraints suffice.

3 Implementation

The code is written in Python 3.7 and is divided into 3 parts:

The input data are in an SQL database. The advantage of using the SQL database as a data storage is that it allows us to define integrity constraints (e.g., not null, unique, foreign key) declaratively and that essentially each data engineer is familiar enough with SQL to deal with eventual integrity violations. The entity-relationship diagram of the input data is in Fig. 6.

The input data, which must be provided by the auditors for a new schedule, include a set of `auditors`, a set of `engagements`, `availability` of the auditors and `engagements` (`auditor_availability` and `engagement_availability`, respectively), a set of `tasks` that the auditors should do, and a set of the past auditor assignments to the engagements as a proxy of the auditors’ familiarity with the engagements (`is_familiar`).

Inputs not listed in Fig. 6 include parameters, like the travel cost in EUR for 1 km, which are used to calculate the cost matrix C from the input data, and a potential past solution X that should be updated to conform to the new requirements.

For optimization, we experimentally compared open-source OR-Tools v7.7.7810 (with CBC v2.10.5 solver in the backend) to commercial Gurobi v9.0.3.

4 Results

The implementation was evaluated on real data with 71 auditors, 47 engagements, 3 phases, 10 levels, 6 indices, and 365 days. The runtime measurements were obtained on a laptop with an Intel Core i5 processor.

An empirical comparison of the computer-generated schedule to the handmade schedule by an experienced planner based on the criteria listed in section 2.2 is presented in Table 1.

The computer-generated schedule is better or equal to the handmade schedule in all regards except the count of level substitutions. To be fair to the experienced planner and to further highlight the difficulty of finding a feasible schedule given the constraints, if we did not permit level substitutions in the computer-generated schedule, a feasible solution (a solution that does not violate any hard

Fig. 6 Entity-relationship diagram of the input data in IE notation

Table 1 Empirical comparison of a handmade schedule to the optimized schedule. When a handmade schedule takes 24 hours to generate, it actually means 3 man-days

Criterion	Handmade	Computer-generated
# Auditor availability violations	4	0
# Engagement availability violations	46	0
# Familiarity violations violations	157	153
# Level substitutions	0	41
# Auditors to recruit	3	2
Time to collect data	38 hours	38 hours
Time to schedule	24 hours	2654 seconds
Time to update schedule	8 hours	236 seconds

constraint) would require at least five new auditors owing to the clustering of engagements over time.

The comparison also includes the time to collect the engagement availability from the clients and the task set from the engagement managers (based on the survey conducted, each manager spends approximately 2 h on this task). We do not use any additional information beyond what has already been collected. The reported runtime of 2654 s includes input data validation & preprocessing, model creation & optimization, and schedule export. This runtime was achieved using the Gurobi solver. When we used OR-Tools, the runtime increased to 55793 s (21-times

Table 2 The effect of hard versus soft constraints

Implementation	Peak memory [MB]	Runtime [s]
Hard constraint	61	2654
Soft constraint	32311	132649

more than with Gurobi). The runtime of the schedule update (e.g., adjustment of the availability of an auditor or an engagement) is dominated by the model creation as the solver is warm-started from the previous solution.

5 Discussion

Since this article is not only about solving the audit scheduling problem exactly but also about solving the problem quickly, we have to discuss implementation details.

Hard versus soft constraints In comparison to the reference ILP formulation of audit scheduling by Balachandran and Zoltner [1], we exclude the illegal assignments from the problem formulation, whereas Balachandran and Zoltner leave them in the problem formulation and only adjust the corresponding cost in c_{as} to a large constant. The advantage of exclusion is a smaller memory footprint and faster optimization (observe Table 2).

Example. When we conduct an exhaustive cross-join between 71 auditors, 47 engagements, 3 phases, 10 levels,

Table 3 The effect of early pruning

Implementation	Peak memory [MB]	Runtime [s]
With early pruning	61	2654
Without early pruning	630	2791

Table 4 The effect of changing what x means

Implementation	Peak memory [MB]	Runtime [s]
Auditor started to work	61	2654
Auditor works	90	16493

6 indices, and 365 days, we obtain 219240900 edges. However, in real data of the same size, only 245455 are needed (an almost 1000-fold reduction).

Early pruning While modern ILP implementations have efficient presolvers that can quickly prune away unused or fixed variables, we found early pruning during the problem formulation highly beneficial, because it leads to smaller matrices and that speeds up all the following data transformations (observe Table 3). Examples of early pruning include exclusion of edges where an auditor cannot finish the workload in time, or exclusion of days when auditors do not work, like national holidays, as presented in section 2.2.

Problem representation A tempting problem representation of the scheduling problem is to use a binary vector x_{aed} :

$$x_{aed} = \begin{cases} 1 & \text{if auditor } a \text{ works on engagement } e \text{ on day } d \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

However, similarly to Chan, Dodin et al. [3, 4, 6, 7], we opted to use a binary vector x_{aed} :

$$x_{aed} = \begin{cases} 1 & \text{if auditor } a \text{ started to work on engagement } e \text{ on day } d, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

because it gives us the continuity constraint of section 2.2 for free without the need to refine the optimization criterion to also minimize the length of the assignments (observe the impact in Table 4).

Mock auditors The advantage of using a large pool of mock auditors is that the solver can find an initial solution quickly. The disadvantage is that it then takes longer to converge to the optimal solution and to prove that the solution is indeed optimal (observe Table 5).

We recommend using just a slightly higher number of mock auditors than is the expected need to minimize the runtime.

Table 5 The effect of the number of mock auditors

Implementation	Peak memory [MB]	Runtime [s]
2 mock auditors	61	2713
3 mock auditors	61	2654
10 mock auditors	62	2701
20 mock auditors	62	2816

Table 6 The effect of dividing the problem into subproblems

Implementation	Peak memory [MB]	Runtime [s]
Whole at once	44316	23868
4 subsets one after another	92	8292

Divide and conquer After the initial run on a subset of 71 auditors and 47 engagements, described in section 4, the auditing company decided to generate the schedule for all 299 auditors and 271 engagements. This led to an increase in memory requirements and runtime. But after dividing auditors and engagements into 4 disjoint subsets, particularly memory requirements decreased significantly (observe Table 6). Since the division was based on distinct types of audit (banking, non-banking) and geographically distinct areas (east, west), it did not lead, at least in this case, to a degradation of the obtained solution.

6 Conclusion

We implemented a schedule optimizer for an auditing firm and published the code at <https://github.com/janmotl/audit-scheduling>. We formulated the scheduling problem as a multi-commodity network flow problem. Our contribution is in adapting a multi-commodity network flow formulation from *shift-based* scheduling with fixed starting and end times, to *task-based* scheduling, where the tasks can be performed at any time in a temporal window.

According to the auditors, the introduction of some automated planners was necessary because the hand-made scheduling started to take an unacceptably long time. Our planner implementation reduced the time to schedule from 3 days to 45 min. As a byproduct, the computer-generated schedule is of higher quality than the handmade schedule. The most appreciated improvement was the reduction in the number of auditors to recruit from three to two, because the wage makes a significant part of all their expenses.

In future work, we plan to extend the planner to consider the flu season when some of the auditors are expected to be unavailable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Balachandran Bala V, Zoltners Andris A (1989) An interactive decision support system. *Account Rev* 56(4):801–812
2. Peter B, Sigrid K (2001) Resource-constrained project scheduling and timetabling. *Lect Notes Comput Sci* 2079(LNCS):277–293
3. Hung CK, Bajis D (1986) Decision staff support system with and for scheduling constraints precedence dates. *Account Rev* 61(4):726–734
4. Hung CK, Lam SF, Shijun C (1998) Audit scheduling and the control of travel costs using an optimization model for multinational and multinational audits. *J Account Audit Financ* 13(1):67–98
5. Dantzig G, Fulkerson R, Johnson S (1954) Solution of a large-scale traveling-salesman problem. *J Oper Res Soc Am* 2(4):393–410
6. Bajis D, Elimam AA (1997) Audit scheduling with overlapping activities and sequence-dependent setup costs. *Eur J Oper Res* 97(1):22–33
7. Bajis D, Huang CK (1991) Application of production scheduling methods to external and internal audit scheduling. *Eur J Oper Res* 52(3):267–279
8. Andreas D, Juergen G (1993) Nonpreemptive multi-mode resource-constrained project scheduling. *Inst Ind Eng* 25(5):74–81
9. El AAA, Mohamed G, Nashat FM (2018) A new formulation and solution for the nurse scheduling problem: a case study in Egypt. *Alexandria Eng J* 57(4):2289–2298
10. Ernst AT, Jiang H, Krishnamoorthy M, Sier D (2004) Staff scheduling and rostering: a review of applications, methods and models. *Eur J Oper Res* 153(1):3–27
11. Johnson SM (1954) Optimal two- and three-stage production schedules with setup times included. *Nav Res Logist Q* 1(1):61–68
12. Lilia KR, Boctor FF (2018) An efficient genetic algorithm to solve the resource-constrained project scheduling problem with transfer times: The single mode case. *Eur J Oper Res* 265(2):454–462
13. Mala K, Sarbjeet S (2015) A review of metaheuristic scheduling techniques in cloud computing. *Egypt Inf J* 16(3):275–295
14. Edwin RH, Romero MD (2000) A class of greedy algorithms for the generalized assignment problem. *Discret Appl Math* 103(1–3):209–235
15. Salewski F (1994) An integrative approach to audit-staff scheduling. Technical report, Institut für Betriebswirtschaftslehre der Universität Kiel
16. Salewski F, Schirmer A, Drexel A (1996) Project scheduling under resource and mode identity constraints. Part II: an application to audit-staff scheduling. Technical report, Institut für Betriebswirtschaftslehre der Universität Kiel
17. Khodakaram S, Sara B (2020) The multicommodity network flow problem: state of the art classification, applications, and solution methods. *Oper Res* pp 1–47
18. Smith Wayne E (1956) Various optimizers for single-stage production. *Nav Res Logist Q* 3(1–2):59–66
19. Sozou Peter D (1998) On hyperbolic discounting and uncertain hazard rates. *Proc R Soc B Biol Sci* 265(1409):2015–2020
20. Summers Edward L (1974) The audit staff assignment problem: a linear programming analysis. *Account Rev* 49(3):575
21. Den Bergh V, Jorne BJ, Philippe DB, Erik D, Liesje DB (2013) Personnel scheduling: a literature review. *Eur J Oper Res* 226(3):367–385
22. Yu Ren Wang and Siang Lin Kong (2012) Applying genetic algorithms for construction quality auditor assignment in public construction projects. *Autom Constr* 22:459–467

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.