

Distributed Services with Foreseen and Unforeseen Tasks: The Mobile Re-allocation Problem

Jose Antonio Larco, Rommert Dekker and Uzay Kaymak

ERIM REPORT SERIES <i>RESEARCH IN MANAGEMENT</i>	
ERIM Report Series reference number	ERS-2007-087-LIS
Publication	December 2007
Number of pages	40
Persistent paper URL	http://hdl.handle.net/1765/10773
Email address corresponding author	jlarco@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS	
Abstract	In this paper we deal with a common problem found in the operations of security and preventive/corrective maintenance services: that of routing a number of mobile resources to serve foreseen and unforeseen tasks during a shift. We define the (Mobile Re-Allocation Problem) MRAP as the problem of devising a routing strategy to maximize the expected weighted number of tasks served on time. For obtaining a solution to the MRAP, we propose to solve successively a multi-objective optimization problem called the stochastic Team Orienteering Problem with Multiple Time Windows (s-TOP-MTW) so as to consider information about known tasks and the arrival process of new unforeseen tasks. Solving successively the s-TOP-MTW we find that considering information about the arrival process of new unforeseen tasks may aid in maximizing the expected proportion of tasks accomplished on time.
Free Keywords	routing, location, reliability, distributed services
Availability	The ERIM Report Series is distributed through the following platforms: Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal Social Science Research Network (SSRN), SSRN ERIM Series Webpage Research Papers in Economics (REPEC), REPEC ERIM Series Webpage
Classifications	The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems: Library of Congress Classification, (LCC) LCC Webpage Journal of Economic Literature, (JEL), JEL Webpage ACM Computing Classification System CCS Webpage Inspec Classification scheme (ICS), ICS Webpage

Distributed services with foreseen and unforeseen tasks: The Mobile Re-allocation Problem

Jose Antonio Larco¹, Rommert Dekker², Uzay Kaymak²

¹RSM Erasmus University, P.O. box 1738, 3000 DR Rotterdam, The Netherlands,

jlarco@rsm.nl (corresponding author) Tlf +31 10 408 25 44 Fax +31 10 408 90 14

²Econometric Institute, Erasmus University Rotterdam P. O. Box 1738, 3000 DR, Rotterdam,

The Netherlands, dekker@few.eur.nl, u.kaymak@ieee.org

Abstract

In this paper we deal with a common problem found in the operations of security and preventive/corrective maintenance services: that of routing a number of mobile resources to serve foreseen and unforeseen tasks during a shift. We define the (Mobile Re-Allocation Problem) MRAP as the problem of devising a routing strategy to maximize the expected weighted number of tasks served on time. For obtaining a solution to the MRAP, we propose to solve successively a multi-objective optimization problem called the stochastic Team Orienteering Problem with Multiple Time Windows (s-TOP-MTW) so as to consider information about known tasks and the arrival process of new unforeseen tasks. Solving successively the s-TOP-MTW we find that considering information about the arrival process of new unforeseen tasks may aid in maximizing the expected proportion of tasks accomplished on time.

Keywords:

Routing, location, reliability, distributed services.

1 Introduction

The domain of geographically distributed security services and preventive/ corrective maintenance services involves the task of routing mobile resources over a number of sites so that various services can be performed. While the clients, as the recipients of the services, are mainly concerned with receiving timely and reliable services so that their security/maintenance problems are solved effectively, the operational managers of the service providers are mainly concerned with offering competitive and profitable services. In a competitive environment, where the firms choose to differentiate themselves from competitors by focusing on service quality, rather than on costs, responsiveness is one of the most important quality dimensions of service operations as recognized in the "Gap model" by Zeithaml et al. (1990).

One of the measurable dimensions of quality of service is that of fulfilling contract specifications. Given that the clients' main concern is that their security/maintenance problem is being dealt with, the main requirement specified in contracts to be fulfilled in the service operations is that the services must be started on site within certain time windows (Larco, 2007).

The determination of these types of time windows depends on the type of service/task to be performed. We distinguish between two kinds of tasks: foreseen tasks which are known in advance at the start of a shift and unforeseen tasks which are only known at the time an alarm is triggered. While foreseen tasks involve routine activities such as verification of alarm systems or preventive maintenance, unforeseen tasks are emergency related activities addressing critical situations such as a robbery or a failure of an important piece of equipment (Larco et al, 2006). As a result, in the case of foreseen tasks, the latest allowable starting times are specified in advance at the request of the client, whereas in the case of unforeseen tasks, these are defined by a given standard response-time which is a service-guarantee that can be easily understood

(for a discussion on service guarantees see Hart, 1988).

Operational managers and/or dispatchers have the task of ensuring that the routing decisions, which are taken before the start of a shift, and the re-routing decisions are such that they allow servicing as many foreseen and unforeseen tasks as possible within the specified time-windows and within the shift's length. Although we acknowledge that there are several sources of uncertainty in the distributed services domain, such as non-deterministic travelling times or task service durations, we restrict our study in this paper to the main source of uncertainty: that of the release of unforeseen tasks. Hence, in this paper we study the challenge of (re-)routing when dealing with (un)foreseen tasks so as to meet specifications of service contracts and offer reliable and timely services.

To this end, our aim is to solve an incomplete information problem, which we call the Mobile Re-allocation Problem (MRAP). In contrast to other models of geographically distributed services where the goal is to minimize response times (see Weintraub et al., 1999) or minimize costs by minimizing total travel distance (see Johns, 1995), in the MRAP, our goal is to maximize the expected (weighted¹) number of tasks that are serviced by mobile resources within the specified contractual deadlines.

To solve the MRAP, we propose to solve successively a static multi-criteria problem which we call the stochastic Team Orienteering Problem with Multiple Time Windows (s-TOP-MTW). The first criterion makes use of complete information of known tasks while the second criterion makes use of incomplete information about the arrival process of new unforeseen tasks. The first criterion is to serve as many known weighted tasks as possible within given time-windows and a given shift using a variant of the Team Orienteering Problem (Golden et al. 1996). The second criterion, recognizes that given that the arrival time of unforeseen tasks is not known a-priori, it is important to know the risk that a certain set of routes of mobile resources implies. To evaluate such a risk, we extend Daskin's Maximum Expected Covering Location Problem

¹The weights provide an indication of the relative task's importance.

(MEXCLP) (Daskin, 1983) where it is possible to identify which sites are covered so as to service potential requests of unforeseen tasks on time.

In the original version of the MEXCLP, the problem is applied to static coverage situations where emergency vehicles (e.g. ambulances) return to their bases immediately after serving *every* unforeseen task, hence to identify which sites are covered only the fixed location of the base is relevant. In contrast, the extended MEXCLP version we propose is applied to dynamic coverage situations where mobile resources (such as security guards or service engineers) have to modify their routes to accommodate new unforeseen tasks without having to return to the base implying that the location of each mobile resource at each point in time is relevant for identifying the sites been covered and must be updated. The location of mobile resources can now be identified using the tracking capabilities of GPS that is already been used in the geographical distributed services domain. By incorporating the proposed extended MECXLP model as the second criterion, we are interested in investigating the added value of the information that the arrival process can bring for the design of more reliable routes. To our knowledge, there is no available literature that deals with the combination of routing and location models where the location refers not to static bases or sites but to the mobile resources themselves.

This paper is organized as follows. In Section 2, we describe the MRAP and its assumptions. In Section 3, we outline the multi-criteria problem that we refer to as the stochastic-Team Orienteering Problem with Multiple Time Windows (s-TOP-MTW). In Section 4, we provide an algorithm to solve the s-TOP-MTW problem. In Section 5, we present two experiments where we first investigate the relationship between routing and coverage considerations and then we test through simulation whether or not including risk-coverage considerations yields better results for solving the MRAP.

2 The Mobile Re-allocation Problem (MRAP)

The Mobile Re-allocation Problem is a dynamic problem with incomplete information and as such requires the use of schedules that are re-evaluated each time new information is available. Thus, in the MRAP, we start the shift with a schedule for foreseen tasks and then this schedule can be modified at any moment as soon as new (unforeseen) tasks become available. In the following, we describe the dynamic characteristics of the MRAP and then provide a precise definition of the MRAP. To conclude, we illustrate the dynamics of the MRAP with an example.

2.1 General Assumptions

We assume that the travelling times between sites are deterministic (the mobile resources travel fixed distances at a constant c speed) and that the durations of both foreseen and unforeseen tasks are known beforehand. Furthermore, we abstract from road networks and assume that all sites may be reached from any other location.

On the other hand, we assume the problem during a shift of length TT where all the mobile resources start at a base and have to return at the end to the base at the end of the shift. Moreover, we assume that it is a single dispatcher who makes a schedule for every mobile resource and that every mobile resource is able to serve any task (i.e. there is no zoning). In this problem setting, the fixed information of the MRAP (i.e. known at any point in time) is given by the following:

- I is the set of sites.
- K is the set of mobile resources.
- $d(i_1, i_2)$ is the Euclidian distance between any pair of sites $i_1, i_2 \in I$.
- δ is the system-wide standard response time indicating the time available to start serving an unforeseen task since its release.

- $1/\lambda$ is the mean of the exponentially distributed interarrival time between unforeseen tasks.
- P_i is the probability that a given unforeseen task is at site $i \in I$.
- μ is the deterministic and fixed duration of any unforeseen tasks generated.
- Δ_n is the deterministic duration of task n , when the task is of unforeseen type $\Delta_n = \mu$; otherwise Δ_n is a fixed value known at the start of the shift.
- TT is the total shift length.

2.2 Tasks

The services delivered by mobile resources are a series of tasks. Every task $n \in N$ has the following tuple associated to it: $\varphi_n = (L(n), w_n, r_n, \Delta_n, e_n, l_n)$ where $L(n) = i : i \in I$ is the site at which the task is to be performed, w_n is the importance of servicing task $n \in N$, r_n is the release time at which the task is known to the dispatcher and e_n and l_n are the earliest and latest allowable starting times respectively.

Note that the release time allows to distinguish between a foreseen and an unforeseen task. If $r_n = 0$ the task is of foreseen type (i.e. $n \in F$) and if $r_n > 0$ the task is of unforeseen type (i.e. $n \in U$). The starting time windows are also derived differently for each type of task. While for foreseen tasks, e_n and l_n are determined a-priori, for unforeseen tasks these are determined with the following relations: $e_n = r_n$ and $l_n = r_n + \delta$. Additionally, note that $\Delta_n = \mu$ for any unforeseen task (i.e. if $r_n > 0$) and that Δ_n is known a-priori for any foreseen task.

However, since the MRAP is an incomplete information problem, not all the tasks of a shift are necessarily known. If at instant $j \in J$ an unforeseen task arises, we define TK_j as the set of known pending tasks; i.e. all foreseen and released unforeseen tasks that have not been serviced before instant j .

2.3 Valid Schedules

In the MRAP, the main task of a central dispatcher is that of defining valid schedules for each of the mobile resources available. However, as these schedules are expected to be modified each time new unforeseen tasks are released, we need to (re-)evaluate schedules at instants $j \in J$. A valid schedule at instant j is then defined as follows.

Definition 1: *A valid schedule for a mobile resource $k \in K$ at a given evaluation instant $j \in J$, is defined as an ordered set of known pending tasks denoted by $\Pi_j^{(k)} \subseteq TK_j$ to be accomplished by the mobile resource that fulfills the following dynamic conditions:*

1. *The mobile resources have a given starting position, which is the home base (common to all mobile resources) at the start of their shift and their location at any time the route is re-evaluated. The mobile resources must end their shift at a unique base; servicing only tasks that allow them to return to the base before the termination of their shift at $t=TT$.*
2. *Each task included in the schedule must start to be executed within the service starting time windows: $[e_n, l_n]$.*
3. *Once a task is executed, the mobile resource re-allocates immediately to either service another task or return to the base. If the mobile resource arrives at a site before the corresponding earliest allowable starting time, it is assumed that it waits just outside the premises of the site until the earliest allowable starting time, e_n .*
4. *If a mobile resource is en-route to service an unforeseen task, such a mobile resource can not be re-allocated to another task until it finishes servicing the unforeseen task.*
5. *If a mobile resource is servicing a foreseen or an unforeseen task, such a mobile resource can not be re-allocated to serve another task until the mobile resource finishes servicing its current task.*

The first two conditions enforce that mobile resources start servicing tasks within the specified time windows so that the mobile resources are able to return to the base before the end of the shift. The third condition is used for simplification purposes where a sequence of tasks fully determines the actual position of mobile resources at given time instants. The last two conditions are preemptive rules that guarantee that a mobile resource servicing a task is not interrupted either while serving a task or while transferring to serve an unforeseen task so that clients do not perceive that a mobile resource is abandoning the service of its task to serve another task.

Given that we assume in the MRAP the time to transfer between sites is pre-determined as the speed and distances between sites are constant, it is possible to know in advance if a released task can be started to be serviced within given time windows. Thus, we can also assume that tasks are not serviced if it is impossible to meet their associated time windows. Indeed, it is often the case in real settings that a dispatcher notifies the client when it is not possible to serve a task on time.

2.4 Information Set

At the start of the shift and when an unforeseen task is released, updated information is available to the dispatcher for it to generate new schedules. The information set, \mathcal{I}_j , at evaluation instant $j \in J$, includes the updated information of the tasks available to be served as well as information about the current situation of every mobile resource $k \in K$.

$$\mathcal{I}_j = (TK_j, \bigcup_k GL_k^{(j)}, \bigcup_k GS_k^{(j)}, \bigcup_k LT_k^{(j)}) \quad (1)$$

where:

- $GL_k^{(j)}$: is the actual location of mobile resource k in the Euclidian space at evaluation instant $j \in J$.

- $GS_k^{(j)}$ identifies the state at which a mobile resource k is at evaluation instant $j \in J$ such

that:

$$GS_k^{(j)} = \begin{cases} 1 & \text{if mobile resource } k \text{ is waiting to serve a foreseen task at its site.} \\ 2 & \text{if mobile resource } k \text{ is currently servicing a (un)foreseen task.} \\ 3 & \text{if mobile resource } k \text{ is travelling to serve a foreseen task.} \\ 4 & \text{if mobile resource } k \text{ is travelling to serve an unforeseen task.} \\ 5 & \text{if mobile resource } k \text{ is travelling back to the base.} \end{cases}$$

- $LT_k^{(j)}$ is the first time at which a dispatcher can reschedule a mobile resource k given the current time at evaluation instant $j \in J$ in accordance with the conditions defined in Definition 1 .

2.5 Strategy

We define a strategy as a mechanism that generates valid schedules for each mobile resource upon the occurrence of each event (i.e. a request to serve a new unforeseen task). More formally, by defining a system schedule as the collection of all valid mobile resources schedules at a given evaluation instant $j \in J$: $\Gamma_j = \bigcup_{k \in K} \Pi_j^{(k)}$, we define a strategy as follows:

Definition 2: *A strategy denoted by s is a function that maps the information set at a given evaluation instant $j \in J$ to a system schedule: $s : \mathcal{I}_j \rightarrow \Gamma_j$.*

2.6 Performance measurement in the MRAP

To evaluate the quality of the schedules generated by a given strategy during the realization of a certain shift, we introduce the concept of weighted fulfilment yield which is in-line with the service operational managers' objective of servicing as many tasks as possible while recognizing that not all the tasks are of equal importance.

Definition 3: *The weighted fulfilment yield, Ψ , is a random variable defined as the weighted proportion of tasks $n \in N$ (weighted by w_n) that are completely serviced within their time-*

Table 1: Initial data for the example considered.

Task index {Site Location}	1{a}	2{b}	3{c}	4{d}	5{e}	6{f}
X-coordinate (km)	-1.0	1.0	-1.0	1.0	-1.0	1.0
Y-coordinate (km)	1.0	1.0	0.0	0.0	-1.0	-1.0
Earliest allowable starting time (hr)	2.0	1.0	1.0	5.0	0.5	2.5
Latest allowable starting time (hr)	5.0	7.0	9.5	8.5	3.5	4.5
Duration shift (hr)	12.0					
Duration of foreseen Δ_n {unforeseen μ } tasks (hr)	1.0{1.2}					
Relative importance of foreseen {unforeseen} tasks	16.7{100}					
Mean rate of unforeseen tasks (# tasks in $[TT - \mu]$)	1					
Speed of mobile resources (km/hr)	1					
Response time standard δ (hr)	1.5					

windows during a shift $[0, TT]$.

As a problem with incomplete information, the objective of the MRAP is to maximize the *expected* weighted fulfilment yield γ over all possible realizations of unforeseen tasks in a shift. Therefore, we define the MRAP as follows:

Definition 4: *The objective of the MRAP is to devise a strategy $s \in S$ for the allocation and re-allocation of mobile resources to tasks so that the strategy maximizes the expected weighted fulfilment yield $E(\Psi)$ during a shift given that unforeseen tasks arise following a Poisson process with intensity λ and that the probability that the unforeseen task is located at a given site is given by P_i for every $i \in I$.*

The constraints of this incomplete information optimization problem are stated implicitly by constraining possible solutions to strategies that generate *only* valid schedules (see the conditions that a valid schedule has to fulfill in Section 2.3 for the constraints of the MRAP problem).

2.7 Risks inherent to routes

To illustrate the MRAP, consider the next example where the information initially available is given in Table 1. There are six sites which are fully connected. Further, assume that two mobile resources are available and that the number of unforeseen tasks that occur in a shift are Poisson distributed with an equal likelihood of being located at any site.

Table 2: Timing information of solution sets.

Solution Set -Mobile resource	1{a}	2{b}	3{c}	4{d}	5{e}	6{f}	Base
Solution Set A: Mobile resource 1	4.41	1.41	6.41	-	-	-	8.41
Solution Set B: Mobile resource 2	-	-	-	6.41	1.41	4.41	8.41
Solution Set A: Mobile resource 1	3.00	6.00	1.00	-	-	-	8.41
Solution Set B: Mobile resource 2	-	-	-	6.41	1.41	4.41	8.41

For illustration purposes, consider two alternative solution sets as an initial schedule for the MRAP. In Solution Set A, mobile resource 1 follows an anticlockwise pattern visiting the sites in the sequence b-a-c, while mobile resource 2 also follows an anticlockwise pattern visiting the sites in the sequence e-f-d. Solution Set B differs from Solution Set A only in the route of mobile resource 1 that visits the sites in a clockwise manner in the sequence c-a-b. Both sets of plans are feasible solutions, fulfilling the corresponding time windows and returning to the base before the end of the shift (see Table 2).

If no alarms occur during the duration of the shift then it is clear that both solution sets serve all the foreseen tasks achieving a 100% weighted fulfilment yield. However, unforeseen tasks do occur and hence it is worth verifying if both sets of plans yield also the same risk.

To assess such a risk, it is useful to track the positioning and state of the mobile resources at different instants. While in Solution Set A (see Figure 1(a)) both mobile resources tend to be positioned in opposite sides at the same instants, in Solution Set B (see Figure 1(b)), the mobile resources are located in the same side (i.e. left or right) of the graph. If we then draw a circle with radius $\delta = 1.5$ hrs (the response time standard) with the position of each mobile resource as a centre, we can identify the sites that are able to be reached on time if an unforeseen task is released and thus have certain insight about the inherent risks to routes.

For example, at time instant $t = 7.50$ hrs Figure 1 shows that while the setting in Solution Set A covers all the sites (Figure 1(c)), the setting in solution set B (Figure 1(d)) does not cover sites a, c and e. If, for example, an alarm occurs at site a at $t = 7.50$ hrs, mobile resource 1 in Solution Set A will be able to serve it returning to the base on time. However, none of the




Figure 1: Coverage of mobile resources for unforeseen tasks.

mobile resources in Solution Set B will be able to serve the alarm because these are located too far from the site as shown in Figure 1(d). In the case only one unforeseen task is released at site a at $t = 7.50$ hrs during the shift, then Solution Set A provides a $\gamma = 100\%$ of fulfilment yield. In contrast, Solution Set B, provides only $\gamma = 50\%$ of fulfilment yield. Thus, given the increased separation of mobile resources and the higher number of sites covered in time of Solution Set A compared to Solution Set B, it is reasonable to choose Solution Set A over Solution Set B on the grounds of being a less risky choice as the mobile resources cover more sites on average.

The example illustrates that for obtaining higher fulfilment yields it is important to design

routes that serve as many known tasks as possible, while maintaining at the same time the capacity of mobile resources to handle unforeseen tasks on time (i.e. identifying which sites are covered). However, it is also important to quantify such capacity over time so as to have an objective criterion that identifies which sets of routes are more risky in terms of handling unknown unforeseen tasks on time. The next section deals with how to address the complete and incomplete information aspects of the MRAP.

3 A valid strategy for the MRAP: the s-TOP-MTW

In this section we present a multi-objective optimization problem to be solved successively each time a new unforeseen task is generated at instants j . We call such multi-objective problem the stochastic Team Orienteering Problem with Time Windows (s-TOP-MTW). The problem is in itself composed of two sub-problems, each with its own objective function.

- The Team Orienteering Problem with Time Windows (TOP-MTW)
- The Time-averaged Maximum Expected Location Problem (TAMEXCLP)

The TOP-MTW addresses the known tasks by maximizing the weighted number of tasks served on time. On the other hand, the TAMEXCLP, addresses the stochastic process of the arrival of new unforeseen tasks by maximizing the capacity of mobile resources routes to attend new unforeseen tasks on time.

Nonetheless, both sub-problems share the same solution structure: an ordered sequence of tasks to be served per mobile resource. In the case of the TOP-MTW, the ordered sequence of tasks directly provides the weighted number of tasks accomplished on time. In the case of the TAMEXCLP, the sequence of tasks influences the objective function in a more indirect way. Namely, given that we assume that the sites are represented in a complete graph in the Euclidean plane, the sequence of tasks determines the locations of mobile resources at certain points in

time, defining which sites are covered and subsequently what is the inherent capacity of mobile resources to serve unforeseen tasks on time.

To aggregate the individual objectives of both sub-problems in one objective function Z we propose to integrate the normalized TOP-MTW and TAMEXCLP (i.e. Z_{S1} and Z_{S2}) in an affine combination in line with a common method of multi-criteria optimization: the point-estimate weighted sums approach (Steuer, 1986). The suitability of choosing a certain weighing value (i.e. α) will be studied later in the experimental section. The general structure of the s-TOP-MTW is then as follows:

$$\max Z = (1 - \alpha)Z_{S1} + \alpha Z_{S2} \quad (3)$$

s.t.

$$\text{TOP-MTW constraints} \quad (4)$$

$$\text{TAMEXCLP and interaction constraints} \quad (5)$$

Our overall approach is to solve the MRAP by solving successively the s-TOP-MTW as follows.

1. *Choose a given α .*
2. *Solve the s-TOP-MTW given the information set \mathcal{I}_j .*
3. *Upon release of unforeseen task, update the information set \mathcal{I}_j adding the newly released unforeseen task to the set of known tasks $n \in TK_j$.*
4. *Re-solve the s-TOP-MTW given the updated information set \mathcal{I}_j .*

We note that since any unforeseen task generated after $TT - \mu$ will not be able to be served in a way that allows the base to be reached before the end of the shift we assume that the release time of the last unforeseen task to be generated is before $TT - \mu$.

In the next sub-sections we present the formulations and advantages of including each sub-problem in the s-TOP-MTW. The formulations are given for the case where the resources start

at the home base. It is easy to change the formulation for more general starting conditions, which is needed for rescheduling during the shift.

3.1 Subproblem 1: The Team Orienteering Problem with Time Windows (TOP-MTW)

To select a suitable routing model that can serve as a basis for dealing with known tasks in the MRAP, we notice the similarity between the objective of the MRAP which is to serve as many tasks on time as possible and that of a problem found in the literature called the Orienteering Problem (OP) that has been classified as part of the TSP with Profits set of problems (Feillet, 2005). The OP (Golden, 1987) is inspired by an outdoor sport where the participants are equipped with a compass and a map and have to visit a number of checkpoints displayed on the map. Each check point has an associated score which is collected when a participant visits the check point. The goal is to maximize the total score collected by visiting a number of checkpoints within a limited time-span. Hence, the OP and the MRAP have certain similarities, where the relative importance of servicing a task is equivalent to a reward and the limited time-span is equivalent to a shift length.

Next, we can further extend the OP to incorporate multiple time windows per site for several tasks per site and multiple resources which we call the Team Orienteering Problem with Multiple Time Windows (TOP-MTW). It is important to observe that if we set the mean number of unforeseen tasks per shift to zero (i.e. $\lambda = 0$) then no unforeseen tasks can be generated in a shift, converting the MRAP into a complete information problem and thus, solving the MRAP is reduced to solving the TOP-MTW.

To extend the OP to the TOP-MTW we review the existing variants of the OP. The multi-resource version of the OP already exists in the literature and is known as the Team Orienteering Problem (TOP) (Chao, 1996). Kantor and Rosenwein (1992) developed a time-windows version (one time-window per site/vertex) for the OP but the time windows refer to the end of the service

time rather than the start service time. Later, Nguyen and Gao (2003) extend the possibility for several time-windows per site. However, both formulations of time windows consider deadlines instead of starting servicing times. Hence we have to adapt existing formulations and develop a new one. The formulation defined for every occurrence j of the TOP-MTW is as follows using the notation introduced in Section 2.

Data

$d(m, n)$: the Euclidean distance between the locations $L(m), L(n)$ of two tasks $m, n \in TK_j$,

$D(m, n)$: the transfer time between the locations $L(m), L(n)$; $D(m, n) = d(m, n)/c$,

$i = 0$: the index value of the base as a site,

$n = 0$: the initial artificial task located at the base such that

$$L(0) = \Delta_0 = e_0 = l_0 = w_0 = 0$$

$n = N$: the last artificial task located at the base such that

$$L(N) = \Delta_N = e_N = w_N = 0, l_N = TT$$

M : a large number,

ζ_n : real numbers associated with task $n \in TK_j$ used to eliminate sub-tours,

Decision variables

$x_{mn}^{(k)}$: 1, if task $n \in TK_j$ is scheduled immediately after task $m \in TK_j$,
by mobile resource $k \in K$ such that $n \neq m$; 0 otherwise,

Derived values from $x_{mn}^{(k)}$

z_n : 1, if task $n \in TK_j$ is scheduled to be served; 0 otherwise,

$y_n^{(k)}$: 1, if task $n \in TK_j$ is scheduled to be served by mobile resource $k \in K$;
0 otherwise,

$s_n^{(k)}$: the starting servicing time of task $n \in TK_j$ by mobile resource $k \in K$.

The TOP-MTW problem to be solved at evaluation instant j is now defined by the following.

$$\max Z_{S1} = \frac{\sum_{n \in TK_j} w_n z_n}{\sum_{n \in TK_j} w_n} \quad (6)$$

s.t.

$$\sum_{k \in K} y_n^{(k)} = z_n \quad \forall n \in TK_j \quad (7)$$

$$\sum_{\substack{n \in TK_e, \\ n \neq m}} x_{mn}^{(k)} = y_m^{(k)} \quad \forall m \in TK_j, \forall k \in K, \quad (8)$$

$$e_n \leq s_n^{(k)} \leq l_n \quad \forall n \in TK_j, \forall k \in K, \quad (9-10)$$

$$M(1 - x_{mn}^{(k)}) \geq s_m^{(k)} + \Delta_m + D(m, n) - s_n^{(k)} \quad \forall m, n \in TK_j : n \neq m, \forall k \in K, \quad (11)$$

$$M(x_{mn}^{(k)} - 1) \leq s_m^{(k)} + \Delta_m + D(m, n) - s_n^{(k)} \quad \forall m, n \in TK_j : n \neq m, \forall k \in K, \quad (12)$$

$$\zeta_m - \zeta_n + |TK_{\hat{e}}| x_{mn}^{(k)} \leq |TK_{\hat{e}}| - 1 \quad \forall m, n \in TK_j, : m \neq 0, n \neq m, \quad (13)$$

$$\sum_{\substack{n \in TK_e, \\ n \neq 0}} x_{0n}^{(k)} = \sum_{\substack{m \in TK_e, \\ m \neq N}} x_{mN}^{(k)} = 1, \quad \forall k \in K, \quad (14-15)$$

$$\sum_{\substack{n \in TK_{\hat{e}}, \\ n \neq N}} x_{Nn}^{(k)} = 0, \quad \forall k \in K, \quad (16)$$

$$s_0^{(k)} = 0, y_0^{(k)} = 1 \quad \forall k \in K, \quad (17-18)$$

$$z_n \in \{0, 1\}, \zeta_n \in \mathbb{R} \quad \forall n \in TK_j, \quad (19-20)$$

$$y_n^{(k)} \in \{0, 1\}, s_n^{(k)} \in [0, TT) \quad \forall n \in T, \forall k \in K, \quad (21-22)$$

$$x_{mn}^{(k)} \in \{0, 1\} \quad \forall m, n \in TK_j, : n \neq m, \quad (23)$$

The normalized objective (6) maximizes the weighted proportion of tasks accomplished on time (note that the denominator is only used to integrate it with the with the second goal of the s-TOPMTW).

Constraints (7) assure that a task will at most be served by one mobile resource as the maximum value of z_n is 1. Similarly, constraints (8) assure that a task m to be served may have at most one task preceding it (i.e. n), because the maximum value of $\hat{y}_n^{(k)}$ is 1.

Constraints (9) and (10) assure that any task n that is serviced by a mobile resource starts to be executed within the respective time windows.

The set of constraints (11) and (12) force that if task m immediately precedes task n (i.e. $y_{mn}^{(k)}$), then the starting time of a task n must be *exactly* equal to the starting time of the preceding task m plus the time taken to service task n and to transfer to task m . Hence,

the relation $s_n = s_m + \Delta_m + D(m, n)$ is forced by making the upper and lower bounds equal. Constraints (11) and (12) allow $s_m^{(k)}$ and $s_n^{(k)}$ take any positive value if task n does not precede task m in mobile resource k route using the "big M " method, reducing the expressions to $-M \leq s_m + \Delta_m + D(m, n) - s_n^{(k)} \leq M$. The formulations of these constraints differ from those in Nguyen and Gao (2003) as in the MRAP we have time windows that refer to the starting times and not the ending times of the respective services. Moreover, we also assume that as soon as one mobile resource finishes serving a task it immediately reallocates to the next task (i.e. Definition 1, condition 3) and thus its starting time of the next served task is fully predetermined by the visiting sequence of tasks. In the case of Nguyen and Gao (2003) and Kantor and Rosenwein (1992), the starting times of a service are not fully predetermined by the visiting sequence of tasks.

Next, constraints (13) eliminate the possibility of sub-tours of tasks located at sites that do not involve the base, while constraints (14), (15) and (16) introduce two fictitious task located at the base so as to ensure that the mobile resources' routes start and finish at the base within the shift interval $[0, TT]$.

Finally, constraints (17) to (23) define the types of the decision variables of the model.

3.1.1 Subproblem 2: The Time-averaged Maximum Expected Location Problem (TAMEXCLP)

In Section 3.1 it was shown that it is possible to identify which sites are covered for servicing potential unforeseen tasks on time. If it is known which sites are covered, it is possible to obtain a measure of the capacity of mobile resources to serve unforeseen tasks on time.

A useful measure of the capacity for servicing potential unforeseen tasks is the probability that if an unforeseen task is released, at least one mobile resource will be able to be service it on time. Multiplying this probability with the associated relative importance of servicing an unforeseen task and the probability that an unforeseen task will occur at such site, it is possible

to obtain a measure of the expected demand covered by the mobile resources at given time instants.

In fact, for a given time instant, the expected demand covered by the mobile resources is equivalent to the objective function of the Maximum Expected Location Problem or MEXCLP (Daskin, 1983), which is used for locating emergency vehicles at certain bases to cover possible sites that request emergency services. The difference with the MRAP lies in the fact that emergency vehicles serve emergency requests and then return to the base while mobile resources in the MRAP follow a route servicing several tasks returning to the base only when there are no more tasks to perform. However, by taking a representative number of "snapshots" of the system at regular intervals of length ϵ it is possible to measure the instantaneous capacity of mobile resources to serve unforeseen tasks. Nonetheless, it is important to note that this is only an approximation of the expected demand covered at each snapshot as we assume that the current routes followed are left unchanged until the "snapshot" where the coverage is evaluated. Furthermore, this method for calculating the expected demand covered is limited for situations in which queuing phenomena are unlikely to occur, such as in the case of practical applications of the MRAP where the emergency events are rare and the standard response times are short enough so that it is unlikely that pending released unforeseen tasks accumulate at the same time (for a discussion on the accuracy of the MEXCLP model see Saydam and Atuğ (2003)).

If we apply the concept of sampling MEXCLP objective values over time to the presented example we obtain a graph (see Figure 2) that clearly indicates that on average Solution Set A is more reliable than Solution Set B by covering more sites during the shift. The MEXCLP measured over time can then be seen as a measurement of the quality of coordination between mobile resources. When mobile resources are located further apart and do not execute tasks simultaneously there is less overlapping of coverage and thus the mobile resources can more efficiently cover sites for unforeseen tasks.




Figure 2: MEXCLP sampling over time.

For the purpose of devising a strategy for the MRAP, we propose to measure the MEXCLP over time and average it to construct a new time-dependent dynamic coverage problem: the Time-Averaged Maximum Expected Location Problem (TAMEXCLP). In the next lines we present the general structure for the formulation of the TAMEXCLP. It must be noted, however, that due to the application of interpolation and Euclidian distances in the formulation it is not possible to model it as a mixed integer program implying that certain solution techniques can not be applied.

We first define the following variables.

Data

- TS : the total number of discrete points in time,
- ϵ : the time between discrete points in time, i.e. $\epsilon = \frac{TT-\mu}{TS}$,
- p : index identifying discrete point in time; $t_p = \epsilon p, t_p = t_1, \dots, t_{TS}$
- w_i : the relative importance associated to each site i , assuming that if an unforeseen task is generated at this site it will have such a weight,

Derived values from $x_{mn}^{(k)}$

- $g_{ip}^{(k)}$: 1, if mobile resource k is covering site i at time t_p ; 0, otherwise,
- $\rho_p^{(k)}$: radius covered (in time units) by mobile resource k at time t_p ,
- $X_p^{(k)}, Y_p^{(k)}$: X and Y coordinates of mobile resource k at point at time t_p ,
- $q_p^{(k)}$: busy fraction; the probability that mobile resource k at point at time t_p can not serve an unforeseen task if it is covered and is generated within the interval (t_p, t_{p+1})
- $q_{ip}^{*(k)}$: the probability that mobile resource k at point at time t_p can serve an unforeseen task if it is generated specifically at site i within interval (t_p, t_{p+1}) .

We can then define the TAMEXCLP as follows.

$$\max Z_{S2} = \frac{1}{FP \sum_{i \in I} w_i P_i} \sum_{p=1}^{p=FP} \sum_{i \in I} P_i w_i \left[1 - \prod_{k=1}^{|K|} q_{ip}^{*(k)} \right] \quad (24)$$

s.t.

$$\text{Determination of } X_p^{(k)}, Y_p^{(k)} \text{ constraints} \quad \forall k \in K, \forall p = 1, \dots, TS \quad (25)$$

$$\text{Determination of } \rho_p^{(k)} \text{ constraints} \quad \forall k \in K, \forall p = 1, \dots, TS \quad (26)$$

$$\text{Determination of } g_{ip}^{(k)} \text{ constraints} \quad \forall i \in I, \forall k \in K, \forall p = 1, \dots, TS \quad (27)$$

$$\text{Determination of } q_p^{(k)} \text{ constraints} \quad \forall k \in K, \forall p = 1, \dots, TS \quad (28)$$

$$\text{Determination of } q_{ip}^{*(k)} \text{ constraints} \quad \forall i \in I, \forall k \in K, \forall p = 1, \dots, TS \quad (29)$$

$$q_{ip}^{*(k)} \in [0, 1], g_{ip}^{(k)} \in \{0, 1\} \quad \forall i \in I, \forall k \in K, \forall p = 0, 1, \dots, TS \quad (30-31)$$

$$\rho_p^{(k)} \in \mathbb{R}, L_{Xp}^{(k)} \in \mathbb{R}, L_{Yp}^{(k)} \in \mathbb{R}, q_p^{(k)} \in [0, 1] \quad \forall k \in K, \forall p = 0, 1, \dots, TS \quad (32-35)$$

All the variables of the TAMEXCLP formulation are derived from the sequence of tasks

scheduled per mobile resource constrained by the TOP-MTW sub-problem. Thus, the TAMEX-CLP is structured in the logical sequence in which these variables should be derived. An overview of this sequence is given as follows.

For calculating the average expected demand covered across time, it is first necessary to track the guards' locations at discrete points in time given by constraints (25). Next, given that the MRAP conditions for a valid strategy defined in Definition 1 imply that the actual area covered per mobile resource is time dependent, the coverage radius must be determined with constraints (26). Having establishing the coverage radius, it is then possible to identify whether a certain mobile resource covers a certain site as defined in constraints (27). When the sites covered are identified, it is then possible to calculate the potential workload due to unforeseen tasks. The workload in turn is a basic input to calculate in constraints (28) the probability can be estimated that a guard will be busy to service an unforeseen task because he is servicing another unforeseen task. Finally, combining the information of constraints (29) and (30), it is then possible to determine $q_{ip}^{*(k)}$ which is the basic input of the objective function.

Note the difference between $q_p^{(k)}$ and $q_{ip}^{*(k)}$ in that the former refers to the availability of the mobile resource to serve unforeseen tasks, while the latter refers to the incapability of a mobile resource to service an unforeseen task in a specific site. In other words, a guard may be available to service unforeseen tasks but may not be able to service an unforeseen task at a given site because it can not be reached on time.

3.1.2 Objective function (24)

The normalized objective function of the TAMEXCLP (24) is based on that of the Maximum Expected Coverage Location Problem (MEXCLP) (Daskin, 1983). In effect, the TAMEXCLP objective function samples repeatedly at each snapshot a modified MEXCLP measure to then be averaged over time. In each snapshot, the objective function calculates the probability that at least one mobile resource is able to serve an unforeseen task generated at a given site and

then weighs such probability by ω_i and P_i to calculate the expected demand covered.

To calculate the probability that at least one mobile resource would be able to serve an unforeseen task generated at site i on time, we consider the complimentary event, namely, the probability that all of the mobile resources would be unable to serve such unforeseen task. Assuming that the probability of a mobile resource k to be unable to serve an unforeseen task (i.e. $q_{ip}^{*(k)}$) is independent from that of other mobile resources, the probability that the complimentary event occurs is simply $1 - \prod_{k=1}^{|K|} q_{ip}^{*(k)}$.

3.1.3 Determination of $X_p^{(k)}, Y_p^{(k)}$, constraints (25)

For tracking the mobile resources positions over time given by the tuples $(X_p^{(k)}, Y_p^{(k)})$, we distinguish between two situations in which a mobile resource can be in (see Figure 3).

1. The mobile resource is at a site, hence the coordinates of the mobile resource correspond to that of the site's.
2. The mobile resource is transferring to serve a task at a site, hence the coordinates are linearly interpolated between origin and destination.

Hence, for calculating $X_p^{(k)}$ we obtain the following relationship for $\forall p = 0, 1, \dots, TS, \forall k \in K$ as follows.

$$X_p^{(k)} = \begin{cases} X_m & \text{if } \begin{aligned} &\exists m, n \in TK_j: n \neq m, \\ &x_{mn}^{(k)} = 1, a_n^{(k)} \leq t_p \leq s_n^{(k)} + \Delta_n \end{aligned} \\ X_m + \frac{(t_p - s_m^{(k)} - \Delta_m)(X_n - X_m)}{D(m, n)} & \text{if } \begin{aligned} &\exists m, n \in TK_j: n \neq m, \\ &x_{mn}^{(k)} = 1, s_m^{(k)} + \Delta_m < t_p \leq a_n, \end{aligned} \end{cases} \quad (36)$$

where :

- X_m, X_n : X coordinates of locations $L(m)$ and $L(n)$, and
 $a_n^{(k)}$: is the arrival time of mobile resource k for servicing task n such that
: $a_n^{(k)} = s_m^{(k)} + \Delta_m + D(m, n)$

The calculation of the Y -coordinate $Y_p^{(k)}$ is analogous to that of the X -coordinate $X_p^{(k)}$.

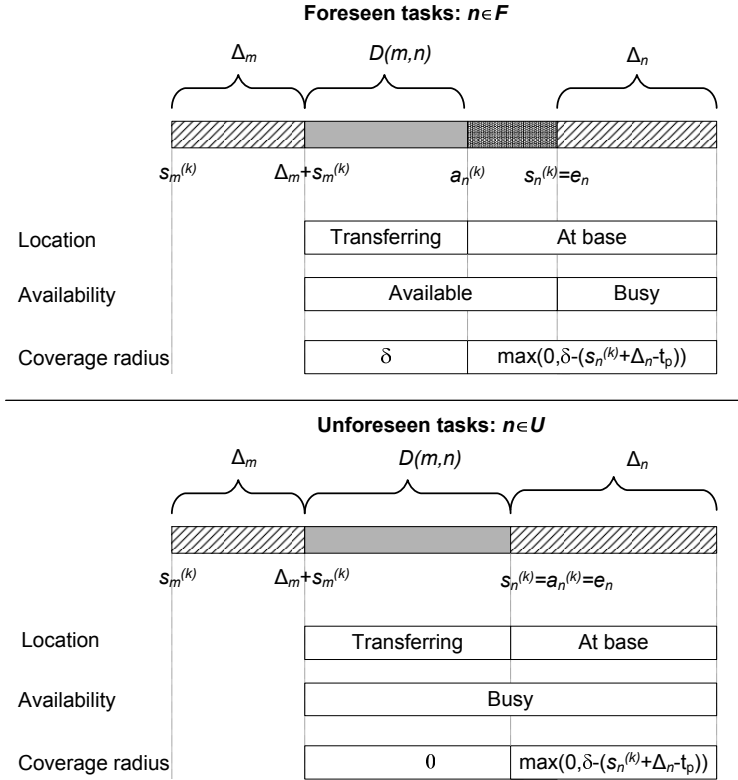


Figure 3: The status and coverage radius of mobile resources according to a time line.

3.2 Determination of $\rho_p^{(k)}$, constraints (26)

Next, to determine the coverage radius $\rho_p^{(k)}$, we identify three possible situations that a mobile resource k at an instant ep may encounter, in accordance with the preemptive rules of the MRAP (see Figure 4):

1. Transferring or waiting to serve a foreseen task $n \in F$: In this case, the mobile resource

can be immediately re-scheduled to serve another task and the coverage radius corresponds to that of the standard response time.

2. Servicing a task (i.e. either foreseen or unforeseen): In this case, the mobile resource can not serve any unforeseen task immediately but after it finishes servicing the task and hence has a reduced coverage area.
3. Transferring to serve an unforeseen task $n \in U$. If we assume for simplicity that $\Delta_n \geq \delta$, then the mobile resource has an effective coverage of zero as it will never finish servicing the unforeseen task early enough to be able to serve the new unforeseen task on time.

Consolidating the three cases described above, we have that $\rho_p^{(k)}$ is defined as follows for $\forall p = 0, 1, \dots, TS, \forall k \in K$.

$$\rho_p^{(k)} = \begin{cases} \delta & \text{if } \begin{aligned} &\exists m, n \in TK_j : n \neq m, x_{mn}^{(k)} = 1, \\ &n \in F, s_m^{(k)} + \Delta_m \leq t_p \leq s_n^{(k)} \end{aligned} \\ \max(0, \delta - (s_n^{(k)} + \Delta_n - t_p)) & \text{if } \begin{aligned} &\exists n \in TK_j : y_n^{(k)} = 1, \\ &s_n^{(k)} \leq t_p \leq s_n^{(k)} + \Delta_n \end{aligned} \\ 0 & \text{if } \begin{aligned} &\exists m, n \in TK_j : n \neq m, x_{mn}^{(k)} = 1, \\ &n \in U, s_m^{(k)} + \Delta_m \leq t_p \leq s_n^{(k)} \end{aligned} \end{cases} \quad (37)$$

3.2.1 Determination of $g_{ip}^{(k)}$, constraints (27)

For determining if a mobile resource is actually covering a site at a given time step (i.e. $g_{ip}^{(k)} = 1$), it is only required to compare the transfer time from the current mobile resource's location to the site's location, with the mobile resource's coverage radius $\rho_p^{(k)}$ of the circle that covers sites for servicing unforeseen tasks on time as seen in Figure 5. If we denote (X_i, Y_i) as the pair

of coordinates of site i then $D((X_p^{(k)}, Y_p^{(k)}), (X_i, Y_i))$ is the transfer time between the current location of the mobile resource and the site of interest i . The calculation of $g_{ip}^{(k)}$ is then as follows for $\forall i \in I, \forall p = 0, 1, \dots, TS, k \in K$.

$$g_{ip}^{(k)} = \begin{cases} 1 & \text{if } \rho_p^{(k)} \geq D((X_p^{(k)}, Y_p^{(k)}), (X_i, Y_i)) \\ 0 & \text{if } \rho_p^{(k)} < D((X_p^{(k)}, Y_p^{(k)}), (X_i, Y_i)) \end{cases} \quad (38)$$

3.2.2 Determination of $q_p^{(k)}$, constraints (28)

For calculating the busy fraction, $q_p^{(k)}$, we observe the definition of Revelle (1989) who states that the busy fraction is the ratio between the time the mobile resource is busy and the shift length available to serve tasks. For the purpose of the MRAP, the busy fraction is the probability that a mobile resource would be unavailable to serve an unforeseen task because that mobile resource is servicing another new unforeseen task. This means that we do not consider the possibility of being busy by servicing a task that is already known to the dispatcher, as in such a case, the decreased availability of the mobile resource is already factored by a reduced coverage radius $\rho_p^{(k)}$. It is important to note that effectively, the busy fraction is adjusting the availability of mobile resources to be able to serve unforeseen tasks due to time-overlapping unforeseen requests in the interval (t_p, t_{p+1}) .

For calculating the busy fraction,, we first calculate the expected workload that an unforeseen task at a site i would generate for a mobile resource k . The total expected workload generated by an unforeseen task at a given site i for a shift is simply the workload generated if the unforeseen task occurs multiplied by λP_i , the expected number of unforeseen tasks generated in such a site during the interval $TT - \mu$. In the case the unforeseen task is generated, then the implied workload is the required time to transfer from the current position of the mobile resource (i.e. $(X_p^{(k)}, Y_p^{(k)})$) to the site where the task is requested plus the time taken to service such a task.

However, with respect to the specific mobile resource, such workload is only really effective if the mobile resource is actually covering such a site, otherwise the workload due to unforeseen

tasks at such site is zero (hence we multiply the expression by the $g_{ip}^{(k)}$ indicator). Moreover, since more than one mobile resource may be covering such a site, then it is reasonable to assume that on average the workload would be distributed evenly among all the mobile resources that cover such a site. Thus, the expected workload generated by a site for mobile resource k at time t_p for a shift is given by:

$$W_{ip}^{(k)} = \frac{\lambda P_i \left(\mu + D((X_p^{(k)}, Y_p^{(k)}), (X_i, Y_i)) \right) g_{ip}^{(k)}}{\sum_{k' \in K} g_{ip}^{(k')}} \quad (39)$$

The total expected workload for a given mobile resource k is simply obtained by summing $W_{ip}^{(k)}$ over all the sites. The busy fraction for a given shift would then only be such total workload divided by the $[TT - \mu]$ time span. Notice that the time span is irrelevant for optimization purposes as it affects the numerator and the denominator of the expression alike. The final expression for $q_p^{(k)}$ is then as follows²:

$$\begin{aligned} q_p^{(k)} &= \frac{1}{TT - \Delta_E} \sum_{i \in I} W_{ip}^{(k)} \\ &= \frac{\lambda}{TT - \Delta_E} \sum_{i \in I} \left[\frac{P_i \left(\mu + D((X_p^{(k)}, Y_p^{(k)}), (X_i, Y_i)) \right) g_{ip}^{(k)}}{\sum_{k' \in K} g_{ip}^{(k')}} \right] \end{aligned} \quad (40)$$

3.2.3 Determination of $q_{ip}^{*(k)}$, constraints (29)

The probability that at time point p a guard k would not be able to serve an unforeseen task generated in the interval (t_p, t_{p+1}) at site i is given by $q_{ip}^{*(k)}$. Such probability is the same as the busy fraction, with the difference that if site i is not covered then the probability should be set to 1 as it is certain that an unforeseen task in site i would not be able to be serviced by guard k . Hence, the expression of $q_{ip}^{*(k)}$ is as follows for $\forall i \in I, \forall p = 0, 1, \dots, FP, \forall k \in K$.

$$q_{ip}^{*(k)} = 1 - (1 - q_p^{(k)})g_{ip}^{(k)} \quad (41)$$

²Note that we assume that unforeseen tasks are rare events and that the sum of the expected workload $\sum_{i \in I} \omega_{ip}^{(k)}$ should not exceed $TT - \mu$ so that $q_p^{(k)} \leq 1$.

Table 3: Example for calculating coverage for Solution Set B at t=7.50 hrs.

λP_i	=	0.167	(for all i)		δ	=	1.500
W_i	=	16.777	(for all i)		TT	=	12.000
Δ_E	=	1.200			$TT - \mu$	=	10.500

Site i	$D_1(\cdot)$	$D_2(\cdot)$	$g_{ip}^{(1)}$	$g_{ip}^{(2)}$	$\sum_{k' \in K} g_{ip}^{(k')}$	$\omega_{ip}^{(1)}$	$\omega_{ip}^{(2)}$	$q_{ip}^{*(1)}$	$q_{ip}^{*(2)}$	Cov $_i$
a	1.684	2.160	0	0	0	0.000	0.000	1.000	1.000	0.000
b	0.500	1.004	1	1	2	0.142	0.184	0.029	0.063	16.637
c	1.769	1.914	0	0	0	0.000	0.000	1.000	1.000	0.000
d	0.737	0.086	1	1	2	0.161	0.107	0.029	0.063	16.637
e	2.328	2.160	0	0	0	0.000	0.000	1.000	1.000	0.000
f	1.684	1.004	0	1	1	0.000	0.367	1.000	0.063	15.622
					\sum	0.303	0.658			Cov = 48.895
					$q_p^{(1)}$	=	0.029			
					$q_p^{(2)}$	=	0.063			

Using the presented example in Section 2.5, and the procedure hereby described in Table 3 we illustrate for Example 1 the calculation for coverage of Solution Set B at $t_p = 7.50$ hrs, $p = 15$ the as shown. Note that $D_k(\cdot)$ is an abbreviation for $D((X_p^{(k)}, Y_p^{(k)}), (X_i, Y_i))$ and that Cov_i is the coverage for a site and an abbreviation for $P_i \cdot \hat{w}_i \left[1 - \prod_{k=1}^{|K|} q_{ip}^{*(k)} \right]$.

4 Solving the s-TOP-MTW

To solve moderately sized s-TOP-MTW problems³ each time a new problem instance is generated we first generate a number of feasible solutions for the s-TOP-MTW (a sequence of tasks to be performed by each mobile resource) calculating the associated TOP-MTW scores for each of them. Next, we track the mobile resources' paths of each solution at regularly spaced discrete points in time to then calculate the TAMEXCLP associated objective function of each proposed solution (see Appendix Algorithm 1 for details). Finally, we select the best path according to the s-TOP-MTW objective function for a given value of α .

³"Moderately" in this context means a problem with less than 4 guards and 15 tasks to be solved in less than 3 minutes on a 1.6 Ghz computer.

The main underlying idea of the presented algorithmic structure is to take advantage of the fact that the TOP-MTW is more constrained than the TAMEXCLP, where a solution for the TOP-MTW satisfies the TAMEXCLP, but a solution for the TOP-MTW does not necessarily satisfies the TAMEXCLP. The procedure for generating feasible solutions is then described as follows.

For solving the TOP-MTW we first have to recognize that the computational complexity of the TOP-MTW is at least as difficult as the OP which is already NP-hard (Laporte, 1990). Hence, most of the approaches available in the literature to solve the OP and its derivatives are based on metaheuristics (see Chao et al., 1996; Tasgetiren, 2002; Liang and Smith, 2006). Furthermore, since the relationship between the TOP-MTW and the TAMEXCLP is poorly understood, it is reasonable to choose for an algorithm that it is fast and simple; that can provide "good" solutions for the TOP-MTW and at the same time a diversity of alternative solutions with a variety of TAMEXCLP reliability measures to choose from.

In this context, we construct "good" feasible plans by extending a stochastic-based algorithm proposed by Tsiligirides for the Orienteering Problem (OP) (Tsiligirides, 1984). Tsiligirides' S-algorithm is based on devising a reasonable measure of the desirability to append a task to an existing route. The desirability measure is the ratio between the reward collected and the extra distance needed to serve such a task. The S-algorithm constructs a sequence of tasks by adding one task at a time and making the probability of adding a certain task in a route dependent on such desirability ratio. Therefore, the algorithm provides a certain directionality in the search for good "solutions" using a greedy approach, but also allows for exploring neighborhood areas in the solution space by introducing stochasticity.

In this paper, we adapt Tsiligirides' S-algorithm for solving the TOP-MTW. Instead of the extra distance required to service a candidate task, we use in the denominator of the fitness function the extra time needed to serve the candidate task. In this way, the time required for

servicing a task and also the time required to wait at a site for an earliest starting time e_n to occur are considered in the desirability measure. Hence, the further in the future the earliest allowable starting time is, the lower the desirability for appending such a task to the current schedule is.

The fitness function is then defined as follows.

$$f(n) = \left[\frac{w_n}{\max(Last_k + D((L(k), L(n)) + \Delta_n, e_n) - Last_k)} \right]^3 \quad (42)$$

where:

- $Last_k$: is the finishing time of last scheduled task by mobile resource $k \in K$,
- $D(a, b)$: is the travel time between two pairs of coordinates a and b ,
- $L(k)$: is the Euclidian coordinates where mobile resource k is located,
- $L(n)$: is the Euclidian coordinates where task n is located,

A power of 3 is used in the fitness function in order to amplify the differences between ratios, and bias further the probability of selecting tasks with higher desirability ratios.

Another extension made to Tsiligrades' S-algorithm is that of dealing with more than one mobile resource. In Chao et al. (1996), Tsiligrades' S-algorithm is extended by scheduling the mobile resources either sequentially or concurrently. The sequential construction of routes involves constructing first a route for the first mobile resource, then for the second and so on. In the concurrent approach, each mobile resource takes turns to add tasks to their routes. We propose to randomize the selection of a mobile resource turn⁴ so as to further increase solution variety.

The logic of the extended algorithm is simple. First it is checked whether a task n can be feasibly added to the route of a mobile resource k . Such feasibility depends on the possibility of meeting the task deadline on time (i.e. start servicing the task before the latest allowable starting time) and to have sufficient time to return back to the base before the shift ends. Next,

⁴Each mobile resource has an equal chance of being selected.

the fitness functions $f(n)$ are computed and then ranked. To limit the running time and similarly to Tsiligirides (1984) only four tasks are considered for inclusion in the mobile resources' route. These tasks are referred to as members of the elite set, ES . Finally, task n^* is selected based on the following probability.

$$P(n^*) = \frac{f(n^*)}{\sum_{n \in ES} f(n)} \quad (43)$$

5 Experimental study

In the previous section, we suggested the possibility of embedding coverage considerations in devising a strategy for solving the MRAP by integrating in an affine manner two different objectives: the TOP-MTW and the TAMEXCLP. However, it is not obvious which α weight to use to integrate both objectives. Furthermore, even a more basic question is whether the two objectives are in conflict. If the two objectives are not in conflict then it is possible to solve in a hierarchical way the TOP-MTW and the TAMEXCLP sub-problems. Otherwise, the suitability for selecting the value of α should be investigated.

In this section, we report the results from two experiments. In Experiment 1, we identify non-dominated solutions of the multi-objective s-TOP-MTW and observe if there is effectively a trade-off between the objectives of both sub-problems. In Experiment 2, we study the suitability of selecting several variables of α in maximizing the fulfilment yield by conducting a simulation study with on-line (re-)planning. This experiment is designed to study whether embedding coverage considerations and measuring the capacity of routes to respond on time to the arrival of new unforeseen tasks is effective in increasing the fulfilment yield beyond solving only the TOP-MTW routing model.

Both experiments have a common setting based on the same basic layout of sites and time-windows for foreseen tasks. Some important design considerations common to the two experiments include the following:

- The purpose of the experiments is to explore the possibility of embedding coverage considerations in the strategies devised for the MRAP. Hence, they are kept computationally simple. Only 10 sites with 1 foreseen task each are considered, and two mobile resources available.
- The time windows of the tasks (e_n, l_n) have been generated randomly such that $e_n = U(0, 0.7TT)$, and $l_n = \min(e_n + U(0, 0.4TT), TT)$, where U stands for the uniform distribution. These parameters are selected in order to have wide enough time windows that allow several feasible solutions s .
- The expected total weight of unforeseen tasks was made equal to the total weight of foreseen tasks so that intuitively the routing of foreseen tasks should be as important as servicing unforeseen tasks, i.e. $\lambda w_U = \sum_{n \in F} w_n$.
- The probability of occurrence of unforeseen tasks is not homogeneous, highlighting the need for coverage in areas of "high risk".
- The mean number of unforeseen tasks per shift, λ is fixed low as is normally the case for emergency situations: only 2 unforeseen tasks per shift compared to 10 foreseen tasks. Moreover, low values of λ are consistent with applying the TAMEXCLP that assumes that the routes remain unchanged for the calculation of the routes' reliability measure

5.1 Experiment 1: Trade-offs between TOP-MTW and TAMEXCLP

In this experiment we were concerned with identifying a trade-off curve of solutions between the TOP-MTW objective function and the TAMEXCLP objective function to assess the need for solving both sub-models hierarchically or in an integrated way. The experiment is considered off-line as no unforeseen tasks are released and the optimization procedure is applied only once. In the experiment, we generated 8000 solutions (some repeated) and measured the TOP-MTW

and the TAMEXCLP scores. Note that of the 8,000 solutions only some will lie in the efficient frontier. Figure 4, shows a plot of the solutions found.

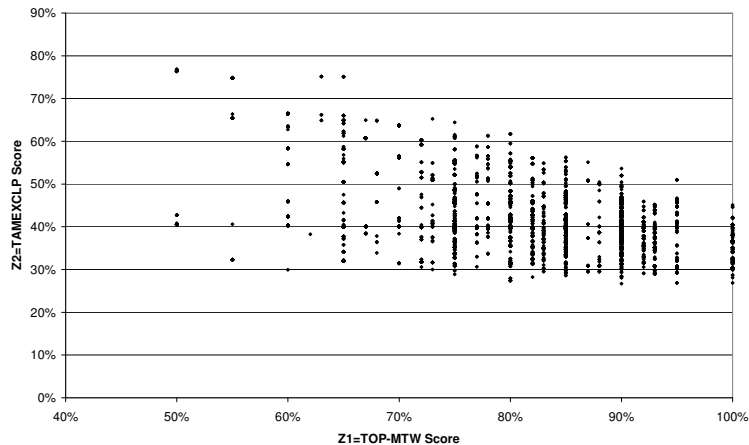


Figure 4: Off-line solutions found for the s-TOP-MTW.

From the chart it can be observed that there exists alternative optimum solutions for the TOP-MTW that provide different objective values of TAMEXCLP. For this particular case it means that solving hierarchically the s-TOP-MTW by first optimizing the TOP-MTW and then selecting the best TAMEXCLP solution we can obtain an increase of up to 19% in the TAMEXCLP reliability measure for $Z_{TOP-MTW} = 100\%$. This result is significant as it implies that a significant increase in routes reliability is possible by simply selecting the best alternative optimum of the TOP-MTW.

Next, by sorting the solutions we may eliminate the dominated solutions and obtain an efficient frontier as Figure 5 shows. Note that the efficient frontier is actually discontinuous due to the dichotomous definition of coverage (i.e. a site is covered if the coverage radius is greater than the distance between the mobile resource and the site). The results of the numerical experiment suggest that a trade-off does exist where an increase of 45% in TOP-MTW score implies a reduction of 35% of TAMEXCLP scores. Several other scenarios were tested yielding similar results. However, it must be stated that the results were sensitive to changes in the standard response time δ . For example, if δ was chosen with a high enough value a trade-off

is no longer found as the response-time is so long that it is always possible to start to service unforeseen tasks within the standard response time (assuming also that the shift length allows for accomplishing all these tasks).




Figure 5: Efficient frontier of TOP-MTW and TAMEXCLP subproblems.

5.2 Experiment 2: On-line MRAP simulation

For the on-line MRAP simulation we generated 75 different series (see Figure 6 for visualization of one simulation run) of unforeseen tasks for an equal number of working shifts and tested different α integration factors (with every re-schedule selecting a solution from 8000 runs as before). As explained in the previous section, the tasks were generated up-front so as to be able to compare for each run how different solution strategies perform with the MRAP (i.e. with different α integration factors).

The α factors were chosen in steps of 0.25 from 0 to 1. An $\alpha = 0$ implies that a simple TOP-MTW problem is solved disregarding the TAMEXCLP problem altogether, while an $\alpha = 1$ implies that the plan with the best TAMEXCLP score is selected. However, as the basic algorithm is based on solving the TOP-MTW it can not be said that $\alpha = 1$ solves the TAMEXCLP to optimality. In addition to the range of α integration factors tested, we introduce two more α values to reflect a hierarchical way of selecting a solution for the MRAP. In this way, an

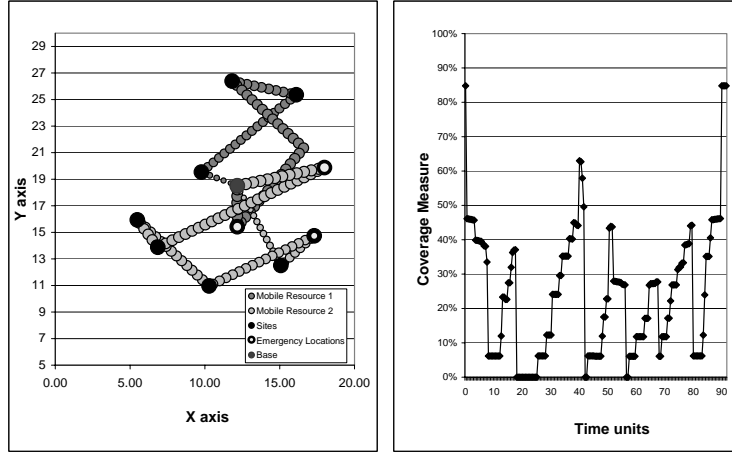


Figure 6: Results of a simulation run: Routes (left) and coverage (right).

$\alpha = 0.1 * 10^{-5}$ has such a low α that in effect forces to select from all the alternative optima for the TOP-MTW the solution that provides the best TAMEXCLP measure.

In addition, we also solve the complete information sub-problem TOP-MTW for each vector of unforeseen tasks in hindsight (a-posteriori) where all the generated unforeseen tasks are known at the start of the shift. Solving each simulation run in hindsight allows us to compare each solution with the highest achievable fulfilment yield.

The results of Experiment 2 are shown in Table 4. Using the TAMEXCLP sub-problem can positively influence the expected % of weighed tasks accomplished in due time. A total improvement of 2.94% of fulfilment yield γ over solving a "plain" TOP-MTW is achieved by solving the problem hierarchically with $\alpha = 10^{-6}$. A low value of $\alpha = 0.25$, also yields an improvement (i.e. 2.64%). To verify the significance of the improvement we conduct a two-sided t-test suggested by Law and Kelton (2000) to compare two different simulation alternatives. The test consists of constructing a confidence interval with an auxiliary variable that is the difference between the fulfilment yields of two alternatives in a same simulation run: $Z_j = \gamma_{1j} - \gamma_{2j}$. Comparing the alternative of solving hierarchically the s-TOP-MTW (i.e. solving with $\alpha = 10^{-6}$) with solving solely the TOP-MTW (i.e. solving with $\alpha = 0$), we find that with a significance of $p < 0.050$ both alternatives are different and thus, conclude with a 95% confidence that solving

Table 4: Experiment 2 Results of fulfillment yields γ for different values of α

α	0	10^{-6}	0.25	0.50	0.75	$1-(10^{-6})$	1	Hindsight
Mean Fulfilment yield $\gamma\%$	79.83	82.78	82.48	78.84	54.47	36.06	36.22	93.24
Standard Deviation %	18.77	18.11	18.91	18.79	16.73	15.20	15.23	10.12

the s-TOP-MTW hierarchically outperforms solving the TOP-MTW.

The significance of the improvement obtained with $\alpha = 10^{-6}$ and $\alpha = 0.25$ should not be overlooked, as a maximum improvement of 13.41% is attainable (i.e. considering the hindsight result as threshold and solving the TOP-MTW as reference). Nonetheless, it is worth noticing that these positive results occur only for low values of α , indicating that reliability is of second importance after that of greedily maximizing the weighed amount of known tasks accomplished in time.

Moreover, for high values of α (i.e. >0.5) the obtained objective values are low with a decrement in the fulfillment yield of up to 43% compared to solving a "pure" TOP-MTW. Such results can be explained by the fact that although high values of α yield a high percentage of coverage, servicing known tasks including released unforeseen tasks is not considered as important as good coverage and thus despite the good positioning of mobile resources to serve unforeseen tasks it is possible that it doesn't take advantage of such positioning to serve known tasks.

For Experiment 2 it was expected that the best value for factor α be $\alpha = 0.5$ as the expected total weight of unforeseen tasks was made equal to the total weight of foreseen tasks so that intuitively the routing of foreseen tasks should be as important as servicing unforeseen tasks. However, the best value was obtained at $\alpha = 10^{-6}$. The underlying reason for this result may be that the reliability measures taken at time steps are calculated assuming that the routes would remain unaffected (so that the mobile resources' routes remain known) which is not strictly true as unforeseen tasks may arrive that change the routes followed by the mobile resources.

6 Conclusions

The MRAP applied to service engineering is a new approach in a focus also on service quality rather than on costs only. Therefore, instrumental, in recognizing the specific needs of clients through the consideration of specific service deadlines and response times. In addition to only information about known tasks, we have also used in this paper incomplete information about possible unforeseen tasks. Moreover, it was shown that by tracking the position and state of each mobile resource it is possible to devise a measure reliability of constructed routes.

We have observed that the improvements in reliability can be made by choosing among alternative equivalent optima of the TOP-MTW problem or at the expense of the performance of the known tasks. Through simulation experiments we observed that moderate considerations of the routes' reliability measure (i.e. the TAMEXCLP objective value), improvements can be achieved making it possible to serve more tasks on time. Nonetheless, further investigation is required to quantify the optimum relative weight that the TAMEXCLP component should have for optimum results.

The MRAP can be further extended by adding the possibilities of delays at sites without immediate re-allocation to serve other tasks and even by being able to modify the exact paths between sites. Increased control on the mobile resources path may add new possibilities to obtain higher route reliabilities. Other possible extensions for the MRAP include the use of individual response times, the recognition of different travelling speeds, stochastic transferring times and the possibility of mobile resource breakdowns.

7 Appendix

Algorithm 1 Calculation of TAMEXCLP

```
for  $p = 1$  to  $TS$  do
  for every mobile resource  $k \in K$  do
    Calculate coverage radius:  $\rho_p^{(k)}$ 
    for every site  $i \in I$  do
      Verify if site  $i \in I$  is covered by mobile resource  $k \in K$ 
      Update numbers coverers of site  $i$ 
    end for
  end for
  for every mobile resource  $k \in K$  do
    for every site  $i \in I$  do
      Calculate workload:  $W_{ip}^{(k)}$ 
    end for
    Calculate busy fraction;  $q_p^{(k)}$ 
  end for
  for every site  $i \in I$  do
    Calculate specific busy fraction per site  $q_{ip}^{(k)}$ 
    Calculate MEXCLP objective value for site  $i$  at step  $p$ 
  end for
  Calculate average MEXCLP objective for step  $p$ 
end for
Calculate TAMEXCLP objective value for feasible solution  $\Gamma_j^{(f)}$ 
```

References

- Chao, I., Golden, B. and Wasil, E. (1996) The team orienteering problem. *European Journal of the Operational Research Society*, 88(3), 464-474.
- Daskin, M., (1983) The maximal expected coverage location model: Formulation, properties and heuristic solution. *International Journal of Production Research*, 17(1): 48-70.
- Feillet, D., Dejax P. and Gendreau, M. (2005) Traveling salesman problems with profits, *Transportation Science*, 39(2), 188-205.
- Gao, M.. and Nguyen, M (2003) Modeling and prototype development of an electronic tourist guide. Master thesis, Centre for Industrial Management, Katholieke Universiteit Leuven, Herverlee, Belgium.
- Golden, B., Levy, L., and Vohra, R. (1987). The orienteering problem, *Naval Research Logistics*, 34(3):307-318.
- Golden, B., Assad, A., (1988). *Vehicle routing: methods and studies*. North Holland, Amsterdam, The Netherlands.

- Hart, C. (1988). The power of unconditional service guarantees. *Harvard Business Review*, 88(4):54-62.
- Johns, S. (1995). Heuristics to schedule service engineers within time windows. *Operational Research Society*, 46(3):339-346.
- Kantor, M. and Rosenwein, M. (1992) The orienteering problem with time windows. *Journal of the Operational Research Society*, 43:629–635.
- Laporte, G., and Martello, S. (1990) The selective traveling salesman problem, *European Journal of the Operational Research Society*, 177(1):1-21.
- Larco, J. Kaymak, U. and Dekker, R. (2006) Multi-objective scheduling for distributed security services. *Proceedings of the 2nd European Symposium on Nature Inspired Smart Information Systems 2006*, Puerto de la Cruz, Tenerife, Spain, 141-148.
- Larco, J., (2007). *The Mobile Reallocation Problem: Embedding coverage considerations in routing problems*. MPhil thesis, Erasmus Research Institute of Management, Erasmus University of Rotterdam, Rotterdam, The Netherlands.
- Liang, Y. and Smith, A. (2006) An ant colony approach to the orienteering problem. *Journal of the Chinese Institute of Industrial Engineers*, 23(5):403-414.
- Revelle, C. (1989) Review, extension and prediction in emergency service siting models. *European Journal of the Operational Research Society*, 40(1):58-69.
- Saydam, C. and Atug C. (2003) Accurate estimation of expected coverage: revisited, *Socio-Economic Planning Sciences* 37(1):69-80.
- Steuer, R., (1986). *Multiple criteria optimization theory, computation and application*. John Willey and Sons, United States.
- Tasgetiren, M. (2002) A genetic algorithm with an adaptive penalty function for the orienteering problem, *Journal of Economic and Social Research*, 4(2):1-27.
- Tsiligrirides, T. (1984) Heuristic methods applied to orienteering, *Journal of the Operations Research Society*, 35(9):797-909.
- Weintraub, A., Aboud, J., Fernandez, C., Laporte, G. and Ramirez, E. (1999) An emergency vehicle dispatching system for an electric utility in Chile, *Journal of the Operational Research Society*. 50(7): 690-696.
- Zeitham, V.A., Parasaurman, A. and L., B., (1990) Delivering quality services: Balancing cus-

tomers' perceptions and expectations. *International Journal of Production Research*. 39(9):1865-1883.

Publications in the Report Series Research * in Management

ERIM Research Program: "Business Processes, Logistics and Information Systems"

2007

India: a Case of Fragile Wireless Service and Technology Adoption?

L-F Pau and J. Motiwalla

ERS-2007-011-LIS

<http://hdl.handle.net/1765/9043>

Some Comments on the Question Whether Co-occurrence Data Should Be Normalized

Ludo Waltman and Nees Jan van Eck

ERS-2007-017-LIS

<http://hdl.handle.net/1765/9401>

Extended Producer Responsibility in the Aviation Sector

Marisa P. de Brito, Erwin A. van der Laan and Brijan D. Irion

ERS-2007-025-LIS

<http://hdl.handle.net/1765/10068>

Logistics Information and Knowledge Management Issues in Humanitarian Aid Organizations

Erwin A. van der Laan, Marisa P. de Brito and S. Vermaesen

ERS-2007-026-LIS

<http://hdl.handle.net/1765/10071>

Bibliometric Mapping of the Computational Intelligence Field

Nees Jan van Eck and Ludo Waltman

ERS-2007-027-LIS

<http://hdl.handle.net/1765/10073>

Approximating the Randomized Hitting Time Distribution of a Non-stationary Gamma Process

J.B.G. Frenk and R.P. Nicolai

ERS-2007-031-LIS

<http://hdl.handle.net/1765/10149>

Application of a General Risk Management Model to Portfolio Optimization Problems with Elliptical Distributed Returns for Risk Neutral and Risk Averse Decision Makers

Bahar Kaynar, S. Ilker Birbil and J.B.G. Frenk

ERS-2007-032-LIS

<http://hdl.handle.net/1765/10151>

Optimal Zone Boundaries for Two-class-based Compact 3D AS/RS

Yugang Yu and M.B.M. de Koster

ERS-2007-034-LIS

<http://hdl.handle.net/1765/10180>

Portfolios of Exchange Relationships: An Empirical Investigation of an Online Marketplace for IT Services

Uladzimir Radkevitch, Eric van Heck and Otto Koppius

ERS-2007-035-LIS

<http://hdl.handle.net/1765/10072>

From Closed-Loop to Sustainable Supply Chains: The WEEE case

J. Quariguasi Frota Neto, G. Walther, J. Bloemhof, J.A.E.E. van Nunen and T. Spengler

ERS-2007-036-LIS

<http://hdl.handle.net/1765/10176>

A Methodology for Assessing Eco-Efficiency in Logistics Networks

J. Quariguasi Frota Neto, G. Walther, J. Bloemhof, J.A.E.E van Nunen and T. Spengler

ERS-2007-037-LIS

<http://hdl.handle.net/1765/10177>

Strategic and Operational Management of Supplier Involvement in New Product Development: a Contingency Perspective

Ferrie E.A. van Echtelt, Finn Wynstra and Arjan J. van Weele

ERS-2007-040-LIS

<http://hdl.handle.net/1765/10456>

How Will Online Affiliate Marketing Networks Impact Search Engine Rankings?

David Janssen and Eric van Heck

ERS-2007-042-LIS

<http://hdl.handle.net/1765/10458>

Modelling and Optimizing Imperfect Maintenance of Coatings on Steel Structures

R.P. Nicolai, J.B.G. Frenk and R. Dekker

ERS-2007-043-LIS

<http://hdl.handle.net/1765/10455>

Human Knowledge Resources and Interorganizational Systems

Mohammed Ibrahim, Pieter Ribbers and Bert Bettonvil

ERS-2007-046-LIS

<http://hdl.handle.net/1765/10457>

Revenue Management and Demand Fulfilment: Matching Applications, Models, and Software

Rainer Quante, Herbert Meyr and Moritz Fleischmann

ERS-2007-050-LIS

<http://hdl.handle.net/1765/10464>

Mass Customization in Wireless Communication Services: Individual Service Bundles and Tariffs

Hong Chen and Louis-Francois Pau

ERS-2007-051-LIS

<http://hdl.handle.net/1765/10515>

Individual Tariffs for Mobile Services: Analysis of Operator Business and Risk Consequences

Hong Chen and Louis-Francois Pau

ERS-2007-052-LIS

<http://hdl.handle.net/1765/10516>

Individual Tariffs for Mobile Services: Theoretical Framework and a Computational Case in Mobile Music

Hong Chen and Louis-Francois Pau

ERS-2007-053-LIS

<http://hdl.handle.net/1765/10517>

Individual Tariffs for Mobile Communication Services

Hong Chen and Louis-Francois Pau

ERS-2007-054-LIS

<http://hdl.handle.net/1765/10518>

Is Management Interdisciplinary? The Evolution of Management as an Interdisciplinary Field of Research and Education in the Netherlands

Peter van Baalen and Luchien Karsten

ERS-2007-047-LIS

<http://hdl.handle.net/1765/10537>

Detecting and Forecasting Economic Regimes in Multi-Agent Automated Exchanges
Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta and Paul Schrater
ERS-2007-065-LIS
<http://hdl.handle.net/1765/10594>

Emergency Messaging to General Public via Public Wireless Networks
P.Simonsen and L-F Pau
ERS-2007-078-LIS
<http://hdl.handle.net/1765/10718>

Flexible Decision Control in an Autonomous Trading Agent
John Collins, Wolfgang Ketter and Maria Gini
ERS-2007-079-LIS
<http://hdl.handle.net/1765/10719>

Discovering the Dynamics of Smart Business Networks
L-F Pau
ERS-2007-081-LIS
<http://hdl.handle.net/1765/10732>

Performance Approximation and Design of Pick-and-Pass Order Picking Systems
Mengfei Yu and René de Koster
ERS-2007-082-LIS
<http://hdl.handle.net/1765/10733>

Distributed Services with Foreseen and Unforeseen Tasks: The Mobile Re-allocation Problem
Jose Antonio Larco, Rommert Dekker and Uzay Kaymak
ERS-2007-087-LIS
<http://hdl.handle.net/1765/10773>

* A complete overview of the ERIM Report Series Research in Management:
<https://ep.eur.nl/handle/1765/1>

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship