



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

A Pilgrim Scheduling Approach to Increase Safety During the Hajj

Knut Haase, Mathias Kasper, Matthes Koch, Sven Müller

To cite this article:

Knut Haase, Mathias Kasper, Matthes Koch, Sven Müller (2019) A Pilgrim Scheduling Approach to Increase Safety During the Hajj. *Operations Research* 67(2):376–406. <https://doi.org/10.1287/opre.2018.1798>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Operations Research*. Copyright © 2019 The Author(s). <https://doi.org/10.1287/opre.2018.1798>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2019, The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

A Pilgrim Scheduling Approach to Increase Safety During the Hajj

Knut Haase,^a Mathias Kasper,^b Matthes Koch,^a Sven Müller^c
^aInstitut für Verkehrswirtschaft, Universität Hamburg, D-20146 Hamburg, Germany; ^bInstitut für Wirtschaft und Verkehr, Technische Universität Dresden, D-01187 Dresden, Germany; ^cInstitute for Transport and Infrastructure, Karlsruhe University of Applied Science, D-76133 Karlsruhe, Germany

Contact: knut.haase@uni-hamburg.de (KH); mathias.kasper@tu-dresden.de (MathK); matthes.koch@uni-hamburg.de (MattK); sven.mueller@hs-karlsruhe.de,  <http://orcid.org/0000-0002-3406-2133> (SM)

Received: May 28, 2015

Revised: April 1, 2015; May 18, 2015; January 24, 2017; April 30, 2018

Accepted: April 30, 2018

Published Online in Articles in Advance: January 25, 2019

Subject Classifications: transportation: scheduling; simulation: system dynamics

Area of Review: OR Practice

<https://doi.org/10.1287/opre.2018.1798>
Copyright: © 2019 The Author(s)

Abstract. The Hajj—the great pilgrimage to Mecca, Saudi Arabia—is one of the five pillars of Islam. Up to four million pilgrims perform the Hajj rituals every year. This makes it one of the largest pedestrian problems in the world. *Ramy al-Jamarat*—the symbolic stoning of the devil—is known to be a particularly crowded ritual. Up until 2006, it was repeatedly overshadowed by severe crowd disasters. To avoid such disasters, Saudi authorities initiated a comprehensive crowd management program. A novel contribution to these efforts was the development of an optimized schedule for the pilgrims performing the stoning ritual. A pilgrim schedule prescribes specific routes and time slots for all registered pilgrim groups. Together, the assigned routes strictly enforce one-way flows toward and from the ritual site. In this paper, we introduce a model and a solution approach to the Pilgrim Scheduling Problem. Our multistage procedure first spatially smooths the utilization of infrastructure capacity to avoid dangerous pedestrian densities in the network. In the next optimization step, it minimizes overall dissatisfaction with the scheduled time slots. We solve the Pilgrim Scheduling Problem by a fix-and-optimize heuristic, and subsequently simulate the results to identify necessary modifications of the scheduling constraints. Our numerical study shows that the approach solves instances with more than 2.3 million variables in less than 10 minutes on average. At the same time, the gap between optimal solution and upper bound never exceeds 0.28%. The scheduling approach was an integral part of the Hajj planning process in 2007–2014 and 2016–2017. No crowd disaster occurred in these years. Our approach was not applied in 2015, when a severe crowd crush happened close to the ritual site. We briefly discuss possible causes and consequences of this accident.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Operations Research. Copyright © 2019 The Author(s). <https://doi.org/10.1287/opre.2018.1798>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This work was partially supported by the Ministry of Municipal and Rural Affairs of the Kingdom of Saudi Arabia.

Keywords: Hajj • crowd management • scheduling • timetabling • pedestrian simulation • crowd disaster • crowd dynamics • pilgrimage

1. Introduction

The number and severity of deadly crowd disasters at public events has risen significantly over the past decades (Helbing et al. 2015). A frequent cause of such disasters is overcrowding. If congestion reaches critical levels, crowd motion transitions into a stop-and-go pattern and eventually causes a phenomenon called *crowd turbulence* (Helbing et al. 2007). Turbulent crowd motion is characterized by random, unintended displacements of groups in all possible directions (mass motion). In this state, individuals who have fallen over are trampled on if they cannot get back on their feet quickly enough and so become obstacles for others. The people following on are pushed forward by the crowd’s motion. Consequently, the area of trampled persons expands

rapidly, although nobody wants to intentionally harm anybody.

Until 2006, this phenomenon was observed at several crowd disasters with hundreds of casualties during the annual great pilgrimage to Mecca (see Table 1).

With up to four million Muslim pilgrims from all over the world, the Hajj is one of the largest recurring mass gatherings in the world (Baxter 2010). At the *Jamarat* pillars, three symbolic representations of the devil, pilgrims repeatedly perform *Ramy al-Jamarat*—the stoning of the devil ritual (Figure 1). During the four-day ritual, dense pedestrian flows approach the ritual site through a road network. Because many pilgrims prefer to follow the ritual times of the prophet Muhammad, the risk of particularly high crowd densities at the ritual site is significant (Johansson et al. 2007, Johansson 2008).

Table 1. Crowd Disasters in Mina, 1990–2015

Year	Location	Casualties
1990	Pedestrian Tunnel	1,426
1994	Jamarat Bridge	266
1997	Eastern entrance	22
1998	Eastern entrance	118
2001	Jamarat Bridge	35
2004	Jamarat Bridge	251
2006	Jamarat Bridge [Eastern entrance]	363
2015	Mina valley—crossing of streets 223/204	769

Notes. Data of 1990–2006 based on Helbing and Johansson (2009). Figure of 2015 is the official figure released by the Saudi Arabian Ministry of Health (Browning 2015). Other sources report higher numbers for 2015 (Gambrell and Ahmed 2015).


In the aftermath of the accident of 2006, the Ministry of Municipal and Rural Affairs of the Kingdom of Saudi Arabia (MOMRA) embarked on a massive endeavor to

prevent crowd disasters at the Jamarat Bridge (Haase et al. 2016). Most significant actions of MOMRA’s reorganization plan are summarized in Table 2.

The developed infrastructure solutions aimed to increase the capacity of the ritual site and to enforce a separation of incoming and outgoing pedestrian flows on the road network. As a result, most roads leading to and from the Jamarat Bridge are now one-way streets (see Figure 2). However, MOMRA recognized that the immense infrastructure investments are necessary but not sufficient to guarantee pilgrims’ safety. In the past, dangerous crowd conditions arose, when the flows of pilgrims between the pilgrim camps and the Jamarat Bridge lacked spatiotemporal coordination.

Consequently, MOMRA initiated a program to route, schedule, predict, and monitor pilgrim flows at the holy sites. Key objectives were the implementation of a safe routing plan and a detailed timetable for the

Figure 1. (Color online) Mena Tent City and Jamarat Bridge



Sources. (a) 21.388718 N, 39.906427 E; Google Earth, July 2015. (b) MOMRA. (c) Zull Must / Shutterstock.com.

Table 2. Jamarat Ritual Reorganization Measures

1	Replacement of the old Jamarat Bridge by a new, simulation-based design with increased capacity
2	Unidirectional flow organization of the street network (see Figure 5)
3	Separate ramps at the Jamarat Bridge for incoming and leaving pilgrims
4	Redesign of Jamarat Plaza to allow redirection of flows between different floors (see Figure 2)
5	Real-time flow monitoring
6	Rerouting in situations of high capacity utilization of certain routes
7	Contingency plans for all kinds of situations (including bad weather conditions)
8	Expansion of tunnel capacity
9	Introduction of a metro line
10	Control tower jointly used by all responsible authorities
11	Improved signage
12	Awareness campaigns including videos, leaflets, guides
13	Enhanced communication systems
14	Tightened protection against unauthorized access to the ritual site

Note. Summary of main reorganization measures for the Ramy al-Jamarat ritual based on Johansson (2008) and Helbing et al. (2015).

stoning ritual. Therefore, we have developed a scheduling approach that consists of two main components. First, we solve a large-scale mixed-integer problem to construct a stoning ritual schedule with route assignments. Because of the size of the problem, we omit dynamic features of pedestrian flows such as queuing and dynamic velocity in the first step. In the second step, we incorporate these first-step omissions into a mesoscopic pedestrian flow simulation. Both components combined


generate computationally feasible and safe solutions to the scheduling problem.

Naturally, we can only provide formal movement schedules for pilgrims who are officially registered with the Saudi authorities. Despite all efforts by Saudi Arabian security forces to prevent unauthorized access, a considerable number of illegal local pilgrims perform Hajj every year.¹ Their presence poses an additional challenge for a safe operation. Figure 3(d) shows the use of required infrastructure by illegal pilgrims. In practice, we adapt our scheduling and simulation program to account for their presence.

2. Problem Setting

Pilgrim accommodation and transport are organized by *establishments*. Currently, there are eight establishments. Each of them represents pilgrim organizations of a greater geographical region. Figure 4 displays the location of camps and their assigned establishments in 2016. An establishment is usually further subdivided into multiple service-offices (*Mutawifs*). A single service-office may be in charge of up to 5,000 pilgrims. It manages the supply of food, information, and spiritual guidance, and it is also a recipient of the schedules.


The tents of the Mena valley are grouped in *blocks*, which are usually surrounded by streets and can include one or multiple *camps* (see Figure 3). Pilgrims of one camp are pooled into *groups* of around 250 pilgrims each. Designated local guides lead their groups from the camp to the Jamarat Bridge and back.

Figure 2. (Color online) Separated Access and Egress Flows

Source. MOMRA 2014.

Notes. Access to and egress from the Jamarat Bridge is based on unidirectional and steady road usage. The picture shows the Jamarat Plaza in front of the Jamarat Bridge and the Mena tent city in the background. Densities are such that the flow of pilgrims is continuous and noncongested.

Figure 3. (Color online) Tents in the Mena Valley Accommodating the Pilgrims



Source. MOMRA.

Pilgrims can access the Jamarat Bridge via the street network and the metro line (Figure 4). Major pedestrian routes through the tent city are designated as *main paths*. These main paths strictly adhere to a one-way flow organization (Figure 5). More specifically, a main path consists of a route toward the Jamarat Bridge and a route leading back into the tent city. Camps are linked to a main path by connection paths.

Because such connections must also adhere to the one-way flow principle, pilgrims might have to take longer routes than in a two-way flow network. In our optimization approach, main paths and connections paths are static and individually determined in collaboration with external crowd management experts. For each camp, a set of such round trips is defined. We refer to each round trip of a camp as one of its possible *paths* (see Figure 6).

We divide each of the four ritual days into discrete time periods. The period length is defined such that a group guide can gather and dispatch at least one group within a single period.

We define a set of resources to represent infrastructure bottlenecks along the paths. Resources correspond to street segments, crossings, and metro stations, as well as ramps,


floors, and the escalators of the Jamarat Bridge. A single street may be represented by multiple resources, and several paths may use the same resource. Figure 5 shows the main resources considered in the scheduling problem.

Each resource has a limited capacity. We define the capacity as the maximum allowed number of pilgrims crossing the bottleneck per period.

A safe schedule avoids variability in densities, speeds, and flows of pilgrims². Hence, we must ensure that a pilgrim schedule yields a smooth utilization of the available infrastructure capacity over time.

Though not a strict requirement, many pilgrims prefer to perform the stoning ritual at the exact times of the Prophet Muhammad. Thus, crowd density is expected to peak at these times. We define these times as set of *peak periods* for each ritual day. Nonregistered pilgrims, in particular, use much of the available capacity on the ground level and the first level of the Jamarat bridge during peak hours.

For safety reasons, we schedule no registered pilgrims to the lower levels of the Jamarat Bridge during the peak periods. However, establishments and service-offices provide *time preferences* of their pilgrims for the off-peak hours. Pilgrim organizations may request different times

Figure 4. (Color online) Pilgrim Tent Camps in the Mena Valley 2016

Note. Camps are colored according to their associated establishment in the season 2016.

based on outside factors such as departure flight schedules, weather conditions, or the health and demographics of their groups. In addition, the preferred times can reflect pilgrims' local cultural traditions. As an example, some pilgrim groups send their female pilgrims in the early hours while the male pilgrims will go in the afternoon.

Currently, pilgrim organizations provide time preference data at different levels of aggregation. We receive preference distributions for blocks, camps, or service-offices. They state the percentage of pilgrims who want to perform the ritual in a particular time interval, for instance, in one or multiple periods. In a preprocessing step described in Appendix B, we transform these inputs to determine one preferred stoning period for each pilgrim group and ritual day.


Pilgrim Scheduling Problem. Based on these definitions, we formulate the following. For each pilgrim group, assign a path and determine a stoning period for each ritual day such that, (i) the capacity limit of each resource is not exceeded, (ii) the resulting utilization of each resource over time is smooth, and (iii) the sum of the weighted deviations from the preferred

stoning periods over all groups and periods is minimized.

3. Related Literature

To the best of our knowledge, the Pilgrim Scheduling Problem (PSP) just defined has not been considered in the literature to date. The problem is related to evacuation planning to some extent (Choi et al. 1988, Church and Cova 2000, Bretschneider and Kimms 2011), mainly because of routing and capacities. However, pilgrim scheduling is not about evacuation—that is, leaving a given location as fast as possible—but rather about coordinating the movement pattern of millions of pilgrims (Helbing et al. 2002). For the PSP there is only one direction of flow and a single destination to which all pilgrims head. In contrast, evacuation planning usually relies on dynamic network flow problems (Hamacher and Tjandra 2001, Stepanov and Smith 2009). In the PSP, multiple flows amalgamate toward the destination: the Jamarat Bridge. This allows us to relax the explicit flow component of the problem. One objective of our problem, to minimize the deviation of scheduled stoning periods

Figure 5. (Color online) One-way System and Locations of Resources Proximate to the Jamarat Bridge




Notes. One-way system and locations of resources near to the Jamarat Bridge. In-flows and out-flows are strictly separated. Each path is designed according to the one-way system. Crossing the Jamarat Plaza takes less than 15 minutes on average.

from preferred stoning periods, is similar to the objective of timetabling and scheduling problems (Nemhauser and Trick 1998, Cayirli and Veral 2003, Ernst et al. 2004, Cardoen et al. 2010, Van den Bergh et al. 2013).


There have been only a few applications of operations research to the challenge of crowd management during Hajj to date. For a discussion see Haase et al.

(2016). Charnes et al. (1989) developed a mathematical model to plan infrastructure at the sacred places in the region of Mecca and Mena such that overcrowding and congestion would be avoided. Feng and Miller-Hooks (2014) propose a network optimization-based methodology to support efficient crowd movement during large public gatherings. The presented bilevel integer program

Figure 6. Assignment Principle of Camps, Paths, and Resources



Notes. Service-offices provide stoning time preferences for their pilgrim groups. Camps accommodate pilgrims. A path consists of an access route and an egress route from the camp to the Jamarat Bridge and back. More than one path might be available for a given camp. Only one path is selected in the solution. A path might contain several resources, and a resource might be used by more than one path.

Figure 7. (Color online) Relationship of Flow, Density, and Velocity

Notes. Shown are fundamental diagrams describing the relationship between density ρ_{kl} and flow F_{kl} (a), and between density ρ_{kl} and velocity V_{kl} (b). Besides the empirical data recorded during the crowd accident at the old Jamarat Bridge in 2006, we show two additional flow–density (velocity–density) curves with flow $F_{kl} = \rho_{kl} V_0 (1 - \exp(-\omega(\rho^{\max}/\rho_{kl} - 1)))$ and $V_{kl} = F_{kl}/\rho_{kl}$ denoting the velocity. $\rho^{\max} = 10$ is the maximum density. V_0 is the desired velocity of pilgrims under free flow conditions. ω is a scale parameter to adjust the curves to the maximum flow ($F^{\max}$). Since V_0 is unknown, we assume various values in the simulation.

accounts for travel time, minimal physical layout, and utility-maximizing routing of individuals. In contrast to both references, we focus entirely on operational issues instead of the infrastructural setup.

Based on Abdelghany et al. (2014) and Abdelghany et al. (2016), Verbas et al. (2016) recently presented an iterative, integrated bi-objective optimization and simulation framework to schedule pilgrims for the stoning ritual. They consider three interrelated subproblems: (i) A network analysis determines a set of paths for each camp to a main path. The optimization component (ii) determines the size of each group, the departure time, and the path to be taken by the pilgrims. Their approach aims to minimize pilgrims' travel time and to match the preferred stoning time distribution to the specified demand. The third component (iii) simulates each pilgrim in the network according to his/her path and departure time determined by (i) and (ii). The mesoscopic simulation computes actual travel times and capacity utilization. These are first submitted to the network analysis and then to the optimization component for the next iteration. Their process was reported to converge after three iterations.

Since the authors of this study essentially tackle the same practical problem that we have considered until 2014, let us highlight the significant differences of our approach. Verbas et al. (2016) tightly constrain the set of available time slots for a block only to the time periods preferred by the service-offices. Therefore, only the allocation of pilgrims *within* the preferred periods may differ from the service-offices' desired allocation. In contrast, we


consider for stoning time allocation not only the preferred time periods, but all feasible periods for a group. Pilgrim preferences are temporally clustered. To derive solutions with smooth capacity utilization, flexibility in stoning time allocation is beneficial for our approach.

Verbas et al. (2016) distinguish between main paths and access paths. Main path and access path together constitute a feasible route for the pilgrims of a camp. In their implementation, several different routes may be assigned to a given camp as the set of possible paths may vary over time. In contrast to our approach, a group of pilgrims might be assigned to different paths over the course of the stoning days.

The main paths described in Verbas et al. (2016) are designed in compliance with the one-way flow concept established in 2007. Access paths, in contrast, connect the camps to main paths and are generated by constrained shortest-path calculations on the road network. The presented algorithm determines a path of minimal travel time that cannot cross nodes or links of any main path. Apart from this constraint, the authors report no further details on whether the one-way-flow concept has been strictly applied in these path calculations and whether the return routes from the Jamarat Bridge to the campsites have been included in their analysis.

In our approach to the PSP, we use for every camp a set of optional paths designed in collaboration with local engineers and international crowd management experts. The design of a path represents a strict implementation of the one-way flow concept that was in

Figure 8. (Color online) GPS Tracking Results



Notes. Sample of 28 pilgrim groups tracked by GPS (one GPS tracker per guide) collected on all days in Hajj season 2012, yielding 352 observations. The mean walking velocity is 0.98 meters per second with a standard deviation of 0.28. (a) Relationship between velocity (ordinate) and distance to the Jamarat Bridge (abscissa). We have fitted various functions to the data (polynomial regression). The velocity declines slightly close to the Jamarat Bridge. As depicted by the box plots in (b), velocities vary mostly between 0.5 and 1.5 meters per second, regardless of the distance to the Jamarat Bridge.

operation from 2007 to 2014 and 2016 to 2017. We restrict the use of main road links and all access road links to strictly their designated flow direction. Hence, our implementation generally avoids crossing of access and return paths to and from the main paths, but may contain routes that are not minimal with respect to the pilgrims travel time. For safety reasons, we assign the same path to a camp on each day of the ritual. By incorporating this restriction, we give up some flexibility when minimizing dissatisfaction with the timetable; but we avoid group guides' having to handle different maps handed out for different days, which might cause confusion. Pilgrim groups may also remember their designated route after the first day of the ritual and tend to follow it again.

Apart from these differences the approach presented by Verbas et al. (2016) seems to be comparable to our approach, in that it integrates large-scale crowd simulation with an integer-programming approach to form a crowd-scheduling framework. In the next section, we will introduce the integer-programming model and its underlying assumptions, as well as the simulation model.

4. Modeling Approach

4.1. Model Assumptions

The definition of our deterministic planning model is based on the key assumptions outlined below.

(1) Flow capacity safety margin. Saudi authorities have defined a target pedestrian density of one pilgrim per square meter on the street network. At this density,

pedestrians can walk uninterrupted on uncongested streets. Empirical data collected during the Hajj season in 2006 and shown in Figure 7 reveal that this density corresponds to a pedestrian flow of one pilgrim per meter and second. In our model, this constitutes a corresponding flow limit (or capacity) of about 3,600 pilgrims per hour for a street segment (resource) of 1 meter's width. To accommodate possible obstacles, irregular road variations, and nonregistered pilgrim flows, we work with a reduced flow capacity of 3,250 registered pilgrims per hour and meter of street width. In comparison, the maximum flow observed by the CCTV system in 2006 corresponded to a capacity of about 6,790 pilgrims per hour and meter at a density of 4.4 pilgrim per square meter in Figure 7. By design, our defined capacity is well below this maximum flow. Hence, we can assume free flow conditions even if our model solution uses a resource to full capacity.

(2) Static group velocity. Given a pilgrim group's stoning time period, its camp location, and its designated path, we calculate the resulting utilization of the resources endogenously in our model. Most of the considered resources are located near the Jamarat Bridge (see Figure 5). A pilgrim group is assumed to use such resources during the assigned stoning period. For more distant resources, we calculate the utilization period based on the assigned stoning period and an estimated travel time. For this calculation, we assume a constant speed of motion, which is typical for groups in an uncongested network (Yamaguchi et al. 2011). Research

on pedestrian behavior indicates that groups are likely to walk at a more constant speed than individuals and that they tend to walk at the pace of its slowest or leading members (Cheng et al. 2014). As stated above, our definition of capacities always enables a free flow on the network segments. On an uncongested road, pilgrim groups can walk at their desired speed. Hence, we assume in the scheduling model that walking velocities are independent of time and of the scheduled resource utilization. We calculate a static offset for the utilization period based on the stoning period. A sample of GPS data collected in 2012 shows that pilgrims' velocity is 1 meter per second on average (Figure 8).

(3) Period length and velocity. In our scheduling program, period corresponds to a time interval of 30 minutes' length. Over- or underestimating the actual group velocity, and hence travel time, likely has a very limited impact on the calculation of resource utilization. Our scheduling model restricts capacity utilization at bottleneck resources for each 30-minute period. Because of the assumption of static average velocity, the realized arrival times of the implemented schedule will differ slightly from the planned arrival times at the bottlenecks. However, as differences are small, the scheduled period will correspond to the actual arrival period, so that the capacity utilization control is still effective. Imagine a pilgrim group walking 600 meters to the Jamarat Bridge. Given a speed of 1 meter per second (3.6 km/h), a person needs about 10 minutes to walk along the street segment. A speed of 0.7 meters per second (2.52 km/h) would result in 14.2 minutes for the same segment, while a speed of 1.3 meters per second (4.7 km/h) would correspond to 7.4 minutes. In all cases, the group would be expected to arrive within its scheduled 30-minute period. The maximum observed speed for Hajj in 2006 was 1.5 meters per second, whereas the lowest velocity, measured during congestion, was around 0.5 meters per second (Helbing et al. 2007). However, measured over many pilgrims, these extreme values are averaged out. We do not expect a high variation of average velocities, because the optimized schedule avoids congestion and produces steady flows. Note, however, that we assume constant velocities over time only in the scheduling model, while possible variations in velocity are taken into account by the dynamic network loading module described in Section 4.3.

(4) Pilgrim dissatisfaction with the timetable. Due to constrained infrastructure capacities, the resulting timetables never fully match the given time preferences of all pilgrim groups. However, we account for preferences in our approach because we reasonably expect that pilgrims' compliance with the timetable increases with decreasing deviation between assigned stoning times and stated preferences. This relation may not

necessarily be linear. Here, we assume that we can maximize overall compliance if we schedule as many groups as possible to the preferred period or adjacent periods. Specifically, we expect that reducing the deviation of the scheduled to the preferred period from 12 hours to 10 hours has a smaller effect on the compliance than an equally long reduction of that deviation from two hours to zero. However reasonable, we cannot directly validate whether this particular assumption holds for Hajj. Nevertheless, empirical studies on commuting behavior demonstrate that deviations from preferred arrival times can result in behavioral changes. For example, empirical studies documented changes in route choices and departure times (Senbil and Kitamura 2004, Saleh and Farrell 2005, Jou et al. 2008, Li et al. 2012). Moreover, Saudi authorities consider a close match of pilgrims' ex ante preferred and ex post realized—that is, scheduled—stoning times as defining the quality of service.

(5) Compliance in the model. In our scheduling model and in the simulation model, we assume 100% compliance with the timetable and the assigned paths, but we ensure sufficiently large, considerable reserves for deviations from the plan.

4.2. Timetabling and Routing

At its heart, the PSP has two objectives. First, to minimize spatiotemporal peaks in infrastructure utilization and, second, to minimize total dissatisfaction with the assigned time periods. These two scheduling objectives, *safety* and *satisfaction*, are conflicting. We do not consider a bi-objective problem formulation because a weighting scheme might be difficult to handle and to communicate. Instead, we decompose the problem into two subproblems.

In the first subproblem, we determine minimal resource utilization levels by solving a lexicographic min-max problem. This preprocessing step is part of the solution approach described in Section 5. In the second, subsequent scheduling problem, we constrain resource utilization to the precomputed levels and minimize total dissatisfaction with the timetable. By this approach, we literally put safety first as we select from all feasible and safe schedules one that fits best to the preferred stoning periods. To formulate a mathematical planning model for this problem, we introduce the following sets:

$\mathcal{T}$ stoning periods in ascending order, index t ;

$\mathcal{G}$ pilgrim groups, index g ;

$\mathcal{M}$ ritual days (movements), index m ;

$\mathcal{S}$ scheduling groups with $\mathcal{S} = \mathcal{G} \times \mathcal{M}$, index s ; and

$\mathcal{T}_s$ available stoning periods for scheduling group s ($\mathcal{T}_s \subset \mathcal{T}$)

and consider the following parameters:

τ_s^* preferred stoning period of scheduling group $s \in \mathcal{S}$,

θ threshold difference for objective function coefficient calculation, and

η scaling parameter for objective function coefficient calculation ($\eta \leq 1$).

We use scheduling groups instead of explicit combinations of groups and days to simplify model notation. A group g may be represented by more than one element in the set $\mathcal{S}$, if the group's ritual is scheduled on multiple days.

Using subsets $\mathcal{T}_s$, we define the available periods for each group. In general, we include all periods of the movement associated with group s in its subset $\mathcal{T}_s$. We may exclude periods due to group-specific time constraints. Such constraints are imposed to accommodate pilgrims' arrival times on the first day of the ritual and their departure times on the final day. They typically arrive and leave via prescheduled shuttle buses from and to the airport. We also use these subsets to model metro operation times and temporary road closures affecting specific subsets of camps.

As discussed in Assumption 4 of Section 4.1, we set up a nonlinear time-preference matching parameter. For each scheduling group $s \in \mathcal{S}$ and period $t \in \mathcal{T}_s$, we compute the *dissatisfaction*.

$$f_{st} = \begin{cases} (t - \tau_s^*)^2 & \text{if } t - \tau_s^* \leq \theta \\ \theta^2 + \eta \cdot (|t - \tau_s^*|) & \text{otherwise.} \end{cases} \quad (1)$$

The dissatisfaction penalty of an assignment of group s to period t increases quadratically as long as the deviation from the preferred period is not greater than the threshold θ . For deviations greater than θ , we set up a linear increase in the penalty. Following Assumption 4 about the compliance with the timetables, we define function (1) to have its steepest decline for absolute deviations from the preferred period not greater than θ .

Now, consider the additional sets:

$\mathcal{C}$ pilgrim camps;

$\mathcal{P}_c$ feasible paths for camp c ;

$\mathcal{P}_s$ feasible paths for group s ;

$\mathcal{S}_c$ scheduling groups in camp c ;

$\mathcal{S}_p$ scheduling groups that can use path p ;

$\mathcal{R}$ resources, index $r \in \mathcal{R}$; and

$\mathcal{Q}$ period-resource combinations, index $(r, t) \in \mathcal{Q}$ with $|\mathcal{Q}| < |\mathcal{R} \times \mathcal{T}|$

and parameters:

a_{rp} period offset between stoning time period and capacity utilization period of resource r on path p ,

b_{rt} capacity of resource r in period t in number of pilgrims,

n_s number of pilgrims of scheduling group s ,

$\bar{u}_{rt}$ relative utilization limit of resource r in period t with $\bar{u}_{rt} \leq 1$, and

σ_r maximum increase or decrease in utilization of resource r between two successive periods.

Now, we define the variables:

$X_{stp} = 1$, if scheduling group s is scheduled to perform stoning in period t and to use path p (0, otherwise);

$Y_{cp} = 1$, if camp c is assigned to path p (0, otherwise);

U_{rt} relative utilization of resource r in period t ; and

DS total dissatisfaction with timetable (objective function value).

We now define the following mixed-integer program **P1**:

Minimize

$$DS = \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}_s} \sum_{p \in \mathcal{P}_s} f_{st} \cdot X_{stp} \quad (2)$$

subject to

$$\sum_{p \in \mathcal{P}_c} Y_{cp} = 1 \quad \forall c \in \mathcal{C}, \quad (3)$$

$$\sum_{t \in \mathcal{T}_s} X_{stp} = Y_{cp} \quad \forall c \in \mathcal{C}, p \in \mathcal{P}_c, s \in \mathcal{S}_c, \quad (4)$$

$$\sum_{p \in \mathcal{P}_r} \sum_{s \in \mathcal{S}_p} n_s \cdot X_{s,t-a_{rp},p} = b_{rt} \cdot U_{rt} \quad \forall r \in \mathcal{R}, t \in \mathcal{T}, \quad (5)$$

$$U_{rt} - U_{r,t-1} \leq \sigma_r \quad \forall (r, t) \in \mathcal{Q}, \quad (6)$$

$$U_{r,t-1} - U_{rt} \leq \sigma_r \quad \forall (r, t) \in \mathcal{Q}, \quad (7)$$

$$0 \leq U_{rt} \leq \bar{u}_{rt} \quad \forall r \in \mathcal{R}, t \in \mathcal{T}, \quad (8)$$

$$X_{stp} \in \{0, 1\} \quad \forall s \in \mathcal{S}, t \in \mathcal{T}_s, p \in \mathcal{P}_s, \quad (9)$$

$$Y_{cp} \in \{0, 1\} \quad \forall c \in \mathcal{C}, p \in \mathcal{P}_c. \quad (10)$$

The decision enforced for each group $s \in \mathcal{S}$ through (3) and (4) is twofold: Assign a time period $t \in \mathcal{T}_s$ and a path $p \in \mathcal{P}_s$ by setting exactly one variable X_{stp} to 1. Constraints (4) enforce a common path choice for all groups of a specific camp. As each group uses the same path on each day of the ritual, pilgrims and their guides might remember the path between their camp and the Jamarat Bridge. According to (5), each positive decision—that is, $X_{stp} = 1$ —increases the capacity utilization of infrastructure resources on the selected path p at the chosen stoning time period t , adjusted by the static time period offset a_{rp} . The offsets are precomputed based on average travel times as described in Assumption 2 in Section 4.1. Equations (5) and (8) together impose capacity constraints on the resources. As outlined in Section 4.1, resource capacities b_{rt} are given such that even with $U_{rt} = 1$, full resource utilization, free flow conditions are retained. To effectively ensure that solutions will feature low utilization levels and a smooth flow of pedestrians, relative utilization limits $\bar{u}_{rt}$ must be set up carefully in the preprocessing step as described in Section 5.

Through (6) and (7) we limit the volatility in the utilization of a given resource over time. These constraints facilitate smooth utilization levels at the resources and a constant velocity of the scheduled pilgrims. Parameter σ_r defines the limit of the volatility of each resource contained in set $\mathcal{Q}$. Low values of σ_r ensure less volatility

in scheduled infrastructure utilization. Parameter σ_r also offers a way to coordinate flows of multiple resources. Pilgrims on a particular return route to their camps might have entered the Jamarat Plaza via different access roads and thus have used different access resources. By constraining the fluctuation of flow on the shared return road resource, our scheduling program effectively coordinates utilization between multiple access roads.

We exclude several infrastructure resources from the domain of smoothing constraints (6) and (7) because fluctuation in their utilization may either be acceptable or cannot be avoided at all. For instance, we exclude all resources related to the metro system, such as ticket gates, ramps, and stations. The train convoys follow a predetermined timetable and transport large groups of 3,000–4,000 passengers per train. This variable arrival rate cannot be smoothed and is dealt with through capacity reserves.

4.3. Dynamic Network Loading

As described in Section 4.1, the scheduling model relies on a number of assumptions about the pedestrian traffic to be computationally feasible. To analyze network performance in detail, we simulate the model solution under the consideration of dynamic aspects of crowd movements. The simulation explicitly implements the relationship between pedestrian density, velocity, and flow shown in Figure 7. In contrast to the scheduling model, the simulation covers not only a subset of key resources, but the entire infrastructure—that is, all road segments, escalators, etc.—and also explicitly includes the estimated flows of nonregistered pilgrims.

We employ a pedestrian flow simulation model (Helbing et al. 2002) based on a discrete version of the Lighthill–Whitham–Richards (LWR) traffic performance model (Treiber and Kersting 2012, pp. 81–126). In this approach, we partition the unidirectional route network into distinct sequentially ordered cells. For each cell, we determine the demand and supply of pedestrians using current densities, flow capacities, and assumptions about the desired walking speed. The flow between two adjacent cells balances supply and demand.

This macroscopic modeling approach, typically applied to study highway traffic, can handle scenarios with a large number of individuals (Daganzo 1995). Macroscopic models do not explicitly simulate each individual but describe traffic flow analogously to fluids in motion. Additionally, our simulation selects at each intersection the correct target cell for every outgoing group to ensure that they travel in accordance with the scheduled paths. This routing requires the model to track the groups throughout the network and to maintain a waiting queue in every cell. Because tracking and queuing are microscopic features incorporated into the macroscopic, discrete LWR model, our approach can

be categorized as a mesoscopic hybrid simulation model (Treiber and Kersting 2012, pp. 57–59).

We consider discrete time steps, indexed l , of $\delta = 10$ seconds duration. The cells, indexed k , are of sufficient length to ensure that pilgrims do not skip one or more cells in a single time step. To guarantee a correct mapping of the street network, the length of intersection cells may slightly deviate from the standard cell length.

Let ϕ_k be the area and μ_k the width of cell k , N_{kl} the simulated number of pilgrims, $\rho_{kl} = N_{kl}/\phi_k$ the density (pilgrims per square meter), and F_{kl} the flow (pilgrims per meter and second) in cell k at time l . Then the supply of cell k at time l is given as

$$S_{kl} = \mu_k \times \delta \times \begin{cases} F_{kl}^{\max} & \text{if } \rho_{kl} \leq \rho^{\text{crit}} \\ F_{kl} & \text{if } \rho_{kl} > \rho^{\text{crit}} \end{cases} \quad (11)$$

and the demand as

$$D_{kl} = \mu_k \times \delta \times \begin{cases} F_{kl}^{\max} & \text{if } \rho_{kl} > \rho^{\text{crit}} \\ F_{kl} & \text{if } \rho_{kl} \leq \rho^{\text{crit}}, \end{cases} \quad (12)$$

where $F^{\max}$ denotes the maximum flow and ρ^{crit} the associated density (Coscia and Canavesio 2008). Based on the empirical data visualized in Figure 7, we set $F^{\max}$ to 1.94 pilgrims per meter and second and ρ^{crit} to 4.4 pilgrims per square meter.

Let us consider one-way flow without any intersection: The flow from cell k' to the succeeding cell k in time step δ is given by the minimum of demand of cell k' and supply of cell k —that is, $\Delta N_{k',k,l} = \min(D_{k',l}, S_{k,l})$. Now, if we denote by k the successive cell of k'' , then the number of pilgrims in cell k at time $l + 1$ is given by

$$N_{k,l+1} = N_{kl} + \Delta N_{k',kl} - \Delta N_{kk',l}. \quad (13)$$

Now, let us consider an intersection with flows from the cells k' and $\tilde{k}$ to their successive cell k : If necessary, the supply of cell k is distributed proportionally to the demand of the cells k' and $\tilde{k}$,

$$\Delta N_{k',k,l} = \min(D_{k',l}, S_{kl} \cdot D_{k',l}/(D_{k',l} + D_{\tilde{k},l})). \quad (14)$$


Thus, in general, we calculate via the formula

$$N_{k,l+1} = N_{kl} + \sum_{k' \in K_k^{\text{Pre}}} \Delta N_{k',kl} - \sum_{k'' \in K_k^{\text{Suc}}} \Delta N_{kk'',l} \quad (15)$$

the number of pilgrims in cell k at time $l + 1$ where K_k^{Pre} are the predecessor cells of cell k and K_k^{Suc} are the successor cells of cell k .

To illustrate the simulation, we consider a part of a street network (Figure 9(a)) that is occupied by eight scheduling groups (Figure 9(b)). At time $l = 1$, 500 pilgrims are in cell $k = 2$. This yields a density of $\rho_{2,1} = 500/200 = 2.5$ pilgrims per square meter. The associated flow toward cell $k = 3$ is $F_{2,1} = 1.75$ pilgrims per meter and second. Taking into account $\mu_2 = 10$ meters

Figure 9. (Color online) Numerical Example of the Simulation Approach



Notes. Cell $k = 3$ is the successive cell of cells $k = 1, 2$. Cells $k = 4, 5$ are the successive cells of cell $k = 3$.

and $\delta = 10$ seconds, the demand of cell $k = 2$ is $D_{2,1} = 1.75 \cdot 10 \cdot 10 = 175$ pilgrims. The supply of cell $k = 3$ is 185. In an analogous manner we obtain $D_{1,1} = 125$ (demand of cell $k = 1$ toward cell $k = 3$) and $S_{3,1} = 185$ (supply of cell $k = 3$). As $125 + 175 > 185$, the supply of cell $k = 3$ is proportionally distributed to the cells $k = 1$ and $k = 2$ —that is, we obtain $\Delta N_{2,3,1} = 185 \cdot 175 / (125 + 175) \approx 108$ and $\Delta N_{1,3,1} = 185 \cdot 125 / (125 + 175) \approx 77$. Consequently, groups may split up within the time step (Figure 9 (c) and (d)). We apply the first-in, first-out (FIFO) principle to process the groups inside the cell queues. If one subsequent cell in a junction is congested and the other is not, the model does not strictly adhere to the FIFO principle. Rather, we allow groups that follow a path to the free cell to overtake waiting groups that follow a path to the congested cell. Since pilgrims of a group tend to stay together, split groups are joined again if all pilgrims of a scheduling group are located in the same cell.

The model can be adapted to location-specific details via parametrization. For example, we can consider cell-specific flow–density diagrams. We calibrate escalators, ramps, and stairs with individually adjusted flow–density diagrams based on GPS samples.

5. Solution Approach

We derive solutions for the PSP by solving four consecutive subproblems. First, we determine group-specific dissatisfaction parameters f_{st} based on aggregated preference distribution data. Second, we compute minimal resource utilization profiles for each resource using a lexicographic minimax approach. These are used to set up utilization bounds $\bar{u}_{rt}$ for the scheduling problem **P1**. In the third step, we solve **P1** by applying a heuristic fix-and-optimize strategy. The schedules computed in this step are then forwarded to the dynamic network loading module. If it indicates critical conditions in the network, we introduce additional capacity constraints

to the optimization model for a further scheduling iteration.

5.1. Parameter Setup

Establishments and service-offices submit for each stoning day one or multiple relative distributions of pilgrim numbers over the day. For example, a service-office that supports pilgrims in three different camps provides up to three different distributions per stoning day.

In Algorithm 3 of Appendix B, we outline how we use this preference information to derive a single preferred stoning period for each scheduling group. If we scheduled each group for the preferred period obtained by our procedure, the preference distributions given by the service-offices would be closely approximated.

The solution of **P1** depends on the values of parameters $\bar{u}_{rt}$ and σ_r . To determine upper bounds on the resource capacities $\bar{u}_{rt}$ we solve an aggregated path assignment problem first. We use the following additional sets:

$\mathcal{R}^{\min}$ set of resources considered,

$\mathcal{W}$ set of considered combinations (r, m) with $r \in \mathcal{R}^{\min}$, $m \in M$,

the following additional parameters:

d_{cm} the number of groups in camp c on day m ,

$\tilde{b}_{rm}$ the aggregated capacity of resource r on day m ,

$\tilde{u}_{rm}^*$ maximum allowed resource utilization for resource r on day m ,

and variables:

$\tilde{U}_{rm}$ aggregated capacity utilization level for resource r and day m , and

$\bar{U}$ maximum of capacity utilization.

The objective of the capacity precalculation problem is to lexicographically minimize the maximum aggregated capacity utilization over all considered resources and stoning days. We exclude several resources from the set $\mathcal{R}^{\min}$. As an example, all resources related to the metro system (ticket gates, train stations) are excluded. The metro train convoys follow specific timetables, which implicitly determine their time-dependent capacity bounds. We use the full available capacity for them in problem **P1**.

Now, we formulate the mixed-integer program **P0** as follows.

$$\text{Minimize } \bar{U} \quad (16)$$

subject to (3), (10),

$$\tilde{U}_{rm} \leq \bar{U} \quad \forall (r, m) \in \mathcal{W}, \quad (17)$$

$$\sum_{p \in P_r} \sum_{c \in C_p} d_{cm} \cdot Y_{cp} = \tilde{b}_{rm} \cdot \tilde{U}_{rm} \quad \forall r \in R, m \in M, \quad (18)$$

$$0 \leq \tilde{U}_{rm} \leq \tilde{u}_{rm}^* \quad \forall r \in R, m \in M, \quad (19)$$

$$0 \leq \bar{U} \leq 1. \quad (20)$$

Constraint (17) is used to calculate the maximum utilization of all resource-movement combinations in set $\mathcal{W}$. Capacity constraint (18) is used to calculate aggregated

utilization levels of the resources for every movement, and (19) and (20) define the variable domains. In this subproblem, binary choices Y_{cp} determine an assignment of camps to available paths and thus to a set of resources. The objective is to determine lowest possible capacity utilization levels given the current pilgrim numbers and their spatial distribution in the Mena Valley. As summarized in Algorithm 1, we accomplish this by repeatedly solving problem **P0** and updating parameters $\tilde{u}_{rm}^*$ and set $\mathcal{W}$.

Algorithm 1 (Capacity-Bound Precalculation)

$\mathcal{W} := \{(r, m) \mid r \in \mathcal{R}^{\min}, m \in M\}$

$\tilde{u}_{rm}^* = 1$

while $|\mathcal{W}| > 0$ **do**

Solve **P0** and obtain $\tilde{U}_{rm}, \bar{U}$

$(r', m') = \arg \max_{(r, m) \in \mathcal{W}} \{\tilde{U}_{rm}\}$

$\tilde{u}_{(r', m')}^* = \bar{U}$

$\mathcal{W} := \mathcal{W} \setminus \{(r', m')\}$

The final values for $\tilde{u}_{rm}^*$ represent the relative utilization for each resource and ritual day. But constraining the relative utilization of each resource and time period to this level would only be possible if the scheduler could distribute the pilgrim groups uniformly among the periods of a ritual day. Such a distribution scheme is usually not feasible for instances of problem **P1** because of the individual time constraints defined in the sets $\mathcal{T}_s$. Feasibility is retained for problem **P1** by gradually increasing the calculated minimum utilization levels. We introduce a scaling parameter λ to control the maximum utilization of resources exceeding their minimum utilization levels ($\lambda \geq 1$). Using $\tilde{u}_{rm}^*$, the right-hand side of capacity constraints (8) is calculated as

$$\bar{u}_{rt} = \tilde{u}_{rm}^* + \frac{1 - \tilde{u}_{rm}^*}{\lambda} \quad \forall r \in R, m \in M, t \in \mathcal{T}_m. \quad (21)$$


External requests to reduce resource utilization during specific time slots can be introduced by further modification of parameter $\bar{u}_{rt}$. If scaling parameter $\lambda = 1$ or $\tilde{u}_{rm}^* = 1$, then $\bar{u}_{rt} = 1$. Since $\bar{u}_{rt}$ declines in λ , larger values of λ are associated with safer schedules.

By construction, $\bar{u}_{rt}$ is more sensitive to λ for low levels of minimum utilization, small values of $\tilde{u}_{rm}^*$, than for high utilization levels. If the minimum utilization is high, close to unity (100%), then large values of λ are needed to achieve a small incline in $\bar{u}_{rt}$ compared with $\tilde{u}_{rm}^*$. Otherwise, smaller values of λ yield $\bar{u}_{rt}$ close to 1 as shown in Figure 10(a). Of course, the lower $\tilde{u}_{rm}^*$ the more control on $\bar{u}_{rt}$ is provided by λ .

5.2. Timetabling and Routing

Required input data for the scheduling and simulation modules is usually finalized only shortly before the beginning of Hajj. To incorporate ad hoc updates of data and yet ensure accurate numerical solutions, simulations, and visualizations, fast solution methods

Figure 10. (Color online) Scheduling Parameters λ and σ_r



Notes. (a) Impact of λ on the utilization limit $\bar{u}_{rt}$ as formally given in Equation (21). (b) Impact of σ_r on utilization volatility. In (b), the variance is nearly 0.6 for $\sigma_r = 1$ and 0.3 for $\sigma_r = 0.5$.

for various data scenarios are crucial. Table 4 shows that despite the reduced number of scheduling groups in 2016, solution times of the mixed-integer program **P1** are still significant. Therefore, we consider a three-stage solution approach to solve the PSP with varying (λ, σ_r) settings heuristically with low computation times.

First, we determine upper bounds on capacity utilization as described by Algorithm 1 and parameter assignment (21). Then we determine the path assignment for every camp. Finally, we select a stoning period for each scheduling group—that is, we determine X_{stp} .

To solve the path assignment problem, we leave out the capacity utilization smoothing constraints (6) and (7). We replace binary constraints (9) with

$$X_{stp} \geq 0 \quad \forall s \in \mathcal{S}, t \in \mathcal{T}_s, p \in \mathcal{P}_s \quad (22)$$

and define problem **P2** as (2) subject to (3)–(5), (8), (10), and (22). The solution to **P2** yields binary path choices Y_{cp}^{*P2} and provides a lower bound to problem **P1**. We fix the path assignment decision by setting

$$Y_{cp} = Y_{cp}^{*P2} \quad \forall c \in \mathcal{C}, p \in \mathcal{P}_c. \quad (23)$$

Now, we introduce the period-wise capacity utilization smoothing constraints (6) and (7) as well as binary constraints (9) to solve the scheduling problem for the remaining X_{stp} .

Algorithm 2 (Solution Procedure for Problem P1)

Initialize λ and σ_r .

Process Algorithm 1, obtain $\bar{u}_{rm}^*$.

$\bar{u}_{rt} := \bar{u}_{rm}^* + \frac{1 - \bar{u}_{rm}^*}{\lambda} \quad \forall r \in \mathcal{R}, m \in \mathcal{M}, t \in \mathcal{T}_m$.

Solve **P2** to obtain Y_{cp}^{*P2} .

Set $Y_{cp} = Y_{cp}^{*P2} \quad \forall c \in \mathcal{C}, p \in \mathcal{P}_c$.

Solve **P1** to obtain final X_{stp} .

The solution process summarized in Algorithm 2 allows us to derive timetables and path assignments from given pilgrim data in only a few minutes of computation time. In this approach, the two scheduling parameters λ and σ_r jointly characterize the expected safety of a schedule. We generate schedules for various values of σ_r and λ , and evaluate them in the dynamic network loading module.

The simulation enables us to evaluate pedestrian densities and fluctuation in the network. We incorporate network segments indicating critical states into the scheduling models, where they are considered as capacity-constrained resources. Further, if we detect high local densities caused by return flows, additional restrictions on the set of available time periods can be imposed for the respective groups in the scheduling models. We reoptimize the scheduling results by combining both measures—that is, by adding spatial and temporal constraints based on the simulation results. We repeat the combination of optimization and simulation until we can verify good conditions for the entire network. In Figure C.1, the reoptimization principle is visualized for one ritual day.

6. Computational Study

The data for the PSP of the year 2016 are summarized in Table 3. We used $\eta = 0.1$ and $\theta = 2$ to calculate objective coefficients f_{st} .

6.1. Study Setup

We implemented all problems of this study using the algebraic modeling system GAMS 24.7 (GAMS 2016). We used the CPLEX 12.6 mixed-integer-programming solver to perform deterministic concurrent optimization. Numerical tests were processed on a workstation with

Table 3. Data Summary of Hajj 2016 (1437H)

Name	Annotation	Quantity
Scheduling groups	$ \mathcal{S} $	27,676
Number of periods	$ \mathcal{T} $	192
Number of camps	$ \mathcal{C} $	868
Number of paths	$ \mathcal{P} $	44
Number of resources	$ \mathcal{R} $	39
Number of network segments (simulation)		2,885
Number of periods (simulation)		14,000
Feasible path–camp combinations	$ \mathcal{P} \times \mathcal{C} $	1,615
Resource–period combinations considered for smoothing in Q	$ \mathcal{R} \times \mathcal{T} $	1,043
Feasible group–period–path combinations	$ \mathcal{S} \times \mathcal{T} \times \mathcal{P} $	2,390,747

Note. Because of large-scale expansion works at the Al-Haram Mosque in Mecca, the number of registered pilgrims declined significantly from 2013, with registrations increasing again for the first time in 2017.

256 GB of memory and two Intel Xeon E5-2667v3 3.2-GHz processing units. We set a maximum solution gap of 0.01% for **P1** and **P2**, and solve **P0** to optimality. Computation time was limited to six hours per study instance.

The instances differ in their parameter combination of $\lambda \in \{1.0, 1.5, 2.0, \dots, 5.0\}$ and $\sigma_r \in \{0.05, 0.10, 0.15, 0.20, 1.0\}$. We assume $\sigma = \sigma_1 = \dots = \sigma_r$ for this study. We defined a period length of 30 minutes. We included $\lambda = 1.0$ to calculate solutions with 100% of the available scheduling capacity b_{rt} . Instances with $\sigma_r = 1.00$ effectively relax smoothing constraints (6) and (7). Instance $(\lambda, \sigma) = (1.0, 1.0)$ represents a schedule with the greatest emphasis on time-preference satisfaction.

We normalize parameter settings for group sizes and capacities in this study. We set $n_s = 1 \forall s \in \mathcal{S}$ and define capacities $b_{rt} = \lfloor \frac{b_r}{250} \rfloor \forall r \in \mathcal{R}, t \in \mathcal{T}$. We consider slightly more scheduling groups with each having 250 pilgrims or fewer, but problem **P2** exhibits integrality for the relaxed X_{stp} .

Each instance is processed twice. First, we compute a lower bound for our fix-and-optimize heuristic (FO) by solving the integer program **P1** to optimality. Because of the problem size of **P1** (Table 3), not all instances were solved to optimality within the six-hour time limit. For instances without an optimal solution, we include the solver's estimate of the best possible solution in the results. Second, we determine a solution with the fix-and-optimize approach. This solution is an upper bound to the PSP. We calculate the solution gap between both solutions relative to the lower bound value.

6.2. Scheduling Results

Table 4 summarizes the computational results of the study. In general, computational effort rises when σ declines and λ increases. Highest computation times are recorded for the instance with maximum availability of capacity and tightest bounds on the volatility of utilization. Here, finding a feasible integer solution

during the branch and bound process for **P1** took more than 80% of overall computation time. The solution space for the relaxed linear program problem **P2** expands with increasing infrastructure capacity (low λ settings). But with decreasing σ —that is, lower allowed volatility of the utilization—the number of feasible integer solutions in the solution space decreases. Setting σ too low—that is, $\sigma = 0$ —requires the resources to be utilized uniformly throughout the ritual days. As discussed in Section 5, this would render the problem infeasible.

Our results confirm that the fix-and-optimize approach can find good solutions quickly. On average, finding the lower bound with the integer-programming model took more than 12 times as long as finding a solution with the fix-and-optimize heuristic for instances that have been solved within the six-hour time limit. The significant difference can be explained by the reduction of the solution space after fixing path choices Y_{cp} . This step effectively eliminates the path dimension from the domain of the X_{stp} variables and reduces problem **P1** to a one-dimensional assignment problem.

Table 4 suggests that the fix-and-optimize heuristic produces solutions of high quality. The gap between optimal and heuristic solution is low with 0.056% on average and a maximum of 0.27%.

A deterioration in the objective function of **P1**^{fx}, the final heuristic solution, compared with **P2** can only be explained by disadvantageous path choice fixation—that is, when adding volatility constraints (6) and (7) to **P2**, different path choices would have been realized. To avoid violating constraints (6) and (7), the set of groups that have been fixed to a path and thus to a set of resources may need to be reassigned to different time periods. If the time preferences do not allow this exchange to be free of cost for the objective function value, the solution deteriorates.

Table 4 shows that the objective function value increases in λ . At the same time, total dissatisfaction

Table 4. Excerpt of Computational Study Results

Instance λ	σ	Computation time (CPU seconds)			Objective value				Goal evaluation	
		P1 ^{fx}	P2	FO	LB	DS ^{FO}	DS ^{LB}	GAP	MT	TSRU
1.00	1.00	50	84	133	315	1.0000	1.0000	0.0000	0.755	1.000
1.00	0.20	66	84	150	1,950	1.0012	1.0006	0.0625	0.755	0.876
1.00	0.15	87	84	171	2,194	1.0016	1.0010	0.0645	0.755	0.873
1.00	0.10	708	84	792	-	1.0031	1.0022 *	0.0880	0.755	0.837
1.00	0.05	5,305	84	5,388	-	1.0107	1.0080 *	0.2715	0.751	0.710
2.00	1.00	55	110	165	2,280	1.1225	1.1225	0.0003	0.720	0.458
2.00	0.20	119	110	229	3,503	1.1229	1.1228	0.0125	0.720	0.397
2.00	0.15	126	110	236	1,542	1.1233	1.1230	0.0300	0.720	0.395
2.00	0.10	285	110	395	5,687	1.1240	1.1235	0.0384	0.720	0.387
2.00	0.05	757	110	867	-	1.1267	1.1257 *	0.0934	0.718	0.363
3.00	1.00	55	174	229	1,400	1.1719	1.1718	0.0045	0.705	0.294
3.00	0.20	108	174	282	1,022	1.1725	1.1720	0.0397	0.705	0.287
3.00	0.15	129	174	303	2,662	1.1729	1.1722	0.0586	0.705	0.288
3.00	0.10	164	174	338	4,087	1.1736	1.1728	0.0664	0.704	0.286
3.00	0.05	573	174	747	-	1.1752	1.1741 *	0.0948	0.703	0.263
4.00	1.00	54	210	264	1,070	1.2033	1.2032	0.0054	0.696	0.242
4.00	0.20	123	210	333	1,587	1.2037	1.2033	0.0364	0.696	0.241
4.00	0.15	149	210	359	2,569	1.2040	1.2035	0.0448	0.696	0.241
4.00	0.10	156	210	366	3,504	1.2045	1.2038	0.0619	0.696	0.235
4.00	0.05	627	210	837	-	1.2065	1.2051 *	0.1182	0.695	0.220
5.00	1.00	56	255	311	3,632	1.2217	1.2216	0.0050	0.690	0.220
5.00	0.20	91	255	346	5,404	1.2220	1.2217	0.0210	0.690	0.218
5.00	0.15	118	255	373	10,618	1.2223	1.2218	0.0363	0.690	0.219
5.00	0.10	147	255	402	18,292	1.2229	1.2221	0.0626	0.690	0.214
5.00	0.05	625	255	880	-	1.2249	1.2233 *	0.1344	0.689	0.203

Notes. Computation time includes the time the solver needed to solve the respective integer problems in cpu seconds. P1^{fx} denotes the problem **P1** with fixed path choices Y_{cp} . Column FO contains combined computation times of P1^{fx} and **P2**. Objective function values are given relative to a column's minimum objective function value. DS^{FO} shows the final objective function values of the fix-and-optimize approach. Column DS^{LB} contains objective values of integer-programming solutions of **P1**. GAP denotes the solution gap between the final fix-and-optimize heuristic solution and the optimal solution of **P1** in percent. Column MT shows the percentage of groups assigned either to their preferred or an adjacent time period—that is, $M = \frac{1}{|\mathcal{G}|} \cdot \sum_{s \in \mathcal{G}} \sum_{p \in \mathcal{P}_s} \sum_{t \in \mathcal{G}_s | t_{st} \leq 1} X_{stp}$. Column TSRU displays the sum of the squared exceedance of utilization over a threshold value of 0.5. It is calculated for resources contained in $\mathcal{Q}$ only and is displayed relative to the column's maximum.

*Instances that could not be solved to optimality within six hours of computation time.

increases with decreasing σ . The more emphasis the planner puts on balanced capacity utilization through λ and σ , the higher is the total deviation from the preferred stoning periods.

This trade-off between *safety* and *satisfaction* objectives is highlighted in the goal evaluation columns of Table 4. Column MT lists the share of groups assigned to either their preferred period or an adjacent period. In the safest solution, more than two-thirds of the groups match this criterion. However, the least safe solution raises this figure by a mere 6.6% points. Given the capacity constraints, the stated time-preference distribution can never be fully matched. The attainable maximum is 75.5%. The column TSRU, total squared resource utilization, is an indicator for balanced and smooth resource utilization in the solution. To obtain this value, we sum the squared exceedance of relative capacity utilization over a threshold value of 0.5 for all resources contained in $\mathcal{Q}$.

Column TSRU displays the indicator relative to the column's maximum value.

Low values of λ correspond to high TSRU values and highest sensitivity for the σ parameter. Given the time preferences (see Section 7.2), Table 4 reveals that low TSRU values can only be achieved by accepting low preference matching rates. The better we match the pilgrim preferences, the higher the magnitude of peak capacity utilization and volatility in the solution.

The results also show the interaction of both parameters. The effect of varying σ on TSRU decreases with rising λ values. The same is true for the objective values. The increase in dissatisfaction when changing from $\sigma = 1$ to $\sigma = 0.05$ ranges between 1.07% points for $\lambda = 1$ and 0.32% points for $\lambda = 5$.

By raising λ , we decrease the gap between minimum utilization levels $\tilde{u}_{rm}^*$ and the maximum allowed utilization levels $\tilde{u}_{rt}$ in **P1**. We have shown in Section 5 that

$\tilde{u}_{rm}^*$ is calculated assuming that pilgrims can be distributed uniformly over each ritual day. Hence, if the $\tilde{u}_{rt}$ are very close to the value of their associated $\tilde{u}_{rm}^*$, the resulting utilization is smooth regardless of the σ setting. Hence, varying σ is less effective in instances with high λ settings.

However, we can set up σ to ensure smooth capacity utilization in scenarios with less strict λ settings. Potential use cases for such settings are less critical road segments. If the authority responsible for crowd management is confident to handle higher peak utilization at a particular segment, the σ -constraints would still ensure that low fluctuation is scheduled for such resource. Figure 11(a) illustrates the effect of σ on scheduled resource utilization. Fluctuation of the utilization is reduced on the lowest σ setting. In this solution, dangerous stop-and-go waves in the scheduled utilization can be avoided.

6.3. Simulation Results

We simulate all instances under the assumption of 100% schedule compliance for the registered pilgrims. Based on CCTV counting at access and exit points around Jamarat Plaza, we estimate arrival times and the distribution of nonregistered groups over the arrival paths. We assume a total of 433,000 nonregistered pilgrims of whom almost 50% perform the ritual at the ground level for this simulation. Registered groups are not scheduled during the peak hours, but nonregistered arrivals are estimated with a beta distribution covering the peak periods.

The variance in cell densities is mostly determined by the underlying schedule. However, the example in Figure 11 shows that we usually detect slight differences between utilization calculated in scheduling model P1 and the respective simulated flows. We can identify three effects to explain these differences. First, at multiple points in the network, return flows toward the campsites merge into flows directed toward Jamarat (see Figure D.3 for an example). Returning groups then walk toward Jamarat until they reach their camp. At the entry points, merging groups temporarily increase local densities. Such small local peaks carry over to adjacent network segments and fade out with increasing distance to the bridge.

Second, nonregistered pilgrim groups are simulated to perform stoning during the peak hours. On their way to Jamarat they also use the tent city's infrastructure in periods adjacent to the peak hours. As a result, scheduled utilization deviates from simulated utilization in these particular periods.

Table 5 reveals the third effect on simulated densities. Using different fundamental diagrams (see Figure 7), we can determine the effect of assuming various desired velocities V_0 . The table displays the network's *level of service* (LOS) as defined by Fruin (1971). LOS assesses the traffic performance on an ordinal scale from A (free


flow) to F (congested). We measure LOS by the available space per person in a crowd (see Still 2018 for practical demonstrations of LOS). Lowest service levels and highest maximum density are recorded for $V_0 = 0.75$ m/s. In this extreme-case scenario, pilgrim groups that performed stoning before the peak hours return very slowly to their campsites. To respect the one-way flow system, they must join the access flows toward Jamarat at some point. If the returning flow moves very slowly, it merges with the high-density flows of the non-registered pilgrims during the peak periods. Table 5 shows that the average velocity of the pilgrim flows scheduled toward Jamarat is close to 1 m/s and varies with V_0 . The results suggest that differences between simulated travel times and the static travel times used to calculate scheduled departures are small for all tested fundamental diagrams.

Results presented in Table 6 show the effect of scheduling parameter λ on the simulated network performance. With increasing λ the share of LOS categories A and B rise while categories C to F show declining numbers. The simulation confirms that by narrowing the available scheduling capacity in the optimization model, the accumulated duration of high densities in the system can be reduced. The highest maximum density is detected for $\lambda = 1$. Here, simulation helps to identify critical network cells that require crowd management.

Figure 11(b) shows the simulated utilization of a resource for different σ_r settings and constant $V_0 = 1.1$ m/s. With rising σ_r , volatility of the utilization increases. The relative standard deviation of the simulated flow increases from 0.78 ($\sigma_r = 0.1$) to 0.86 ($\sigma_r = 0.3$). The progression of the scheduled utilization in Figure 11 differs from the simulated utilization because of the effects of return flow merging described above and the non-registered pilgrim flows. Because, in this study, the capacity constraints are employed using unweighted group counting—that is, $n_s = 1$ —the scheduled utilization slightly overestimates real utilization.

We evaluate two additional schedules in this study: the actual schedule of Hajj season 2016 and a schedule generated without capacity constraints—that is 100% time-preference matching—to highlight possibly dangerous segments and periods. In this unconstrained schedule, we assign paths to camps by shortest walking distance. Figure 12 highlights the importance of spatiotemporal coordination of pilgrim flows by visualizing the disparity between the two scenarios in crowd density, Figure 12 (a) and (b), and maximum density, Figure 12 (c) and (d). Congestion is expected to occur at densities above $\rho^{crit} = 4.4$ P/m². Panels (c) and (d) of Figure 12 illustrate areas of expected congestion in red. For the unrestricted schedule, panels (b) and (d) of Figure 12 suggest that the highest densities would occur along the northern return path leading into the

Figure 11. Scheduled and Simulated Utilization of a Network Resource Over 1 Day



Notes. Utilization profiles for northern access resource King Fahd Road on Day 11. Each line represents a different setting of σ_r . Parameter λ is kept constant. Parameter σ influences the volatility in the scheduled utilization. The lower σ , the lower the allowed volatility of utilization between succeeding time periods. No regular pilgrim groups are scheduled between 6:00 a.m. and 10:00 a.m. to avoid potential overcrowding with nonregistered pilgrims.

tent city because multiple arriving flows converge into it after the stoning ritual. The simulation identifies multiple hot spots of congestion directly at Jamarat Plaza but also deep inside of the tent city and on the southern return paths. For the actual schedule of 2016 (Figure 12(a)), densities never exceeded $1P/m^2$ for more than 99% of the time. Maximum densities in Figure 12(c) indicate no congestion. Because maximum densities are below ρ^{crit} (see Figure 7 for reference), the schedule features

free flow conditions at any time, low travel times, and avoids congestion. In contrast, the instance with complete preference matching (no schedule) produces massive congestion as well as disastrous maximum densities surpassing $8P/m^2$. In this scenario, density levels are higher, especially within and around the peak periods, because congestion cannot be dissolved early enough and overlaps with the nonregistered pilgrim flows during the peak hours.

Table 5. Simulated Level of Service for Varying Desired Speed V_0

V_0	LOS A	LOS B	LOS C	LOS D	LOS E	LOS F	$\max \rho$	$\emptyset V$	$\emptyset \text{GAP}_{\text{TT}}$
$[\frac{m}{s}]$	$\geq 3.3 [\frac{m^2}{p}]$	$2.3-3.3 [\frac{m^2}{p}]$	$1.4-2.3 [\frac{m^2}{p}]$	$0.9-1.4 [\frac{m^2}{p}]$	$0.5-0.9 [\frac{m^2}{p}]$	$\leq 0.5 [\frac{m^2}{p}]$	$\frac{p}{m^2}$	$[\frac{m}{s}]$	[min]
1.50 (emp.)	96.83	1.76	1.34	0.07	< 0.01	0.00	1.64	1.28	-08.24
1.50	97.18	1.71	1.06	0.04	< 0.01	0.00	1.49	1.36	-09.70
1.25	96.35	2.04	1.54	0.07	< 0.01	0.00	1.57	1.14	-04.82
1.10	95.38	2.54	1.96	0.11	< 0.01	0.00	1.66	1.00	-02.03
1.00	94.48	3.02	2.30	0.19	0.01	0.00	1.74	0.91	+03.26
0.75	91.51	3.88	3.39	1.11	0.10	0.01	2.51	0.69	+15.44

Notes. Level of service (LOS) as defined by Fruin (1971). For each network segment and time step, we determine the LOS. Columns LOS A–LOS F show the relative frequency of the respective category over all segments and time steps. Each row displays simulation results based on a fundamental diagram with desired speed V_0 . The first row contains results for the fundamental diagram of 2006. Results are given for Day 11 and scheduling instance ($\lambda = 5$, $\sigma = 0.05$). Average velocity, $\emptyset V$, calculated over all registered scheduling groups on their way from camp to the Jamarat Bridge. $\max \rho$ is the highest density found. $\emptyset \text{GAP}_{\text{TT}}$ displays the average difference between simulated travel time and assumed travel time of a group in minutes.

Because the velocity of a pilgrim group moving through the network depends on the crowd density of the utilized segments, traveling through systems of low density should not only be safer, but also faster. Even though the path choice logic in the scheduling model does not consider travel times, we have observed a decline in total travel time with increasing safety settings. In our extreme-case example, accumulated travel time over all groups in the unrestricted scenario increased by more than 16% compared with the schedule implemented in 2016. Travel times increased by more than one hour for individual groups compared with the regular schedule. In the scorching heat of the Mena Valley, low travel times and low densities may affect the travel experience even more than in a regular urban transport scenario. The results verify the initially stated trade-off: pilgrims' travel experience on the streets could be improved if lower stoning time-preference matching were accepted.

Using the simulation approach, we can also evaluate planned infrastructure projects. For instance, simulation results showed that the ramp planned after the 2015 accident (Section 8) yields high performance improvements especially at the intersection of streets 223 and 204. This infrastructure project was successfully implemented in 2016. However, infrastructure is not modified very frequently, and only minor updates to the scheduling resources are required from year to

year. Most scheduling resources are located around Jamarat Plaza. Figure 13 provides an explanation for this pattern; it illustrates the distribution of simulated densities on network segments at different distances to the Jamarat Bridge. We observe that densities decline with increasing distance to the Jamarat Bridge. Close to Jamarat Plaza, multiple flows merge to produce the highest densities. If densities at these resources enable a free flow of pedestrians, it is reasonable to assume that resources upstream are uncongested as well.

7. Implementation and Challenges

We implemented the approach described in Section 5 for the Hajj scheduling process of 2016 and 2017. Variants of it were applied between 2007 and 2014. Following the death of King Abdullah bin Abdulaziz Al Saud in early 2015, MOMRA underwent political reorganization. The ministry, whose primary mission is to develop infrastructure and large-scale construction projects, has withdrawn from the pilgrim scheduling operations ever since. As a result of these changes, the authors of this study were not involved in the scheduling project in 2015. However, the approach presented here has been reapplied for Hajj seasons 2016 and 2017.


In the interaction with the planning authorities, we have noticed a growing awareness and openness toward the analytical approach over the years. Even aspects that appear secondary to the academic problem,

Table 6. Simulated Level of Service for Different Settings of λ

λ	LOS A	LOS B	LOS C	LOS D	LOS E	LOS F	$\max \rho$
	$\geq 3.3 [\frac{m^2}{p}]$	$2.3-3.3 [\frac{m^2}{p}]$	$1.4-2.3 [\frac{m^2}{p}]$	$0.9-1.4 [\frac{m^2}{p}]$	$0.5-0.9 [\frac{m^2}{p}]$	$\leq 0.5 [\frac{m^2}{p}]$	$[\frac{p}{m^2}]$
1.00	94.37	2.73	2.35	0.54	0.01	< 0.01	2.91
2.00	94.64	2.76	2.34	0.26	0.01	0.00	1.60
3.00	94.77	2.86	2.18	0.18	0.01	0.00	1.65
4.00	95.04	2.66	2.15	0.15	0.01	0.00	1.65
5.00	95.22	2.57	2.08	0.11	0.01	0.00	1.65

Notes. Each row represents a schedule with a λ setting as shown in the first column. The σ parameter is fixed to 1. $\max \rho$ denotes the highest density found.

Figure 12. (Color online) Simulation Results for Hajj in 2016—Day 10



Notes. Comparison of the actual schedule of 2016 and a schedule without capacity restrictions. Maps (a) and (b) visualize the total time a network cell exhibits densities above $1P/m^2$. Maps (c) and (d) illustrate maximum densities of network cells.

such as the requirement to prepare detailed input data for the scheduling models, contribute to establishing a quantitative, analytical viewpoint to supplement religious and traditional aspects in discussions. Today, mathematical planning, simulation, and monitoring systems are known and established tools in the authorities' approach to crowd management. Detailed scheduling reports can provide a starting point for the authorities to implement a comprehensive operation plan and to begin constructive discussions with the pilgrim organizations.


7.1. Implementation Process

The timeline in Figure 14 summarizes the timing of the planning process. Before receiving required input data for our scheduling module, organizational processes have to be completed. Among these are the formation of the service offices, negotiations about pilgrim contingents and tent zone allocations, means of transport, and others. This process involves numerous government authorities and contractors who organize pilgrim data

individually. We consolidate their inputs using several data sources, such as spreadsheet data, (hand-) drawn maps, and telephone calls. During this phase, the planner incorporates pilgrim-group-specific preferences and constraints on times and paths into the data set. Once a data set is considered final, we employ the fix-and-optimize approach described in Section 5 to obtain preliminary solutions.

The Ministry of Hajj and the associated pilgrim organizations review the material and provide feedback to fine-tune the schedules. Such fine-tuning may include the modification of camp routes—that is, in the case of temporary roadblocks, time slot assignment fixations, restrictions on the available time slots for camps, or updates on the pilgrim numbers and time preferences. Each of these changes is made in close cooperation with the local security forces, external crowd management experts, and the pilgrim establishments. Finally, we incorporate feedback into our data set and calculate schedules by solving the integer problem P1 to optimality.

Figure 13. (Color online) Relationship of Density and Distance to the Jamarat Bridge



Notes. The box plots summarize simulated densities at resources on access paths. We assumed $V_0 = 1.1$ m/s and simulated the actual schedule of 2016. A box plot summarizes the densities of all resources within a certain distance interval (0–500 m, for example) over all time steps l of a given day. The horizontal position of the box indicates the average distance of all resources within the respective distance interval. A linear fit to the data additionally indicates the relationship of density and distance.


The distribution of timetables and route maps to the service offices is a time-critical operational challenge year by year. The Ministry of Hajj and the pilgrim establishments must distribute the individual scheduling reports, shown in Appendix D, to more than 800 entities within 48 hours. Within this time span, the respective group guides must be trained to apply the timetable and routing instructions. All relevant authorities can access localized versions of the schedules and routing instructions via a web application (Figures D.1–D.3).

An ongoing challenge in every Hajj season is schedule compliance. The guides, employed by the Ministry of Hajj and the service offices, are responsible for keeping up with the scheduled departure period and the assigned path of their group. At present, there exists no effective control system that verifies compliance with the schedule directly at the campsite.

7.2. Aggregated Camera Counting

The Jamarat Bridge is equipped with a video-based counting system (Johansson 2008). This system can

Figure 14. Timeline of Implementation



Notes. The timeline shows generalized and approximated timings of the planning process. Each tick mark corresponds to a week. During Ramadan, planning activities are significantly reduced. Establishments determine the service-offices from a pool of potential Mutawifs. The Mutawifs are assigned to pilgrim groups and camps in the tent city. This process is very dynamic, with frequent changes in the input data. Saudi authorities organize several workshops to discuss the state of the preparations with local and international experts. We process frequent ad hoc updates until the beginning of the Hajj.


provide aggregated flow statistics for the different floors of the building. Figure 15 shows the simulated flows based on the distributed schedule of 2016, video count data, and expected utilization based on preferences as provided by the pilgrim organizations.

Pilgrim groups without access to the metro system can state time preferences only within the off-peak intervals. As a result, stated preferences peak before

and after the unavailable peak time periods while the revealed preferences indicate that many pilgrims enter Jamarat during these hours. Most eye-catching is a peak on Day 10 at Ground Level, where the majority of the unregistered pilgrims perform stoning.

In response to the crowd density, security agents at Jamarat Plaza redirect flows to balance the utilization of Level 1 and Ground Level. The data shows that this

Figure 15. (Color online) Utilizations and Counting Data



Notes. Utilization of the Jamarat Bridge based on video counting data (*Counted*), stated preferences (*Preferences*), and simulated schedules. Each level has an operating capacity of 120,000 pilgrims per hour (i.e., 60,000 pilgrims per 30 minutes). We consider an operative capacity of 480,000 pilgrims per hour.

operational task was not successful in 2016 because the peak on Level 1 is at around 35,000 pilgrims per 30 minutes and the peak on the ground floor is at more than 100,000 pilgrims per 30 minutes. Security might have expected way more pilgrims to enter Level 1 because of the stated preferences and scheduled pilgrim groups. In the preceding Hajj season, Level 1 was more utilized than the ground floor, and crowd management might have redirected some flows based on experience.

Comparing the ground floor and Level 1, a mismatch becomes apparent between stated preferences and utilization based on the video counting system. Again, an operational redirection could be a reason for this pattern as well as data uncertainty—in both preference data and video count data. The video counting system, which has parts that have been in operation for more than 10 years, operates in a very challenging environment. The time-preferences data, on the other hand, may not cover 100% of the offices—especially when there were many last-minute changes in the assignments. As a result, we currently have very accurate preference data for many pilgrim offices that cooperate closely with the ministry while we have to assume estimated default preference values if the pilgrim data are less clear. An automated system to increase the coverage is currently under development and may be piloted in the coming Hajj seasons.

Only a few pilgrims are assigned to Level 2 because pilgrims enter it from the eastern access ways and not from the tent city. Here, crowd management mostly expects unregistered pilgrims.

Level 4 is almost exclusively used by metro ticket holders. Only a few other pilgrims approach Level 4 via the escalators embedded in the Jamarat Bridge. Thus, the metro train schedules have a significant influence on the aggregated flow pattern. The schedule compliance for that level is usually better than on the lower levels of Jamarat.

Level 3 shows a typical pattern of the preferred stoning times, with highest demand right before and after the peak periods. The counting data follows this pattern, except for Day 10.

The schedule, restricted by capacity bounding and smoothing, mirrors the stated stoning time preferences well. However, we recognize mismatches between the desired and actual behavior of scheduled and nonscheduled pilgrims as detected by the aggregated counting system. By anecdotal evidence, we have learned that assigned time periods may not always be accepted by a group's spiritual leader. In some cases, ex ante unknown logistical reasons (food supply times, shuttle bus schedules, etc.) might hinder group guides' dispatching the group in time. For those guides, it is very challenging to keep a big group of people mostly unknown to them together inside of a large pedestrian mass. At present, the actual work of the group guides is hardly surveyed. Thus, it is not clear whether each camp

has an adequate number of group guides for the dispatching operation.

To tackle these issues, we recommend a tailored dispatching controlling system to identify noncompliant camps and their reasons. Further, we suggest investigating the causes for the actual peak on the morning of Day 11. Maybe service-offices were not aware of the weather conditions at the time of the stating of stoning time preferences, hence demanding stoning times close to the peak periods. However, these periods exhibit high temperatures, and therefore pilgrims might have preferred to perform stoning in the very early morning before sunrise.

8. The 2015 Crowd Crush

On September 24, 2015, the first day of the stoning ritual, a severe crowd accident happened in the Mena Valley (Bucks 2015, Graham-Harrison et al. 2015, Almuhtar and Watkins 2016). According to media reports, two streams of pilgrims intersected at the crossing of streets 204 and 223 (Figure 16(a)). The two flows collided at around 9:00 a.m. Security forces rapidly stopped any additional inflow of pilgrims from street 204 by blocking the road with service vehicles, but the crush had already left many pilgrims dead. The pilgrims who approached the accident site through street 223 had evidently finished the stoning ritual before the accident and were on the way back to their camps. The pedestrian traffic encountered on street 204 was directed toward the Jamarat Bridge, where the pilgrims wanted to perform the stoning ritual.

The Saudi government established a committee to investigate the accident. The authors were not involved in the scheduling and routing process of Hajj season 2015 and have no insight into the confidential results of the investigation committee. However, it is known that the accident happened during a peak period. In previous years, no registered pilgrims were scheduled to perform the ritual during that period. According to our own routing concept described in this paper, street 204 would have been operated strictly as a one-way road—that is, a flow from street 223 in the opposite direction would not have been allowed. It is known that a majority of registered and nonregistered pilgrims prefer to perform the stoning rituals during the peak period between 6:00 a.m. and 10:30 a.m.; thus, the dense pedestrian flow on street 204 was not unexpected. In the authors' opinion, the main cause of the crowd disaster was likely the counter flow from street 223, which must eventually have stopped or considerably obstructed both flows. When the density reached critical levels, the victims collapsed under the pressure and the extreme weather conditions. However, it is unknown to us why the groups' paths crossed in the first place.

In 2015–2016, the Saudi authorities made several improvements to the pilgrim routes to avoid future accidents in this area. Figure 16(b) illustrates two

essential enhancements. First, street 206 is now a one-way route leading directly to the Jamarat Bridge—that is, no pilgrims from street 206 can switch to street 204 anymore. Second, the connection to street 223 is strictly blocked—at least during peak periods.

Despite all efforts, however, it may never be possible to eliminate all risks entirely at complex mass gatherings. Even though significant investments in the infrastructure were made and sophisticated information systems in and around Mecca were set up, human factors, such as the behavior of unregistered pilgrims and the compliance of registered pilgrims with their schedules and dedicated paths, are still difficult to control.

9. Final Remarks

In this paper, we introduce the Pilgrim Scheduling Problem to the operations research literature. We provide a comprehensive problem statement and propose an integer-programming model in which we integrate central results of the crowd research literature. In our model, we implement strict separation of flows, and coordinate access and egress flows by imposing flow level and volatility constraints. Our mesoscopic simulation module provides critical feedback to the planner, which we also use to reoptimize the schedules.


In the computational study, we have shown that our approach solves practical instances with more than 2.3 million variables in less than 10 minutes on average. At the same time, the maximum gap between optimal solution and upper bound never exceeded 0.28%. The study illustrates the central trade-off between the *safety* and *satisfaction* objectives in our model: the more emphasis a planner puts on scheduling smooth and

well-balanced pedestrian flows, the higher the violation of the pilgrims' stoning time preferences. We have demonstrated how the level and the volatility of infrastructure utilization in a solution can be controlled, and highlighted interaction effects between the scheduling parameters. Results presented for the simulation approach verify the safety of the schedule and the effectiveness of the scheduling parameters.

Our approach, among several other measures to improve safety, has been applied in the years 2007–2014 and 2016–2017. In these years, Hajj was a safe mass gathering. To external observers, it might be “[...] surprising [...] that the Hajj got through eight consecutive quiet years” (Benedictus 2015). But in fact, technological progress in crowd management has contributed to the “[...] achievement that most years pass without a major incident at the Hajj [...]” (Plackett 2016). The Hajj season of 2007 was a particularly remarkable success. Despite ongoing construction work at the Jamarat Bridge and significantly higher pilgrim numbers than in the disastrous previous year, the pilgrimage was free of crowd accidents. In that year, a scheduling approach based on mathematical optimization was applied for the first time in the history of Hajj. Today, Saudi authorities consider the scheduling and routing of pilgrims as “[...] the most important part of the operational programme [...]” (The Guardian 2016).

The pilgrimage is a yearly stress test of the Saudi authorities' ability to protect pilgrims' safety. Thousands of trained professionals, soldiers, firefighters, engineers, medical personnel, metro staff, pilgrim guides, and others are mobilized in each Hajj season. Over the recent decades, Hajj management has accumulated experience

Figure 16. (Color online) Accident Location in 2015 and Infrastructure Modifications for 2016



Accident area in 2015 at the intersection of street 223 and street 204.

Modifications of path system for 2016: street 223 is blocked and street 206 has been extended to the Jamarat Bridge.

in coping with most severe safety threats, including epidemic diseases, extreme weather conditions, dense pedestrian crowds, political tensions between participating groups, and even terrorist activities (possible attacks/provoked accidents). In the authors' opinion, it is key to preserve this knowledge for the future. In each area of crowd management, staff fluctuation and changes in competencies might lead to a loss of critical knowledge. This includes operational planning as well as infrastructure development. Any future infrastructure expansion must adhere to the established traffic flow concept.

Considering that pilgrim numbers are expected to rise significantly in the coming years (Kassens-Noor et al. 2015, Müller 2015, Al Arabiya 2017), it is crucial not only to preserve the progress already made, but also to expand the digital infrastructure of Hajj based on the presented concepts. The increasing use of mobile devices, for instance, introduces both new opportunities as well as subtle new risks for crowd management. As an example, generic smartphone navigation services are, to the best of our knowledge, currently not aware of the one-way flow system established to avoid dangerous counterflows. As a result, smartphone navigation and third-party Hajj guide applications using such services might mislead disoriented pilgrim groups into counterflows. On the other hand, the technology could

allow the authorities to broadcast official information to the pilgrims just in time.

We would like to make scheduling information available via mobile devices to all camps and connect the data to an effective real-time dispatching control system. Based on such a system, we further encourage the development of an incentive-based approach to increase compliance with the assigned dispatching times. Moreover, we intend to enhance our approach to allow real-time evaluation of dispatching postponement strategies and ad hoc redirection plans in the case of unacceptably high densities.

We hope that the presented approaches, embedded in a well-designed, comprehensive safety concept, will contribute to keeping the Hajj safe in the future.

Acknowledgments

Valuable suggestions from Habib Zein Al-Abideen, Salim Al-Bosta, Sonic Chan, and Dirk Serwill, as well as many Saudi Arabian experts, are gratefully acknowledged. Further, the authors are very grateful to Anders Johansson for providing pilgrim count data. Most of all, they thank Dirk Helbing (ETH Zurich) for his many helpful suggestions that improved the paper. Finally, the authors thank an associated editor and an anonymous reviewer for their very helpful comments on earlier versions of the paper.

Appendix A. Glossary

Block	Geographical subunit of an establishment
Camp	A camp contains tents (accommodating pilgrims); located in the Mena Valley
Capacity	Maximum number of pilgrims occupying a resource per period such that a safe flow of pilgrims is ensured
Dissatisfaction	Squared deviation of scheduled stoning period and preferred stoning period
Establishment	Organization responsible for transport and accommodation of pilgrims. Establishments are distinguished according to regions of the world
Group	Aggregation of 250 pilgrims, all affiliated to the same camp. Sometimes, the size of a group differs from 250 (i.e., the number of pilgrims in a camp is not always strictly a multiple of 250)
MOMRA	Ministry of Municipal and Rural Affairs, Kingdom of Saudi Arabia
Path	A route (access and egress) connecting a camp with (a certain level of) the Jamarat Bridge
Peak period	Periods when most of the pilgrims prefer to perform the stoning ritual
Period	30 minutes
Pilgrims	Individuals who perform religious rituals. We distinguish between registered and nonregistered pilgrims. Registered pilgrims use the service of establishments—i.e., the number of registered pilgrims is more or less known in advance. Concerning nonregistered pilgrims, there is no information regarding their total numbers, locations, etc.
Preferred stoning period	Period within an interval of preferred stoning periods provided by establishments
Resource	Infrastructure of limited capacity used by pilgrims (Jamarat Bridge and Shoaibain Tunnel, for example)
Service office	Organizational subdivision of establishments
Service office camp	A distinct combination of service office and a camp
Scheduling group	A distinct combination of a group and a day (10th day, 11th day, or 12th day)
Utilization	Relative capacity demand per period and resource

Appendix B. Time-Preference Allocation Example

This section contains an algorithm used to compute the preferred stoning times of scheduling groups from aggregate information about the preferred stoning times. Aggregate information may stem from the establishments, plots, or camps, for example. Here, we outline the procedure for plot information. Since we only obtain the preferred intervals of stoning periods for each plot, we generate the preferred stoning periods for each group. Consider the following example: On the 10th day, 40% of the pilgrims of a block may prefer to perform stoning from period 12 to period 32 (i.e., 3:00 a.m. to 8:00 a.m.), and 60% from period 24 to period 42 (i.e., 6:00 p.m. to 10:30 p.m.). The first interval covers 20 periods and the second 18 periods. Note that we have assumed that the percentage distribution of intervals of a service office camp is in accordance with the distribution of the corresponding plot. Now consider a service office camp of the plot that contains six groups. We assign two groups (~40%) to the first interval and four groups (~60%) to the second interval. Usually, it is not possible to perfectly break down the groups to the percentages. Therefore, we assign a group to an interval such that the resulting deviation is as small as possible. Further, the number of groups of a block assigned to a certain interval should be evenly distributed over the periods of the interval.

Now, consider the sets:

Π set of blocks, index $\pi \in \Pi$;

$\mathbb{O}$ set of service office camps, index $o \in \mathbb{O}$;

$\mathbb{O}_\pi$ set of service office camps assigned to block π ($\mathbb{O}_\pi \subset \mathbb{O}$);
and

$\mathcal{G}_{om}$ set of pilgrim groups of service office camp o for day m ($\mathcal{G}_{om} \subset \mathcal{G}$).

Then, let

$\beta_{\pi mt}$ be the fraction of pilgrims of block π who prefer stoning up to period t on day m ,

$\tilde{\beta}_{omt}$ be the fraction of pilgrims of service office camp o who prefer stoning up to period t on day m . If service office camp o belongs to block π , then $\tilde{\beta}_{omt} = \beta_{\pi mt}$,

γ_π the number of pilgrim groups in block π , and

$\tilde{\gamma}_o$ the number of pilgrim groups in service office camp o .

By Algorithm 3 we obtain for each group g and each day m the preferred stoning period t_{gm}^* . The results of an example with five office camps, 10 periods, two time intervals, and 30 groups are summarized in Table B.1.

Algorithm 3 (Calculate Preferred Stoning Period t_{gm}^* for Each Group g and Day m)

for all $m \in \mathcal{M}$ **do**

for all $\pi \in \Pi$ **do**

$j_1 = 1$

for all $o \in \mathbb{O}_\pi$ **do**

$t \leftarrow 1$

$i \leftarrow 1$

for all $g \in \mathcal{G}_{om}$ **do**

while ($j_t < \text{round}(\beta_{\pi mt} \gamma_\pi)$) or ($i < \text{round}(\tilde{\beta}_{omt} \tilde{\gamma}_o)$) **do**

$t \leftarrow t + 1$

$j_t \leftarrow j_t + 1$

$i \leftarrow i + 1$

$t_{gm}^* \leftarrow t$.

Some establishments provide only a few preferred stoning periods per block; others provide very detailed information. The results of the scheduling problem and the simulation are presented to the establishments. Usually, the establishments then reconsider the preferred stoning times.

Table B.1. Stoning Time Preference Calculation: Example


	60%				40%					
	$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$	$t = 6$	$t = 7$	$t = 8$	$t = 9$	$t = 10$
Service office camp										
$o = 1$	2	1	2	1	1		1	1		1
$o = 2$	2	1	2	1	1		1	1		1
$o = 3$	1	1		1		1			1	
$o = 4$		1		1				1		
$o = 5$				1			1			
Target	5	4	4	5	2	2	2	2	2	2
Actual	5	4	4	5	2	1	3	3	1	2

Notes. Example of the derivation of preferred stoning periods from preferred stoning intervals. Each cell displays the quantity of groups of a service office camp (row) to be best scheduled in a given period (column). The first interval covers the periods $t \in \{1, \dots, 4\}$ and the second the periods $t \in \{5, \dots, 10\}$. Overall, 60% of the groups prefer to be scheduled in the first time interval and 40% in the second time interval. In each row of a service office camp, we see the numbers of groups to be best scheduled in a period. For example, two groups of service office camp $o = 2$ should be scheduled in period $t = 3$.

Appendix C. Reoptimization


Maximum simulated densities for three scenarios on Day 11. If critical densities are discovered in the simulation (see panels (a) and (b) of Figure C.1), we incorporate respective segments into the scheduling model and impose capacity constraints. Results in panel (c) of Figure C.1 indicate no critical densities after reoptimization.

Figure C.1. (Color online) Reoptimization by Increasing the Set of Scheduling Resources



Appendix D. Hajj Information System

Figure D.1. (Color online) Hajj Schedule Information System



Notes. Download requests logged for Hajj season 2016 (1437H). The eighth Day of Dhu al-Hijjah 1437 corresponds to September 9, 2016. The first day of the stoning ritual in 2016 was on September 11. Establishments, service-offices, and other authorities download maps and timetables for distribution to the guides.

Figure D.2. (Color online) Excerpt of a Camp's Timetable for Distribution to the Guides, English and Arabic Version





Figure D.3. (Color online) Map Handed Out to Support Group Leaders

Note. The concrete route is depicted for a given camp.

Endnotes

¹Newspapers report of more than 180,000 people without Hajj permits who were denied entry at several checkpoints around Mecca during the Hajj of 2016 (Al Arabiya 2016).

²Helbing (1992) demonstrated that pedestrian movements could be optimized by reducing variances in velocity. The larger the velocity variance of the pedestrians, the larger the space demanded by each one. As an illustrative example, consider dancers on a dance floor, who occupy more space than surrounding spectators. Analysis of pedestrian flows in the Hajj season of 2006 confirms these observations: stop-and-go flows can be avoided in a system featuring low variance in velocity (Helbing et al. 2007).

References

- Abdelghany A, Abdelghany K, Mahmassani H (2016) A hybrid simulation-assignment modeling framework for crowd dynamics in large-scale pedestrian facilities. *Transportation Res. Part A: Policy Practice* 86:159–176.
- Abdelghany A, Abdelghany K, Mahmassani H, Alhalabi W (2014) Modeling framework for optimal evacuation of large-scale crowded pedestrian facilities. *Eur. J. Oper. Res.* 237(3): 1105–1118.
- Al Arabiya (2016) Illegal Hajj pilgrims to be tried, sentenced within 24 hours. (September 7), <http://english.alarabiya.net/en/News/middle-east/2016/09/07/Illegal-hajj-pilgrims-to-be-tried-sentenced-within-24-hours.html>.
- Al Arabiya (2017) Saudi king approves Hajj capacity increase. (January 5), <http://english.alarabiya.net/en/News/gulf/2017/01/05/Saudi-King-approves-increasing-Hajj-capacity.html>.
- Almukhtar S, Watkins D (2016) How one of the deadliest Hajj accidents unfolded. *New York Times* (September 6), <http://www.nytimes.com/interactive/2016/09/06/world/middleeast/2015-hajj-stampede.html>.
- Baxter E (2010) 15 percent increase in Saudi pilgrims in 2010. *Arabian Bus.* (March 4), <http://www.arabianbusiness.com/582928-15-increase-in-saudi-pilgrims-in-2010>.
- Benedictus L (2015) Hajj crush: How crowd disasters happen, and how they can be avoided. *The Guardian* (December 23), <https://www.theguardian.com/world/2015/oct/03/hajj-crush-how-crowd-disasters-happen-and-how-they-can-be-avoided>.
- Bretschneider S, Kimms A (2011) A basic mathematical model for evacuation problems in urban areas. *Transportation Res. Part A: Policy Practice* 45(6):523–539.
- Browning N (2015) Haj death toll rises to 769, Iran denounces ‘crime.’ *Reuters* (September 26), <http://www.reuters.com/article/us-saudi-haj-toll-idUSKCN0RQ0FB20150926>.
- Bucks J (2015) The Hajj crush: It was the closest thing to hell on earth. *The Guardian* (December 23), <https://www.theguardian.com/news/2015/dec/23/hajj-crush-pilgrimage-mecca-stampede-saudi-arabia-mina-valley>.
- Cardoen B, Demeulemeester E, Belien J (2010) Operating room planning and scheduling: A literature review. *Eur. J. Oper. Res.* 201(3):921–932.

- Cayirli T, Veral E (2003) Outpatient scheduling in health care: A review of literature. *Production Oper. Management* 12(4):519–549.
- Charnes A, Duffuaa S, Yafi A (1989) A non-linear congestion network model for planning internal movement in the Hajj. *Eur. J. Oper. Res.* 40(3):292–298.
- Cheng L, Yarlagadda R, Fookes C, Yarlagadda PKDV (2014) A review of pedestrian group dynamics and methodologies in modelling pedestrian group behaviours. *World* 1(1) 002–013.
- Choi W, Hamacher H, Tufekci S (1988) Modeling of building evacuation problems by network flows with side constraints. *Eur. J. Oper. Res.* 35(1):98–110.
- Church RL, Cova TJ (2000) Mapping evacuation risk on transportation networks using a spatial optimization model. *Transportation Res. Part C: Emerging Tech.* 8(1–6):321–336.
- Coscia V, Canavesio C (2008) First-order macroscopic modelling of human crowd dynamics. *Math. Models Methods Appl. Sci.* 18(supp01):1217–1247.
- Daganzo CF (1995) The cell transmission model, part II: Network traffic. *Transportation Res. Part B: Methodological* 29(2):79–93.
- 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.
- Feng L, Miller-Hooks E (2014) A network optimization-based approach for crowd management in large public gatherings. *Transportation Res. Part C: Emerging Tech.* 42:182–199.
- Fruin JJ (1971) Pedestrian planning and design. Technical report, Metropolitan Association of Urban Designers and Environmental Planners, New York.
- Gambrell J, Ahmed B (2015) Hajj stampede in September killed over 2,400, new count finds. *New York Times* (December 10) <http://www.nytimes.com/aponline/2015/12/10/world/middleeast/ap-ml-saudi-hajj.html>.
- GAMS, Development Corporation (2016) General Algebraic Modeling System (GAMS), release 24.7.1. GAMS Development Corporation, Washington, DC. <http://www.gams.com>.
- Graham-Harrison E, Dehghan SK, Salih ZM (2015) Hajj pilgrimage: More than 700 dead in crush near Mecca. *The Guardian* (September 24), <https://www.theguardian.com/world/2015/sep/24/mecca-crush-during-hajj-kills-at-least-100-saudi-state-tv>.
- Guardian, The (2016) Saudi Arabia tightens up hajj planning to avoid repeat of 2015 disaster. (September 4), <https://www.theguardian.com/world/2016/sep/04/saudi-arabia-tightens-up-hajj-planning-avoid-repeat-2015-disaster>.
- Haase K, Al Abideen H, Al-Bosta S, Kasper M, Koch M, Müller S, Helbing D (2016) Improving pilgrims safety during Hajj: An analytical and operational research approach. *Interfaces* 46(1): 74–90.
- Hamacher H, Tjandra S (2001) Mathematical modelling of evacuation problems: A state of art. Technical Report 24, Fraunhofer Institut Techno- und Wirtschaftsmathematik, Kaiserslautern, Germany.
- Helbing D (1992) A fluid-dynamic model for the movement of pedestrians. *Complex Systems* 6:391–415.
- Helbing D, Johansson A (2009) Pedestrian, crowd and evacuation dynamics. Meyers R, ed. *Encyclopedia of Complexity and Systems Science* (Springer, New York), 6476–6495.
- Helbing D, Johansson A, Al-Abideen HZ (2007) Dynamics of crowd disasters: An empirical study. *Phys. Rev. E* 75(4):046109.
- Helbing D, Farkas I, Molnar P, Vicsek T (2002) Simulation of pedestrian crowds in normal and evacuation situations. Schreckenberg M, Sharma S, eds. *Pedestrian and Evacuation Dynamics* (Springer, Berlin).
- Helbing D, Brockmann D, Chadeaux T, Donnay K, Blanke U, Woolley-Meza O, Moussaid M, et al. (2015) Saving human lives: What complexity science and information systems can contribute. *J. Statist. Phys.* 158(3):735–781.
- Johansson A (2008) Data-driven modeling of pedestrian crowds. PhD thesis, Faculty of Traffic Sciences, Technische Universität, Dresden, Germany. <http://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa-20900>.
- Johansson A, Helbing D, Shukla P (2007) Specification of the social force pedestrian model by evolutionary adjustment to video tracking data. *Adv. Complex Systems* 10(suppl 2):271–288.
- Jou R-C, Kitamura R, Weng M-C, Chen C-C (2008) Dynamic commuter departure time choice under uncertainty. *Transportation Res. Part A: Policy Practice* 42(5):774–783.
- Kassens-Noor E, Wilson M, Müller S, Maharaj B, Huntoon L (2015) Towards a mega-event legacy framework. *Leisure Stud.* 34(6): 665–671.
- Li Z, Tirachini A, Hensher DA (2012) Embedding risk attitudes in a scheduling model: Application to the study of commuting departure time. *Transportation Sci.* 46(2):170–188.
- Müller S (2015) Spaces of rites and locations of risk: The great pilgrimage to Mecca. Brunn SD, ed. *The Changing World Religion Map: Sacred Places, Identities, Practices and Politics* (Springer, Dordrecht, Netherlands), 841–853.
- Nemhauser GL, Trick MA (1998) Scheduling a major college basketball conference. *Oper. Res.* 46(1):1–8.
- Plackett B (2016) To stone the devil, safely. *Inside Science* (September 8), <https://www.insidescience.org/news/stone-devil-safely>.
- Saleh W, Farrell S (2005) Implications of congestion charging for departure time choice: Work and non-work schedule flexibility. *Transportation Res. Part A: Policy Practice* 39(7–9):773–791.
- Senbil M, Kitamura R (2004) Reference points in commuter departure time choice: A prospect theoretic test of alternative decision frames. *J. Intelligent Transportation Systems* 8(1):19–31.
- Stepanov A, Smith JMG (2009) Multi-objective evacuation routing in transportation networks. *Eur. J. Oper. Res.* 198(2):435–446.
- Still K (2013) Fruin - levels of service. Online course support, Crowd Risk Analysis and Crowd Safety, May 29, <http://www.gkstill.com/Support/crowd-flow/fruin/LoS.html>.
- Treiber M, Kersting A (2012) *Traffic Flow Dynamics: Data, Models and Simulation* (Springer, Berlin).
- Van den Bergh J, Belien J, De Bruecker P, Demeulemeester E, De Boeck L (2013) Personnel scheduling: A literature review. *Eur. J. Oper. Res.* 226(3):367–385.
- Verbas Ö, Abdelghany A, Mahmassani H, Elfar A (2016) Integrated optimization and simulation framework for large-scale crowd management application. *Transportation Res. Record: J. Transportation Res. Board* 2560(1):57–66. <https://doi.org/10.3141/2560-07>.
- Yamaguchi K, Berg AC, Ortiz LE, Berg TL (2011) Who are you with and where are you going? 2011 IEEE Conf. Comput. Vision Pattern Recognition (CVPR 2011) (IEEE, Washington, DC), 1345–1352.

Knut Haase is the director of the Institute for Transportation at the University of Hamburg, Hamburg, Germany. He received his diploma and doctoral degree and his lecture qualification in business administration at the University of Kiel, Kiel, Germany. His research topics are focused on optimization approaches for solving large-scale problems with applications in logistics and public transport.

Mathias Kasper is a research and teaching assistant at the Dresden University of Technology. He holds a master's degree in transport economics from University of Technology Dresden. His research topics are focused on crowd dynamics and optimization approaches for solving large-scale problems.

Matthes Koch studied economics of transport and logistics at the Dresden University of Technology. His research includes work on simultaneous vehicle and crew scheduling

approaches for transport companies. He is a PhD student at the University of Hamburg and works on the technical integration of scheduling, controlling, and data management for the Hajj crowd management system.

Sven Müller is professor of transport business economics at the Institute for Transport and Infrastructure at the University

of Applied Sciences Karlsruhe, Karlsruhe, Germany. He holds a PhD in transport economics and logistics from the University of Technology Dresden, Dresden, Germany. His research focuses on the integration of predictive and prescriptive analytics with applications to transport, logistics, energy, retail, and healthcare.