Vehicle Routing with Endogenous Learning: Application to Offshore Plug and Abandonment Campaign Planning

Steffen J. Bakker[a,*], Akang Wang[b,c], Chrysanthos E. Gounaris[b,c]

[a] *Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Alfred Getz veg 3, NO-7491 Trondheim, Norway*
[b] *Department of Chemical Engineering, Carnegie Mellon University, 5000 Forbes Avenue, 15213 Pittsburgh PA, USA*
[c] *Center for Advanced Process Decision-making, Carnegie Mellon University, 5000 Forbes Avenue, 15213 Pittsburgh PA, USA*

Abstract

When a particular service is performed many times, the duration of the service might reduce due to the effect of learning from similar tasks that have been performed before. In this article, we present an approach to account for such learning effects that arise in the context of vehicle routing operations. Our approach enables the consideration of endogenous learning, where the service times are dependent on the experience that is to be gained in the same routing horizon. We apply our approach to the problem of planning an offshore plug and abandonment campaign, where different vessels are being used to perform plugging operations on offshore oil and gas wells. We extend existing instances for this problem with observed learning data and investigate the effects of learning and cooperation. Results show that the inclusion of an endogenous learning effect leads to different and significantly better solutions compared to those that are found when the learning effect is neglected.

Keywords: Routing, Endogenous Learning, Plug and Abandonment Operations, Logistics, OR in Maritime Industry

*Corresponding author
Email address: `steffen.bakker@ntnu.no` (Steffen J. Bakker)

1. Introduction

The Vehicle Routing Problem (VRP) literature is rich and there exist myriad variants of the VRP (Toth & Vigo, 2002). Besides the routing part, most VRP variants also consider the performance of some service at the customer nodes. When this service entails a repetitive task, a learning effect might arise for the routed asset; for example, the duration of the service provided by a particular routed asset shall reduce when the latter performed similar tasks in the past. Since such a learning effect occurs within the routing horizon of the problem, it should be modeled in an endogenous way. In this article, we present a methodology to incorporate such a learning effect in VRP models, and we refer to this variant as the Vehicle Routing Problem with Endogenous Learning.

Our interest in this topic arises from our work on plug and abandonment (P&A) campaigns. When an oil or gas well reaches the end of its lifetime, it must be permanently plugged and abandoned (P&A'd) (Vrålstad et al., 2019). When several wells are P&A'd together, making use of one or several available specialized vessels, we call this a P&A campaign. Even though P&A operations have been conducted for a long time (Calvert & Smith, 1994), the focus on P&A has increased during recent years due to the large number of offshore wells that are approaching the end of their life-time in established areas such as the North Sea and the Gulf of Mexico (Khalifeh et al., 2013; Kaiser, 2017). As P&A operations can be very time-consuming and costly, it is of interest to optimize the P&A process as much as possible.

Bakker et al. (2019) developed an extension of an uncapacitated Vehicle Routing Problem with Time Windows (uVRPTW) model that can be used for the planning of P&A campaigns. However, this formulation does not allow for a learning effect. In line with the case of drilling of wells where the existence of a learning effect that has been observed (Brett & Millheim, 1986), recent experience from operators suggests that a significant effect is also present in the execution of P&A operations (Straume, 2018). The learning effect in P&A operations occurs due to various reasons. The primary one is that there is significant diversity and unpredictability regading the status of each well and/or the field formation, and as a result, information is constantly gained about these characteristics throughout the duration of the plugging campaign. A second reason contributing to a learning curve is the general gaining of experience by the crew in performing such operations. It should be noted that seldom does the industry maintain dedicated crews to perform P&A. The main activity of crews is most often related to exploration and production, with P&A operations being a secondary assignment.

Hence, a crew's initial experience with P&A might not be very high when the campaign begins, leading to a significant learning effect. We note that, when dedicated P&A crews are being used, the learning effect might become less pronounced, but it is generally still observable.

We also highlight that learning effects are often present in many other VRP settings as well. For example, other variants of the VRP in which learning might occur due to the performance of repetitive services include the Technician Routing Problem (Chen et al., 2016), the Workover Rig Routing Problem (Aloise et al., 2006; Ribeiro et al., 2012), the Maintenance Routing and Scheduling Problem (Irawan et al., 2017), and the (multiple) Traveling Repairman Problem (Luo et al., 2014). To that end, methodologies to incorporate learning effects in the context of routing models can have broader applicability beyond the offshore oil and gas industry.

The main contributions of this article include: (i) the development of a method to incorporate an endogenous learning effect in a standard VRP setting by means of a linearization approach that does not introduce any additional binary variables, (ii) the compilation of a suite of realistic benchmark instances for the problem of planning a plug and abandonment campaign under learning, and (iii) the elucidation of the benefits from modeling such a learning effect as well as the quantification of the value of cooperation for operators.

The structure of this paper is as follows. Section 2 reviews relevant VRP literature as well as literature on learning effects. Section 3 presents the uVRPTW as the base model for our work and its extension to account for the learning effect. In Section 4, we introduce the problem of P&A campaign planning, which serves as an application of the VRP with endogenous learning. In addition, we develop a clustering-based solution approach that efficiently reduces the combinatorial complexity introduced by the routing. Finally, Section 5 presents our computational studies, before we conclude in Section 6.

2. Literature Review

2.1. Vehicle Routing Problems

The problem studied in Bakker et al. (2019) arises from a real world problem and can be referred to as a rich VRP (Lahyani et al., 2015), which extends classical VRPs with issues arising in real world applications. At its base lies the Vehicle Routing Problem with Time Windows (VRPTW), which is one of the most important generalizations of the classical VRP (Cordeau et al., 2007).

We focus in particular on the uncapacitated VRP with Time Windows and Precedence Constraints (uVRPTWPC), which can also be referred to as the multiple Traveling Salesman Problem with Time Windows and Precedence Constraints (mTSPTWPC) (Balas et al., 1995; Ascheuer et al., 2001). An overview of formulations and solution procedures for these problems has been given by Bektas (2006).

There exist multiple alternative modeling avenues when formulating VRPs and TSPs (Bektas, 2006). For example, one may utilize any of a class of assignment-based formulations, with the well-known Miller-Tucker-Zemlin (MTZ) constraints being the most notable example (Miller et al., 1960). Another traditional approach is to use single-commodity flow (SCF) formulations (Langevin et al., 1990; Gouveia & Pires, 1999; Öncan et al., 2009). The main advantage of SCF and/or MTZ models is that they exclude subtours using a polynomially-sized formulation that can in principle be monolithically solved by generic mixed-integer linear optimization solvers, making them suitable for adoption by practitioners who have access to only off-the-shelf software. When a custom-built branch-and-cut implementation can be considered, one may instead resort to a vehicle-flow model (Laporte & Nobert, 1983). In this case, starting with a bare-bones formulation capturing only the node degrees, one separates and dynamically adds to the formulation cuts deriving from the well-known subtour elimination constraints (Dantzig et al., 1954), or in the context of a capacitated routing problem, rounded capacity inequalities (Laporte & Nobert, 1983). In this work, we follow the SCF paradigm, where arrival times and experience levels are modeled as flows. Based on preliminary results, we had determined that the SCF approach is a reasonable choice for our case, as it was able to address satisfactorily the problem benchmarks of interest, while yielding better performance than the MTZ counterpart in this context. In fact, the commodity flow constraints in this model are satisfactorily tight to eliminate the formation of most subtours in the linear programming relaxations, to the extent that strengthening the model with subtour elimination constraints was not computationally favorable. At the interest of brevity, we have chosen to omit the details of those preliminary investigations from the manuscript.

A fundamentally different way to formulate VRPs is by means of a set partitioning formulation, where binary variables are associated with feasible routes (Baldacci et al., 2008). Due to the existence of exponentially many feasible vehicle routes, the resulting formulation is tackled by a branch-and-price algorithm. In particular, only a subset of routes are initially considered at each branch-and-bound node, and a column generation step is called iteratively to introduce additional routes with the potential to reduce the objective value, until no more such routes exist. The best performing exact algorithms for capacitated VRPs combine elements of cut separation and column generation and are known as branch-price-and-cut (BPC) approaches. Notable contributions within this field include the works by Jepsen et al. (2008); Baldacci et al. (2008, 2011); Pecin et al. (2017); Pessoa et al. (2019). Nevertheless, adapting these algorithms to the VRP variant discussed in this paper is not straightforward. The main challenge lies in the fact that the VRP variant that we consider here includes requirements for precedence between nodes, considering that P&A operations at each well have to be performed in sequence. As these nodes can be visited by different vehicles, there is a need for synchronization between routes for different vehicles, leading to challenges when implementing a BPC algorithm. Indeed, how to best address such challenges would be a promising area for future research.

2.2. The Learning Effect

The term *learning* has often been used in the literature to refer to the impact of experience on service or production times (Chen et al., 2016). This is motivated by the fact that, when engaged in repetitive tasks, workers tend to use less time to perform the later tasks due to their familiarity with the operation. Mathematical representations of this process are referred to as learning curves. There exists an extensive literature on the learning effect and corresponding learning curves. Detailed discussions of various learning curves and their applications are available in Anzanello & Fogliatto (2011).

The first quantification of a learning curve is given by Wright (1936). He observed that assembly costs of airplanes decreased as repetitions were performed. The Wright's model is now known as the Power, or Log-linear, model of learning. Many other learning models have since been proposed, in an effort to represent the learning effect more realistically in various contexts. Notable classes of learning models are, for example, the Stanford-B, DeJong, S-curve, plateau and the exponential model (Nembhard & Uzumeri, 2000).

In regards to the oil and gas industry, Brett & Millheim (1986) were the first to use a learning curve to assess drilling performance (in terms of completion time) for a series of similar wells that have to be drilled. They used an exponential model to specify the learning effect and their approach is still the standard when considering learning curves in offshore drilling operations. A visualization is given in Figure 1.

Plugging operations are similar to drilling operations in that they require much of the same equipment and procedures. Hellström (2010) studied learning curves in drilling and well operations for a Norwegian operator, investigating among others the application of the Brett & Millheim (1986)

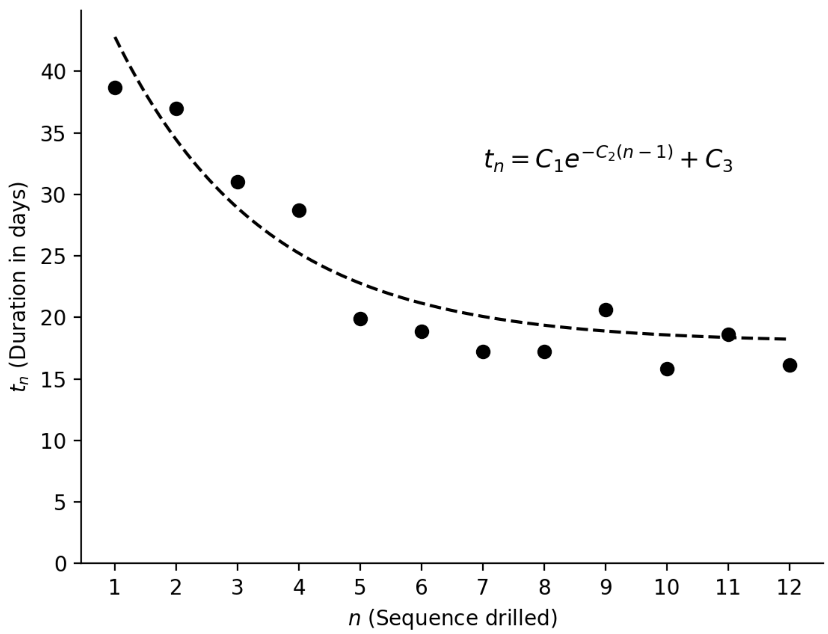

Figure 1: Visualization of a typical learning curve for the drilling of oil wells adapted from (Brett & Millheim, 1986).

model to P&A operations. More recently, Moeinikia et al. (2014b) evaluated cost efficiency of rigless P&A for subsea multi-well campaigns using Monte Carlo simulations. The learning effect, modeled by means of the Brett and Millheim model, was a key element in this analysis. In addition, data on a significant observed learning effect in a recent real-life plugging campaign has been presented by a Norwegian operator (Straume, 2018).

Whereas the learning effect has been extensively discussed in the context of manufacturing and machine/project/workforce scheduling (Biskup, 2008; Azzouz et al., 2018), it has received less attention in the VRP community. When learning is defined in the context of familiarity of vehicles with certain customers or areas, it is referred to as VRP with driver learning, driver familiarity or driver-specific travel time information. The first attempt to incorporate a learning effect targeting travel times is from Zhong et al. (2007). They consider a multi-day vehicle routing problem, where driver learning, or familiarity, results from visiting service areas repeatedly. With increased familiarity, driver performance increases due to ease of finding addresses and locations. The model is solved heuristically over a 30-day planning period simulation. The learning effect is then taken into account by updating the parameters between the model runs for each day. In a similar fashion, Kunkel & Schwind (2012) considers a multi-day VRP with Driver Learning, which again is solved heuristically. Contrary to the previous approaches, Schneider (2016) assumes that driver learning already has taken place in the past, leading to different familiarity levels. That problem can then be viewed as a variant of a heterogeneous fleet VRP. The authors note that the inclusion of a learning model that describes the reduction of travel and service times in dependence of the number of visits to each customer presents an interesting opportunity for future research.

The work of Chen et al. (2016) and Chen et al. (2017) focuses on a learning effect for service times, considering the Multi-Period Technician Routing Problem with Experienced-based Service Times. Their model is formulated based on a Markov decision process, and it is solved using a rolling-horizon procedure based on a heuristic. The service time parameters are then updated between the different periods, according to the attained increase in experience and defined learning effects.

The above presented VRP literature includes learning that tends to be on an operational level and is used for daily planning. The learning effect is treated in an *exogenous* way. That is, the learning effect is considered outside of the main model. These routing models are solved in a rolling horizon fashion, where the learning effect is taken into account by iteratively updating the parameters between the model runs. In contrast, in this work we consider a problem with a long time horizon at the strategic level. Learning occurs within the time horizon of the problem and directly depends on the decisions to be made. Hence, in the VRP with Endogenous Learning, we incorporate the learning effect directly into the model.

3. Mathematical Model

In this section, we present a commodity-flow formulation for the uVRPTW and show how to extend this with endogenous learning.We explain the notation (sets, indices, parameters and variables) used in the model and we provide the mathematical formulation of the constraints and objective function.

3.1. Model Formulation

In the uVRPTW, the objective is to find minimum-cost routes for a set of vehicles, $\mathcal{K}$, such that all customers, gathered in the node-set $\mathcal{N}$, are being served. We consider a heterogeneous fleet in which vehicles are not necessarily compatible with all customers. The vehicles start and finish at a depot, typically modeled as two locations, denoted by $o(k)$ and $d(k)$, respectively. The union of the depots and customers is denoted by the vertex set $\mathcal{V}$. Each vehicle $k \in \mathcal{K}$ has its own node- and vertex-set, denoted by $\mathcal{N}_k$ and $\mathcal{V}_k$ respectively. In addition, we associate arc sets $\mathcal{A}_k = \{(i,j) : i,j \in \mathcal{V}_k \wedge i \neq j\}$ with each vehicle k. Moreover, given vertex i, $\delta_k^+(i)$ is defined as the set of vertices j such that arc $(i,j) \in \mathcal{A}_k$. Similarly, given vertex i, $\delta_k^-(i)$ is defined as the set of vertices j such that $(j,i) \in \mathcal{A}_k$.

Each vehicle k has travel times T_{ijk}^{TR} for all $(i,j) \in \mathcal{A}_k$. Moreover, each customer i has a time window, $\left[\underline{T}_i, \overline{T}_i\right]$, when it may accept service. The service times at customers are given by the continuous variables τ_{ik}^{SE}. In addition, we define binary flow variables x_{ijk}, for each vehicle $k \in \mathcal{K}$ and arc $(i,j) \in \mathcal{A}_k$, such that x_{ijk} equals 1 if vehicle k uses arc (i,j) in the optimal solution, and 0 otherwise. The cost of traversing arc (i,j) for vehicle k is given by parameters c_{ijk}. We also define continuous variables t_{ik} and w_{ik}, for each $k \in \mathcal{K}, i \in \mathcal{V}_k$, representing the arrival time from and waiting time at customer i by vehicle k, respectively. When vehicle k does not visit customer i, these variables equal zero. The cost parameters corresponding to service times and waiting times are given by d_{ik} and e_{ik}, respectively. Note that we opt to model the waiting times explicitly because the cost of waiting might be different from the cost of serving customers, in general. In addition, let the continuous variables $\tilde{t}_{ijk}$ be defined as follows:

$$\tilde{t}_{ijk} = \begin{cases} t_{jk}, & \text{if } x_{ijk} = 1, \\ 0, & \text{if } x_{ijk} = 0, \end{cases}$$

where $k \in \mathcal{K}$ and $(i,j) \in \mathcal{A}_k$.

A commodity-flow formulation of the uVRPTW is given below:

$$\min \quad \sum_{k \in \mathcal{K}} \left(\sum_{(i,j) \in \mathcal{A}_k} c_{ijk} x_{ijk} + \sum_{i \in \mathcal{V}_k} (d_{ik} \tau_{ik}^{SE} + e_{ik} w_{ik}) \right) \tag{1}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{K}} \sum_{j \in \delta_k^+(i)} x_{ijk} = 1 \qquad i \in \mathcal{N} \tag{2}$$

$$\sum_{j \in \delta_k^+(o(k))} x_{o(k)jk} = 1 \qquad k \in \mathcal{K} \tag{3}$$

$$\sum_{i \in \delta_k^-(j)} x_{ijk} - \sum_{i \in \delta_k^+(j)} x_{jik} = 0 \qquad j \in \mathcal{N}_k, k \in \mathcal{K} \tag{4}$$

$$t_{ik} + \tau_{ik}^{SE} + w_{ik} = \sum_{j \in \delta_k^+(i)} \left(\tilde{t}_{ijk} - T_{ijk}^{TR} x_{ijk}\right) \qquad i \in \mathcal{N}_k, k \in \mathcal{K} \tag{5}$$

$$t_{ik} = \sum_{l \in \delta_k^-(i)} \tilde{t}_{lik} \qquad i \in \mathcal{N}_k \cup \{d(k)\}, k \in \mathcal{K} \tag{6}$$

$$t_{ik} = \sum_{j \in \delta_k^+(i)} \left(\tilde{t}_{ijk} - T_{ijk}^{TR} x_{ijk}\right) \qquad i = o(k), k \in \mathcal{K} \tag{7}$$

$$\underline{T}_j x_{ijk} \leq \tilde{t}_{ijk} \leq \overline{T}_j x_{ijk} \qquad (i,j) \in \mathcal{A}_k, k \in \mathcal{K} \tag{8}$$

$$x_{ijk} \in \{0,1\}, \tilde{t}_{ijk} \in \mathbb{R}_+ \qquad (i,j) \in \mathcal{A}_k, k \in \mathcal{K} \tag{9}$$

$$t_{ik}, \tau_{ik}^{SE}, w_{ik} \in \mathbb{R}_+ \qquad i \in \mathcal{N}_k, k \in \mathcal{K} \tag{10}$$

The objective function (1) minimizes the routing costs and/or costs associated with time usage. Constraints (2)–(4) are the degree constraints. Constraints (2) require that all customers must be visited by exactly one vehicle, constraints (3) require that the routes start at the depot, and constraints (4) ensure that, when a vehicle arrives at a customer, it also leaves that customer. Constraints (5) are known as the commodity flow constraints. If a vehicle travels between two customers, then they ensure correct accounting of travel time and arrival time for each of the visits. Together with the degree constraints (2)–(4), the commodity flow constraints (5) are also responsible for eliminating subtours. Constraints (6) and (7) link the different departure time variables, while constraints (8) impose time windows on the times when the customers can be visited. Finally, the domains of the variables are defined in (9) and (10).

We observe that the service time variables τ_{ik}^{SE} only appear in constraints (5). When a learning effect is not considered, the following substitution can be made:

$$\tau_{ik}^{SE} = \sum_{j \in \delta_k^-(i)} T_{ik}^{SE} x_{jik}, \quad i \in \mathcal{N}_k, k \in \mathcal{K}. \tag{11}$$

Here, T_{ik}^{SE} are the deterministic service times. The uVRPTW now consists of constraints (1)–(11).

However, when there exist dependencies and/or restrictions between the start times of service at the different customers, then we can extend the formulation for the uVRPTW with generalized precedence constraints of the form

$$\sum_{k \in \mathcal{K}} (t_{ik} + \delta_{ijk}) \leq \sum_{k \in \mathcal{K}} t_{jk}, \quad (i,j) \in \Delta, \tag{12}$$

where the parameters δ_{ijk} specify the minimum difference in time between when customers i and j are being serviced, and the set Δ defines all customer pairs (i, j) for which a temporal dependency exists. This constraint captures all types of temporal dependencies between customers, such as, for example, synchronization, overlap or precedence. The resulting model is generally referred to as the VRPTW with Temporal Dependencies (Dohn et al., 2011).

3.2. Endogenous Learning

When services have to be performed at the customer nodes, a learning effect might arise; that is, the service times reduce as a function of the number of times this task has been performed before. Under this setting, the constraints (11) from the uVRPTW model would no longer be valid, as they assume a fixed service time for each customer. In the following, we describe a way to allow for an endogenous learning effect in the context of a commodity-flow uVRPTW model.

3.2.1. Experience Level

A learning effect arises when a particular task is being performed repeatedly. As different services might have to be done at different customers, we define the set $\mathcal{S}$ to contain the services for which a learning effect exists. To account for such an effect, we have to keep track of the *experience level* of the vehicles for these different services.

We define nonnegative continuous variables z_{isk}, $i \in \mathcal{N}_k, s \in \mathcal{S}_k$ and $k \in \mathcal{K}$, that measure the experience level. More specifically, if vehicle k performs service s at customer i, then z_{isk} will represent the number of times vehicle k will have performed a service s, after having visited customer i; it will equal zero otherwise. Moreover, we define nonnegative continuous variables $\tilde{z}_{ijsk}$ for $(i, j) \in \mathcal{A}_k, s \in \mathcal{S}_k$ and $k \in \mathcal{K}$. These variables are flow variables that keep track of the experience level. Finally, $\sigma(\cdot) : \mathcal{N} \to \mathcal{S}$ is a function that maps customers to the service they require.

The relationship between the experience variables is defined in the model in the following way:

$$z_{isk} = \sum_{j \in \delta_k^+(i)} \tilde{z}_{ijsk}, \quad i \in \mathcal{N}_k, s \in \mathcal{S}_k, k \in \mathcal{K}, \tag{13}$$

and

$$z_{jsk} = \sum_{i \in \delta_k^-(j)} \left(\tilde{z}_{ijsk} + \mathbb{1}_{\{\sigma(j)=s\}} x_{ijk} \right), \quad j \in \mathcal{N}_k, s \in \mathcal{S}_k, k \in \mathcal{K}, \tag{14}$$

where $\mathbb{1}_{\{\sigma(j)=s\}}$ is the indicator function, equaling one if the service required at customer j equals s, and equaling zero otherwise. Here, constraints (14) state that the flow of the experience level out of a certain node should equal the incoming flow, increased with the possible gain in experience level when executing that particular operation. In addition, the experience flow variables should only be allowed to be positive when their corresponding x-variables equal one, namely

$$\tilde{z}_{ijsk} \leq \tilde{M}_s x_{ijk}, \quad (i,j) \in \mathcal{A}_k, s \in \mathcal{S}_k, k \in \mathcal{K}, \tag{15}$$

where $\tilde{M}_s$ represents the maximum number of customers that require service s.

Finally, one might encounter a reduction in the experience level in a certain node. This can be the result of a change in the agent(s) performing these services. For a plugging campaign this means, for example, that the crew is being refreshed during a harbour visit. Let j' and k' represent the node and vehicle for which this is the case. We can then reset the experience level to zero, or to any other affine combination of the incoming experience level,

$$z_{j'sk'} = \alpha_{j'sk'} + \beta_{j'sk'} \left(\sum_{i \in \delta_{k'}^-(j')} \tilde{z}_{ij'sk'} \right), \quad s \in \mathcal{S}_{k'}. \tag{16}$$

3.2.2. Learning Effect Representation

Anzanello & Fogliatto (2011) present a wide range of learning curves based on different mathematical relationships. The most popular such relationships can be categorized as power models or exponential models. The typical learning curve in Figure 1 is an example of an exponential model. It is important to highlight that, even though these models of learning are non-linear, they are generally convex. It is also important to observe that we are interested in the values of the learning curve over a discrete range; that is, we are interested in the service time as a function of the number of times a similar task has been performed before. This implies that we can represent the (non-linear) learning curve in an exact manner (i.e., without any approximation error) using a continuous piecewise-linear function. Hence, moving forward, we can assume that the learning curves ($\tilde{f} : \mathbb{R} \to \mathbb{R}$) are of the following form:

$$\tilde{f}(z) := \max_{i \in \{1,\dots,N\}} f_i(z), \tag{17}$$

where z is the experience level (or sequence number), $f_i(z) := a_i z + b_i$, and the parameters a_i and b_i can be deduced from the original learning curve relationship. Figure 2 shows a typical learning curve, together with its linear representation.

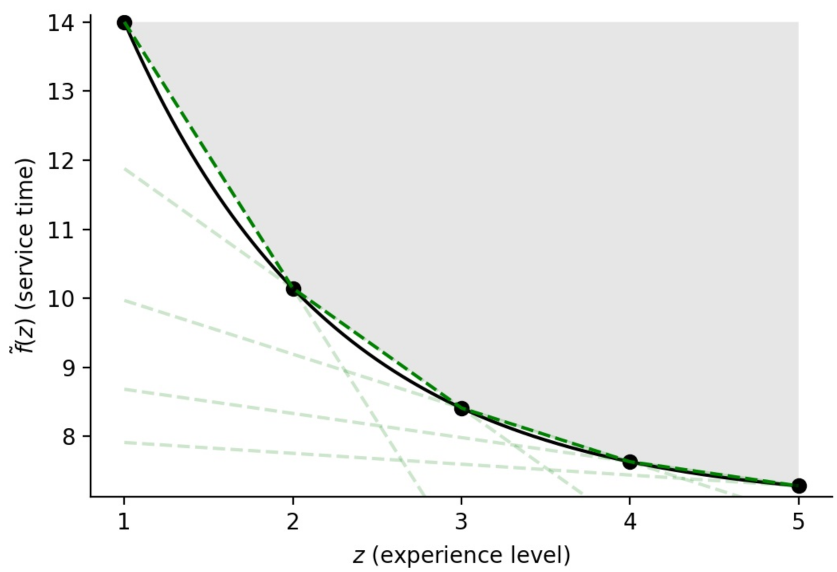

Figure 2: Visualization of a convex learning curve (solid line) and its representation using the epigraph of a piecewise linear function (dashed lines)

We remark that, since our intentions in the VRP model are to minimize makespan/task durations, there will exist an incentive by the optimizer to maximize the experience level and minimize the service time. This implies that we only have to bound f from below. Consequently, in order to account for the learning effect, we need only define the following equations of the service time variables τ_{ik}^{SE}:

$$\tau_{ik}^{SE} + b_{ikl}\left(1 - \sum_{j\in\delta_k^-(i)} x_{jik}\right) \geq a_{ikl} z_{i\sigma(i)k} + b_{ikl}, \tag{18}$$

for $l \in \left\{1, ..., \tilde{M}_s\right\}, i \in \mathcal{N}_k, k \in \mathcal{K}$. Neglecting momentarily the second term on the left hand side, these equations impose that that the service times with learning should be larger than the piecewise linear functions evaluated at $z_{i\rho(i)k}$. This is represented in Figure 2. But as the service times are minimized, these functions will be bounding and the linear approximation will be exact. However, we should be careful not to apply such lower bounds when the vehicle k does not service customer i (i.e., when $z_{i\rho(i)k} = 0$). In this case, we must simply allow that $\tau_{ik}^{SE} \geq 0$. To that end, the second term of the left hand side of equation (18) adds b_{ikl} when vehicle k does not service customer i (i.e., when $\sum_{j\in\delta_k^-} x_{jik} = 0$) so as to appropriately relax these constraints. Finally, we can tighten the overall formulation by enforcing:

$$\tau_{ik}^{SE} \leq \bar{b}_{ik} z_{i\sigma(i)k}, \quad i \in \mathcal{N}_k, k \in \mathcal{K}, \tag{19}$$

where $\bar{b}_{ik} = \max_{l\in\left\{1,...,\tilde{M}_s\right\}} b_{ikl}$.

4. P&A Campaign Planning

4.1. Problem Description

When an offshore oil or gas well has reached the end of its productive lifetime, it has to be plugged and abandoned to prevent leakages from or into the well. While platform wells can be plugged and abandoned with the existing drilling rig at the platform, subsea wells require dedicated vessels, referred to as mobile offshore units (MOU). When several subsea wells are plugged and abandoned together, making use of one or several MOUs, the process is referred to as a *P&A campaign* (Bakker et al., 2019). We use the problem of planning such a P&A campaign as an application of the Vehicle Routing Problem with Endogenous Learning.

When plugging a well permanently, several operations have to be performed. Based on the work by Oil & Gas UK (2015) and Moeinikia et al. (2014a), these operations can be divided into four phases. Phase 0 ("preparatory work") includes pre-P&A work such as stopping the flow from the well, logging the tubing quality and establishing temporary barriers. Phase 1 ("reservoir abandonment") and phase 2 ("intermediate abandonment") include the setting of barriers towards the reservoir, possible barriers in the overburden, and establishment of a surface plug. These two phases are typically performed consecutively, as a single service. Finally, phase 3 ("wellhead and conductor removal") includes the cutting of casing and conductor strings as well as retrieval of the wellhead (Vrålstad et al., 2019).

Conventionally, P&A operations are performed by semi-submersible rigs (SSR) with high spread rates. Current available technology still requires an SSR to perform phase 1 and 2 operations. Among other functions, the rig provides capacity to handle fluids returns as well as heavy lifting and cutting operations. However, more recently, lighter vessels are being used to perform simple P&A operations (Saasen et al., 2013; Sørheim et al., 2011; Valdal, 2013). This includes Riserless Light Well Intervention (RLWI) Vessels and Light Construction Vessels (LCVs).

P&A operations are in general not time critical. This means that when there are no integrity issues with the well and the operator maintains satisfactory control of the well, then the various phases can be executed at different times by different vessels. However, due to regulations or well conditions, wells might have to be plugged and abandoned within a certain time window.

Shut-down decisions are usually taken on a field level by the responsible operator/license holders. This implies that the scope of a plugging campaign usually is restricted to a single field. Nonetheless, plugging campaigns can be planned across multiple fields and licenses. As Bakker et al. (2019) shows, such large scale campaigns can lead to cost-savings. On a field, subsea wells can be found at different locations on the seabed. They may be located on their own as single satellites, or clustered on templates. As a result, the MOUs must move between the wells to perform the plugging operations. However, when performing operations on multi-well templates, a vessel does not have to be relocated.

The problem of planning a plugging campaign can now be defined as follows. A given number of subsea wells, possibly located on different fields, has to be plugged and abandonded within certain time windows. To plug a well, certain operations have to be performed in a strictly ordered sequence, but not directly after each other and different vessels can be used to perform these operations. The objective in a plugging campaign is then to find optimal routes and schedules for a fleet of MOUs, such that all plugging operations are performed.

Figure 3 visualizes the problem and a possible plan for the plugging of two offshore fields, making use of three vessels (SSR, RLWI and LCV). Each field contains three wells on which three operations have to be performed, related to the three different phases. In the first field, the RLWI vessel is used to perform all three phase 0 ($p0$) operations as well as the phase 3 ($p3$) operation on the first well, while in the second field, the LCV is used for the $p3$ operations. For all other operations, the SSR is being used.

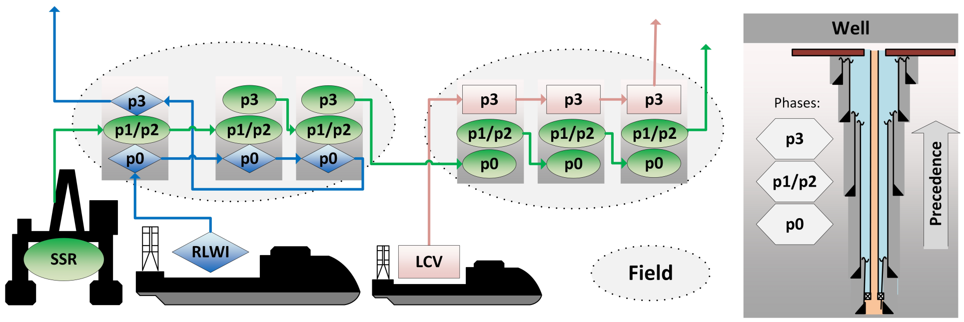

Figure 3: Visualization of a possible solution for the problem of planning a plugging campaign

4.2. Formulation

The problem of planning a P&A campaign can be formulated as a uVRPTW with Precedence Constraints. An extensive treatment is given in Bakker et al. (2019), but a short interpretation is provided here for completeness. Customer nodes represent operations, which have to be performed on wells to be plugged and abandoned. The vehicles now represent vessels that can perform these operations. On each well, three operations have to be performed. These operations are categorized into phase 0 ($p0$), phase 1 and 2 ($p1/p2$), and phase 3 ($p3$), which represent the three different services that can be performed.

We base our model on the formulation for the uVRPTW that was presented in Section 3. The objective (1) is to minimize total rental costs, which is a function of the durations the vessels are being used and their corresponding daily rates. This can be obtained by setting $c_{ijk} = T^{TR}_{ijk} C^{DAY}_{k}$ and $d_{ik} = e_{ik} = C^{DAY}_{k}$, where C^{DAY}_{k} represents a uniform day rate for vessel k. Constraints (2)–(11) can be interpreted in the standard way. In addition, we require precedence constraints that capture the fact that, on each well, the corresponding operations have to be performed in a strictly ordered sequence, starting with the $p0$ operation and finishing with the $p3$ operation. Constraints (12) capture this when setting $\delta_{ijk} := \tau^{SE}_{ik}$ and letting Δ consist of all pairs of operations (i, j) for which a precedence relationship exists. Besides time windows for the operations, we also impose restrictions on when the vessels can be used. This can be due to other planned activities or restricted operability during the winter months. For each vessel, we introduce an extra node to represent a harbour. This makes it possible for a vessel to return to the harbour during the winter months, where it does not incur any rental costs. We prescribe a learning effect for each of the different operation types. Consequently, we have that $\mathcal{S} = \{p0, p1/p2, p3\}$.

4.3. Learning Curve

We make use of the model developed by Brett & Millheim (1986) as a representation of the learning effect, since this model is being widely being used by the Oil & Gas Industry. However, we note that other specifications of the learning effect also may also be used as an alternative. A mathematical representation of the Brett and Millheim learning curve that was visualized in Figure 1 is given by:

$$t_n = C_1 e^{-C_2(n-1)} + C_3, \tag{20}$$

where t_n is the time required to perform the n-th operation in a sequence, C_1 is a constant reflecting how much longer the initial operation takes to perform than the idealized operation, C_2 is referred to as the learning rate and reflects the speed with which the operator reaches the minimum execution time for an operation, and C_3 is a constant that reflects the idealized minimum execution time for an operation. We note that $C_2 \in \mathbb{R}^+$ and $\frac{\partial t_n}{\partial C_2} < 0$. This means that a high learning rate (C_2) leads to a shorter execution time compared to low learning rates.

4.4. Node Clustering

The computational study from Bakker et al. (2019) was based on ten realistically sized instances. In that study, only the three smaller instances could be solved to optimality within the given time limit, with the remaining seven larger instances having an integrality gap of up to 2.5% after one hour. When extending this model with an endogenous learning effect, the computational complexity of the problem increases even more. This led us to investigate different solution approaches, in order to obtain better computational performance across the board.

More specifically, in the context of offshore logistics, locations tend to be clustered together. For example, oil and/or gas wells tend to be located relatively close to each other within the same field. In offshore windmill parks, wind turbines are positioned along narrow wind paths and the distances between them are relatively small. When servicing such clustered locations, the time it takes to travel between nodes within a cluster constitutes only a small fraction of the time required to travel across multiple clusters and to complete the whole route. Hence, when a cluster is serviced, the order in which the corresponding nodes are visited is insignificant for the evaluation of the total distance traveled or time consumed. This feature creates a certain kind of combinatorial hierarchy and makes the problems computationally difficult. In order to take into account the clustered nature of our datasets, we follow a solution approach as described below. The effectiveness of this approach is specifically tested later, in Section 5.2.1.

First, we define clusters of wells that meet the following two criteria:

1. the wells are located on the same field and are within a certain distance (application dependent) of each other, and
2. the wells have the same time windows.

Then, within a cluster, we define a fixed sequence in which the wells have to be plugged. This order can be chosen in different ways, such as based on the solution of a traveling salesman problem, or based on the complexity of servicing each specific well. The former would lead to shorter travel times within clusters, while the latter reaps the benefits from a learning effect. Specifically, when the wells have different complexities, one would start plugging the least complex wells first in order to accumulate experience for the more complex jobs scheduled for later. In this way, the learning gains can be more substantial. For reference, the solution visualized in Figure 3 abides to the proposed approach, where the fixed order of visiting wells within a cluster/field is defined to proceed from left to right. We note that a vessel is allowed to go back to an earlier well in the sequence so as to perform a different operation/phase, as is illustrated by means of the route of the RLWI vessel in that figure.

To incorporate such a fixed sequence within a cluster in the model, we simply reduce the graph on which the problem is defined, without the need to add any extra constraints. First, we remove all arcs going into a cluster that do not end in the first well in the sequence. Then, within a cluster, we remove the arcs that do not abide to the order of the sequence that has been set. As an example, for a cluster with two wells, W_1 and W_2, we could set the fixed sequence to be (W_1, W_2). In this case, we would remove the arcs leading to W_2 from outside the cluster, as well as the arcs that lead from W_2 to W_1.

Finally, we remark that the above described clustering approach bears resemblance to the Generalized VRP (GVRP) (Baldacci et al., 2010; Pop et al., 2012) as well as the Clustered VRP (CluVRP) (Battarra et al., 2014). In both these problems, customers are grouped into clusters, each of which is being served by exactly one vehicle, while each cluster can only be visited once. The main difference between the two is that, whereas in the GVRP exactly one customer is visited in each cluster, in the CluVRP all customers have to be visited. So, our approach is related to the CluVRP, with the exception that we allow customers within a cluster to be serviced by multiple vehicles. We highlight that, even though we partially fix the order in which wells have to be visited in a cluster, we do not fix the assignment of vessels to operations.

5. Computational Studies

5.1. Data

We apply the learning effect methodology to the problem of planning a P&A campaign. To test this effect, we make use of the instances that were defined in Bakker et al. (2019). These consist of synthetically constructed subsea fields based upon realistic data and well locations resembling typical Norwegian subsea fields. Each of these instances contains data on the number of wells, well complexities, templates, and operations. An overview of the dataset is given in Table 1. Moreover, Bakker et al. (2019) present information about the different vessels that are available for these campaigns, including their traveling speeds, day rates, operability restrictions, and (de-)mobilization times.

Table 1: Overview of the data instances

Instance	1	2	3	4	5	6	7	8	9	10
Number of Operations	12	15	15	24	24	33	33	42	42	48
Number of Templates	4	5	5	8	8	11	11	14	14	16
Number of Wells										
Total	8	14	18	13	25	29	32	32	33	44
Low complexity	2	9	0	5	2	15	17	12	7	16
Medium complexity	6	5	14	6	19	6	5	10	20	22
High complexity	0	0	4	2	4	8	10	10	6	6

The characterization of the wells in terms of their complexity are determined based on each well's match with the capabilities and restrictions of the lighter vessels, and are taken taken from Øia et al. (2018). For a well of low complexity, only a single zone has to be sealed off, the production tubing can be left in-hole, and no casing has to be retrieved. For a well of medium complexity, there is a zone for which the cement integrity is unknown and the production tubing has to be (partially) retrieved to log cement bonding. Moreover, it can include additional zones that need to be sealed off, if the casing need not be removed for this zone. Finally, for highly complex wells, the cement integrity is unknown, the production tubing needs to be retrieved fully, and the casing string(s) have to be removed, while multiple zones may exist.

We extend these instances to include a learning effect in the model. For this, we calibrate the learning curve from Equation (20) using data about the learning rates (C_2), minimum execution times for operations (C_3) and maximum execution times ($C_1 + C_3$). Brett & Millheim (1986) categorizes the values of C_2 in four groups, namely excellent, good, average and poor performers. Operators are assumed to work with a learning rate value corresponding to an average performer. This implies a value of C_2 between 0.25 and 0.45. In our analyses, we fixed $C_2 = 0.35$, which has been found to be the industry average for the drilling of wells. To obtain values for C_1 and C_3, we make use of the data presented in Øia et al. (2018). They provide a thorough description of operational procedures for both SSR and RLWI vessels, as well as they present duration estimates for subsea wells of all three complexities. These duration estimates are provided as minimal, expected and maximal values. On our end, we aggregated the data from Øia et al. (2018) at a phase level. That is, we determined the duration of a phase by summing up the durations of all the operations that are included in this phase. Finally, as we lack data on durations of performing services using LCVs, and since LCV and RLWI vessels have similar capabilities, we assumed that the durations of phase $p3$ operations are equal for these two vessel types. A summary of the resulting data is presented in Table 2.

5.2. Results

In this section, computational results are presented and discussed for the ten different instances defined above. We focus on the inclusion of a learning effect as well as the performance of the clustering approach. The model was implemented in Python 3.5.3, formulated using Pyomo 5.1.1 and is solved with CPLEX version 12.7. The analyses have been carried out on an HP EliteBook 820 G2 computer with an Intel Core i5-5200U CPU, 2.2 GHz processor, 16Gb RAM, running Windows 10 and using up to eight threads.

Table 2: Durations (in days) of the different phases when performed by SSR or RLWI vessels, for wells of different complexities, and categorized by minimum, expected and maximum value (data based on Øia et al. (2018)).

Complexity	Low			Medium			High		
	Min	Exp.	Max	Min	Exp.	Max	Min	Exp.	Max
SSR									
$p0$	3.75	5.29	6.88	3.65	4.71	6.19	3.58	4.58	5.79
$p1/p2$	6.04	8.75	12.50	7.94	9.52	12.71	11.33	14.21	18.17
$p3$	1.08	1.38	1.75	1.08	1.38	1.75	0.58	0.88	1.17
RLWI									
$p0$	2.58	3.33	4.50	4.06	4.81	6.08	6.58	8.33	10.92
$p3$	1.08	1.38	1.75	0.69	0.96	1.38	1.08	1.38	1.75

5.2.1. Clustering Approach Validation

We shall first focus on quantifying the performance of the clustering based solution approach. We start by comparing the results from the original approach as, presented in Bakker et al. (2019), with the results obtained here using clustering. As we do not yet consider a learning effect, we fix the sequences within the different clusters using the shortest route. The instances are run with a time limit of one hour. Table 3 compares the original and the clustering approach in two ways. The left side of the table presents results regarding CPU times, objective function values, and (for cases when the time limit was hit) optimality gaps. The right side of the table compares the structure of the best solutions stemming from each approach.

Firstly, from a computational tractability perspective, we observe that the solution times decrease drastically when making use of clusters, which can be attributed to the combinatorial reduction of the size of the feasible region, as compared to the original model. In fact, for nine out of ten instances, we can solve the reduced problems to zero optimality gap in mere seconds. In contrast, the majority of the problems solved with the original approach timed out after an hour, with a residual–albeit small–optimality gap. Considering the quality of the solutions obtained when using clustering, we observe that the objective function values for most of the instances are nearly identical. The worst deviation arises in the fourth instance, where we observe a 0.15% upwards change in objective value as compared to the full (non-clustering based) search. In contrast, for instances seven through nine, we manage to get a small improvement in the objective function value due to the original approach not having converged to zero gap.

Table 3: Performance of the clustering approach as compared to the original one. The first part of the table compares the two approaches in terms of CPU time (in seconds), best known solution value (in million dollars) and MIP gap (if time limit), on a per-instance basis. The second part of the table sketches the structure of the best solutions in terms of the number of times a particular vessel is used to perform a certain phase*.

	Computational performance							Solution structure					
	Original approach			Clustering approach				Original approach			Clustering approach		
								RLWI		LCV	RLWI		LCV
Inst.	CPU Time (s)	BKS (MM$)	MIP gap (%)	CPU Time (s)	BKS (MM$)	MIP gap (%)	ΔBKS (%)	$p0$	$p3$	$p3$	$p0$	$p3$	$p3$
1	14.2	36.88	(opt.)	0.2	36.88	(opt.)	0.01 %	4	1	0	4	1	0
2	316.7	56.29	(opt.)	0.2	56.31	(opt.)	0.02 %	5	2	0	5	2	0
3	21.0	83.42	(opt.)	0.2	83.42	(opt.)	0.01 %	5	3	0	3	0	3
4	3,600	56.59	(0.65)	0.4	56.67	(opt.)	0.15 %	6	3	0	6	1	0
5	3,600	109.73	(0.19)	1.0	109.75	(opt.)	0.02 %	6	0	6	6	0	6
6	3,600	133.60	(0.28)	0.4	133.61	(opt.)	0.01 %	9	0	9	9	0	9
7	3,600	140.67	(1.19)	8.6	140.51	(opt.)	-0.11 %	8	0	7	8	0	8
8	3,600	143.95	(2.49)	1.1	143.54	(opt.)	-0.28 %	11**	0	11	11**	0	8
9	3,600	143.11	(1.00)	205.9	142.55	(opt.)	-0.39 %	10	0	10**	10	4	6
10	3,600	186.14	(0.88)	3,600	186.21	(0.50)	0.04 %	12	2	10	12	2	10

* A rig always performs $p1/p2$, so a plan/structure is characterized by the allocation of vessels to $p0$ and $p3$ operations.

** These operations are seperated by a winter period.

Furthermore, with respect to the structure of the solutions, we observe that every identified optimal plan requires an SSR to perform $p1/p2$ operations, while the other operations are to be performed by the lighter vessels. Hence, a plan is mainly differentiated by the usage of these lighter vessels and assignment to $p0$ and $p3$ operations at the different wells. The second part of Table 3 summarizes this by presenting the assignment of vessels to $p0$ and $p3$ operations. We observe that the clustering approach provides solutions with similar structures as those provided by the original appeoach, as the assignment of vessels to operations tends to be nearly identical. The above empirical evidence suggests that clusters constructed in this way make sense for the real problem setting and lead to solutions that do not significantly sacrifice optimality.

5.2.2. Learning Effect

In order to judge the impact of the inclusion of an endogenous learning effect, we take the model including learning and solve it in two different ways. First, we take the best known plan that results from the model using the expected execution times without considering learning. We then solve the learning model with the obtained routing variables fixed. This gives us a measure of how the plan that does not consider learning would perform in the real setting. Second, we solve the learning model without any restrictions, resulting in the optimal plan for that case. Table 4 presents a summary of this comparison for all instances.

Firstly, we note that the *no learning* plan is not feasible in the learning model for instance 6. In addition, we observe that the plans that do not consider learning perform significantly worse than the plans that do. More specifically, we can achieve reductions in the objective function values between 3% and 20% This reduction in objective function values results from significant changes in the plans. Indeed, as the second part of Table 4 summarizes, the solutions differ in terms of the distribution of vessels that perform the $p0$ and $p3$ operations. More specifically, when not taking learning into account, both the RLWI vessel and LCV are being used to perform $p3$ operations. However, in this way, the benefits from learning are not optimally utilized. When we consider learning, we see that, in the optimal plans, the solution utilizes the RLWI solely for $p0$ operations and the LCV for $p3$ operations. Therefore, in general,

we can argue that the inclusion of learning leads to the use of fewer vessels to perform certain operations. This effect is illustrated in Figure 4, which focuses on dataset 9, as a representative instance. We observe that, when not considering learning, the optimal plan calls for utilizing all three vessels.

Table 4: Comparison of the plans resulting from the *no learning* (*NL*) and *learning* (*L*) model, evaluated in a learning setting. We present the best-known solution, CPU time and solution structure for each approach.

Instance	Computational performance				Solution Structure					
	NL model	*L* model			*NL* model			*L* model		
	BKS ($MM)		Change	CPU time (s)	RLWI		LCV	RLWI		LCV
					*p*0	*p*3	*p*3	*p*0	*p*3	*p*3
1	58.04	52.32	-10 %	0.2	4	1	0	0	0	0
2	73.53	71.46	-3 %	0.2	5	2	0	3	0	5
3	118.95	105.63	-11 %	0.2	4	0	4	0	0	5
4	76.11	64.08	-16 %	0.7	6	2	0	0	0	8
5	168.26	134.11	-20 %	0.8	7	0	7	0	0	8
6	-	130.59	-	1.6	-	-	-	8	0	11
7	146.73	135.98	-7 %	75.5	8	0	8	9	0	12
8	152.24	145.00	-5 %	655.4	11*	0	11	7	0	14
9	178.50	153.65	-14 %	289.4	11	4	7	0	0	14*
10	217.91	197.62	-9 %	3,600.0	13	2	10	6	0	12

*Operations performed over two disjoint periods, separated by the winter period.

However, when learning is considered, only the LCV and SSR vessels are being used, in which case the LCV performs all *p*3 operations.

5.2.3. Benefits of learning in large campaigns

In previous work, Bakker et al. (2019) have quantified potential benefits in running large plugging campaigns in lieu instead of several small ones. They refer to this as the *value of cooperation*. The presence of learning effects should make these benefits even more pronounced. To test this hypothesis, we focus on instances 5, 7, 8 and 10, for which the wells could be located on fields belonging to two different operators. Note that the results that we have obtained so far assume that these operators cooperated in one big campaign. We now consider the execution of two separate campaigns for the different operators. Table 5 compares the total costs of these two campaigns with the cumulative cost of the single cooperative campaign. Overall, we observe that cooperation in the planning of these campaigns leads to cost savings in the $11-13\%$ range, approximately. This is significantly more than the $3-4\%$ cost savings that Bakker et al. (2019) calculated when not considering a learning effect. Finally, Figure 5 depicts the cost savings when operators cooperate instead of planning separately for different values of the learning rate, C_2. We observe that the benefits of cooperation peak somewhere at a learning rate between 0.2 and 0.4, resulting in gains between 11% and 20%.

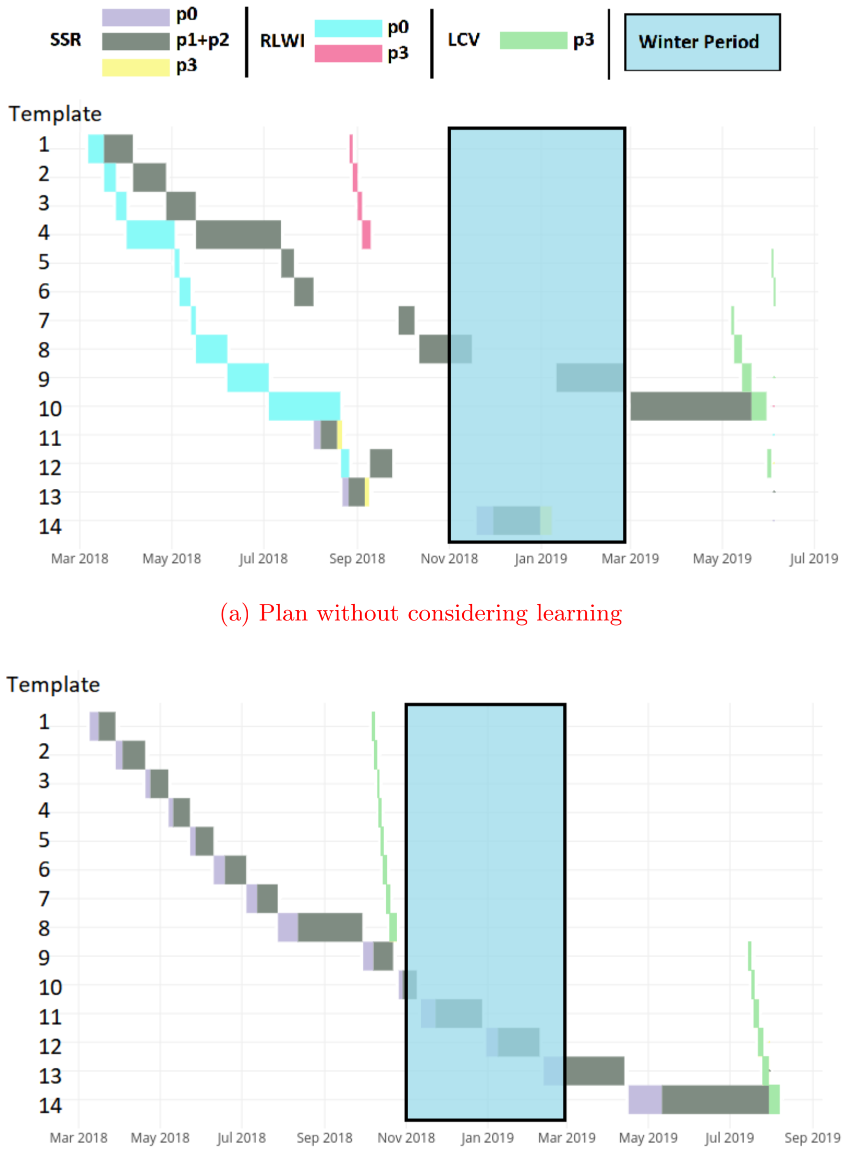

(a) Plan without considering learning

(b) Plan considering learning

Figure 4: Optimal Gantt charts for the representative instance 9, with and without considering the effect of learning.

Table 5: Total P&A campaign costs (in million dollars) for the two operators ("Oper. 1" and "Oper. 2"). The last column provides the percentage cost savings when cooperating ("Coop.") instead of running separate campaigns.

	P&A Campaign Costs ($MM)				Cost Savings
	Separate campaigns			Coop.	
Inst.	Oper. 1	Oper. 2	Total	Total	
5	45.59	105.63	151.22	134.11	11.3 %
7	71.46	84.07	155.53	135.98	12.6 %
9	105.63	71.48	177.11	153.65	13.2 %
10	143.98	77.51	221.49	197.62	10.8 %

After these peaks, the gains slowly decrease to a steady value of around $7 - 12\%$. Interestingly, these results suggest that there are significant gains for relatively slow learning rates.

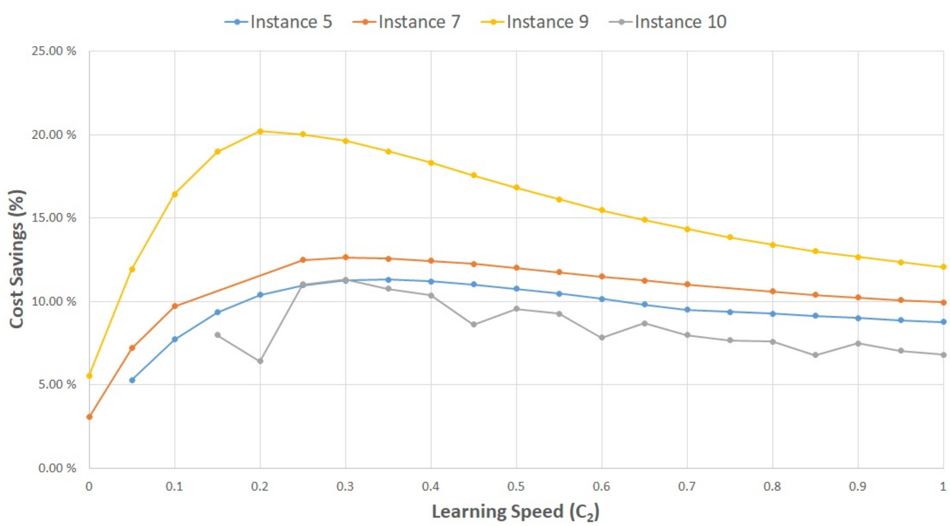

Figure 5: P&A cost savings when operators cooperate instead of planning separately, as a function of learning rate.

5.2.4. "What-if" Analysis

The parameters of the learning curves have to be estimated based on the results of previous plugging campaigns. However, the properties of a new campaign are likely to be somewhat different, and as a result, we might observe a different learning effect than the one postulated. We are therefore interested in the *robustness* of the optimal solutions we obtained against various scenarios about the learning curve parameters, as well as the difference in performance of these solutions from what would have otherwise been the optimal solutions, if we had *a-priori perfect information* about the prevailing values of these parameters.

We start by evaluating the effect of having misspecified the learning rate (C_2). We choose this parameter because it might vary significantly between different campaigns, and it can arguably be the most difficult to estimate out of the three parameters defining the learning curve. More specifically, we first determine the optimal plan using the nominal value of $C_2 = 0.35$ (average performer), and we fix the obtained routing variables. We then consider six different possibilities for the realized learning rate that span the learning rate categorization of Brett & Millheim (1986), namely the values $C_2 = 0.05$ and $C_2 = 0.20$ (poor performers), $C_2 = 0.50$ and $C_2 = 0.65$ (good performers) and $C_2 = 0.80$ and $C_2 = 0.95$ (excellent performers), and we solve the partially fixed model to determine actual timings and costs in each case. Table 6 presents the percentage cost differences in objective function values compared to the optimal solution for the postulated value of $C_2 = 0.35$, which provides an indication of the robustness of this solution to changes in the learning rate parameter.

Table 6: Performance of the optimal plan corresponding to the nominal learning rate value ($C_2 = 0.35$), under different realizations of this parameter.

Inst.	Change in realized cost of the nominal solution under various C_2 scenarios					
	$C_2 = 0.05$	$C_2 = 0.20$	$C_2 = 0.50$	$C_2 = 0.65$	$C_2 = 0.80$	$C_2 = 0.95$
1	13.1 %	5.5 %	-4.0 %	-6.9 %	-9.1 %	-10.7 %
2	infeas.	6.6 %	-4.5 %	-7.6 %	-9.9 %	-11.5 %
3	21.5 %	8.2 %	-5.2 %	-8.6 %	-10.9 %	-12.5 %
4	infeas.	8.6 %	-4.2 %	-6.4 %	-7.7 %	-8.5 %
5	26.4 %	8.3 %	-4.2 %	-6.5 %	-7.9 %	-8.8 %
6	26.2 %	7.7 %	-3.7 %	-5.8 %	-7.1 %	-8.0 %
7	infeas.	infeas.	-4.0 %	-6.2 %	-7.6 %	-8.5 %
8	infeas.	8.4 %	-4.3 %	-6.7 %	-8.2 %	-9.1 %
9	infeas.	6.7 %	-2.4 %	-3.4 %	-4.0 %	-4.3 %
10	infeas.	7.0 %	-2.9 %	-4.3 %	-5.2 %	-5.7 %

We observe that the realized learning rate strongly affects the feasibility and costs of a plugging campaign. In fact, when the realized learning rate is lower than the anticipated (nominal) rate, the overall duration of the campaign increases. This causes the planned routes to become infeasible in many instances, while even when this is not the case, the campaign costs increase between 5.5% and 26.2%. On the other hand, for a realized learning rate that is higher than the nominal rate, the campaigns turn out to be between 2.4% and 12.5% cheaper than anticipated.

Despite the sensitivity of the total costs on the learning rate, we highlight that advance knowledge of the exact learning rate does not necessarily help. To showcase this, we conduct an alternative analysis where we judge the quality of the obtained plan under the nominal learning rate. Again, we fix the routing variables obtained from the nominal case, and resolve the model for the cases with different learning rates. Subsequently, we solve the model without any restrictions for the different learning rates, to obtain the optimal plans in each case. Table 7 now presents the percentage cost increase of the objective function value of the nominal plan compared to the optimal plan, as the latter is evaluated under the different realizations of the learning rate. This can be considered to be a measure of the value of perfect information.

Table 7: Comparison of the plan obtained from the nominal learning rate ($C_2 = 0.35$) and the optimal plan obtained under different settings for this parameter.

Inst.	Difference in optimal costs using nominal solution as reference					
	$C_2 = 0.05$	$C_2 = 0.20$	$C_2 = 0.50$	$C_2 = 0.65$	$C_2 = 0.80$	$C_2 = 0.95$
1	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
2	infeas.	0.0 %	0.0 %	0.1 %	0.1 %	0.1 %
3	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
4	infeas.	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
5	0.0 %	0.0 %	0.0 %	0.0 %	0.4 %	0.8 %
6	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
7	infeas.	infeas.	0.0 %	0.0 %	0.0 %	0.0 %
8	infeas.	0.7 %	0.0 %	0.0 %	0.0 %	0.0 %
9	infeas.	0.0 %	0.1 %	0.1 %	0.1 %	0.1 %
10	infeas.	0.0 %	0.2 %	0.5 %	0.2 %	0.5 %

We observe that, as long as the plan corresponding to the nominal learning rate remains feasible, then this plan tends to be (nearly) optimal. This means that the value of perfect information is relatively low. In other words, knowing the realization of C_2 a priori would not affect the plan that would be generated from solving the optimization model. This finding follows from the fact that the optimal plans that are generated when considering learning tend to have a structure as described in Section 5.2.2. Moreover, these structures tend to be similar for different realizations of the learning rate.

With regards to parameters C_1 and C_3, it can be argued that these parameters can be estimated reasonably well, as they are mainly determined by technical restrictions. For example, Øia et al. (2018) presents estimates of the minimum, most likely, and maximum time it takes to perform P&A operations based on the analysis of operational sequences, where it was assumed that the methods and associated equipment used consist of known practice and technology. Nevertheless, as there is always a possibility that an operator is unaware, or misspecifies, the values of C_1 and C_3, we conducted a similar "what-if" analysis around those parameters as well, varying them by $\pm 25\%$ of their nominal values in 5% increments. The detailed results are deferred to the Appendix. In summary, our findings were very similar to the case of C_2. First and foremost, the change in realized cost can be substantial, which can be attributed to the fact that any shift in the duration of operations directly affects the objective function through the rental costs. By and large, this change was found to be linear in the deviation of these parameters off their nominal values. Moreover, whereas the nominal solutions tend to remain optimal for the different C_1 parameter values we considered, they might become infeasible for C_3 increases of 10% and upwards, as was the case in few of our instances. Regardless, the nominal solutions tend to perform very well for all other scenarios, leading us to conclude that the value of perfect C_1 and C_3 information is also fairly low.

6. Conclusions

In this article, we presented an approach that allows for the inclusion of an endogenous learning effect in the setting of the uncapacitated Vehicle Routing Problem with Time Windows. This approach consisted of the definition of continuous experience variables as well as the formulation of (possibly non-linear) learning curves using piecewise-linear functions. To evaluate the effects of the endogenous learning effect, we applied the methodology to the problem of planning a Plugging and Abandonment campaign in the context of the offshore oil and gas industry. For this application, we developed a solution approach based on clustering that manages to solve the majority of real-life instances in seconds. Moreover, we extended existing instances for this problem with additional data on the learning effect. We observe that the inclusion of a learning effect leads to significantly different optimal plans than when neglecting the learning part. In general, we see that the optimal plans try to reap the benefits of learning by utilizing the vessels with most experience. The consideration of learning in the planning of plugging operations might lead to savings in the order of $3 - 20\%$. In addition, we showed that there exists significant value in cooperation between operators in terms of planning campaigns together, as a result of learning effects. This effect occurs even for very slow learning rates. We also tested the robustness of the obtained solutions for possible deviations in the learning curves, and we showed that deviations in the realized learning rate strongly affect the feasibility and costs of the campaign. However, we found that the value of perfect information is very low, and hence the nominal plan would perform equally well under different realizations of the learning rate. Only when the learning effect is much smaller than anticipated, the original plan might become infeasible, due to an increase in time usage. A possible direction for future work can be to investigate this challenge by means of an appropriate technique that deals with decision making under uncertainty. Overall, we conclude that the implications of a learning effect on VRP solutions can be significant and should therefore be explicitly incorporated in the decision-making process, whenever such effects are applicable.

Acknowledgments

This paper was prepared as a part of the project "Economic Analysis of Coordinated Plug and Abandonment Operations" (ECOPA), funded by the Research Council of Norway through the PETROSAM2 and PETROMAKS2 programs (p-nr: 247589).

References

Aloise, D. J., Aloise, D., Rocha, C. T., Ribeiro, C. C., Ribeiro Filho, J. C., & Moura, L. S. (2006). Scheduling workover rigs for onshore oil production. *Discrete Applied Mathematics*, *154*, 695–702. doi:10.1016/j.dam.2004.09.021.

Anzanello, M. J., & Fogliatto, F. S. (2011). Learning curve models and applications: Literature review and research directions. *International Journal of Industrial Ergonomics*, *41*, 573–583. doi:10.1016/j.ergon.2011.05.001.

Ascheuer, N., Fischetti, M., & Grötschel, M. (2001). Solving the asymmetric travelling salesman problem with time windows by branch-and-cut. *Mathematical Programming*, *90*, 475–506. doi:10.1007/PL00011432.

Azzouz, A., Ennigrou, M., & Ben Said, L. (2018). Scheduling problems under learning effects: classification and cartography. *International Journal of Production Research*, *56*, 1642–1661. doi:10.1080/00207543.2017.1355576.

Bakker, S., Vrålstad, T., & Tomasgard, A. (2019). An optimization model for the planning of offshore plug and abandonment campaigns. *Journal of Petroleum Science and Engineering*, *180*, 369–379. doi:10.1016/j.petrol.2019.05.042.

Balas, E., Fischetti, M., & Pulleyblank, W. R. (1995). The precedence-constrained asymmetric traveling salesman polytope. *Mathematical Programming*, *68*, 241–265. doi:10.1007/BF01585767.

Baldacci, R., Bartolini, E., & Laporte, G. (2010). Some applications of the generalized vehicle routing problem. *Journal of the Operational Research Society*, *61*, 1072–1077. doi:10.1057/jors.2009.51.

Baldacci, R., Christofides, N., & Mingozzi, A. (2008). An exact algorithm for the vehicle routing problem based on the set partitioning formulation with additional cuts. *Mathematical Programming*, *115*, 351–385. doi:10.1007/s10107-007-0178-5.

Baldacci, R., Mingozzi, A., & Roberti, R. (2011). New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research*, *59*, 1269–1283. doi:10.1287/opre.1110.0975.

Battarra, M., Erdogan, G., & Vigo, D. (2014). Exact algorithms for the clustered vehicle routing problem. *Operations Research*, *62*, 58–71. doi:10.1287/opre.2013.1227.

Bektas, T. (2006). The multiple traveling salesman problem: An overview of formulations and solution procedures. *Omega*, *34*, 209–219. doi:10.1016/j.omega.2004.10.004.

Biskup, D. (2008). A state-of-the-art review on scheduling with learning effects. *European Journal of Operational Research*, *188*, 315–329. doi:10.1016/j.ejor.2007.05.040.

Brett, J., & Millheim, K. (1986). The drilling performance curve: A yardstick for judging drilling performance. In *SPE Annual Technical Conference and Exhibition, 5-8 October, New Orleans, Louisiana*. Society of Petroleum Engineers. doi:10.2118/15362-MS.

Calvert, D. G., & Smith, D. K. (1994). Issues and techniques of plugging and abandonment of oil and gas wells. In *Proceedings - SPE Annual Technical Conference and Exhibition* (pp. 507–518). doi:10.2523/28349-ms.

Chen, X., Thomas, B. W., & Hewitt, M. (2016). The technician routing problem with experience-based service times. *Omega*, *61*, 49–61. doi:10.1016/j.omega.2015.07.006.

Chen, X., Thomas, B. W., & Hewitt, M. (2017). Multi-period technician scheduling with experience-based service times and stochastic customers. *Computers and Operations Research*, *82*, 1–14. doi:10.1016/j.cor.2016.12.026.

Cordeau, J.-F., Laporte, G., Savelsbergh, M. W., & Vigo, D. (2007). Chapter 6 Vehicle Routing. *Handbooks in Operations Research and Management Science*, *14*, 367–428. doi:10.1016/S0927-0507(06)14006-2.

Dantzig, G., Fulkerson, R., & Johnson, S. (1954). Solution of a large-scale traveling-salesman problem. *Journal of the Operations Research Society of America*, *2*, 393–410. URL: https://www.jstor.org/stable/166695.

Dohn, A., Rasmussen, M. S., & Larsen, J. (2011). The vehicle routing problem with time windows and temporal dependencies. *Networks*, *58*, 273–289. doi:10.1002/net.20472.

Gouveia, L., & Pires, J. M. (1999). The asymmetric travelling salesman problem and a reformulation of the Miller-Tucker-Zemlin constraints. *European Journal of Operational Research*, *112*, 134–146. doi:10.1016/S0377-2217(97)00358-5.

Hellström, A. H. (2010). *Statoil drilling and well learning curves, experience and theory: is there a learning curve from drilling the first well with a new rig and onwards?*. Master's thesis University of Stavanger. URL: http://hdl.handle.net/11250/182806.

Irawan, C. A., Ouelhadj, D., Jones, D., Stålhane, M., & Sperstad, I. B. (2017). Optimisation of maintenance routing and scheduling for offshore wind farms. *European Journal of Operational Research*, *256*, 76–89. doi:10.1016/j.ejor.2016.05.059.

Jepsen, M., Petersen, B., Spoorendonk, S., & Pisinger, D. (2008). Subset-row inequalities applied to the vehicle-routing problem with time windows. *Operations Research*, *56*, 497–511. doi:10.1287/opre.1070.0449.

Kaiser, M. J. (2017). Rigless well abandonment remediation in the shallow water U.S. Gulf of Mexico. *Journal of Petroleum Science and Engineering*, *151*, 94–115. doi:10.1016/j.petrol.2017.01.004.

Khalifeh, M., Saasen, A., Hodne, H., & Vralstad, T. (2013). Techniques and materials for North Sea plug and abandonment operations. In *Offshore Technology Conference 2013*. Houston, Texas, USA. doi:doi.org/10.4043/23915-MS.

Kunkel, M., & Schwind, M. (2012). Vehicle routing with driver learning for real world CEP problems. In *Proceedings of the Annual Hawaii International Conference on System Sciences* (pp. 1315–1322). doi:10.1109/HICSS.2012.633.

Lahyani, R., Khemakhem, M., & Semet, F. (2015). Rich vehicle routing problems: From a taxonomy to a definition. *European Journal of Operational Research*, *241*, 1–14. doi:10.1016/j.ejor.2014.07.048.

Langevin, A., Soumis, F., & Desrosiers, J. (1990). Classification of travelling salesman problem formulations. *Operations Research Letters*, *9*, 127–132. doi:10.1016/0167-6377(90)90052-7.

Laporte, G., & Nobert, Y. (1983). A branch and bound algorithm for the capacitated vehicle routing problem. *OR Spektrum*, *5*, 77–85. doi:10.1007/BF01720015.

Luo, Z., Qin, H., & Lim, A. (2014). Branch-and-price-and-cut for the multiple traveling repairman problem with distance constraints. *European Journal of Operational Research*, *234*, 49–60. doi:10.1016/j.ejor.2013.09.014.

Miller, C. E., Zemlin, R. A., & Tucker, A. W. (1960). Integer programming formulation of traveling salesman problems. *Journal of the ACM*, *7*, 326–329. doi:10.1145/321043.321046.

Moeinikia, F., Fjelde, K. K., Saasen, A., & Vralstad, T. (2014a). An investigation of different approaches for probabilistic cost and time estimation of rigless P&A in subsea multi-well campaign. In *SPE Bergen One Day Seminar*. doi:10.2118/169203-MS.

Moeinikia, F., Fjelde, K. K., Saasen, A., Vralstad, T., & Arild, O. (2014b). Evaluating cost efficiency of rigless P&A for subsea multiwell campaign. In *IADC/SPE Drilling Conference and Exhibition*. URL: 10.2118/167923-MS.

Nembhard, D. A., & Uzumeri, M. V. (2000). An individual-based description of learning within an organization. *IEEE Transactions on Engineering Management*, *47*, 370–378. doi:10.1109/17.865905.

Øia, T. M., Aarlott, M. M., & Vrålstad, T. (2018). Innovative approaches for full subsea P&A create new opportunities and cost benefits. In *SPE Norway One Day Seminar*. Bergen, Norway. doi:10.2118/191315-MS.

Oil & Gas UK (2015). *Guidelines for the abandonment of wells.* issue 5.

Öncan, T., Altinel, I. K., & Laporte, G. (2009). A comparative analysis of several asymmetric traveling salesman problem formulations. *Computers and Operations Research*, *36*, 637–654. doi:10.1016/j.cor.2007.11.008.

Pecin, D., Pessoa, A., Poggi, M., & Uchoa, E. (2017). Improved branch-cut-and-price for capacitated vehicle routing. *Mathematical Programming Computation*, *9*, 61–100. doi:10.1007/s12532-016-0108-8.

Pessoa, A., Sadykov, R., Uchoa, E., & Vanderbeck, F. (2019). A generic exact solver for vehicle routing and related problems. *Lecture Notes in Computer Science, Springer*, *11480*, 354–369. doi:10.1007/978-3-030-17953-3{_}27.

Pop, P. C., Kara, I., & Marc, A. H. (2012). New mathematical models of the generalized vehicle routing problem and extensions. *Applied Mathematical Modelling*, *36*, 97–107. doi:10.1016/j.apm.2011.05.037.

Ribeiro, G. M., Laporte, G., & Mauri, G. R. (2012). A comparison of three metaheuristics for the workover rig routing problem. *European Journal of Operational Research*, *220*, 28–36. doi:10.1016/j.ejor.2012.01.031.

Saasen, A., Moeinikia, F., Raksagati, S., Fjelde, K. K., & Vralstad, T. (2013). Plug and abandonment of offshore exploration wells. In *Offshore Technology Conference*. Houston, Texas. doi:10.4043/23909-MS.

Schneider, M. (2016). The vehicle-routing problem with time windows and driver-specific times. *European Journal of Operational Research*, *250*, 101–119. doi:10.1016/j.ejor.2015.09.015.

Sørheim, O., Ribesen, B., Sivertsen, T., Saasen, A., & Kanestrøm, O. (2011). Abandonment of offshore exploration wells using a vessel deployed system for cutting and retrieval of wellheads. In *SPE Arctic and Extreme Environments Conference and Exhibition* (pp. 69–81). Moscow, Russia. doi:10.2118/148859-MS.

Straume, M. (2018). Improved P&A performance on Valhall. Presentation at the Plug & Abandonment Seminar 2018, Sola, Norway. URL: `https://www.norskoljeoggass.no/contentassets/f60cf93f2c9c4c129b0a73303ff8081c/11---improved-pa-performance-on-valhall.pdf` Accessed: 8 June 2020.

Toth, P., & Vigo, D. (Eds.) (2002). *The vehicle routing problem*. Society for Industrial and Applied Mathematics.

Valdal, M. (2013). *Plug and abandonment operations performed riserless using a light well intervention vessel*. Master's thesis University of Stavanger. URL: `http://hdl.handle.net/11250/183553`.

Vrålstad, T., Saasen, A., Fjær, E., Øia, T., Ytrehus, J. D., & Khalifeh, M. (2019). Plug & abandonment of offshore wells: ensuring long-term well integrity and cost-efficiency. *Journal of Petroleum Science and Engineering*, *173*, 478–491. doi:`10.1016/j.petrol.2018.10.049`.

Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, *3*, 122–128. doi:`10.2514/8.155`.

Zhong, H., Hall, R. W., & Dessouky, M. (2007). Territory planning and vehicle dispatching with driver learning. *Transportation Science*, *41*, 74–89. doi:`10.1287/trsc.1060.0167`.

Appendix

Table 8: Performance of the optimal plan corresponding to the nominal value of C_1 (normalized here to 100%), under different realizations of this parameter.

Inst.	Change in realized cost using nominal solution as reference under various C_1 scenarios									
	C_1 =75 %	80 %	85 %	90 %	95 %	105 %	110 %	115 %	120 %	125 %
1	-6.76 %	-5.41 %	-4.06 %	-2.71 %	-1.35 %	1.35 %	2.71 %	4.06 %	5.41 %	6.76 %
2	-6.33 %	-5.07 %	-3.80 %	-2.53 %	-1.27 %	1.27 %	2.53 %	3.80 %	5.07 %	6.33 %
3	-5.00 %	-4.00 %	-3.00 %	-2.00 %	-1.00 %	1.00 %	2.00 %	3.00 %	4.00 %	5.00 %
4	-3.79 %	-3.03 %	-2.27 %	-1.52 %	-0.76 %	0.76 %	1.52 %	2.27 %	3.03 %	3.79 %
5	-3.89 %	-3.11 %	-2.33 %	-1.56 %	-0.78 %	0.78 %	1.56 %	2.33 %	3.11 %	3.89 %
6	-3.33 %	-2.66 %	-2.00 %	-1.33 %	-0.67 %	0.67 %	1.33 %	2.00 %	2.66 %	3.33 %
7	-3.91 %	-3.13 %	-2.35 %	-1.56 %	-0.78 %	0.78 %	1.56 %	2.35 %	3.13 %	3.91 %
8	-3.45 %	-2.76 %	-2.07 %	-1.38 %	-0.69 %	0.69 %	1.38 %	2.07 %	2.76 %	3.45 %
9	-1.78 %	-1.42 %	-1.07 %	-0.71 %	-0.36 %	0.36 %	0.71 %	1.07 %	1.42 %	1.78 %
10	-2.15 %	-1.72 %	-1.29 %	-0.86 %	-0.43 %	0.43 %	0.86 %	1.29 %	1.72 %	2.15 %

Table 9: Comparison of the plan obtained from the nominal value of C_1 (normalized here to 100%) and the optimal plan obtained under different settings of this parameter.

Inst.	Difference in optimal costs using nominal solution as reference, as a function of C_1									
	C_1 =75 %	80 %	85 %	90 %	95 %	105 %	110 %	115 %	120 %	125 %
1	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
2	0.07 %	0.05 %	0.04 %	0.02 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
3	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
4	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
5	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
6	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
7	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.03 %	0.06 %	0.10 %	0.08 %	0.11 %
8	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.02 %	0.00 %	0.00 %
9	0.05 %	0.05 %	0.05 %	0.05 %	0.05 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
10	0.07 %	0.05 %	0.08 %	0.09 %	0.09 %	0.10 %	0.10 %	0.11 %	0.12 %	0.10 %

Table 10: Performance of the optimal plan corresponding to the nominal value of C_3 (normalized here to 100%), under different realizations of this parameter.

Inst.	Change in realized cost using nominal solution as reference under various C_3 scenarios									
	C_3 =75 %	80 %	85 %	90 %	95 %	105 %	110 %	115 %	120 %	125 %
1	-15.71 %	-12.57 %	-9.43 %	-6.28 %	-3.14 %	3.14 %	6.28 %	9.43 %	12.57 %	15.71 %
2	-16.20 %	-12.96 %	-9.72 %	-6.48 %	-3.24 %	3.24 %	6.48 %	9.72 %	infeas.	infeas.
3	-17.93 %	-14.35 %	-10.76 %	-7.17 %	-3.59 %	3.59 %	7.17 %	10.76 %	14.35 %	17.94 %
4	-19.23 %	-15.39 %	-11.54 %	-7.69 %	-3.85 %	3.85 %	7.69 %	infeas.	infeas.	infeas.
5	-19.33 %	-15.47 %	-11.60 %	-7.73 %	-3.87 %	3.87 %	7.73 %	11.60 %	15.47 %	19.33 %
6	-18.47 %	-14.78 %	-11.08 %	-7.39 %	-3.69 %	3.69 %	7.39 %	11.08 %	14.78 %	18.48 %
7	-19.31 %	-15.45 %	-11.59 %	-7.73 %	-3.86 %	3.86 %	infeas.	infeas.	infeas.	infeas.
8	-19.87 %	-15.90 %	-11.92 %	-7.95 %	-3.97 %	3.97 %	7.95 %	11.93 %	15.90 %	19.88 %
9	-21.38 %	-17.10 %	-12.83 %	-8.55 %	-4.28 %	4.28 %	8.55 %	15.40 %	19.81 %	24.21 %
10	-20.82 %	-16.66 %	-12.49 %	-8.33 %	-4.16 %	4.16 %	8.33 %	infeas.	infeas.	infeas.

Table 11: Comparison of the plan obtained from the nominal value of C_3 (normalized here to 100%) and the optimal plan obtained under different settings of this parameter.

Inst.	Difference in optimal costs using nominal solution as reference, as a function of C_3									
	C_3 =75 %	80 %	85 %	90 %	95 %	105 %	110 %	115 %	120 %	125 %
1	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	25.16 %	24.29 %	24.12 %
2	0.47 %	0.36 %	0.26 %	0.16 %	0.07 %	0.00 %	0.00 %	0.00 %	infeas.	infeas.
3	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
4	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	infeas.	infeas.	infeas.
5	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
6	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
7	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %	infeas.	infeas.	infeas.	infeas.
8	5.77 %	5.11 %	4.30 %	3.81 %	2.52 %	0.00 %	0.00 %	0.00 %	0.00 %	0.00 %
9	0.07 %	0.06 %	0.06 %	0.06 %	0.05 %	0.00 %	0.00 %	1.42 %	1.44 %	1.57 %
10	0.79 %	0.79 %	0.75 %	0.51 %	0.10 %	0.09 %	0.09 %	infeas.	infeas.	infeas.