Öhman, Mikael; Hiltunen, Markus; Virtanen, Kai; Holmström, Jan

Frontlog scheduling in aircraft line maintenance

Published in:
Journal of Operations Management

DOI:
10.1002/joom.1108

Published: 01/03/2021

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

Received: 29 October 2018 | Revised: 28 May 2020 | Accepted: 15 June 2020
DOI: 10.1002/joom.1108

RESEARCH ARTICLE

WILEY

Frontlog scheduling in aircraft line maintenance: From explorative solution design to theoretical insight into buffer management

Mikael Öhman[1] | Markus Hiltunen[2] | Kai Virtanen[2] | Jan Holmström[1]

[1]Department of Industrial Engineering and Management, School of Science, Aalto University, Aalto, Finland

[2]Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, School of Science, Aalto University, Aalto, Finland

Correspondence
Jan Holmström, Department of Industrial Engineering and Management, School of Science, Aalto University, POB15500, 00076 Aalto, Finland.
Email: jan.holmstrom@aalto.fi

Funding information
Academy of Finland, Grant/Award Number: Direct operations/ 323831; Tekes

Handling Editor: Aravind Chandrasekaran

Abstract

Based on a multi-year research engagement with practice, we present a novel solution design for *frontlog* scheduling in aircraft line maintenance and offer theoretical insights into buffer management in operations. The field problem of the case airline was to improve departure reliability for long-haul aircraft without increasing maintenance resources, and without using backup aircraft. Frontlog scheduling is the purposeful introduction of over-maintenance as a buffer of maintenance tasks that can be opportunistically postponed. A detailed simulation of the solution introduced in the airline's operations indicates a performance frontier shift, concurrently improving departure reliability, and reducing maintenance cost. We position the novel practice in operations and maintenance management literature, arguing that the frontlog creates a new type of time buffer, available in contexts where capacity serves predictable as well as unpredictable demand. Further theoretical elaboration leads us to reconceptualize buffer management along time and capacity dimensions, reducing inventory to a special case of time buffering.

KEYWORDS

aircraft line maintenance, buffer management, design science, frontlog scheduling, performance frontier, simulation

1 | INTRODUCTION

Basic research is the source of novelty and innovation in the established, linear approach to research and development (Stokes, 2011). However, researchers in operations management (OM) can potentially reverse this linear process of creating knowledge through engaging with practice in design exploration (Holmström, Ketokivi, & Hameri, 2009) and interventions (Oliva, 2019). In Pasteur's quadrant of knowledge creation (Stokes, 2011), these types of engagements reveal theoretically surprising outcomes, spurring researchers to develop theory and establish new research directions. In this article, we describe an explorative design science research engagement, which resulted in the development of a new practice for aircraft line maintenance, *frontlog* scheduling. Developing the frontlog transcended practical problem solving, resulting in novel insights into buffer management at the core of our theoretical understanding of just-in-time (JIT), lean, and service operations. These insights provide a foundation for future research on responsive and dynamic scheduling (cf. Biçer & Seifert, 2017), enabled by digitalization of operations (Holmström, Holweg, Lawson, Pil, & Wagner, 2019).

Frontlog is a maintenance scheduling practice that creates opportunities for dynamic rescheduling of planned tasks to free up capacity for unplanned tasks. This rescheduling option is especially attractive in aircraft line maintenance, where unplanned tasks caused by random technical failures may lead to (costly) grounding of the aircraft. To illustrate the practice, consider an aircraft that arrives at a hub location for scheduled line maintenance. There will be several recurring planned maintenance tasks, each with a due date based on when it was last performed. Current scheduling practice typically seeks to maximize maintenance intervals (e.g., Başdere & Bilge, 2014; Sarac, Batta, & Rump, 2006). Thus, this will be the last maintenance opportunity to perform most of the planned tasks. Any delay will result in a backlog that grounds the aircraft and disrupts airline operations. However, with a frontlog, there will also be tasks that have a due date that allows postponement to a future maintenance opportunity, without overshooting the due dates and grounding the aircraft. Should the aircraft arrive with a technical fault, the frontlog serves as a buffer of tasks that can be rescheduled, freeing up maintenance capacity for the emergent demand. In our design, the share of planned tasks that can be rescheduled without violating due dates becomes a decision variable for OM, making it an alternative to maintaining slack capacity or delaying the departure of the aircraft (backlog). However, in aircraft line maintenance, the frontlog comes with a cost: The due dates of planned tasks depend on when they were last performed. Thus, performing tasks before the last opportunity to do so results in over-maintenance, which increases the total planned workload.

In this study, to investigate the implications of the frontlog for practice and theory, we combine explorative design science (Holmström et al., 2009) and empirically grounded analytics (Chandrasekaran, Linderman, & Sting, 2018). For practice, the question is: What are the trade-offs among over-maintenance, departure reliability, and slack capacity when frontlog scheduling is implemented? For theory, the questions are: What is a frontlog? In what ways is it novel, and what are the possible theoretical implications beyond the context of aircraft line maintenance? We approached these questions by studying the balancing of different types of buffers as an explicit management decision in a Nordic airline's line maintenance organization. We were initially invited into the maintenance organization to help improve the departure reliability of the long-haul fleet. Over the course of the initial problem-solving project, we became aware of the link between poor departure reliability and the maintenance organization's practices for coping with unplanned repairs. This awareness prompted further investigation of how the organization created and managed buffers against demand variability. As we questioned the current practice of maximizing the maintenance interval, an unconventional proposal took form (the frontlog) which entailed using deliberate over-maintenance as a buffer against demand uncertainty. To investigate the proposed practice, we constructed an empirically grounded discrete event simulation model. Through the model, we explored the trade-off between over-maintenance introduced by the frontlog and its leveling effect on the workload. We found an opportunity for a significant improvement in both cost efficiency and departure reliability, implying a performance frontier shift. Nordic Airline (a pseudonym for the case company) is currently implementing the frontlog as part of a wider effort to digitalize operations, uncovering new design challenges in making the frontlog buffer visible to production, and closing the digital feedback loop from production to planning.

In Section 2, we position the frontlog in the established body of knowledge. In Section 3, we present the research process combining exploratory design science and empirically grounded modeling. In Sections 4 and 5, we detail our engagement with practice. We focus on how the design emerged and how we constructed the empirically grounded simulation model to test the design. In Section 6, we report the four-stage evaluation process. We also describe how Nordic Airline reacted to the design and is proceeding toward implementing it. Finally, based on our theoretical positioning of the frontlog, in Section 7, we elaborate on the significance of the concept to theory, along with implications for practice and research methodology.

2 | LITERATURE REVIEW: POSITIONING AND CONCEPTUALIZING FRONTLOG

The increasing digitalization of operations is prompting researchers to explore dynamic planning and responsive rescheduling as feasible alternatives for responding to demand uncertainty (cf. Biçer & Seifert, 2017). However, we could not find previous research that described a practice similar to frontlog scheduling. Consequently, we build our theoretical positioning and conceptualization on several related practices and concepts. Table 1 provides an overview of the frontlog from different research perspectives and identifies related practices and concepts. In the literature review, we discuss the frontlog from each perspective, starting with a review of maintenance literature, followed by buffer management, and ending

TABLE 1 What is a frontlog? Perspectives, related practices, and concepts

Research perspective	Frontlog	Related practices and concepts
Maintenance	*Opportunity creating maintenance*: Timing of planned tasks in aircraft maintenance creates opportunities for later rescheduling.	Maintenance timing; opportunity seizing maintenance; maintenance strategy
Buffer management	*Time buffer*: Manage the trade-off between timing of work and capacity flexibility in aircraft line maintenance	Flexibility; types of buffers; buffer trade-offs
Scheduling	*Dynamic scheduling*: Schedule over-maintenance in preparation for later rescheduling to cope with unplanned tasks	Uncertainty; disruption; rescheduling; digitalization

with scheduling. We conceptualize the frontlog as the practice of scheduling over-maintenance in preparation for later rescheduling to cope with unplanned additional workload. Thus, the frontlog can be positioned as part of the uncertainty-focused and dynamic research stream in scheduling. However, a frontlog is more than a scheduling practice. It introduces a new type of time buffer in airline maintenance operations. This novel time buffer leads to theoretical insights, which we elaborate in the discussion.

2.1 | Frontlog as maintenance

When we introduced the frontlog, we conceptualized it as a maintenance scheduling practice that creates opportunities for rescheduling. To position the frontlog as a potentially novel practice in the context of maintenance, we review research on maintenance timing, which is a decision variable in operations research (OR) and OM literature. Maintenance timing forms the basis for most taxonomies of maintenance practices (cf. Kelly, 2006; Khazraei & Deuse, 2011; Swanson, 2001). It is typically referred to as policy, highlighting that timing is a maintenance system design decision (Kelly, 2006). The baseline question every maintenance system designer asks is: Can this piece of equipment be allowed to fail? (Waeyenbergh & Pintelon, 2002). The answer distinguishes reactive (run-to-failure) maintenance policies from preventive (fix-before-failure) maintenance policies. Preventive maintenance can be divided into three types of policies: predetermined, predictive, and proactive (Khazraei & Deuse, 2011).

A failure can be prevented only if it is anticipated, which is reflected in how preventive policies are defined—based on how they anticipate failure. Predetermined policies (Khazraei & Deuse, 2011) include age-, time-, and use-based maintenance, in which maintenance timing is based on respective proxies of the equipment's condition. These maintenance policies result in periodic maintenance (Blakeley, Argüello, Cao, Hall, & Knolmajer, 2003; Grigoriev, van de Klundert, & Spieksma, 2006): The periods or intervals are fixed and predetermined by equipment manufacturers or regulations (Öhman, Finne, & Holmström, 2015). The most common predictive policy (Khazraei & Deuse, 2011) is condition-based maintenance. Maintenance timing is based on a variable that is directly related to the deterioration of the equipment's condition (Jardine, Lin, & Banjevic, 2006; Lee et al., 2014; Veldman, Wortmann, & Klingenberg, 2011). Proactive maintenance denotes the situation when the root cause of the failure can be removed through redesign (Khazraei & Deuse, 2011), removing the need for maintenance altogether.

Frontlog maintenance prepares for dealing with unanticipated failures in a predominantly predetermined maintenance system. However, is it a new type of maintenance policy? Having argued that timing is the key question in maintenance, and that maintenance research views timing in relation to expected technical system failure, we focus on the exception found in literature. A maintenance policy increasingly mentioned alongside other policies, which has evaded broader academic attention, is opportunity or opportunistic maintenance (Ab-Samat & Kamaruddin, 2014). This policy is, in our line of argumentation, fundamentally different from the ones noted above. The timing decision is not driven by the equipment's condition but by the operational circumstances of the technical system. In the words of Ab-Samat and Kamaruddin (2014, p. 99): "With this policy, maintenance is to be performed on a given part, at a given time, depending on the state of the rest of the system." The *opportunity* in opportunity maintenance denotes technical system downtime and is caused by an external (to maintenance) factor, such as technical failure (Sherwin, 1999) or a production stoppage (Ab-Samat & Kamaruddin, 2014). During the opportunity, preventive maintenance tasks that require downtime are performed, removing the need for future downtime caused by the tasks performed in the present. A close relative of opportunity maintenance is the block replacement policy (Dekker & Smeitink, 1991). Under this policy, maintenance tasks with the same (laborious) setup process are bundled, in effect implying that the timing decision for maintenance of all parts is determined by the failed or soon-to-fail part.

Despite a focus on availability rather than reliability, previous research on opportunity maintenance[1] tended to model the operational system as having a constant cost (cf. Dekker & Dijkstra, 1992; Dekker & Smeitink, 1994, 1991; Rust, 1987) and/or utility (cf. Radner and Jorgenson's (1963, p. 73). In contrast, frontlog maintenance builds on the idea that the availability of the technical system might be more valuable (as part of a wider operational system) at certain points of time. This idea questions the (sole) objective of maximizing availability. Furthermore, frontlog maintenance differs from opportunity maintenance in that it is not about seizing opportunities, but about creating opportunities for later use. Thus, the maintenance timing *creates opportunity*, in contrast to *seizing opportunity* (cf. Ab-Samat & Kamaruddin, 2014), which is found in the literature. We found a similar opportunity-creating practice in two healthcare papers (Bowers & Mould, 2002; Helm, AhmadBeygi, & van Oyen, 2011). In the first healthcare paper, elective patients were deferred on short notice to reduce waiting time for unplanned nonelective patients. In the second, the problem of meeting unplanned demand was addressed by introducing expedited scheduled demand, shortening the waiting time for more urgent elective patients.

Maintenance strategies (Waeyenbergh & Pintelon, 2002) relate maintenance operations to the overall business objectives (Pinjala, Pintelon, & Vereecke, 2006). Well-known examples of maintenance strategies are reliability-centered maintenance (RCM), which aims for uninterrupted operations while reducing maintenance cost (Nowlan & Heap, 1978; Rausand, 1998), and total productive maintenance (TPM) (Ahuja & Khamba, 2008; McKone, Schroeder, & Cua, 1999), which aims to improve productivity as part of lean manufacturing (McKone, Schroeder, & Cua, 2001). Examples of more recent maintenance strategies are business centered maintenance (Kelly, 2006), output-based maintenance (Ahmad & Kamaruddin, 2013), and value driven maintenance (Rosqvist, Laakso, & Reunanen, 2009; Stenström, Parida, Kumar, & Galar, 2013), with their respective objectives reflected in their names. These maintenance strategies consider maintenance timing as an instrument for manipulating the condition of the technical system in accordance with their respective business objectives. Thus, frontlog is not a new maintenance strategy, but a new timing practice among a bundle of maintenance practices (cf. Shah & Ward, 2003) that the maintenance organization may use when implementing its maintenance strategy.

2.2 | Frontlog as buffer management

As a means of creating opportunities for responding to future maintenance demand, frontlog can also be conceptualized as a time buffer, complementing capacity buffering in aircraft line maintenance. We approach the concept of buffers in operations through the overarching concept of flexibility (Cousens, Szwejczewski, & Sweeney, 2009; Gerwin, 1993). Buffering is a built-in characteristic of the operational system design, distinct from the related concepts of agility and responsiveness (Bernardes & Hanna, 2009). The variability and demand uncertainty that unplanned repairs cause can be seen as natural (Hopp & Spearman, 2004), implying that they “cannot, or only to a limited extent, be influenced or controlled” (Roemeling, Land, & Ahaus, 2017, p. 1231). Production systems cope with variability either through inventory buffers, which stabilize (Bernardes & Hanna, 2009) and protect (Hopp & Spearman, 2004) the production system, or through flexible production resources (D'Souza & Williams, 2000; Zhang, Vonderembse, & Lim, 2003).

In manufacturing, there are three types of buffers: inventory, capacity, and time (Hopp & Spearman, 2008; Newman, Hanna, & Maffei, 1993). These buffers differ, meaning that buffer management practices and the research related to them are generally discussed in terms of one type, or at most two types, of buffers. In manufacturing research, inventory management practices (cf. Williams & Tokar, 2012) have attracted the lion's share of scholarly interest, with the oldest (still relevant) models originating in the early 20th century (Hopp & Spearman, 2008). In contrast, research on capacity management practices is predominant in the context of services, where the characteristics of services (cf. Moeller, 2010) make inventory buffers inapplicable (Akkermans & Voss, 2013). Without inventories, natural demand variability, to the extent it cannot be contained through demand management (cf. Klassen & Rohleder, 2002), is absorbed either by a capacity buffer or a backlog (time, according to Hopp & Spearman, 2008), where demand waits for available capacity (Akkermans & Vos, 2003). In manufacturing, the time buffer is typically discussed in terms of lead time (Caputo, 1996) and order books (e.g., Hedenstierna et al., 2019). With a few exceptions (cf. Ray & Jewkes, 2003; Zijm & Buitenhek, 1996), time tends to figure as the dependent variable, or as a constraint in inventory and capacity decisions, in manufacturing as well as in services (Anderson, Morrice, & Lundeen, 2005). Despite the name, the substantial literature on time-based manufacturing (TBM) is no exception, as it is contextualized in a fast and responsive manufacturing system (Rondeau, Vonderembse, & Ragu-Nathan, 2000) with associated practices (cf. Koufteros, Vonderembse, & Doll, 1998; Tu, Vonderembse, Ragu-Nathan, & Sharkey, 2006). In such a system time is seen as something that must be minimized, rather than being a buffer to be managed.

Research on buffer management typically concentrates on finding the optimal size and location of one type of buffer, such as inventory (cf. Askin & Krishnan, 2009; Mertins & Lewandrowski, 1999; Yang, Hsieh, & Cheng, 2011). A sub-stream of research seeks to balance two buffers with respect to a third. An interesting aspect of this stream is that the third buffer is typically portrayed as rigid (Rappold & Yoho, 2008), "given" (Betts & Johnston, 2005), or inapplicable (Roemeling et al., 2017). Such proclamations are no doubt well-grounded in their contexts. For example, in healthcare, inventories are often deemed inapplicable, "since patients are themselves transformed in the healthcare process, it is impossible to stock the transformed resource, the patients" (Roemeling et al., 2017, p. 4803).

Research on aircraft line maintenance focuses on capacity buffers. In coping with uncertainty while pursuing effectiveness, airlines combine robustness of planning (Ahmed & Poojari, 2008), and responsiveness of execution (Callewaert, Verhagen, & Curran, 2018). Buffer management is almost exclusively viewed as a resource allocation or capacity problem. Further, buffer management in this context is characterized as challenging due to the inherent complexity and criticality (implying high costs for a backlog buffer) of aircraft maintenance (Gupta & Lulli, 2014). As exceptions, we note Dijkstra, Kroon, Salomon, van Nunen, and Van Wassenhove (1994), who mentioned workload smoothing as part of coping with demand variability, but did not identify any related practices. Beliën, Cardoen, and Demeulemeester (2012) considered workload smoothing, but only within a maintenance window (usually defined as the time between an aircraft's arrival and departure). Thus, to the best of our knowledge, the (over-) reliance on capacity buffering has not been questioned in aircraft line maintenance research. Arguments for relying on capacity buffering boil down to maximizing aircraft availability (and by extension, the potential for revenue generation), advocating just-in-time maintenance. A good example can be found in the work of Boeret (1977), who described a simulation model with a scheduling heuristic to alleviate the maintenance planning work (which, at the time, was cumbersome and time-consuming) to reduce the number of *flying hours lost* because of maintenance that was performed before it was due.

To conclude, positioning the frontlog as a buffer and a scheduling practice to manage the trade-offs between time and capacity in aircraft line maintenance is novel. It responds to largely unheeded calls for OM research that explores the interrelationship between different types of buffers (Hopp & Spearman, 2004; Thürer, Tomašević, & Stevenson, 2017).

2.3 | Frontlog as scheduling

For an in-depth understanding of the role of a frontlog in managing operational buffer trade-offs, we turn to scheduling. From the perspective of scheduling, we conceptualize a frontlog as intentional scheduling of over-maintenance (slack) in preparation for later dynamic rescheduling for coping with an unplanned additional workload. Scheduling has been defined as "allocating a set of resources over time to perform a set of tasks" (McKay & Wiers, 1999, p. 242). This makes scheduling the nexus of buffer management. Scheduling is also the activity through which shifts in buffer trade-offs are operationally achieved, in manufacturing (Berglund & Karltun, 2007) and services (cf. Klassen & Rohleder, 1996; Laganga, 2011). Most scheduling research ignores uncertainty, typically in favor of mathematical tractability (McKay & Wiers, 1999). However, a stream of scheduling research focuses explicitly on uncertainty, and how it can be mitigated (Akkan, 2015; Black & McKay, 2012). With a few exceptions (cf. Biçer & Seifert, 2017), recent research in this stream has focused on disruptions (Akkan, 2015) or disturbances (Cho & Lazaro, 2010). Previous contributions distinguished between internal and external uncertainties (cf. Bean, Birge, Mittenthal, & Noon, 1991; Mehta & Uzsoy, 1998; Smith, Ow, Potvin, Muscettola, & Matthys, 1990).

In the manufacturing stream of research, as in the context where frontlog was developed, external uncertainty is manifested as an unplanned additional workload—additional to the planned workload of the given (and finite) capacity. Two distinct research themes for coping with such uncertainty have emerged: development of robust schedules, which are less sensitive to disruptions, and development of approaches for rescheduling in response to disruptions (Akkan, 2015, p. 199). Robust schedules typically involve introducing slack, in terms of capacity (Bourland & Yano, 1994) and time (Lu, Cui, & Han, 2015; Siedlak, Pinon, Robertson, & Mavris, 2018), through which the schedule can recover from disruptions (Bean et al., 1991). However, research on how to determine the amount of slack and how it should be introduced in schedules remains scarce (Akkan, 2015). Further, we note recent developments in dynamic scheduling, for example, in response to an evolving demand forecast (Biçer & Seifert, 2017). This stream of research sees scheduling as a continuous (automated) activity, which is done whenever relevant new information becomes available, raising its prospects in the increasing digitalization of operations (cf. Holmström et al., 2019).

Maintenance is challenging from a scheduling perspective, due to inherent uncertainties and complex dependencies (Paz & Leigh, 1994). There are two basic approaches to scheduling maintenance. The first approach is to schedule planned maintenance, while reserving enough slack capacity to accommodate unplanned demand (cf. Alfares, 1999). In research related to this approach, the use of slack capacity is typically reflected in assumptions such as "an ample supply of space and labor are usually available" (Sriram & Haghani, 2003, p. 38). The second approach is to introduce unplanned demand through dynamic list scheduling (Paz & Leigh, 1994), where planned and unplanned work is scheduled based on prioritization rules (cf. Safaei, Banjevic, & Jardine, 2011). This implies the use of a time buffer. Naturally, the latter approach is more challenging if planned work cannot be readily postponed (cf. Angalakudati et al., 2014). As in manufacturing, further research on scheduling under uncertainties and dynamic scheduling has been called for in maintenance (Froger, Gendreau, Mendoza, Pinson, & Rousseau, 2016; Van den Bergh, De Bruecker, Beliën, & Peeters, 2013) and in healthcare (Helm et al., 2011).

The prevailing capacity-focused buffer management mind-set in aircraft line maintenance leaves dynamic approaches largely unexplored. Van den Bergh et al. (2013) noted that although maintenance scheduling problems in military and commercial aviation have received academic attention, researchers have focused on (deterministic) analytical models rather than on simulation. The literature acknowledges uncertainties, such as delayed flight arrivals, failure rates (emergent demand for maintenance), extended repair times, and workforce availability. However, these inherently dynamic uncertainties are typically not considered in the models. Some exceptions can be found in military aviation (Adamides, Stamboulis, & Varelis, 2004; Guarnieri, Johnson, & Swartz, 2005; Mattila & Virtanen, 2014; Mattila, Virtanen, & Raivio, 2008), engine maintenance (Gatland, Yang, & Buxton, 1997), and part repairs (Cobb, 1995; Kilpi, Töyli, & Vepsäläinen, 2009).

3 | METHODOLOGY

In this section, we describe the explorative design science (Holmström et al., 2009) research process that created the frontlog practice, including how we anticipate the implementation outcomes and evaluate our design in the case setting. Normal design science research (van Aken, Chandrasekaran, & Halman, 2016) is focused on improving practice, emphasizing field testing, and implementation. Explorative design science shares the interest in improving practice, however, the focus is on reframing a field problem (e.g., Groop, Ketokivi, Gupta, & Holmström, 2017), or on exploring alternative ways of operation, for example, through novel technologies (e.g., Hedenstierna et al., 2019). In explorative design science research, modeling and simulation play important roles, taking the place of implementation in exploring outcomes (Hedenstierna et al., 2019). Furthermore, theory and theorizing are methods for anticipating how a solution that is not yet field tested should best be tested and further developed for transferability and application to settings beyond where it was designed.

The explorative research process of this article is presented in Figure 1. The research started from the field problem of the case company, Nordic Airline. Through design exploration and modeling, the effects of changing the scheduling principles of line maintenance could be explored, resulting in the proposal of a new type of buffer, of interest to OM theory. The major milestones and events in the research process consisted of reframing the departure reliability problem as a rescheduling problem (shifting focus from the engagement in practice to design exploration and modeling); articulating the solution as a frontlog buffer in contrast to a backlog (shifting focus from design and modeling to theoretical search and theorizing); and results from a detailed simulation that indicated the potential to both improve departure reliability and maintenance cost efficiency (shifting attention back to implementation and overcoming obstacles to implementation).

The research engagement with Nordic Airline began in 2014 as part of a large strategic research initiative bringing together academic researchers and industry funded by the funding agency for technology and innovation. The management of the maintenance organization was seeking ways to improve the departure reliability of their long-haul fleet and invited the academic researchers to join the search. The airline operates a hub-and-spoke network with a geographic advantage in connecting Asia and Europe, emphasizing the strategic importance of the long-haul fleet and, by extension, departure reliability. The geographic advantage of the airline's hub translates

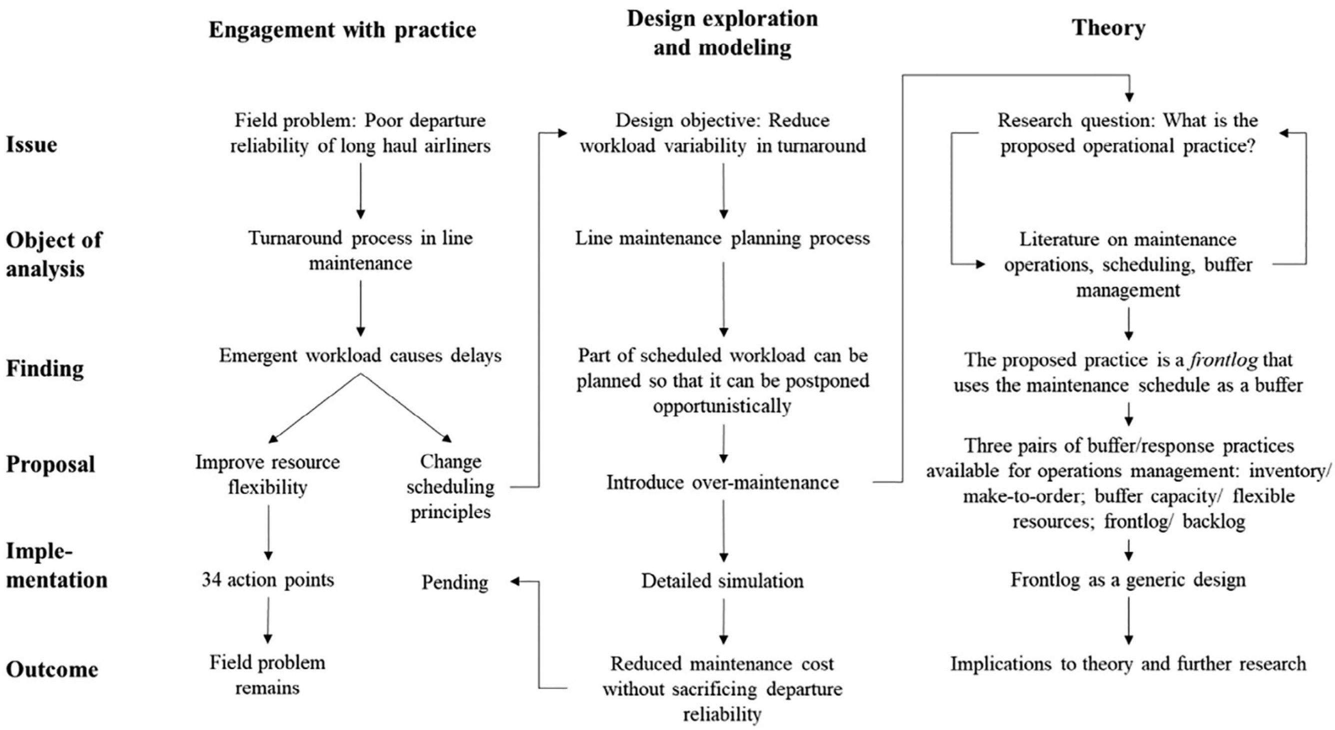

FIGURE 1 Explorative design research process

to one of the highest fleet utilizations for long-haul aircraft, where the aircraft spend more time in the air than on the ground. Because of the high fleet utilization, the typical maintenance window at the hub is limited to a few hours per day, between when the aircraft arrives at the hub and heads back to its Asian destination. The airline refers to this as turnaround maintenance, and increasing resource flexibility in this operation was the initial focus of the research.

By the end of the initial engagement in 2015, through exploring contradictory goals using thinking process tools (Groop et al., 2017), the possibility of change in maintenance scheduling was identified as a potential solution. The possibility was first investigated in the context of the existing line maintenance planning process, with the purpose of reducing workload variability in turnaround maintenance operations. The finding was that, for the long-haul fleet, part of the scheduled workload could be planned so that it later could be postponed opportunistically, in effect purposefully introducing over-maintenance. Elaborating the effect of this change on maintenance cost was the third phase of the research process, the detailed simulation, which was initiated in May 2016. The fourth stage of the research process was to identify what this new maintenance practice was, from a theoretical perspective, and conducted for writing this research article, which began in early 2017. This theorizing consisted of contrasting both our design and our findings against a wide body of academic operations and maintenance management literature. Currently, the introduction of the proposed solution in practice is progressing as part of a wider digitalization effort, with ongoing negotiation of a potential role of the researchers in the project. The data collected for the research through engagement in practice and for modeling are summarized in Table 2.

Engaging with the case company to address their initial field problem initiated a process of framing and reframing (Simon, 1996) that presented an opportunity to innovate aircraft line maintenance operations. Being attentive to frames held by issue stakeholders (Coghlan & Brannick, 2001) creates the opportunity to reframe field problems in novel ways for the purpose of explorative design (Groop et al., 2017). The modeling phase of the research first studied the relationship between over-maintenance and workload variance in a deterministic setting, focusing solely on scheduled maintenance (Öhman, Laine & Holmström, 2016. For the initial study, we conducted five semi-structured interviews (each lasting 1–2 hr) with maintenance planners for long-haul and short-haul fleets and operations, production, and resource planning managers. Through these interviews, we gained an in-depth understanding of the case company's maintenance planning function. We also uncovered heuristics and principles not visible in the operational documentation. The study showed promising results: A

TABLE 2 Data collected throughout the research process

Data	Engagement with practice	Design exploration and modeling
Semi-structured interviews	11 interviews, each lasting 0.5–2.5 hr (e.g., head of line maintenance, head of maintenance planning, production manager, duty manager, and lead mechanic)	5 interviews, each lasting 1–2 hr (e.g., resource planning manager, production planning manager, long-haul planner, short-haul planner)
Observation	8 observations, a total of 33 hr (Maintenance Control Center, long-haul turnaround crew, short-haul turnaround crew)	
Documentation	Maintenance process descriptions Regulatory compliance documentation Reports (fleet performance, workforce, delays) Excerpts of documents used in operations (planning sheets, work instructions, etc.)	Fleet composition
Database extracts	Technical flight delays (9 months)	Aircraft maintenance programs Maintenance event records (12 months) Traffic program data (12 months)
Other	Notes from workshops Notes from informal discussions	Ad hoc discussions to clarify issues that emerged during modeling Notes from meetings and results presentations

1% increase in the total planned workload could result in up to a 6% reduction in workload variance. However, we concluded that the real impact could not be evaluated while the natural demand variability caused by the emergent workload was neglected.

Through the engagement with practice and the initial modeling phase, we developed a rich, in-depth understanding of the case company context (cf. Forrester, 1992), based on both qualitative and quantitative data (Table 2). This understanding was crucial in empirical grounding of the simulation model presented here, and it allowed us to formulate the *good problem*, which is the key to success in any simulation project (Law, 2003). The simulation model then allowed us to explore the effect of the proposed buffer management approach in a stochastic setting.

4 | ENGAGEMENT WITH PRACTICE

Reflecting management's vision of how the process *should* work, the research project was initially named "pit-stop." Two months into the project a workshop was conducted with the maintenance organization of the airline, during which the first impressions and initial problem framings were discussed. At that point, the analysis was still unstructured and consisted of observations on recurring themes and apparent conflicting views in the data. The agenda of the researcher-facilitated workshop was to present and discuss alternative representations of what caused the field problem of poor departure reliability in the long-haul fleet, complementing and possibly challenging the initial case company framing. What had initially been framed (by management) as limited visibility into, and control of turnaround work, was now also explored as an issue of unclear roles within the organization, as well as problematic short-term scheduling principles. Further, the initial research had raised doubts about the extent to which problems with the departure reliability were attributable to maintenance. Based on the workshop, a second round of data gathering, aimed at getting more in-depth insights into the daily work of the turnaround crews and into the reasons for delays in the long-haul fleet, was initiated.

The second phase of data gathering lasted approximately a month, during which data were gathered through observation of the turnaround crews. Further, the case company provided extensive records of flight delays, enabling quantitative analysis of the causes. The data were analyzed using thinking process tools, based on the theory of constraints (Davies, Mabin, & Balderstone, 2005; Goldratt, 1990) with the purpose of framing and reframing to guide explorative design (Groop et al., 2017). We constructed a current reality tree, where the primary undesirable effect was sub-par departure reliability. The delay data also enabled a relative comparison of most of the effects leading to a delay, based on which the impact of the suggested improvements could be evaluated. Further, the problems of

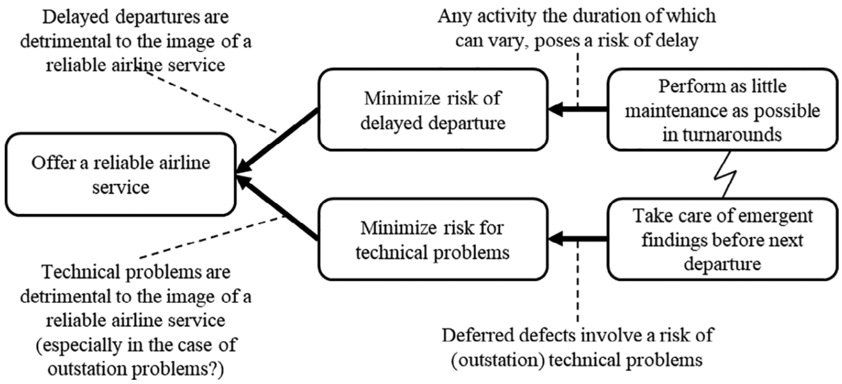

FIGURE 2 Evaporating cloud analysis of the mind-set conflict

short-term scheduling principles were found to be grounded in conflicting mind-sets within the organization. This conflict was approached through an evaporating cloud analysis (included in Figure 2). The history of the maintenance organization offered at least a partial explanation for the conflict. At the time of the study, the maintenance organization was recovering from a recent restructuring that had included considerable downsizing of the maintenance technician workforce. In our interpretation, maintenance resources, which had been relatively abundant, were suddenly scarce. The issue of unclear roles, mentioned above, which further complicated the situation, was rooted in responsibility for turnaround workload management and resource management being housed in different parts of the organization. One visible symptom of the situation was that the 2-week operational planning horizon tended to show a good match of workload and resources, apart from the upcoming 48 hr, for which the workload overshot the planned capacity, which was a situation that persisted. Another symptom was recorded in an early research log entry following a tour of operations: "... just about everyone I've talked to has been busy counting resources—by hand, on paper!"

The bottom line of the analysis was that unpredictable turnaround workload, combined with the tight maintenance windows of the long-haul fleet, was the root cause of the sub-par departure reliability. The results of the analysis were discussed at the workshop which concluded the second cycle of engagement with practice. The management of the maintenance organization insisted that the primary means for remedying the problem should still be through improving resource flexibility and agility, instead of systematically (albeit selectively) postponing (deferring) repairs. The preference was to improve the capacity flexibility over changing the scheduling. At that point in time, the solution based on over-maintenance and rescheduling was not taken into consideration, as the high fleet utilization seemed to imply that any increase in the maintenance workload would only worsen the situation. Thus, it seems natural that the organization would focus its development efforts on increasing resource flexibility and agility, that is, the responsiveness of the capacity buffer.

At that stage of the research, a set of action points, ranging from partial reorganization of the turnaround work to introduction of technology (aimed at improving resource flexibility and information flow in the turnaround process), was proposed. Most of these points were implemented in some form during or after the project.

However, we, the researchers, were still bothered by the case company seeing workload variability as something that "you just have to live with," reminding us of how the bullwhip effect was perceived before solutions for containing it emerged (Geary, Disney, & Towill, 2006). Therefore, we asked whether resource-focused buffering really was the only cost-effective alternative in the case company context. The arguments for avoiding a backlog were sound: Postponing scheduled work typically requires explicit approval from the aviation authorities, and postponing emergent work (even if it could be done within regulatory limits) typically implied increased risks and/or aircraft operating costs. Against this backdrop, the idea of scheduling work earlier than necessary, as part of deliberate active buffer management, emerged.

5 | DESIGN EXPLORATION AND MODELING

The initial design idea for design exploration was that a portion of planned maintenance tasks could be deliberately scheduled so that there would be at least one upcoming maintenance window during which the work could be performed without violating its regulated due date. Thus, if additional workload emerged during a turnaround, many tasks could be postponed without additional effort, risk, or cost. Technically, this could be characterized as a frontlog of tasks rather than a backlog, as this entails doing work before demand is imminent,

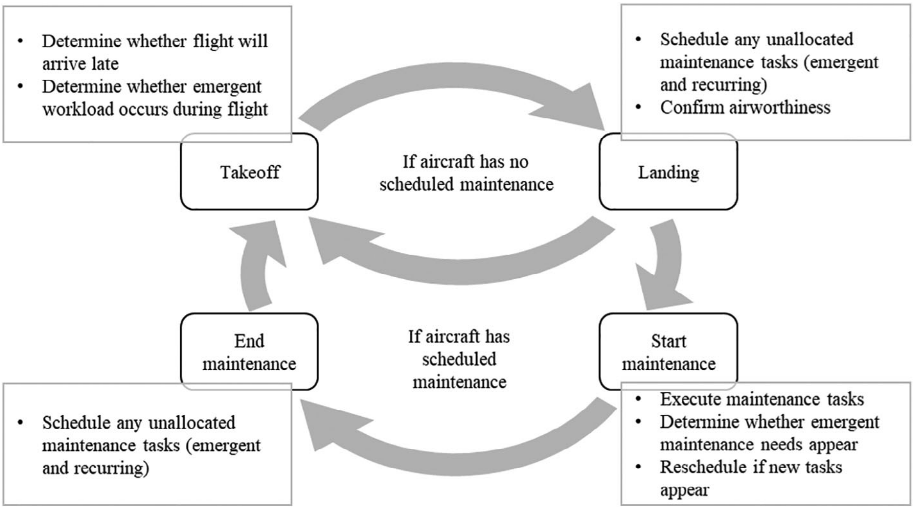

FIGURE 3 Life of a single aircraft as discrete events

TABLE 3 Air fleet operations and scheduling events

Event	Description
Take-off	The time when the flight is scheduled to depart (if the aircraft is delayed, take-off is postponed accordingly). The flight hours are recorded for the aircraft until landing. Emergent failures are recorded during the flight.
Landing	The time when the flight lands (including possible delays). At this time, the scheduling process considers every maintenance need of the aircraft, schedules unallocated recurring maintenance tasks, and checks whether the aircraft is fit for the next flight. If not, rescheduling of maintenance is triggered.
Start maintenance	This is when scheduled maintenance tasks (recurring and emergent) are executed with the corresponding changes made to the aircraft's maintenance status (latest execution updated for recurring tasks). Further emergent maintenance tasks can be found during the maintenance and inspection tasks. Any emergent maintenance tasks trigger a rescheduling procedure for the finding. If feasible, the emergent task is completed right away in the ongoing maintenance event, and if not, the maintenance scheduling algorithm is triggered. This continues until there are no more scheduled tasks.
End maintenance	An event that completes maintenance work. Here, the maintenance scheduling algorithm is always triggered: Planning future maintenance, including the frontlog.

rather than tending to it after it is due. As the due date of recurring work is defined based on the last time the work was performed, we were also dealing with a trade-off: The benefits of an excessive frontlog buffer are outweighed by the costs of the increased average workload.

The relevant events for exploring and modeling the solution design—the frontlog—are the departure and landing of an aircraft and the beginning and ending of a maintenance work package (Figure 3), which need to be processed concurrently for all aircraft in chronological order. All operations concerning flights, maintenance, and changes in aircraft states take place during these events. From the perspective of a single aircraft, the chain of events is typically consecutive take-off and landing events that are occasionally interrupted by maintenance events. A more detailed description of what happens during these events is provided in Table 3.

For every aircraft, a maintenance schedule is needed that contains the next planned execution time for every recurring and pending emergent maintenance task in the future. The maintenance schedule is created (and recreated) by the maintenance scheduling algorithm (Appendix C), which also specifies the decision variable that determines the target frontlog buffer. The frontlog buffer is then created through scheduling recurring maintenance tasks earlier than the last opportunity to perform them.

5.1 | Detailed simulation

The discrete event simulation model describes aircraft operation and maintenance (Figure 4) in a hub-and-spoke context. The focus of the model is maintenance scheduling and rescheduling, as this is where the frontlog buffer is introduced into the maintenance operations. In effect, this includes simulating maintenance routing (known as *tail assignment* in the maintenance organization), identifying suitable ground times for maintenance work packages (bundles of maintenance tasks to be performed during a given maintenance event), and population of these packages with maintenance tasks.[2] The empirically grounded model reflects the empirical context of the study. Five aircraft fleets (which we call E90, A32s, A330, A340, and A350) operate from a central hub, from which the aircraft make roundtrips to their destinations. Maintenance operations are performed only at the hub, reflecting operations at Nordic Airline (at the time of the research). Outstation maintenance is exceptional and typically occurs only due to critical outstation failures.

The model was programmed in Python using *NumPy* (www.numpy.org) in addition to standard libraries. In the remainder of this section, we describe the function of the model: how the model is initialized (Section 5.2) and how the simulation progresses (Section 5.3). Further, we present the measures taken to validate the model (Section 5.4). The results are presented in Section 5.5, in which the relationship between the simulation output (i.e., the workload variance) and the decision variable (i.e., the frontlog buffer) is presented and evaluated based on cost data provided by the airline.

5.2 | Input data and model initialization

Model initialization begins by loading data (Table 4) and building the required flight and maintenance schedules. This brings the simulation into its beginning state, in which each aircraft exists as an individual entity, with its own planned flights and assigned future recurring maintenance tasks. Maintenance routing is naturally nontrivial with respect to detailed maintenance scheduling, as it determines the available maintenance opportunities.

In the studied company, the traffic plan contains predesignated maintenance blocks, which ensure that the plan (at the fleet level) can accommodate recurring maintenance tasks that require exceptionally long ground times. Further, if an engine needs to be changed, the required ground time is jointly planned by traffic and maintenance planning. For the remaining workload, however, maintenance planning treats the flight schedule as given. In practice, changes (apart from the above) requested by maintenance are extremely rare. Thus, maintenance routing produces a feasible flight schedule when all planned flights can be completed without

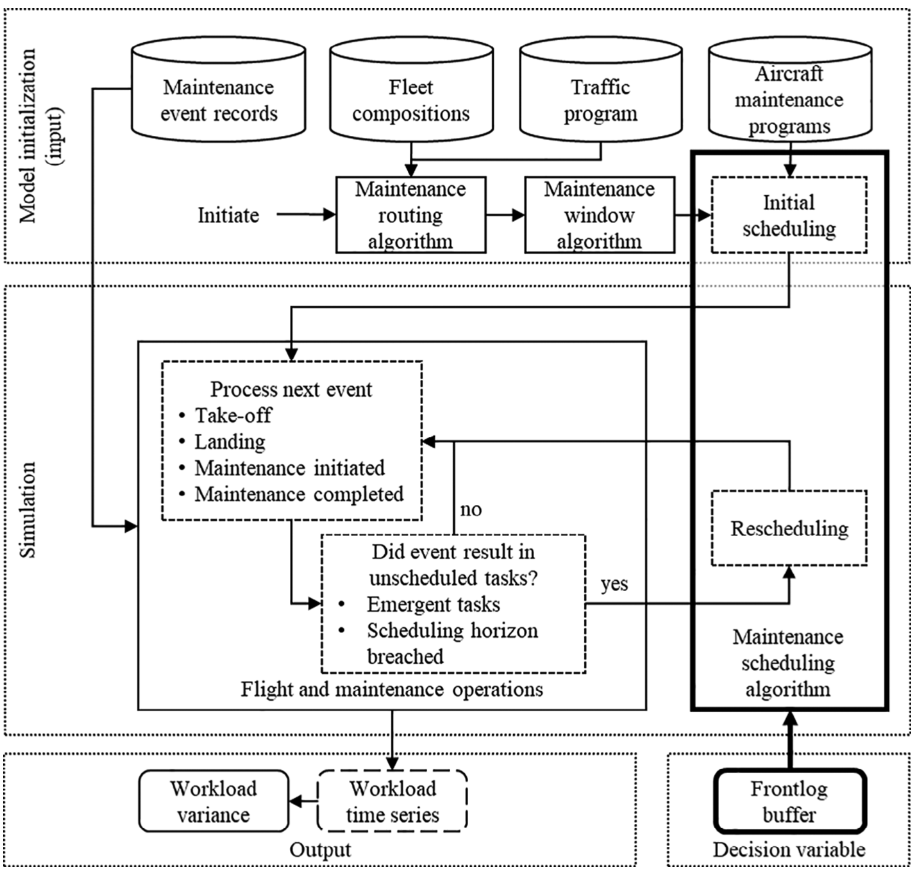

FIGURE 4 Basic structure of the simulation model

TABLE 4 Simulation input data

Data source	Extracted information	Used for
Traffic program	Departure stations, destination stations, take-off time, arrival time, assigned fleet, and seat configuration	Creation of the flight schedule which serves as input to the maintenance routing process
Fleet composition	Aircraft type, seat configurations, registration number	Creating the fleets of aircraft which are assigned to the flight schedule during the maintenance routing process
Aircraft maintenance programs	Aircraft type, maintenance task recurrence (time and/or flight hours/cycles), and expected duration of the maintenance task	Establishing (recurring) aircraft maintenance needs for populating maintenance work packages and monitoring airworthiness
Maintenance event records	Aircraft type, originating maintenance events for emergent findings, distribution for occurrence of emergent failures, and distribution for duration of repairs of emergent failures	Generating emergent failures in the course of the simulation

violating scheduled maintenance constraints. The simulation model allows for aircraft to be swapped between flights during the simulation. However, as this rarely occurs in reality, it should occur rarely in the simulation.

The maintenance routing algorithm (Appendix A) processes each flight of the flight schedule in chronological order and assigns an aircraft to each flight. In addition to the planned flight origin and destination, the flight schedule has an assigned fleet and seat configuration for each flight. The maintenance routing algorithm considers the constraints set by the flight schedule while ensuring that no recurring maintenance task will expire during the flight or during the return flight if the aircraft is destined for an outstation. Further, the algorithm ensures that each aircraft departs from the station the aircraft has landed at and that there is sufficient time (set at 1 hr) between landing and the next departure. During the assignment process, the maintenance routing algorithm also assigns potential maintenance windows where it detects sufficient ground time at the hub. These potential maintenance windows comprise all the positions in which the aircraft can possibly undergo maintenance. Most are later removed. Once maintenance routing has produced a viable flight schedule, the initialization process continues with the maintenance window algorithm.

The maintenance window algorithm (Appendix B) reduces the potential maintenance windows to actual maintenance windows. The algorithm also ensures that the maintenance windows (from a fleet perspective, as the fleets share maintenance resources) are evenly distributed over time, which is a prerequisite for creating a uniform workload in the task scheduling phase. The designated maintenance windows (referred to as work packages) are yet to be populated with recurring maintenance tasks by the subsequent maintenance scheduling algorithm (Appendix C).

Each type of aircraft has several hundred different recurring maintenance tasks that must be executed periodically at different intervals, defined by the time, flight hours, or flight cycles elapsed since the last execution. To generate the initial state of the recurring maintenance tasks, the maintenance scheduling algorithm iterates through all recurring maintenance tasks of an aircraft. The algorithm sets a random previous execution time so that the due date of the recurring task is some time after the first planned maintenance window. The maintenance scheduling algorithm then calculates the projected distribution of work hours for the planning period (taking into account task recurrence). Then, the algorithm re-randomizes the previous maintenance execution dates until the resulting daily workload variance (of the recurring maintenance tasks) for the planning period is less than 20%. (This was found to represent a near-minimum value for workload variance, with the given flight schedule, which was certain to result in a valid maintenance schedule.) Together, the maintenance routing, maintenance window, and maintenance scheduling algorithms create a unique and realistic initial state for the simulation, with a complete flight and maintenance plan.

5.3 | Simulation

Based on the successful initialization of the simulation, every flight has an aircraft assigned to it, and every aircraft has designated maintenance windows that are populated with recurring maintenance tasks for up to 2 months of simulation time (reflecting the practice at the Nordic Airline). The simulation then starts executing flights and performing maintenance in the order they were scheduled. Based on distributions for occurrence, duration, and criticality (derived from actual maintenance records), the simulation also starts generating emergent maintenance tasks that (a) may be accommodated by the current plan, (b) may require rescheduling of planned work, or (c) may lead to the use of more expensive overtime work and cause a departure delay.

The frontlog buffer is then created through scheduling recurring maintenance tasks earlier than the last opportunity to do so. This represents the trade-off where the workload that can be readily postponed is created at the expense of a slight increase in the total average workload. Owing to this trade-off, we determined that the frontlog buffer (that is, the decision variable) was best expressed through the percentage of additional average workload the buffer creates.

While constructing the simulation model, we determined that the best recurring maintenance tasks with which to build the frontlog buffer were those with long and predictable intervals. The reason was quite obvious: The portion of unused maintenance intervals was the smallest for these tasks. Further, the long intervals meant that the tasks had more possible work packages in which the tasks could be executed, as their repeat intervals spanned several work packages. Short interval tasks, in some cases, would have to be executed in every or every other work package. Although we favored longer, predictable interval tasks when building the frontlog, this was not a constraining factor in our model, that is, we made sure to have designated enough buffer tasks for us to be able to test a sufficient range of different frontlogs (ranging from ~0% to 17%, measured as additional total workload).

Scheduling a recurring maintenance task requires knowledge of the task's due date, based on which the next execution of the task is planned to the last work package before it expires. Depending on aircraft utilization (which varies between fleets), these work packages reflect maintenance windows ranging from short apron services lasting a few hours to longer hangar services. Considering task-related constraints (e.g., that some tasks cannot be performed on the apron, and some tasks cannot be performed in the same work package due to regulations), the last possible work package is appointed for each recurring maintenance task. The maintenance scheduling algorithm then applies the sought frontlog buffer by moving recurring maintenance tasks to earlier work packages, while minding constraints, until the total workload in the schedule has risen by the amount set by the decision variable.

Emergent maintenance tasks are scheduled in a similar vein. However, they typically need to be addressed immediately, thus adding work to ongoing or upcoming work packages. The frontlog buffer is used whenever the additional workload exceeds the average additional workload (in relative terms). Whenever an emergent maintenance task pushes the additional workload above the average, a recurring maintenance task designated for the frontlog buffer is postponed. This brings the relative amount of the additional workload back below the average. The frontlog buffer is consumed until no more emergent workload occurs or until the buffer is depleted. If further emergent workload appears, it is absorbed by what could be characterized as overcapacity, as shifts are able to accommodate up to 15% additional workload compared to the planned workload (determined based on maintenance event records). However, any emergent workload beyond this will delay the affected aircraft, basically utilizing a backlog buffer and potentially an overtime capacity buffer (as the workload capacity on a delayed aircraft will be arranged if the capacity is not available).

5.4 | Model validation

During the iterative process of constructing the simulation model, we validated it several times against the available maintenance records and against the airline management's perception (cf. Den Hengst, De Vreede, & Maghnouji, 2007). As for the initialization of the model, the airline management verified that the resulting schedule and maintenance windows with respective work packages corresponded well with the maintenance routing performed by the organization. The main form of validation was achieved by creating an operational planning view (Figure 5) provided by the simulation. The result was recognized by the managers as identical to what they saw in their daily work.

As the airline operates a hub-and-spoke network with long-haul Asian flights feeding short-haul flights to European destinations, the maintenance workload has a distinct daily profile (Figure 6). Despite limited resolution, the simulation produced a daily profile with a distinct workload spike during the afternoon when the long-haul flights from Asian destinations land for apron maintenance before returning to their Asian destinations. Further, the maintenance of the short-haul European feeding traffic is concentrated during evenings and nights, which is also visible in the daily profile.

Finally, the yearly workload produced by the simulation was compared to the operational maintenance data (extending over a 1-year period) supplied by the airline (Figure 7). The simulation displayed a mean error of 2.14% for the total yearly workload. The substantial error is explained by changes in fleet composition, reflected in the operational data, as the A340 fleet was being replaced

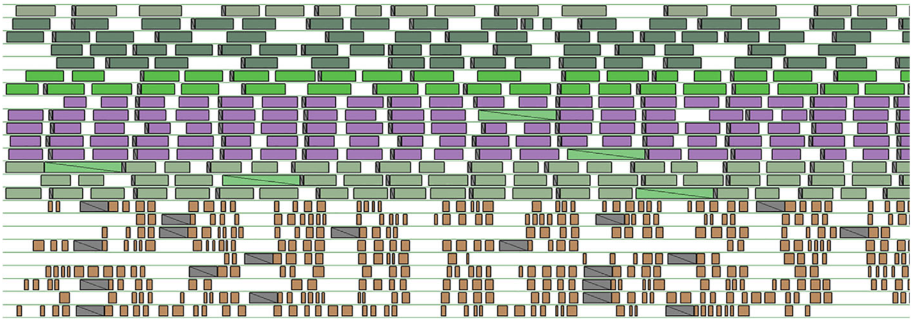

FIGURE 5 Operational planning Gantt chart excerpt generated by the simulation, where each row represents the flights (empty blocks) and maintenance (blocks with the diagonal line) assigned to a specific aircraft

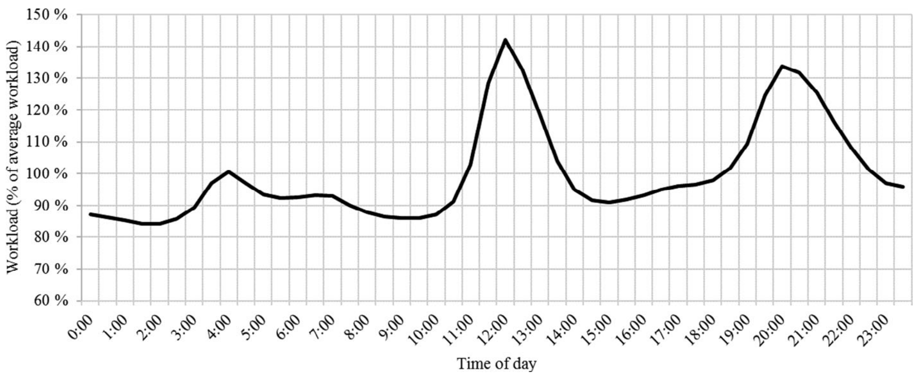

FIGURE 6 Daily workload profile

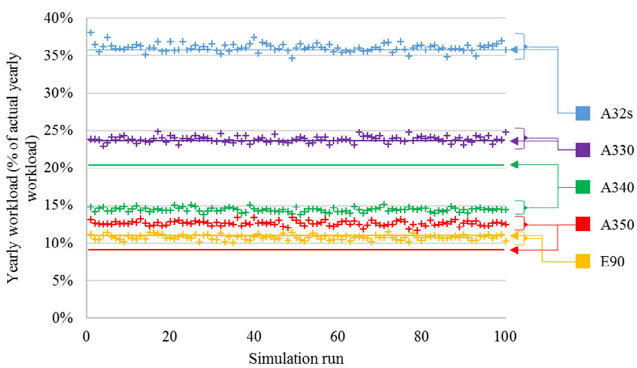

FIGURE 7 Simulated yearly workload (cross) and the actual workload extracted from 1 year of operational data (line)

by the A350 fleet. In practice, this implied a reduction in the number of A340 aircraft and a corresponding increase in the number of A350 aircraft. However, the simulation used stable fleets (reflecting the situation after the time period described by the operational data). When only the fleets that were also stable in the operational data (A32s, A330, and E90) were considered, the mean error is reduced to 0.38%.

5.5 | Simulation results

The impact of utilizing the frontlog buffer on the total workload variance is depicted in Figure 8. Each data point represents 1 year of simulated aircraft operation and maintenance. The simulation gave slightly different buffer variance profiles for different fleets of aircraft (Appendix D), which was expected due to the different operational profiles. The higher the utilization of the aircraft, the stronger the effect of introducing a frontlog buffer.

As expected, utilizing the frontlog buffer had a considerable reducing effect on the workload variance. This means that the frontlog buffer could be expected to improve departure reliability through functioning as a shock absorber against the emergent workload—in accordance with the airline's original objective. The question was, however, at what cost. The frontlog buffer and the capacity buffer (required to cope with any remaining workload variance) had costs, which invited further analysis to arrive at an actionable managerial insight. The cost of utilizing the frontlog buffer could be readily evaluated in terms of the hours of work the buffer added on average. Potential additional spare part costs due to the introduction of the frontlog buffer were assumed to be negligible, as they can be affected by appropriate selection of buffer tasks, and many recurring maintenance tasks do not require spare parts. Further, a cost for "lost availability" was not motivated, as introducing the frontlog buffer did not affect flight plan feasibility. All planned flights could be flown (while minding operational constraints) with all simulated buffer values.

The capacity buffering costs for coping with the remaining workload variance were more challenging to

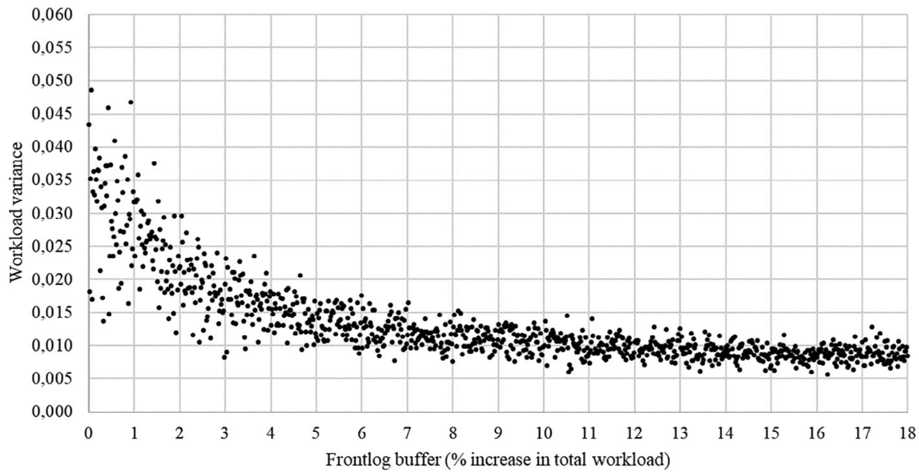

FIGURE 8 The impact of the frontlog buffer on workload variance, where a point represents a (1 year) simulation run

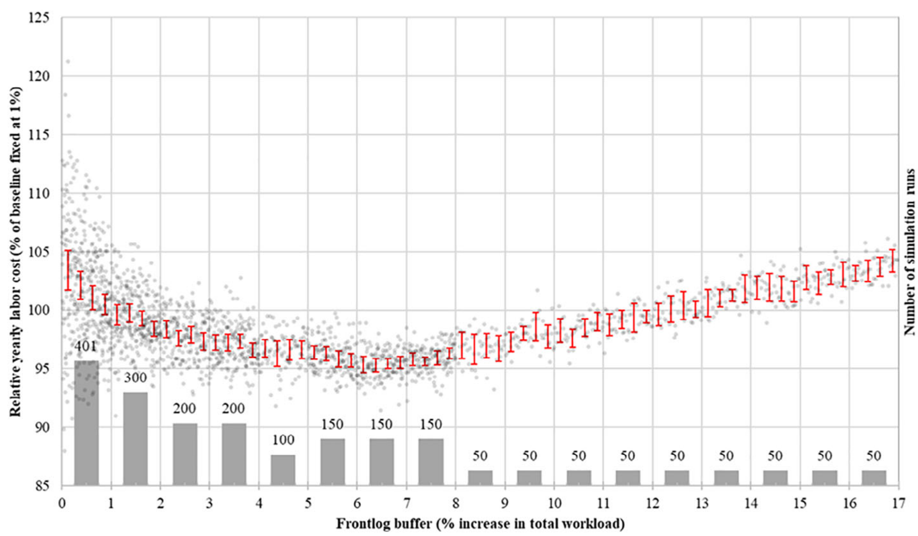

FIGURE 9 The impact of utilizing the frontlog buffer on yearly labor costs, where a point represents a (1 year) simulation run and a bar indicates the number of simulation runs within the buffer interval, spread evenly throughout the interval. The red error bars indicate the 99% confidence interval of the mean simulated yearly labor cost

evaluate. To simplify the analysis (risking a conservative estimate), the cost of coping with the remaining variance was reduced to two components: ad hoc capacity buffering and structural capacity buffering. Ad hoc capacity buffering was reduced to overtime work hours that could be determined from the simulation. According to the airline, overtime work was twice as expensive compared to regularly scheduled work. Structural capacity buffering included all contractual and permanent work organizational arrangements that were intended to concentrate work, to accommodate spikes in the workload. The cost of structural capacity buffering was assessed by the airline managers as adding a few percentage points[3] to the total workload costs and was expected to decrease in proportion to the increase in the buffer when considering its effect on the workload variance.

Through analyzing the maintenance event data supplied by the airline, we determined the current (unmanaged) frontlog, that is, the real amount of work performed before the last opportunity to do it. When compared to the amounts produced by the simulation, we determined that the airline was performing maintenance that corresponded to a buffer of 1%. With the cost of capacity buffering fixed (at a baseline of 100%) at the derived 1% buffer, we simulated 2,101 years of aircraft maintenance and operation with different frontlog buffer values, resulting in the relative (to baseline) yearly labor cost impact depicted in Figure 9.

Based on the overlaps of the 99% confidence intervals for the different buffer values in Figure 9, we concluded that the cost-optimal frontlog buffer would be in the range of the 5–8% of additional average total workload. Increasing the frontlog buffer from the current 1% to the cost-optimal 5–8% range would imply a 4–5% reduction of yearly labor costs (seen as the deviation from the 100% baseline). Further, increasing the frontlog buffer could also be expected to lead to a significant improvement in yearly labor cost predictability (for a 1% buffer, the cost varies by approximately 94–106%; for a 5–8% buffer, the cost varies by approximately 93–98%).

Based on these results, we conclude that introducing a *Frontlog* buffer not only improves departure reliability but also leads to a significant reduction in labor costs. The cost effect presented above should be considered conservative, as we expect that introducing the *Frontlog* buffer also results in (a) a reduction in potentially costly delays (based on the simulation) and (b) a reduction in resource management and the planning workload (based on the engagement with practice).

6 | EVALUATION AND IMPLEMENTATION OF THE FRONTLOG AT NORDIC AIRLINE

The design was evaluated in four phases (evaluations 1–4 in Figure 10). First the focus was on whether frontlog buffering can be implemented at all, then whether it should be done, and finally, whether Nordic Airline can implement frontlog in their current operational system. Although the design work was done in a single context, and the design has not been fully implemented, we can draw a parallel between alpha- and beta-testing (cf. van Aken, 2004). The design first underwent an internal evaluation mainly within the design team, which included the researchers and the head of maintenance planning of Nordic Airline. After the internal evaluation, the design underwent an external (in context) evaluation, in which the design was presented to and evaluated by company stakeholders who had not actively participated in the design process. The internal evaluation was conducted throughout the process of building the simulation model, which resulted in iterative improvement in how well the simulation model represented the context. The design also developed during the internal evaluation, as we were

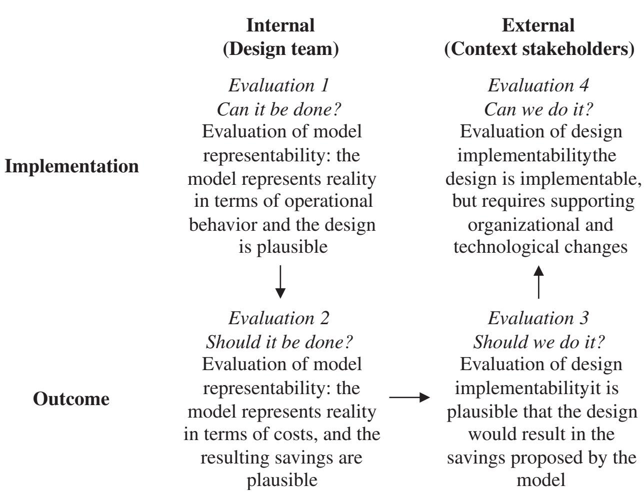

FIGURE 10 Evaluation of the frontlog design

forced to answer questions such as "What tasks should be used for frontlog scheduling?," "Should the frontlog buffer be aircraft or fleet specific?," and "How should we measure the frontlog buffer?"

The main forms of evaluation in Phase 1 are described in Section 5.4 as validation of the simulation model. Apart from several occasions when the model was evaluated within the design team, Evaluation 1 included an evaluation meeting at which the behavior and output of the simulation model were discussed with a maintenance planning manager and an operations manager of the case airline. Evaluation 2 culminated in an evaluation meeting where the cost effect was presented and discussed among the design team, the head of line maintenance, and an operations manager. The cost effect was evaluated as plausible and significant, and only minor corrections were made in the model based on the meeting. A summary of the evaluation meetings and their key takeaways are included in Appendix E.

Evaluation 3 culminated in a meeting where results were presented to key organizational stakeholders. The results were, in general, not questioned, with the exception of the vice president (VP) of technical operations, who noted that aircraft generate revenue when in the air. Thus, focusing on labor costs in the analysis may not provide a complete picture. In response, we emphasized that all planned flights in the current traffic program (which was an input to the simulation model) were successfully executed with all buffer values. In a broader sense, we interpreted this comment as an expression of the mindset emphasizing capacity buffering, which we also found in aircraft line maintenance literature.

Evaluation 4 was initiated after the head of maintenance planning—whom we considered part of the design team—moved to a new position within Nordic Airline. In addition to discussions with airline representatives, we collected evaluation data at three distinct events where implementation challenges were discussed. The first was a meeting (in September 2017), where we presented the design to the new head of maintenance planning and several members of the maintenance planning team. The second was a written implementation update (in December 2018), provided by the new head of maintenance planning. Finally, in November 2019, we met with the new head of maintenance planning, the business development manager, and the process development (lean) manager to gain a more in-depth understanding of the reasons for the slow progress in implementation. These reasons can be summarized in terms of two aspects. First, after the research was conducted, Nordic Airline launched a major digitalization initiative, along with changes in the organization and operational IT systems, which had been business development priorities. Second,

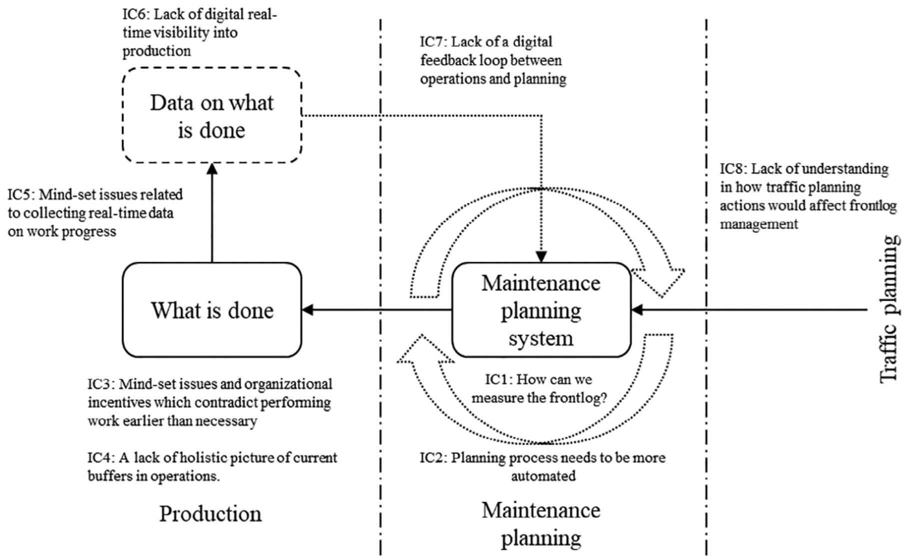

FIGURE 11 Implementation challenges

although the initiative and changes had addressed some of the challenges related to frontlog implementation, Nordic Airline still saw several challenges with introducing frontlog, both operational and organizational. We discuss these challenges in the following subsection.

Based on Evaluation 4, at the time this article was written, eight implementation challenges (IC) must still be addressed by the airline. As indicated in Figure 11, these challenges are related to the production planning system, and to closing the digital feedback loop between planning and production. As a frontlog measure is part of the simulation model, the lack of such a measure in the planning system (IC1) may seem trivial. However, the frontlog measure in the simulation was constructed to analyze its cost effect in terms of buffer trade-offs. Thus, the measure was not designed for decision making in planning. As a managed frontlog would make replanning frequent, Nordic Airline also saw the need to increase automation in planning (IC2), implying the introduction of a form of a dynamic scheduling solution (cf. Biçer & Seifert, 2017). However, automation is difficult, as rescheduling must consider many nontechnical constraints and principles. For which the airline currently relies on the maintenance planners' tacit knowledge.

Challenges 3 and 4 are related to how production is managed, after the maintenance plan is released to production. The mind-sets of the organizational units involved in production management complicate maintaining a planned frontlog. Currently, with an unplanned frontlog, production management postpones many non-due tasks on light grounds (IC3). The fourth challenge is related to the third, and could be described as a current lack of holistic understanding of how different parts of production (and planning) hedge against (natural and artificial) variability (IC4). An example is a task that normally takes 1 hr to complete, but under specific (but infrequent) circumstances takes 3 hr. When such a task is always scheduled for a duration of 3 hr, it indicates the lack of a big picture view on what buffers are available and for what type of variability they are used.

Challenges 5–7 are related to the digital feedback loop from operations to planning. Nordic Airline has made considerable advances in tracking task progress as part of their digitalization initiative. However, the airline is still far from real-time visibility (IC6). Visibility into operations is seen as a prerequisite for frontlog management at the fleet level. Working toward real-time visibility includes addressing negative attitudes toward collecting information on how work is progressing (IC5). Further, closing the feedback loop requires systems integration (IC7). Finally, throughout the fourth evaluation phase there were heated discussions on how the behavior of traffic planning affects frontlog implementation (IC8). Traffic planning decisions affect available maintenance windows and aircraft usage. Therefore, a change in fleet routing might move a planned frontlog from where it is likely to be needed to where it is less likely to be needed.

7 | DISCUSSION

Aircraft line maintenance is a context where technical failures create additional demand on short notice, and failure to meet that demand results in costly delays. It is also a context where demand uncertainty is predominantly addressed through resource flexibility. These were the circumstances for which we found that the proposed frontlog buffer would not only improve departure reliability but also reduce maintenance costs. Next, we discuss how these circumstances were also fortunate for theoretical conceptualization and elaboration of buffer management.

7.1 | Conceptual and theoretical elaboration

The conventional approach to aviation maintenance in practice and academia (cf. Başdere & Bilge, 2014; Boeret, 1977; Sarac et al., 2006) can be described as JIT maintenance, where tasks should not be performed too late or too early. However, with random failures creating additional tasks, the conventional approach leads to over-reliance on capacity buffering for coping with the unplanned demand. In the aviation context, introducing frontlog scheduling creates a time buffer, which can complement and substitute for the capacity buffer in dealing with unplanned demand. Based on the impact of introducing the additional buffer in the airline context, we also arrive at novel theoretical insights into the management of uncertainty in operations.

To effectively cope with demand uncertainty, buffer trade-offs can be reduced to two dimensions with respect to demand: time and capacity. Where demand is predictable and has a due date, production can be performed before demand is imminent, resulting in a (perishable) frontlog buffer. Performing production after demand results in a backlog. To advance the research on buffer trade-offs, we point out that there is no need for the conceptual introduction of a third dimension, such as inventory. Instead, we can view inventories in manufacturing as a special case of a speculative time buffer, which we have denoted as the frontlog. Where demand can be expected, and the product is not immediately perishable, slack capacity can be used ahead of that demand to create a buffer. Just as in our case, the use of slack capacity to build the inventory removes the need to dedicate capacity for the point in time when demand is realized. This freed-up capacity can then be used to meet uncertain, urgent demand (cf. Biçer & Seifert, 2017). If no uncertain, urgent demand appears, then the capacity can again be used ahead of demand to buffer against future demand uncertainty. In other words, scheduling becomes a continuous activity to maintain a future rescheduling option, which we call the frontlog buffer. Reconceptualizing physical inventory as a frontlog time buffer, we make the conceptual argument (in contrast to Hopp and Spearman (2008) that there are only two types of buffers, time and capacity.

Through this reconceptualization, we see temporal buffering as a continuum, where the frontlog (temporal slack) crosses over to the backlog (temporal flexibility). Further, we view capacity buffering as a continuum ranging from overcapacity (capacity slack) to overutilization (capacity flexibility). This way, we can arrive at a generic, demand-centric two-dimensional representation of buffers in operations, against which specific situations can be mapped (Figure 12), corresponding to the alternatives upheld and available to specific operations. Mapping the current and suggested buffer management approaches of Nordic Airline on the chart illustrates the strategic nature of the proposed shift between buffers leading to the improvement observed in the simulation

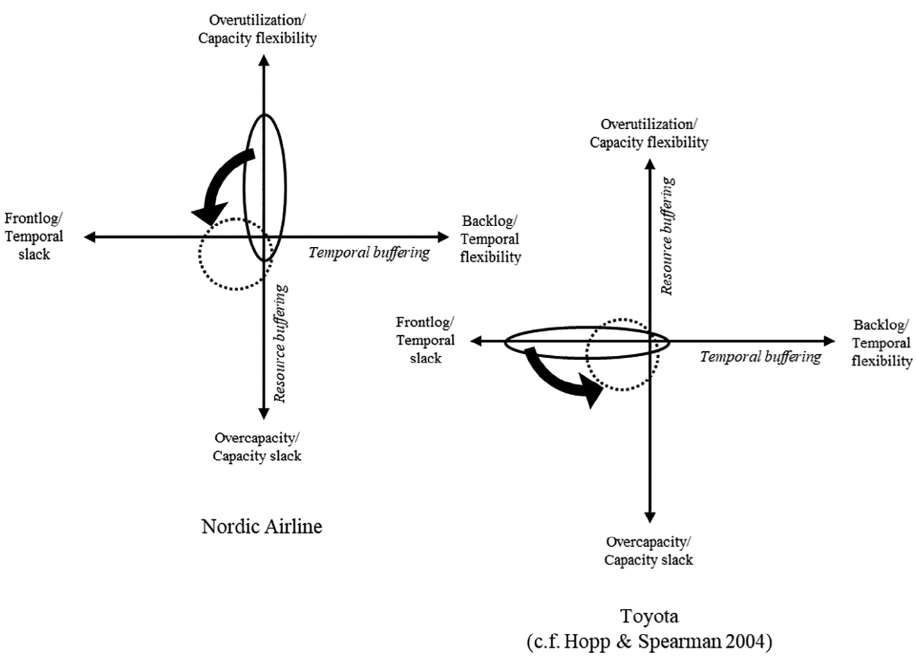

FIGURE 12 Conceptualization of buffer management options in operations

study. This shift is compared to Toyota's strategic shift from (over-)emphasis on frontlog (inventory) buffering to a balanced approach (Hopp & Spearman, 2008). "At a time when automotive plants generally ran three shifts a day, Toyota went to a two-shift schedule separated by 2-hr preventive maintenance (PM) periods. These PM periods served as capacity buffers to allow shifts to make up any shortfalls on their production quotas" (Hopp & Spearman, 2004, p. 145). As Toyota and Nordic Airline have different types of production systems, the companies' buffer shifts are not directly comparable. However, the two examples present a radical change in buffering, and these examples stand in stark contrast to the industry norm. We also note that Toyota's shift would not have been possible without several other (at the time) groundbreaking, variability-reducing developments in how operations were managed (cf. Hopp & Spearman, 2008, p. 310). These developments resemble the implementation challenges that Nordic Airline is addressing as part of the company's digitalization effort.

The two examples suggest that our reconceptualization provides a structure for further research in management of buffer trade-offs and balancing. Seeming distinctions between manufacturing and services (Akkermans & Voss, 2013)—in terms of inventory being available or not for buffer management—could be treated as context-specific attributes and constraints within the same conceptualization. Elaborating further on this conceptualization, we note that in terms of temporal buffering, any given act of production (resulting in one unit of product or service) happens before, after, or exactly at the time that demand arises. Respectively, in terms of capacity buffering, for any given time, we can determine whether production leaves slack capacity, whether overtime is used, or whether capacity is perfectly utilized. By knowing the state of the capacity buffer and by having an average figure for the temporal buffer, we can express an operation's buffering approach at a given point in time, as a single point on the resource/temporal buffering conceptualization plane (cf. Figure 12). Aggregating the position of this single point over a period of time would reveal something akin to the organization's buffering policy (illustrated by the area of the ovals and circles in Figure 12).

In terms of buffer management, JIT manufacturing is a temporal buffering policy that ideally uses neither a frontlog nor a backlog, as the capacities of the entire manufacturing system are operated to exactly fulfill actual customer demand. However, the ideal of lean also drives the reduction in capacity buffering, leading many firms to shun the use of buffering altogether (cf. Kroes, Manikas, & Gattiker, 2018). In this respect, firms that push the limits of their operations through a myopic reduction of both time and capacity buffers, often find their production systems worse off than they were to begin with (cf. Rappold & Yoho, 2008). Instead of shunning buffers, a firm must find an appropriate combination to cope effectively with uncertainty (Pagell et al., 2000). Shifting to a different type of buffer and creating new operational practices that rebalance the use of available time and capacity buffers present opportunities to enact possibly significant shifts in the performance frontier of an operation (Schmenner & Swink, 1998). The well-known performance outcomes of Toyota's development and implementation of JIT manufacturing and the results of this frontlog simulation for Nordic Airline are examples of moving the performance frontier through rebalancing the use of available time and capacity buffers.

7.2 | Design generalizability

As we identified designs similar to the frontlog in healthcare (Bowers & Mould, 2002; Helm et al., 2011), we can discuss what makes this design feasible in the particular contexts of maintenance and healthcare (Denyer, Tranfield, & van Aken, 2008). Through this discussion, we explore the conditions for transferring the design to other contexts, outlining design generalizability based on our available knowledge and theoretical understanding.

We highlight five contextual factors related to frontlog feasibility. First, a common denominator in the example contexts is that the same resources tend to planned and unplanned demand, creating a combination of demand uncertainty (implying that some form of buffer is needed) and demand predictability (implying that scheduling, that is, time buffering, is applicable). Second, related to the first point, capacity is largely generic, which is a precondition for being able to reactively shift capacity from planned work to unplanned work. This factor was reflected in Nordic Airline's effort to increase cross-training and certification of technicians and attempt to standardize the skillsets available in different shifts. The more generic the capacity with respect to demand, the better a capacity option works. Another approach with the same effect is to limit variety in demand which, in the aviation context, translates into operating a limited number of different types of aircraft. Third, the setup costs for (planned) work are low. The higher the setup costs for planned work, the higher the sunken costs that are potentially lost when planned work is postponed. Fourth, the ratio between total available capacity and work task capacity requirements is big enough. The greater the total workload handled, the smaller the proportional effect of an additional workload, assuming random occurrence. This was reflected in Nordic Airline representatives seeing the frontlog buffer as more valuable to small- and medium-sized carriers. This view implies that large carriers (minding the second contextual factor) could be expected to benefit most from smaller frontlog buffers. Fifth, the cost of a backlog is prohibitive for all work, making operations time-critical. If only the unplanned work is critical (as in many other maintenance settings), then the planned work can simply be postponed to a backlog. However, when considerable expense is associated with a planned work backlog (e.g., due to the regulation grounding planes in aviation or human suffering from queuing for treatment in healthcare), frontlog scheduling becomes a feasible alternative.

Finally, although not a direct requirement considering feasibility, the recurring nature of planned work, as in the study context, makes estimating the cost effect of the frontlog easier, compared with such contexts as healthcare (cf. Bowers & Mould, 2002). In addition, the predictable recurrence and duration of tasks make managing the frontlog buffer more approachable for management in aircraft line maintenance than in healthcare.

7.3 | Practical implications

Through the empirically grounded simulation model, we showed how utilizing a frontlog buffer in aircraft line maintenance can have a cost-reducing effect. A frontlog buffer is counterintuitive in the sense that it entails planning (and doing) more work, the cost of which is offset, on average, by additional operational resilience against unplanned work. For the operational context, we conservatively estimated a 4–5% cost reduction for maintenance operations. This estimate excluded likely indirect cost reductions in resource planning and cost-reducing effects of improved departure reliability and reduced major delays. For the operational context specificity of the results, we observed that the near-optimal frontlog buffer and its cost effect varied within the operational context, depending on fleet utilization (Appendix D). Lower fleet utilization rates imply more abundant maintenance windows. However, even for the lower utilization fleets (E90 and A32s), there was a business case for deliberate over-maintenance. This indicates that relying only on capacity buffering should be questioned in any airline operating context, if the capacity buffering costs are comparable to the study context.

In Nordic Airline, the current (unplanned) frontlog buffer was determined to account for 1% of the average total workload. This highlights that one can expect to find a frontlog buffer in any aircraft line maintenance operation, as performing all maintenance at the last possible opportunity is challenging in practice. Consequently, we cannot claim that a frontlog buffer is new per se. However, to address external uncertainty, the frontlog must be purposefully created and managed, making the time buffer explicit.

The results of the empirically grounded simulation model indicated a clear financial incentive for managing the frontlog buffer. However, the simulation results tell less about how the frontlog buffer should be designed and implemented. In our discussions with Nordic Airline, closing the digital feedback loop between production and planning emerged as the main practical challenge. Once this challenge is solved, planning, and dynamic rescheduling could be automated to the extent required for effective frontlog scheduling.

7.4 | Methodological implications

Design and implementation of the frontlog buffer are necessary steps for realizing the expected benefits. However, the theoretical insight offered by the frontlog does not depend on implementation. Based on this work, we emphasize that empirically grounded simulation (Chandrasekaran et al., 2018) is especially useful for theory development through design exploration. Where the design represents a profound change in the way operations are managed (making implementation a long-term project), or where the outcomes of the design are probabilistic (making evaluation a long-term project), a simulation offers a sufficiently rigorous route to pragmatic (pre-) validation, in accordance with the design science ethos (Romme, 2003). Nevertheless, simulations are only simplifications of reality. Novel designs, when implemented, can be expected to reveal unintended consequences (Holmström et al., 2009), both positive and negative, providing opportunities for further design science research.

We also argue that combining design exploration and empirically grounded simulation can spur innovation through academic OM research, by questioning what is rigid (Rappold & Yoho, 2008), given (Betts & Johnston, 2005), and inapplicable (Roemeling et al., 2017). With technology such as the Internet of Things and machine learning changing operational boundary conditions (Holmström et al., 2019), engaging in design exploration with practice provides an opportunity for relevant and theoretically novel research, pushing the state-of-the-art. In effect, it is a golden opportunity to move OM research to Pasteur's quadrant of practical relevance and theoretical insight (Stokes, 2011), as highlighted in the introduction.

7.5 | Limitations and future research

The results of this study should be seen as indicative of further practical and theoretical contributions from research on buffer balancing and trade-offs. Further research on buffer management in different contexts is needed—unearthing and questioning buffer management mind-sets, exploring and developing new practices and approaches for digitalized buffer balancing enabled by dynamic scheduling, and consolidating evidence of outcomes from strategic buffer shifts. The representation of operational buffering alternatives along capacity and time dimensions provides a theoretical foundation on which this further research can be built. The representation itself should also be further developed through strengthening its anchoring in literature on established means of buffering, beyond the initial conceptualization we offered. For example, we may consider demand management (Klassen & Rohleder, 2002) as a potential equivalent to a frontlog. However, as demand management implies a temporal shift in demand rather than in production, there might be a case for developing a dual representation of buffering, one demand-centric (as the one presented here) and the other production-centric (yet to be developed). Further, we considered demand uncertainty solely in terms of volume. However, mix is an equally important consideration, especially in manufacturing. At the very least, nonetheless, this research indicates a need to revisit the OM underpinnings of the supposed divide between manufacturing and services.

As for limitations and further research on aircraft line maintenance, we examined a hub-and-spoke context in which maintenance is always performed at the hub, and where the workload is concentrated in distinct peaks during the day, with implications for capacity buffering costs. As hub-and-spoke networks can be argued to simplify some aviation maintenance problems (Barnhart, Belobaba, & Odoni, 2003), further research is needed to study the effects of utilizing a frontlog buffer in other aviation settings, such as in point-to-point operations and military aviation. Future research could also explore the implications of considering frontlog buffering in connection with maintenance routing, as the latter determines the available maintenance windows utilized by the former.

The direct applicability of inventory management principles in frontlog buffer management cannot be assumed, despite our reconceptualization of inventory, as a special case of frontlog. However, the design of a frontlog could draw on previous research related to containing the bullwhip effect. We approached the uncertainty problem through a variance lens (Towill, Zhou, & Disney, 2007). This view indicates control theory (cf. Ortega & Lin, 2004) could offer insights for detailed implementation and design improvement. This idea is further fueled by our observation that the airline experienced persistent gradual build-up (or amplification) of planned workload within the upcoming 3–4 days.

Based on the present results, we expect the optimal frontlog buffer to depend on at least fleet size, fleet utilization, aircraft reliability, costs of capacity buffering, and capacity sharing among fleets. In addition, as fleet utilization exerts seasonal variation, introducing a frontlog buffer in Nordic Airline will likely have strategic, tactical, and operational dimensions. Considering the purpose of this study, a fixed frontlog was sufficient to show the merits of introducing a deliberately managed frontlog. Against this backdrop, further research could develop a dynamic solution (cf. Biçer & Seifert, 2017) to this problem, for example, through using reinforcement learning or approximate dynamic programming approaches (cf. Mattila & Virtanen, 2011).

8 | CONCLUSION

We set out to explore solutions to a practical problem faced by Nordic Airline. Based on an empirically grounded simulation model, we here propose a counterintuitive solution that represents a new way of doing (aircraft) maintenance. As a result of our study, we have come to question assumptions widely held both in practice and academia. We found that, in aircraft maintenance, sole reliance on capacity buffering is likely not a cost-effective approach for handling uncertainty. We also found that the frontlog buffer shares properties with manufacturing inventory, allowing us to conceptualize inventory as a special case of frontlog time buffering. Based on this reconceptualization, we propose a generic framework of how operations are buffered against demand uncertainty. We argue that time and capacity are the two principal buffers in operations, and that the latter can be considered as a continuum ranging from overcapacity to overutilization, while the former can be considered as a continuum ranging from frontlog to backlog. These two dimensions meet at the perfectly aligned production and demand, any deviation from which implies costs of buffering. We argue that this conceptualization is a long overdue advance for research to understand how managers can balance buffers in operations to effectively cope with uncertainty.

ORCID

Mikael Öhman https://orcid.org/0000-0001-5168-7066

Jan Holmström https://orcid.org/0000-0002-2596-0337

ENDNOTES

[1]This research has its roots in the early '60s. McCall (1963) described the operating characteristics of opportunistic inspection and replacement, and Radner and Jorgenson (1963) provided proof of the optimality of the policy. Rust (1987) provided an analytical solution to the opportunistic replacement problem, supported by empirical data for bus engine replacement decisions.

[2]*Maintenance routing* is a planning activity where the maintenance organization assigns physical aircraft to flights, which are designated by traffic planning as origin, destination, departure time, arrival time, and aircraft type. A *work package* is a bundle of maintenance tasks to be performed in a maintenance window. A *maintenance task* is a sequence of work and inspections performed by the technician.

[3]The exact number is omitted here, as the airline considers it confidential.

REFERENCES

Ab-Samat, H., & Kamaruddin, S. (2014). Opportunistic maintenance (OM) as a new advancement in maintenance approaches: A review. *Journal of Quality in Maintenance Engineering*, *20*, 98–121.

Adamides, E., Stamboulis, Y., & Varelis, A. (2004). Model-based assessment of military aircraft engine maintenance systems. *The Journal of the Operational Research Society*, *55*, 957–967. https://doi.org/10.1057/palgrave.jors.2601756

Ahmad, R., & Kamaruddin, S. (2013). Maintenance decision-making process for a multi-component production unit using output-based maintenance technique: A case study for non-repairable two serial components' unit. *International Journal of Performability Engineering*, *9*, 305–319.

Ahmed, A. H., & Poojari, C. A. (2008). An overview of the issues in the airline industry and the role of optimization models and algorithms. *The Journal of the Operational Research Society*, *59*, 267–277. https://doi.org/10.1057/palgrave.jors.2602350

Ahuja, I. P. S., & Khamba, J. S. (2008). Total productive maintenance: Literature review and directions. *International Journal of Quality & Reliability Managemen*, *25*, 709–756.

Akkan, C. (2015). Improving schedule stability in single-machine rescheduling for new operation insertion. *Computers and Operations Research*, *64*, 198–209. https://doi.org/10.1016/j.cor.2015.05.015

Akkermans, H., & Vos, B. (2003). Amplification in service supply chains: An exploratory case study from the telecom industry. *Production and Operations Management*, *12*, 204–223.

Akkermans, H., & Voss, C. (2013). The service bullwhip effect. *International Journal of Operations & Production Management*, *33*, 765–788. https://doi.org/10.1108/IJOPM-10-2012-0402

Alfares, H. K. (1999). Aircraft maintenance workforce scheduling: A case study. *Journal of Quality in Maintenance Engineering*, *5*, 78–89.

Anderson, E. G., Morrice, D. J., & Lundeen, G. (2005). The "physics" of capacity and backlog management in service and custom manufacturing supply chains. *System Dynamics Review*, *21*, 217–247. https://doi.org/10.1002/sdr.319

Angalakudati, M., Balwani, S., Calzada, J., Chatterjee, B., Perakis, G., Raad, N., & Uichanco, J. (2014). Business analytics for flexible resource allocation under random emergencies. *Management Science*, *60*, 1552–1573. https://doi.org/10.1287/mnsc.2014.1919

Askin, R. G., & Krishnan, S. (2009). Defining inventory control points in multiproduct stochastic pull systems. *International Journal of Production Economics*, *120*, 418–429. https://doi.org/10.1016/j.ijpe.2008.11.020

Barnhart, C., Belobaba, P., & Odoni, A. R. (2003). Applications of operations research in the air transport industry. *Transportation Science*, *37*, 368–391. https://doi.org/10.1287/trsc.37.4.368.23276

Başdere, M., & Bilge, Ü. (2014). Operational aircraft maintenance routing problem with remaining time consideration. *European Journal of Operational Research*, *235*, 315–328. https://doi.org/10.1016/j.ejor.2013.10.066

Bean, J. C., Birge, J. R., Mittenthal, J., & Noon, C. E. (1991). Matchup scheduling with multiple resources, release dates and disruptions. *Operations Research*, *39*, 470–483.

Beliën, J., Cardoen, B., & Demeulemeester, E. (2012). Improving workforce scheduling of aircraft line maintenance at Sabena technics. *Interfaces (Providence).*, *42*, 352–364. https://doi.org/10.1287/inte.1110.0585

Berglund, M., & Karltun, J. (2007). Human, technological and organizational aspects influencing the production scheduling process. *International Journal of Production Economics*, *110*, 160–174. https://doi.org/10.1016/j.ijpe.2007.02.024

Bernardes, E. S., & Hanna, M. D. (2009). A theoretical review of flexibility, agility and responsiveness in the operations management literature: Toward a conceptual definition of customer responsiveness. *International Journal of Operations & Production Management*, *29*, 30–53.

Betts, J. M., & Johnston, R. B. (2005). Just-in-time component replenishment decisions for assemble-to-order manufacturing under capital constraint and stochastic demand. *International Journal of Production Economics*, *95*, 51–70. https://doi.org/10.1016/j.ijpe.2003.10.020

Biçer, I., & Seifert, R. W. (2017). Optimal dynamic order scheduling under capacity constraints given demand-forecast evolution. *Production and Operations Management*, *26*, 2266–2286. https://doi.org/10.1111/poms.12759

Black, G. W., & McKay, K. N. (2012). A parallel machine extension to aversion dynamics scheduling. *International Journal of Industrial Engineering Computations*, *3*, 525–534. https://doi.org/10.5267/j.ijiec.2012.03.002

Blakeley, F., Argüello, B., Cao, B., Hall, W., & Knolmajer, J. (2003). Optimizing periodic maintenance operations for Schindler elevator corporation. *Interfaces (Providence).*, *33*, 67–79. https://doi.org/10.1287/inte.33.1.67.12722

Boeret, N. J. (1977). Air Canada saves with aircraft maintenance scheduling. *Interfaces (Providence).*, *7*, 1–13.

Bourland, K. E., & Yano, C. A. (1994). The strategic use of capacity slack in the economic lot scheduling problem with random demand. *Management Science*, *40*, 1690–1704.

Bowers, J., & Mould, G. (2002). The deferrable elective patient a means of reducing waiting-lists in orthopaedics. *Journal of Management in Medicine*, *16*, 150–158. https://doi.org/10.1108/02689230210434899

Callewaert, P., Verhagen, W. J. C., & Curran, R. (2018). Integrating maintenance work progress monitoring into aircraft maintenance planning decision support. *Transportation Research Procedia*, *29*, 58–69. https://doi.org/10.1016/j.trpro.2018.02.006

Caputo, M. (1996). Uncertainty, flexibility and buffers in the management of the firm operating system. *Production Planning and Control*, *7*, 518–528.

Chandrasekaran, A., Linderman, K., & Sting, F. J. (2018). Avoiding epistemological silos and empirical elephants in OM: How to combine empirical and simulation methods? *Journal of Operations Management*, *63*, 1–5. https://doi.org/10.1016/j.jom.2018.11.003

Cho, S., & Lazaro, A. (2010). Control theoretic model using PID controller for just-in-time production scheduling. *International Journal of Advanced Manufacturing Technology*, *51*, 699–709. https://doi.org/10.1007/s00170-010-2639-x

Cobb, R. (1995). Modeling aircraft repair turntime - simulation supports maintenance marketing efforts. *Journal of Air Transport Management*, *2*, 25–32. https://doi.org/10.1016/0969-6997(95)00024-6

Coghlan, D., & Brannick, T. (2001). *Doing action research in your own organization*. London: Sage.

Cousens, A., Szwejczewski, M., & Sweeney, M. (2009). A process for managing manufacturing flexibility. *International Journal of Operations & Production Management*, *29*, 357–385.

D'Souza, D. E., & Williams, F. P. (2000). Toward a taxonomy of manufacturing flexibility dimensions. *Journal of Operations Management*, *18*, 577–593. https://doi.org/10.1016/S0272-6963(00)00036-X

Davies, J., Mabin, V. J., & Balderstone, S. J. (2005). The theory of constraints: A methodology apart? - a comparison with selected OR/MS methodologies. *Omega*, *33*, 506–524. https://doi.org/10.1016/j.omega.2004.07.015

Dekker, R., & Dijkstra, M. C. (1992). Opportunity-based age replacement: Exponentially distributed times between opportunities. *Naval Research Logistics*, *39*, 175–190.

Dekker, R., & Smeitink, E. (1991). Opportunity-based block replacement. *European Journal of Operational Research*, *53*, 46–63.

Dekker, R., & Smeitink, E. (1994). Preventive maintenance at opportunities of restricted duration. *Naval Research Logistics*, *41*, 335–353.

den Hengst, M., de Vreede, G. J., & Maghnouji, R. (2007). Using soft OR principles for collaborative simulation: A case study in the Dutch airline industry. *The Journal of the Operational Research Society*, *58*, 669–682. https://doi.org/10.1057/palgrave.jors.2602353

Denyer, D., Tranfield, D., & van Aken, J. E. (2008). Developing design propositions through research synthesis. *Organization Studies*, *29*, 393–413. https://doi.org/10.1177/0170840607088020

Dijkstra, M. C., Kroon, L. G., Salomon, M., van Nunen, J. A. E. E., & Van Wassenhove, L. N. (1994). Planning the size and organization of KLM's aircraft maintenance personnel. *Interfaces (Providence)*., *24*, 47–58. https://doi.org/10.1287/inte.24.6.47

Forrester, J. W. (1992). Policies, decisions and information sources for modeling. *European Journal of Operational Research*, *59*, 42–63. https://doi.org/10.1016/0377-2217(92)90006-U

Froger, A., Gendreau, M., Mendoza, J. E., Pinson, É., & Rousseau, L. M. (2016). Maintenance scheduling in the electricity industry: A literature review. *European Journal of Operational Research*, *251*, 695–706. https://doi.org/10.1016/j.ejor.2015.08.045

Gatland, R., Yang, E., & Buxton, K. (1997). Solving engine maintenance capacity problems with simulation. In S. Andradottir, K. J. Healy, D. H. Withers, & B. L. Nelson (Eds.), *Proceedings of the Winter Simulation Conference* (pp. 892–899). Picataway, NJ: IEEE.

Geary, S., Disney, S. M., & Towill, D. R. (2006). On bullwhip in supply chains—Historical review, present practice and expected future impact. *International Journal of Production Economics*, *101*, 2–18. https://doi.org/10.1016/j.ijpe.2005.05.009

Gerwin, D. (1993). Manufacturing flexibility: A strategic perspective. *Management Science*, *39*, 395–410. https://doi.org/10.1007/sl0869-007-9037-x

Goldratt, E. M. (1990). *What is this thing called theory of constraints and how should it be implemented?* Croton-on-Hudson, NY: North River Press.

Grigoriev, A., van de Klundert, J., & Spieksma, F. C. R. (2006). Modeling and solving the periodic maintenance problem. *European Journal of Operational Research*, *172*, 783–797. https://doi.org/10.1016/j.ejor.2004.11.013

Groop, J., Ketokivi, M., Gupta, M., & Holmström, J. (2017). Improving home care: Knowledge creation through engagement and design. *Journal of Operations Management*, *53–56*, 9–22. https://doi.org/10.1016/j.jom.2017.11.001

Guarnieri, J., Johnson, A. W., & Swartz, S. M. (2005). A maintenance resources capacity estimator. *The Journal of the Operational Research Society*, *57*, 1188–1196. https://doi.org/10.1057/palgrave.jors.2602101

Gupta, S., & Lulli, G. (2014). Dynamic resource allocation: A flexible and tractable modeling framework. *European Journal of Operational Research*, *236*, 14–26. https://doi.org/10.1016/j.ejor.2013.10.063

Hedenstierna, C. P. T., Disney, S. M., Eyers, D. R., Holmström, J., Syntetos, A. A., & Wang, X. (2019). Economies of collaboration in build-to-model operations. *Journal of Operations Management*, *1014*, 753–773. https://doi.org/10.1002/joom.1014

Helm, J. E., AhmadBeygi, S., & van Oyen, M. P. (2011). Design and analysis of hospital admission control for operational effectiveness. *Production and Operations Management*, *20*, 359–374.

Holmström, J., Holweg, M., Lawson, B., Pil, F. K., & Wagner, S. M. (2019). The digitalization of operations and supply chain management: Theoretical and methodological implications. *Journal of Operations Management*, *65*, 728–734. https://doi.org/10.1002/joom.1073

Holmström, J., Ketokivi, M., & Hameri, A.-P. (2009). Bridging practice and theory: A design science approach. *Decision Sciences*, *40*, 65–87. https://doi.org/10.1111/j.1540-5915.2008.00221.x

Hopp, W. J., & Spearman, M. L. (2004). To pull or not to pull: What is the question? *Manufacturing and Service Operations Management*, *6*, 133–148. https://doi.org/10.1287/msom.1030.0028

Hopp, W. J., & Spearman, M. L. (2008). *Factory physics* (3rd ed.). Long Grove, IL: Waveland Press.

Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, *20*, 1483–1510. https://doi.org/10.1016/j.ymssp.2005.09.012

Kelly, A. (2006). *Strategic maintenance planning*. Jordan Hill, GBR: Butterworth-Heineman.

Khazraei, K., & Deuse, J. (2011). A strategic standpoint on maintenance taxonomy. *Journal of Facilities Management*, *9*, 96–113.

Kilpi, J., Töyli, J., & Vepsäläinen, A. (2009). Cooperative strategies for the availability service of repairable aircraft components. *International Journal of Production Economics*, *117*, 360–370. https://doi.org/10.1016/j.ijpe.2008.12.001

Klassen, K. J., & Rohleder, T. R. (1996). Scheduling outpatient appointments in a dynamic environment. *Journal of Operations Management*, *14*, 83–101. https://doi.org/10.1016/0272-6963(95)00044-5

Klassen, K. J., & Rohleder, T. R. (2002). Demand and capacity management decisions in services: How they impact on one another. *International Journal of Operations & Production Management*, *22*, 527–548. https://doi.org/10.1108/MRR-09-2015-0216

Koufteros, X. A., Vonderembse, M. A., & Doll, W. J. (1998). Developing measures of time-based manufacturing. *Journal of Operations Management*, *16*, 21–41. https://doi.org/10.1016/S0272-6963(97)00027-2

Kroes, J. R., Manikas, A. S., & Gattiker, T. F. (2018). Operational leanness and retail firm performance since 1980. *International Journal of Production Economics*, *197*, 262–274. https://doi.org/10.1016/j.ijpe.2018.01.006

Laganga, L. R. (2011). Lean service operations: Reflections and new directions for capacity expansion in outpatient clinics. *Journal of Operations Management*, *29*, 422–433. https://doi.org/10.1016/j.jom.2010.12.005

Law, A. M. (2003). How to conduct a successful simulation study. In S. Chick, P. J. Sánches, D. Ferrin, & D. J. Morrice (Eds.), *Proceedings of the 2003 Winter Simulation Conference* (pp. 66–70). Picataway, NJ: IEEE.

Lee, J., Wu, F., Zhao, W., Ghaffari, M., Liao, L., & Siegel, D. (2014). Prognostics and health management design for rotary machinery systems—Reviews, methodology and applications. *Mechanical Systems and Signal Processing*, *42*, 314–334. https://doi.org/10.1016/j.ymssp.2013.06.004

Lu, Z., Cui, W., & Han, X. (2015). Integrated production and preventive maintenance scheduling for a single machine with failure uncertainty. *Computers and Industrial Engineering*, *80*, 236–244. https://doi.org/10.1016/j.cie.2014.12.017

Mattila, V., & Virtanen, K. (2011). Scheduling fighter aircraft maintenance with reinforcement learning. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, & M. Fu (Eds.), *Proceedings of the 2011 Winter Simulation Conference* (pp. 2535–2546). Picataway, NJ: IEEE.

Mattila, V., & Virtanen, K. (2014). Maintenance scheduling of a fleet of fighter aircraft through multi-objective simulation-optimization. *Simulation*, *90*, 1023–1040. https://doi.org/10.1177/0037549714540008

Mattila, V., Virtanen, K., & Raivio, T. (2008). Improving maintenance decision making in the Finnish air force through simulation. *Interfaces (Providence).*, *38*, 187–201. https://doi.org/10.1287/inte.1080.0349

McCall, J. J. (1963). Operating characteristics of opportunistic replacement and inspection policies. *Management Science*, *10*, 85–97. https://doi.org/10.1287/mnsc.10.1.85

McKay, K. N., & Wiers, V. C. S. (1999). Unifying the theory and practice of production scheduling. *Journal of Manufacturing Systems*, *18*, 241–255. https://doi.org/10.1016/S0278-6125(00)86628-5

McKone, K. E., Schroeder, R. G., & Cua, K. O. (1999). Total productive maintenance: A contextual view. *Journal of Operations Management*, *17*, 123–144. https://doi.org/10.1016/S0272-6963(98)00039-4

McKone, K. E., Schroeder, R. G., & Cua, K. O. (2001). The impact of total productive maintenance practices on manufacturing performance. *Journal of Operations Management*, *19*, 39–58. https://doi.org/10.1016/S0272-6963(00)00030-9

Mehta, S. V., & Uzsoy, R. M. (1998). Predictable scheduling of a job shop subject to breakdowns. *IEEE Transactions on Robotics and Automation*, *14*, 365–378. https://doi.org/10.1109/70.678447

Mertins, K., & Lewandrowski, U. (1999). Inventory safety stocks of Kanban control systems. *Production Planning and Control*, *10*, 520–529. https://doi.org/10.1080/095372899232812

Moeller, S. (2010). Characteristics of services—A new approach uncovers their value. *Journal of Services Marketing*, *24*, 359–368. https://doi.org/10.1108/08876041011060468

Newman, W. R., Hanna, M., & Maffei, M. J. (1993). Dealing with the uncertainties of manufacturing: Flexibility, buffers and integration. *International Journal of Operations & Production Management*, *13*, 19–34. https://doi.org/10.1108/01443579310023972

Nowlan, F. S., & Heap, H. F. (1978). *Reliability-centered maintenance*. San Francisco, CA: Dolby, Access Press.

Öhman, M., Finne, M., & Holmström, J. (2015). Measuring service outcomes for adaptive preventive maintenance. *International Journal of Production Economics*, *170*, 457–467. https://doi.org/10.1016/j.ijpe.2015.06.020

Öhman, M., Laine, M., & Holmström, J. (2016). Aircraft fleet maintenance scheduling: heuristics for workload balancing. *19th International Working Seminar on Production Economics*, *3*, 399–410.

Oliva, R. (2019). Intervention as a research strategy. *Journal of Operations Management*, *65*, 710–724. https://doi.org/10.1002/joom.1065

Ortega, M., & Lin, L. (2004). Control theory applications to the production–inventory problem: A review. *International Journal of Production Research*, *42*, 2303–2322. https://doi.org/10.1080/00207540410001666260

Pagell, M., Newman, W. R., Hanna, M., Krause, D., Mark, D., & Daniel, R. (2000). Uncertainty, flexibility, and buffers: Three case studies. *Production and Inventory Management Journal*, *41*, 35–43.

Paz, N. M., & Leigh, W. (1994). Maintenance scheduling: Issues, results and research needs. *International Journal of Operations & Production Management*, *14*, 47–69.

Pinjala, S. K., Pintelon, L., & Vereecke, A. (2006). An empirical investigation on the relationship between business and maintenance strategies. *International Journal of Production Economics*, *104*, 214–229. https://doi.org/10.1016/j.ijpe.2004.12.024

Radner, R., & Jorgenson, D. W. (1963). Opportunistic replacement of a single part in the presence of several monitored parts. *Management Science*, *10*, 70–84.

Rappold, J. A., & Yoho, K. D. (2008). A model for level-loading production in the process industries when demand is stochastic. *Production Planning and Control*, *19*, 686–701. https://doi.org/10.1080/09537280802573726

Rausand, M. (1998). Reliability centered maintenance. *Reliability Engineering and System Safety*, *60*, 121–132.

Ray, S., & Jewkes, E. M. (2003). Customer lead time management when both demand and price are lead time sensitive. *European Journal of Operational Research*, *153*, 769–781. https://doi.org/10.1016/S0377-2217(02)00655-0

Roemeling, O. P., Land, M., & Ahaus, K. (2017). Does lean cure variability in health care? *International Journal of Operations &*

Production Management, *37*, 1229–1245. https://doi.org/10.1108/MRR-09-2015-0216

Romme, A. G. L. (2003). Making a difference: Organization as design. *Organization Science*, *14*, 558–573. https://doi.org/10.1287/orsc.14.5.558.16769

Rondeau, P. J., Vonderembse, M. A., & Ragu-Nathan, T. S. (2000). Exploring work system practices for time-based manufacturers: Their impact on competitive capabilities. *Journal of Operations Management*, *18*, 509–529. https://doi.org/10.1016/S0272-6963(00)00037-1

Rosqvist, T., Laakso, K., & Reunanen, M. (2009). Value-driven maintenance planning for a production plant. *Reliability Engineering and System Safety*, *94*, 97–110. https://doi.org/10.1016/j.ress.2007.03.018

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, *55*, 999–1033.

Safaei, N., Banjevic, D., & Jardine, A. K. S. (2011). Bi-objective workforce-constrained maintenance scheduling: A case study. *The Journal of the Operational Research Society*, *62*, 1005–1018. https://doi.org/10.1057/jors.2010.51

Sarac, A., Batta, R., & Rump, C. M. (2006). A branch-and-price approach for operational aircraft maintenance routing. *European Journal of Operational Research*, *175*, 1850–1869. https://doi.org/10.1016/j.ejor.2004.10.033

Schmenner, R. W., & Swink, M. L. (1998). On theory in operations management. *Journal of Operations Management*, *17*, 97–113. https://doi.org/10.1016/S0272-6963(98)00028-X

Shah, R., & Ward, P. T. (2003). Lean manufacturing: Context, practice bundles, and performance. *Journal of Operations Management*, *21*, 129–149. https://doi.org/10.1016/S0272-6963(02)00108-0

Sherwin, D. J. (1999). Age-based opportunity maintenance. *Journal of Quality in Maintenance Engineering*, *5*, 221–235.

Siedlak, D. J. L., Pinon, O. J., Robertson, B. E., & Mavris, D. N. (2018). Robust simulation-based scheduling methodology to reduce the impact of manual installation tasks on low-volume aerospace production flows. *Journal of Manufacturing Systems*, *46*, 193–207. https://doi.org/10.1016/j.jmsy.2017.12.006

Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). Cambridge, MA: MIT Press.

Smith, S. F., Ow, P. S., Potvin, J.-Y., Muscettola, N., & Matthys, D. C. (1990). An integrated framework for generating and revising factory schedules. *The Journal of the Operational Research Society*, *41*, 539–552.

Sriram, C., & Haghani, A. (2003). An optimization model for aircraft maintenance scheduling and re-assignment. *Transportation Research: Part A Policy and Practice*, *37*, 29–48. https://doi.org/10.1016/S0965-8564(02)00004-6

Stenström, C., Parida, A., Kumar, U., & Galar, D. (2013). Performance indicators and terminology for value driven maintenance. *Journal of Quality in Maintenance Engineering*, *19*, 222–232.

Stokes, D. E. (2011). *Pasteur's quadrant: Basic science and technological innovation*. Washington, DC: Brookings Institution Press.

Swanson, L. (2001). Linking maintenance strategies to performance. *International Journal of Production Economics*, *70*, 237–244.

Thürer, M., Tomašević, I., & Stevenson, M. (2017). On the meaning of 'waste': Review and definition. *Production Planning and Control*, *28*, 244–255. https://doi.org/10.1080/09537287.2016.1264640

Towill, D. R., Zhou, L., & Disney, S. M. (2007). Reducing the bullwhip effect: Looking through the appropriate lens. *International Journal of Production Economics*, *108*, 444–453. https://doi.org/10.1016/j.ijpe.2006.12.024

Tu, Q., Vonderembse, M. A., Ragu-Nathan, T. S., & Sharkey, T. W. (2006). Absorptive capacity: Enhancing the assimilation of time-based manufacturing practices. *Journal of Operations Management*, *24*, 692–710. https://doi.org/10.1016/j.jom.2005.05.004

van Aken, J. E. (2004). Management research based on the paradigm of the design sciences: The quest for field-tested and grounded technological rules. *Journal of Management Studies*, *41*, 219–246. https://doi.org/10.1111/j.1467-6486.2004.00430.x

van Aken, J. E., Chandrasekaran, A., & Halman, J. (2016). Conducting and publishing design science research: Inaugural essay of the design science department of the Journal of operations management. *Journal of Operations Management*, *47–48*, 1–8. https://doi.org/10.1016/j.jom.2016.06.004

van den Bergh, J., de Bruecker, P., Beliën, J., & Peeters, J. (2013). *Aircraft maintenance operations: state of the art (No. 2013/09)*. Brussels: HUBrussel.

Veldman, J., Wortmann, H., & Klingenberg, W. (2011). Typology of condition based maintenance. *Journal of Quality in Maintenance Engineering*, *17*, 183–202. https://doi.org/10.1108/13552511111134600

Waeyenbergh, G., & Pintelon, L. (2002). A framework for maintenance concept development. *International Journal of Production Economics*, *77*, 299–313.

Williams, B. D., & Tokar, T. (2012). A review of inventory management research in major logistics journals: Themes and future directions. *International Journal of Logistics Management*, *19*, 212–232.

Yang, T., Hsieh, C. H., & Cheng, B. Y. (2011). Lean-pull strategy in a re-entrant manufacturing environment: A pilot study for TFT-LCD array manufacturing. *International Journal of Production Research*, *49*, 1511–1529. https://doi.org/10.1080/00207540903567333

Zhang, Q., Vonderembse, M. A., & Lim, J. S. (2003). Manufacturing flexibility: Defining and analyzing relationships among competence, capability, and customer satisfaction. *Journal of Operations Management*, *21*, 173–191. https://doi.org/10.1016/S0272-6963(02)00067-0

Zijm, W. H. M., & Buitenhek, R. (1996). Capacity planning and lead time management. *International Journal of Production Economics*, *46–47*, 165–179. https://doi.org/10.1016/0925-5273(95)00161-1

How to cite this article: Öhman M, Hiltunen M, Virtanen K, Holmström J. Frontlog scheduling in aircraft line maintenance: From explorative solution design to theoretical insight into buffer management. *J Oper Manag*. 2020;1–32. https://doi.org/10.1002/joom.1108

APPENDIX A

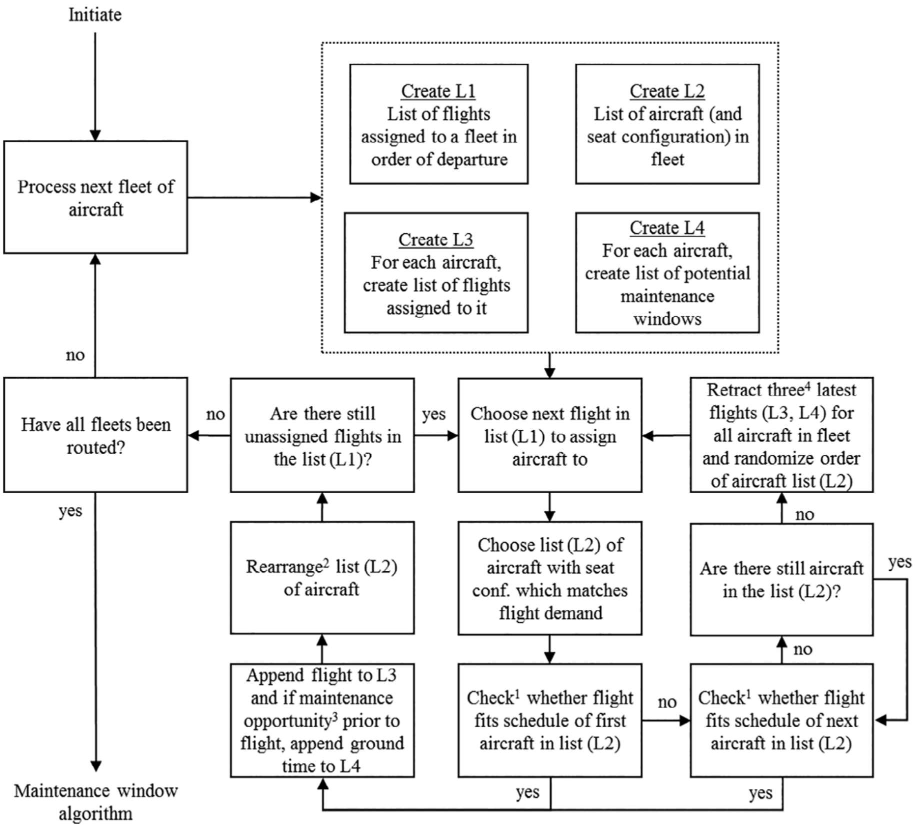

FIGURE A1 Maintenance routing algorithm. [1]Checks whether (1) the aircraft is at the station from which the flight departs (as the model is initialized, no flights are assigned, which means that the aircraft accepts any departure station); (2) enough time (1 hr) has passed since the aircraft landed after its previous flight; and (3) no maintenance expires during the flight (or the return flight if the destination is not the hub). [2]Rearranges the list (L2) so that the aircraft located at an outstation are stacked on top and those at the hub are stacked below, both in ascending order, based on the maintenance expiration. This is done to ensure that if maintenance windows are scarce, they are divided evenly among the aircraft, and the aircraft spend as little time at outstations as possible. [3]Checks whether (1) aircraft are at the hub and (2) whether there is sufficient time for maintenance. [4]If this does not lead to a feasible solution, the number of flights retracted is gradually increased

APPENDIX B

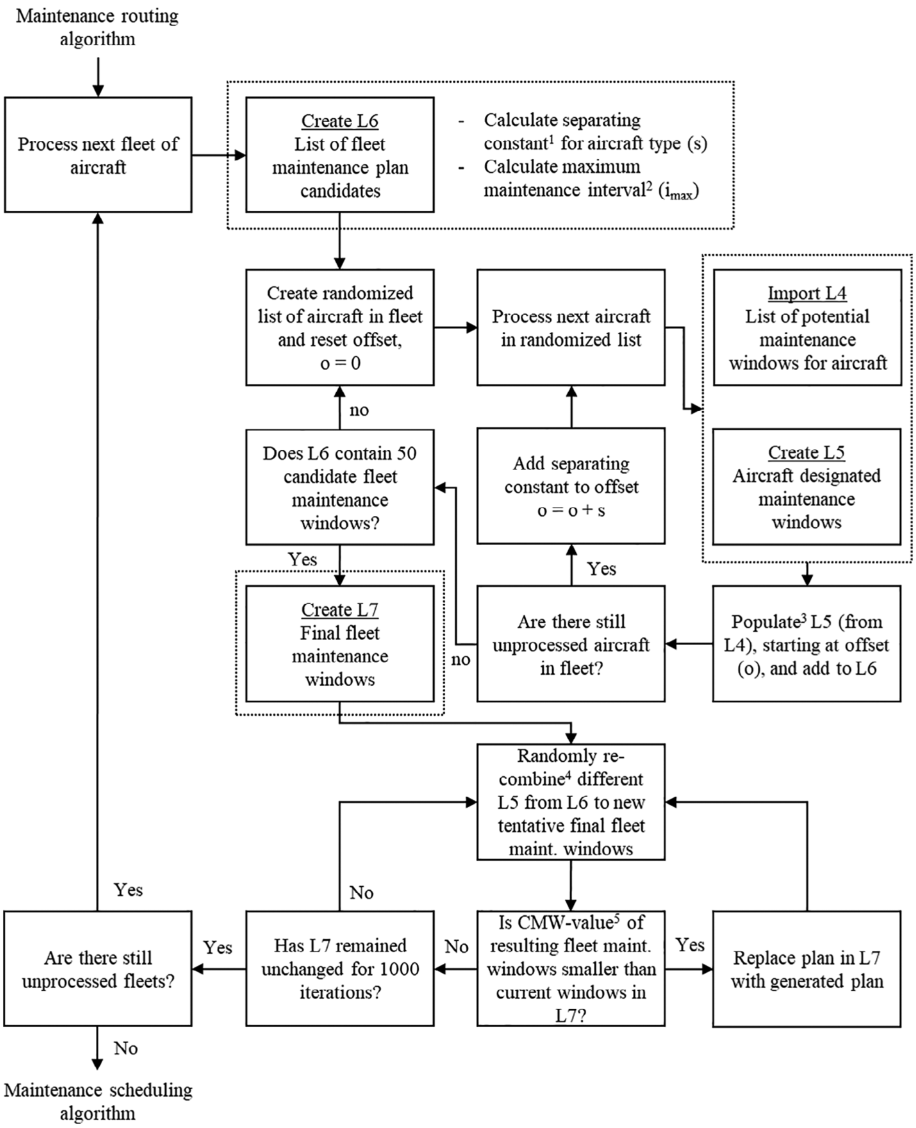

FIGURE B1 Maintenance window algorithm. [1]The *separating constant* is derived from the first processed aircraft in the fleet and refined as other aircraft of the fleet are processed for the first time. In practice, the separating constant is the average time between two maintenance windows, divided by the number of aircraft in the fleet. The resulting value is subsequently used (in the remaining 49 candidate iterations for the fleet) in ensuring that the windows are evenly distributed from the fleet perspective. [2]The maximum maintenance interval corresponds to the shortest recurring maintenance task interval, that is, the longest an aircraft can operate without being maintained. [3]The maintenance plan is populated through reducing potential maintenance windows to actual maintenance windows. This is done by choosing the first maintenance window according to the offset. Then, successive potential maintenance windows are removed until the time between the last designated maintenance window and the next potential maintenance window exceeds the maximum maintenance interval. At this point, the model jumps over one maintenance window (designating it as actual) and continues in a similar vein until the maximum interval again would be exceeded. This procedure leaves close to the minimum number of actual designated maintenance windows. For the next aircraft of the same type, the algorithm adds a separating constant to the previous starting offset, which acts as an approximation of spreading the maintenance windows evenly across time. [4]Randomly constructing different combinations of individual aircraft maintenance plans derived from the 50 fleet maintenance plan candidates. [5]The Concurrent Maintenance Windows (CMWs) value is calculated based on a time series, which expresses the number of concurrent maintenance windows with a 30-min resolution. The CMW value is the sum of the cubed data points of the time series

APPENDIX C

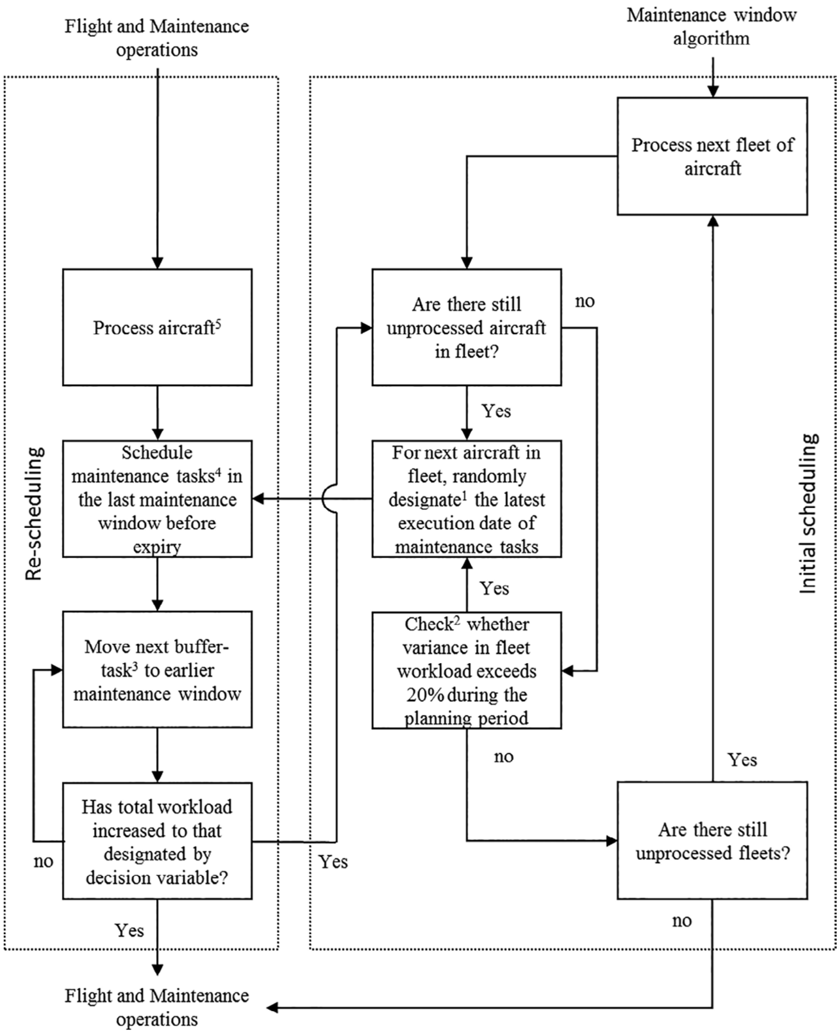

FIGURE C1 Maintenance scheduling algorithm. [1]Randomly designated so that the first due date of the recurring maintenance task falls somewhere at or after the first designated maintenance window. [2]Checks whether the workload variance exceeds 20% within the planning period (2 months), based on a 30-min resolution, taking into account task recurrence. This is done to ensure a realistic and sufficiently uniform planned workload for the initial state of the simulation. [3]Recurring maintenance tasks are initially scheduled to the last possible designated maintenance window. The *frontlog* buffer is applied by moving planned recurring tasks (favoring those with long maintenance intervals) to earlier work packages one by one while updating task recurrences. This procedure is performed until the total workload in the schedule has risen by the amount set by the decision variable. [4]In the simulation, any emergent maintenance tasks are considered with recurring tasks, with the due date of emergent tasks determined by their criticality. Further, the scheduling procedure considers task-related constraints, such as mutually exclusive tasks (e.g., some engine-related tasks on twin-engine aircraft), and whether the work can be performed on apron or not. [5]During flight and maintenance operations, the algorithm is initiated whenever a scheduled maintenance work package is completed or due to the occurrence of emergent maintenance task(s)

APPENDIX D

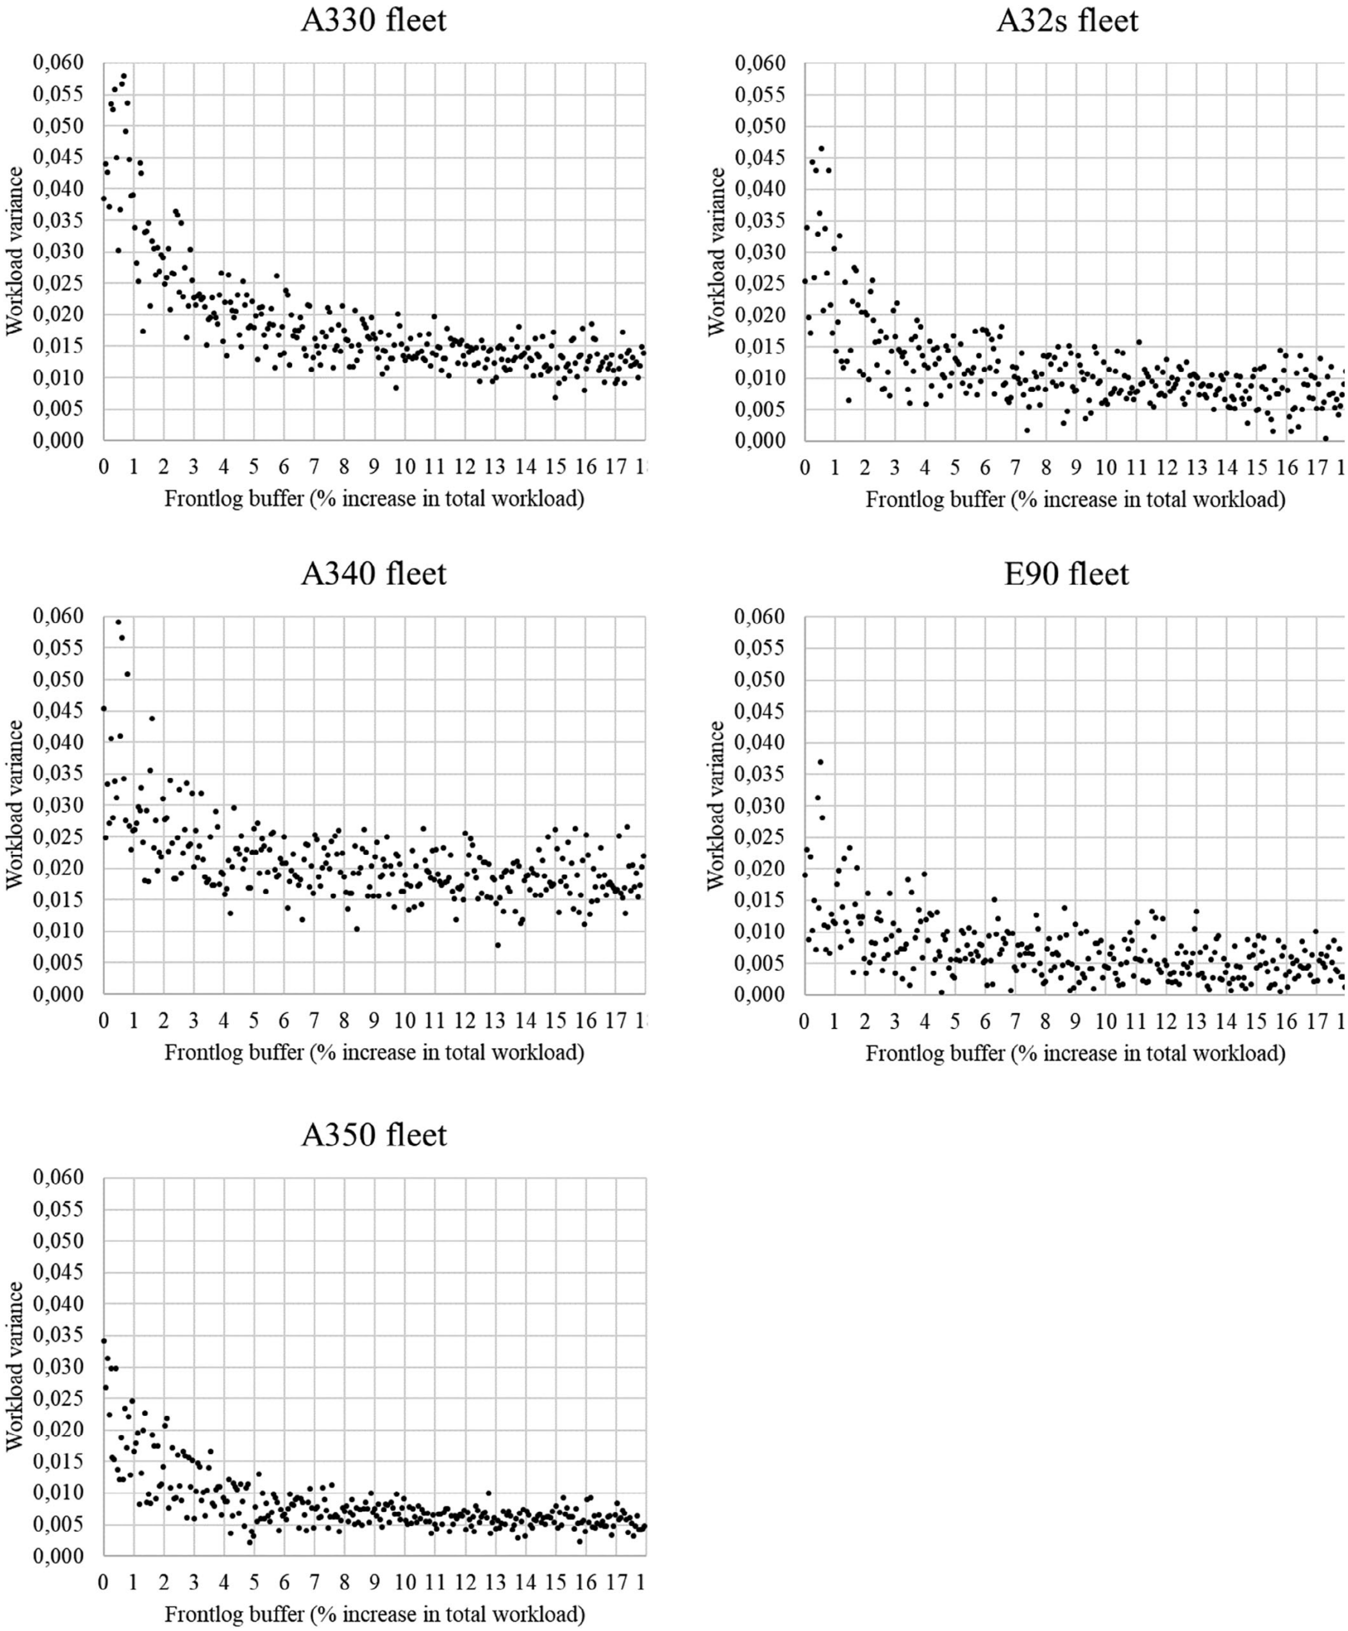

FIGURE D1 Fleet-specific workload variance as a function of the *frontlog* buffer. The figures reflect the different operational profiles of the fleets. The A330, A340, and A350 fleets fly long-haul routes to Asian destinations, and the A32s and E90 fleets fly short-haul connecting/feeding routes. Higher fleet utilization (e.g., in the A330 fleet) leads to increased workload variance, presumably due to the scarcity of maintenance opportunities, which, in turn, leads to a stronger effect of increasing the *frontlog* buffer

APPENDIX E

TABLE E1 Evaluation events

Date	Evaluation event	Evaluation phase 1	2	3	4	Participants (where A* = * author)	Takeaways
June 17, 2016	Meeting—evaluation of conceptual structure of simulation	X				A1, A2, A3, head of maintenance planning	- Simplified view of resource planning can be used without a significant negative impact on representability
September 9, 2016	Meeting—evaluation of first simulation results and model behavior	X				A1, A2, A3, head of maintenance planning, planning manager, operations manager	- On a general level simulation output seems representative of operations - Need to check whether the maintenance routing algorithm results in sufficient ground time for individual aircraft - 24 hr man-hour variation and total fleet man-hours needed for validation - A32s- and E-fleets have spare aircraft, but re-routing should happen infrequently - Buffer impact on departure reliability is interesting, and could be studied directly through the model
December 8, 2016	Meeting—evaluation of simulation model validation	X				A1, A2, A3, head of maintenance planning	- Simulation man hour output is in minutes, which needs to be considered in internal communication - Workload corresponds to reality, deviations in average workload attributed to changes in fleet compositions reflected in the data - Model is not representative in terms of delays, can be fixed through resource re-allocation constraints
February 28, 2017	Meeting - evaluation of cost savings potential indicated by model		X			A1, A2, head of maintenance planning, operations manager, head of line maintenance	- Potential cost savings are plausible and significant - Costs for long term labor flexibility cannot be assumed to approach zero with larger buffer values, since ground time will still be unevenly distributed, to be fixed
April 21, 2017	Results presentation—evaluation of research results by operational stakeholders			X		A1, A2, A3, head of maintenance planning, head of line maintenance, VP of technical operations, head of engineering, head of maintenance control, head of resource planning, head of analytics development, planning manager, operations manager, analyst	- On a general level model results were not questioned. - Workload mismatch was observed by the audience, and explained by the presenters - VP of technical operations noted that with the buffer increasing the total workload, there is a risk that the increased downtime translates to lost flying hours and would hence be away from “generating revenue”. It was highlighted that all

TABLE E1 (Continued)

Date	Evaluation event	Evaluation phase 1	2	3	4	Participants (where A* = * author)	Takeaways
							planned flights in the current traffic program (which was an input to the simulation model) were successfully executed with all buffer values.
September 26, 2017	Implementation meeting—evaluation of efforts needed to implement the design				X	A1, A2, A4, (new) head of maintenance planning, planning manager, maintenance planner, maintenance planner	- Planners strive not to make too tight plans, implying that they are currently introducing a buffer, but it is not managed - Planning systems need to be able to measure the buffer if it is to be managed - Buffer consumption happens in operations, implying that sound principles for when the buffer should be used are needed - Also effect of traffic planning decisions on available ground times, and by extension prerequisites for buffering needs to be better understood
December 28, 2018	Personal communication—recap of implementation status				X	A1, (new) head of maintenance planning	- Buffer would require automated replanning of work, which current planning system is unable to perform, possibilities to remedy this has to be explored - There is a negative mind-set related to not being able to perform all planned work, which could pose a challenge when introducing the buffer - Changes in resource planning cause uncertainty at the moment, which is connected to buffer implementation - Tracking of tasks in operations is improving, which is advantageous considering buffer implementation
November 22, 2019	Interview on implementation challenges				X	A1, A4, (new) head of maintenance planning, process development (lean) manager, business development manager	• Many of the earlier challenges were acknowledged and confirmed as still relevant • There has been considerable progress with task tracking, but there is still work to do before reaching real-time visibility into operations. • Resource planning principles are changing, showing as more stable resource plans. • There are now regular maintenance activities also at a number of outstations for both NB and WB fleets.

(Continues)

TABLE E1 (Continued)

Date	Evaluation event	Evaluation phase 1	2	3	4	Participants (where A* = * author)	Takeaways
							• Buffer would require a digital feedback loop from operations to planning • - Incentives, organizational structure and mind-set in production can be expected to pose challenges related to building and consuming the frontlog.