

Multi-Agent System Support for Scheduling Aircraft De-icing

Adriaan ter Mors

Almende B.V./Technische Universiteit Delft
adriaan@almende.com

Xiaoyu Mao

Almende B.V./Universiteit Maastricht
xiaoyu@almende.com

Nico Roos

MICC/IKAT, Universiteit Maastricht
roos@cs.unimaas.nl

Cees Witteveen

EWI, Technische Universiteit Delft
c.witteveen@tudelft.nl

Alfons Salden

Almende B.V
alfons@almende.com

ABSTRACT

Results from disaster research suggest that methods for coordination between individual emergency responders and organizations should recognize the independence and autonomy of these actors. These actor features are key factors in effective adaptation and improvisation of response to emergency situations which are inherently uncertain. Autonomy and adaptability are also well-known aspects of a multi-agent system (MAS). In this paper we present two MAS strategies that can effectively handle aircraft deicing incidents. These MAS strategies help improve to prevent and reduce e.g. airplane delays at deicing stations due to changing weather conditions or incidents at the station, where aircraft agents adopting pre-made plans that would act on behalf of aircraft pilots or companies, would only create havoc. Herein each agent using its own decision mechanism deliberates about the uncertainty in the problem domain and the preferences (or priorities) of the agents. Furthermore, taking both these issues into account each proposed MAS strategy outperforms a naive first-come, first-served coordination strategy. The simulation results help pilots and companies taking decisions with respect to the scheduling of the aircraft for deicing when unexpected incidents occur: they provide insights in the impacts and means for robust selection of incident-specific strategies on e.g. deicing station delays of (individual) aircraft.

Keywords

Incident management, multi-agent systems, coordination

INTRODUCTION: INCIDENT MANAGEMENT AND MULTI-AGENT SYSTEMS

On April 6, 2005, a large disaster response exercise, the so-called “Oefening Bonfire”) took place in the Netherlands. It involved many different organizations and independent actors, which had to respond to the threat of a terrorist attack. One of the main challenges appeared to be to establish effective coordination between these different actors and organizations. One would be inclined to try to organize the coordination of all parties from one central point; coordination would then be a top-down activity in which high-level decisions are taken on the basis of a complete overview of the situation. Indeed, the eight-o'clock evening news reported that coordination was organized according to the military model of Command and Control (C & C), in which a strong, hierarchical organization directs emergency responders.

However, analyses of past disasters show that the application of the C & C model can be more of an impediment than an enabler of good coordination between different organizations [16]. A tacit assumption of the C & C model is

that individual emergency responders and organizations, when left to their own devices, are unable to act appropriately in the unforeseen circumstances of the emergency; also it is assumed that at a higher level in the chain of command, people *do* know what to do [2, 3, 4]. Both these assumptions have been disproved in practice. Emergency responders and bystanders alike often show remarkable second resilience in the event of a disaster [12]. Second, the *situation awareness* required to make informed decisions at the highest level for emergency responders in the field, is often lacking. In fact, researchers have pointed out that creating this situation awareness for the emergency responders in the field is one of the main challenges of emergency response [5].

If emergency responders in the field as well as at the highest levels are to be left with sufficient autonomy to perform their tasks well, and utilize their improvisational skills, then information systems supporting them should also facilitate the autonomy of the individual actor in taking decision about taking actions. These information systems and consequently those actors can in turn be enabled and represented by *multi-agent systems*, respectively. In a multi-agent system, a set of autonomous agents representing actors each have their own set of tasks to sense, reason and act, but the agents have dependencies between one another which means that for successful task completion they need to communicate and interact. The reasons for agents to coordinate their actions can have a number of different origins. For example, agent coordination exists because the agents have a *common goal* (like saving people from a burning building), to which each of the actors has to make its own contribution. Another cause for agent coordination is the occurrence of a sequence of incidents asking for improvisation of response support to the actors through those agents. For the incident resolving power of a multi-agent system, however, it can be advantageous that each agent pursues its own situational context-specified local goals, without striving for any dictated common goal. For instance, in a traffic system, each agent on a free highway may subject itself to its own driving style, but at a crossing of different roads the agents need to coordinate their behaviour although probably only locally.

In multi-agent systems research, there are at least two main reasons to stress the importance of the autonomy of individual agents, as opposed to handing over all decision-making to a central authority. First of all, agents rarely or could have complete knowledge about the environment (especially in emergency response the environment is changing rapidly) or about the tasks and actions of other agents. Instead, information about the environment and the changes therein can often only be obtained locally, and updates on the intentions of other agents is acquired through direct communication with other (nearby, or related, or friendly, etc.) agents. Hence, deciding on a course of actions for one agent usually involves only a limited number of other agents. A second reason to maintain, augment and enhance individual autonomy is that agents and the actors they represent are *self-interested*, that is, they are only concerned about the goals of itself (or the human it represents), and it does not necessarily care about the welfare of other agents, in particular when aligning of those local (incident-specific) agent goals to those goals of 'distant' agents does not contribute to the individual agent fitness to the changing local environment or incidents. For example, my personal agent that performs bidding in eBay auctions gets me the best price for the items I want, without considering whether these prices are fair to the seller. Additionally, agents may be cautious to share strategically important plans and information with other agents, especially if commercial relations exist between them. In that case, agents not only wish to maintain decision-making authority, they would also use coordination mechanisms in which communication with other agents is minimized.

At first glance, the focus on self-interested agent behavior may seem needlessly restrictive for the emergency response domain, where there *is* a global goal, and (many) actors may be assumed to prefer this global goal above any personal interests. However, in situations where communication lines may break down and agents get disconnected, a coordination mechanism that performs well at a minimum of inter-agent communication and negotiation can prove invaluable. Hence, *decoupling* self-interested but situational context aware agents during task coordination can be a useful attribute in particular in emergency response situations. Once partly decoupled in their limited sub-problem domain, agents and actors have the freedom to act and improvise without the danger of interfering with the tasks of others existing in other sub-problem domains.

In this paper we present two multi-agent coordination mechanisms in which agents have to schedule the use of a shared resource. The domain that inspired our research is airport departure planning under severe weather conditions that necessitate the *de-icing* of aircraft (de-icing is the removal of snow or frost from wings and fuselage that would otherwise endanger safe take-off). This domain is similar to the emergency response domain in the sense that wintry conditions at airports can cause many unforeseen incidents that require the actors to quickly revise their initial plans. A difference may seem that one *aircraft agent* may have little interest in the punctual departure of other aircraft, i.e., the agents are self-interested. The scarce resource they are competing for is namely the *de-icing station*. Other airport actors or representative agents, such as the *de-icing coordinator* and/or its agent that manage the de-icing

station, however, are concerned about the throughput of all the aircraft agents. They can influence the behavior of the individual aircraft agents by just communicating the slots reserved by others. Furthermore, they can impose decommitment penalties to further adapt the local behavior of self-interested agents. Thus in a natural and intuitive way agent coordination is realized by situational context aware but self-interested agents.

The organization of this paper is as follows. In the next section, we will discuss in some greater detail the background of the airport de-icing problem, and we discuss the relevant multi-agent scheduling literature. We will then proceed to present a formal description of our problem, followed by a specification of two different coordination mechanisms: one which explicitly reasons about the uncertainty in the environment, another which primarily takes into account the relative *priorities* of the different aircraft agents, and how this priority affects the competition for the use of a scarce resource. In the section following the specification of the mechanisms we will describe the results of the experiments we conducted using these mechanisms. In the concluding section, we will reflect on the relevance of general multi-agent coordination mechanisms such as the ones presented in this paper for specific domains such as emergency response, and how they can be applied to support human decision making.

BACKGROUND AND RELATED WORK

Aircraft de-icing and anti-icing is required in winter time whenever frost, snow, and ice form on the wings and fuselage of an aircraft. Such a layer of frost or ice on aircraft surfaces influences the aircraft's aerodynamic properties which may cause a loss of lift that could result in a crash. An aircraft may have several alternatives to receive de-icing treatment: most commonly, it will taxi to one of the de-icing stations, located (hopefully) at strategic positions around the airport; second, it can also be deiced at the gate, in which case de-icing vehicle(s) will drive to the gate at which the aircraft is docked. Since the total de-icing capacity at an airport is usually limited, careful planning and scheduling of these resources is of crucial importance to efficient departure planning.

Different multi-agent systems (MAS) have been designed to tackle specific real-world scheduling problems, from patient scheduling in the hospital (cf. [11, 18]) to more general job shop scheduling problems (cf. [10]). In these works, different coordination mechanisms have been designed to coordinate agents' plans, from more cooperative agents (coalition formation) to more competitive agents (market-based mechanisms). Liu and Sycara [10] developed a MAS for job shop scheduling problems in which standard operating procedures are combined with a look-ahead coordination mechanism that should prevent 'decision myopia' on part of the agents. Using their approach, system performance is said to improve in tightly-coupled, real-time job-shop scheduling environments. However, their coordination mechanism is not appropriate for competitive, self-interested agents, which makes it an undesirable choice for coordination in a de-icing setting. Vermeulen et al. [18] developed a Pareto-optimal appointment exchanging algorithm in a patient scheduling problem. The objective is to improve upon the initial first come, first served schedule by letting patient agents exchange their slots. It is quite similar to the work of Paulussen [11] where the agent coordination mechanism is a dynamic schedule-repair affair that can be classified as an after-scheduling coordination mechanism. Although Vermeulen's slot swapping mechanism may be a valuable optimization tool in a dynamic schedule repair context, there is still a need for a coordination mechanism that finds a satisfying initial schedule.

In the following, we present and compare two coordination mechanisms for obtaining an initial schedule: the first is based on an auction for selling de-icing slots; the second is based on decommitment penalties. In previous research, decommitment has been primarily used to enable agents to explore new opportunities from the domain or from other agents [7, 13, 15]; an example might be a package-delivery agent that decommits the contract for one package so that it is able to accept a more profitable package to deliver [7]. Another use of decommitment penalties is that it allows agents to speculate on future events. For example, in [13], agents accept a contract and hope to find other agents to sub-contract parts of the contract. If it fails to do so, it has to pay a decommitment penalty if the agent is unable to meet the original contract. We propose that the concept of decommitment penalties can also be used to coordinate agents, by associating a penalty with the occurrence of an agent decommitting from a slot because it could not make the agreed time. In this sense, the decommitment mechanism curbs the greedy tendency of agents to grab the de-icing station resource as early as possible, before other agents have a chance to take it. Now, every agent gets that chance, but it has to suffer the consequences if it miscalculates its ability to make its slot.

FORMAL MODELLING

In this section we will present a formal model of the aircraft de-icing scheduling problem. The model presented below is a simplified version of the problem described in the previous section, as we leave out other airport planning problems as runway planning and gate allocation, and as a result certain constraints between different planning problems are not taken into account (the *holdover* time is an example of such a constraint: the maximum time allowed between the end of de-icing and the moment of take off). The following model has also been used as the basis of the experiments, which are described in the section “Improvisation support”.

Definition 2 (Aircraft De-icing Scheduling Problem). *The aircraft de-icing scheduling problem is a tuple $\langle A, D, c, \tau, p, p_d, l \rangle$ where*

- A is a set of n aircraft agents,
- D is a set of m de-icing station resources,
- $c : D \rightarrow N$ is a capacity function specifying the number of aircraft that can simultaneously be serviced at the de-icing station (i.e., the number of de-icing bays),
- $\tau : A \rightarrow N$ is a function associating an Target Off-Block Time t_{TOBT} with every agent,
- $p : A \rightarrow N$ is function that specifies the de-icing process duration for a certain aircraft,
- $p_d : N \times A \rightarrow N$ is a function indicating the probability that an incident may happen to an aircraft during its on-block time,
- $l : N \times A \rightarrow N$ is a function that assigns a cost to the delay of an aircraft

The target off-block time $\tau(a_i)$ defined in Definition 2 is the target de-icing start time for aircraft agent a_i , which is in fact the time when all other ground services for this agent are assumed to be finished. The incident probability $p_d(t, a_i)$ indicates the probability that an incident (e.g. passenger boarding delay, ground handling delay, etc.) will occur in the interval $[t, \tau(a_i)]$, i.e., the time during which the aircraft agent receives ground services at the gate. This occurrence of such an incident may delay the target off-block time, and therefore, rescheduling will be needed for an aircraft having a de-icing slot right after τ . The aircraft delay function $l : N \times A \rightarrow N$ maps delay to cost that is specific to each aircraft, reflecting the fact that different agents may hypothetically have different value systems.

A solution to an instance $\langle A, D, c, \tau, p, p_d, l \rangle$ of the de-icing problem is a multi-agent schedule given by the vector $S = \langle (d_1, I_1), \dots, (d_n, I_n) \rangle$ where (d_i, I_i) is a tuple in which d_i is the de-icing station assigned to agent a_i during interval I_i such that

$$I_i = [s_i, s_i + p(a_i)] \wedge s_i \geq \tau(a_i) \quad (1)$$

where s_i is the start time of deicing service for aircraft a_i . A feasible schedule satisfies the following resource constraints: at every point in time t , the de-icing resource utilization for each resource agent does not exceed the resource capacities. We have:

$$\forall t \forall d \in D \left\{ \left| \{ a_j \in A \mid (d, I_j) \in S \wedge t \in I_j \} \right| \leq c(d) \right. \quad (2)$$

Given a target off-block time for each aircraft agent a_i , two optimization criterions can be defined: the first is to minimize the overall delay cost of all aircraft, which reflects the global benefit (cf. formula (3).a); another criterion is to minimize the standard deviation of each aircraft’s delay cost (cf. formula (3).b), which reflects the fairness for each aircraft at the airport.

$$\min \sum_{a_i \in A} l_i = \sum_{a_i \in A} l(s_i - \tau(a_i), a_i) \quad (3.a)$$

$$\min \sqrt{\frac{1}{N-1} \sum_{i=1}^N (l_i - \bar{l})^2} \quad (3.b)$$

where $s_i - \tau(a_i)$ is the delay of aircraft a_i and the \bar{l} is the mean value of each delay cost l_i .

COORDINATION MECHANISMS

In this section we will describe two coordination mechanisms: coordination through decommitment penalties, and coordination using a Vickrey auction to sell de-icing slots to the highest bidder. In both cases, we will make the simplifying assumption that the de-icing duration (the function p in the previous section) is the same for all aircraft. Differences in de-icing duration are relevant when trying to create an efficient de-icing schedule; however, to simplify the experiments, we will assume that it is equal for all agents.

Another simplifying assumption that we will make for both coordination mechanisms is that agents cannot make use of slots that have been freed up because of an incident involving the aircraft agent that initially held that reservation. This assumption can partially be supported by the fact that aircraft only report their inability to make a slot at the very last minute, trying as they are to still make that slot. However, the availability of freed-up slots is relevant to both coordination mechanisms (indeed, we can even imagine the height of the decommitment penalty being dependent on the time at which decommitment occurs), so we consider the relaxation of this simplifying assumption to be (near-) future work.

A final simplifying assumption in both coordination settings is that we assume only a single de-icing station having a single de-icing bay. Again, having multiple de-icing stations makes the problem more interesting from a combinatorial optimization point of view, but it is not especially relevant to our investigation into the relative merits of auctioning and decommitment.

Taking Uncertainty into Account

We define an incident as an unexpected event that causes a disruption to one aircraft's plan regarding its de-icing activities. This means that the occurrence of an incident will force an aircraft to decommit from its current de-icing slot, and find a new one. Although the list of things that can go wrong in airport de-icing operations is too extensive to fit into an elegant model of agent reasoning with uncertainty, observations from real and simulated de-icing operations¹ lead us to conclude that many incidents are concentrated in the ground servicing of the aircraft. For example, if the apron in front of an aircraft accumulates too much snow, it becomes difficult for ground servicing vehicles like baggage carts to reach the aircraft, and push-back vehicles cannot find the grip required to tow an aircraft away from the gate. Hence, there is a great deal of uncertainty surrounding the Target Off-Block Times (TOBT) of an aircraft. If an aircraft agent is considering at time t_1 whether to reserve (or bid for) the de-icing slot starting at time t_2 , then two factors are relevant:

1. $\delta_1 = \tau(a) - t_1$: If δ_1 is large (and greater than 0, obviously), then there are many ground servicing tasks that still need to be performed, in which case the probability that something will delay the TOBT is considerable.
2. $\delta_2 = t_2 - \tau(a)$: If the reserved slot is very far away from the target off-block time, then a small delay during ground handling will not necessarily mean that the de-icing slot will be missed.

¹ For the past two years, the Dutch Aerospace Laboratory NLR has organized a large-scale simulation of Air Traffic Management in winter conditions at Amsterdam Airport Schiphol.

In this paper, we will base the probability-of-decommitment function only on the first factor. Hence, we assume that this function has the following form:

$$p_d(t_1, \tau(a)) = \begin{cases} 0 & \tau(a) < t_1 \\ \min(c, \alpha \cdot (\tau(a) - t_1)) & \text{otherwise} \end{cases} \quad (4)$$

where c is a constant between 0 and 1, and α is a constant greater than 0. Thus, if $p_d(t_1, \tau(a)) = 0.5$, then there is a 50% chance that the aircraft will not be ready in time to make its de-icing slot. In Formula (4), the constant c provides an upper bound on the probability of having to decommit, even if t_1 is an arbitrarily early time of requesting the de-icing slot. The constant α regulates the rate of incident-occurrence; if α very large, then even when requesting a de-icing slot close to the off-block time, there is a high probability of having to decommit.

Vickrey Auction Mechanism

Bidding for a (de-icing) slot is a straightforward way of distributing the scarce slots over the self-interested agents. The idea is that the agents with the highest need (or the biggest clout) get the best slots. In the airport scheduling case, the different preferences of the agents can be the result of, for example, the number of passengers aboard an aircraft, or the level-of-service that an airline wishes to maintain. If we assume that an agent may not sell on a slot to another agent in case it has to decommit, then the value of the slot is a private value. In private value auctions all auction types give the same result according to the revenue equivalence theorem. Therefore, we choose the Vickrey auction (a closed-bid, second price auction), because of its property that (rational) agents are encouraged to bid their true value. Hence, de-icing of aircraft should occur in the order of agents who are willing to pay the most. We will now describe how we set up the auction.

The de-icing station agent(s) will initiate a new auction when the start of the next free de-icing slot (starting at $t_{nextslot}$) is approaching, e.g. half an hour before $t_{nextslot}$. In each auction, the de-icing station agent auctions off the next available de-icing slot (alternative auction schemes like accepting bids for multiple de-icing slots are less appropriate given the dynamic nature of the setting). To determine its value for a certain slot, an agent should first check whether the start time of this slot $t_{nextslot}$ is greater than its target off-block time $\tau(a)$; if it is not, then the agent can't make use of this slot. In case $\tau(a) < t_{nextslot}$, an agent needs to estimate the delay it will incur by not obtaining the current slot. If there are m other aircraft in the system that also need de-icing, then the value of the $(m+1)$ -th slot is 0, because all competing agents can be served before this time. Then, the private value of the slot starting at $t_{nextslot}$ for agent a is:

$$pv(a) = l(t_{m+1} - t_{nextslot}, a) \quad (5)$$

However, not all agents in the system will be able to compete for the next slot, in case their estimated off-block times are greater than $t_{nextslot}$. Therefore, the number m may be smaller than the total number of agents (left) in the system. At the same, we cannot simply equate m to the number k of direct competitors —agents having $\tau < t_{nextslot}$ — because after the first k aircraft have been serviced, more agents will be ready for de-icing. Therefore, the set A_c of competitors for slot $t_{nextslot}$ is the smallest subset $A_c \subseteq A$ such that

$$A_c = \{a \in A \mid \tau(a) < (t_{nextslot} + |A_c| \cdot p(a))\}$$

Finding the set A_c can be done simply by extending, agent by agent, the set of k direct competitors for a slot. Note that in case of insufficient de-icing capacity, the set A_c will quickly equal the set of all agents that have not yet received de-icing. Formula 5 ignores the possibility that incidents can occur during other ground services that will cause an agent to miss its reserved slot. Taking into account the decommitment probability $p_d(t_{nextslot}, \tau(a))$, we get the following private value:

$$pv(a) = l(t_{m+1} - t_{nextslot}, a) \cdot (1 - p_d(t_{nextslot}, \tau(a))) \quad (6)$$

Equation 6 thus expresses that an agent's private value of a slot decreases as the probability increases that it will not make that slot. In the next section, we will introduce an alternative coordination mechanism that focuses not so much on agent preferences, but more on the effects of decommitting on the schedule of an agent.

Strategy 1 (Agent Bidding Strategy). When the auctioneer offers the next de-icing slot, an agent bids its private value for the slot, specified by Equation 6.

In the next section, we will introduce an alternative coordination mechanism that focuses not so much on agent preferences, but more on the effects of decommitting on the schedule of an agent.

Decommitment Penalty Mechanism

When an aircraft agent reserves a particular time slot at a resource such as a de-icing station, it will commit to turn up at that de-icing station at the specified time. If the aircraft fails to show up, it could have to pay a decommitment penalty to the de-icing station². Hence, with the introduction of decommitment penalties, agents have an incentive to reserve as late as possible; after all, if it reserves a slot five minutes from now, it will be fairly certain it can make this slot. On the other hand, the old incentive for scheduling as early as possible remains: agents will still want to schedule slots before other agents get them.

Our approach to coordination using decommitment penalties can be described as follows. An agent can reserve any free slot at a de-icing station, as the de-icing station will accept all requests. However, with a certain probability incidents occur that make it impossible for the aircraft to be present at the de-icing station at the agreed time. When such an event occurs, it must decommit and pay a decommitment penalty, which we assume to be an airport-wide constant δ . We assume that the agent can see when the first available slot is at all de-icing stations (we will refer to this time as $t_{nextslot}$). As an aircraft agent can see the earliest available slot — and we assume that it will want to schedule the earliest available slot³ — it has to solve the following decision problem: Do I reserve the currently available first slot, or do I reserve a slot at a later time? To answer this question, the agent a has to be able to evaluate his two different options. To judge whether the decision to reserve now has any merit, the agent needs to estimate the probability it will have to decommit from the slot. For this, we use Equation 4. Judging the option of reserving a slot at a later time is more difficult, as it needs to predict the availability of de-icing slots in the future. This availability depends on at least the following factors:

1. the passage of time; if a slot is available 11 minutes from now, then, if no-one else takes it, there will be a slot 10 minutes in the future one minute from now,
2. other agents reserve slots.

Trying to incorporate all these factors into a realistic model is a formidable task, especially as the slot-reserving behaviour of agents may be subject to their perception (and prediction) of other agents' behaviour. Therefore, we will make the following simplifying assumptions to make the task of foretelling the future a more tractable one:

- If an agent has to decommit from a slot, then it will have to find a new slot. Apart from the time lost in decommitment, we assume that the number of aircraft needing de-icing per hour stays constant throughout the day. Hence, an agent will not suddenly find itself in a departure peak, after having to decommit.

² It may seem unfair to put the burden of decommitment penalties completely on the airline companies, but this is just because we focus on the decision processes of the aircraft agents; certainly, they can reason about the incidents for which they bear responsibility. Other incidents, for which e.g. the airport authorities will be held accountable, are of lesser significance to the decommitment-penalty reasoning processes of the aircraft.

³ Although it would make sense to reserve a slot later than the target available slot in case an aircraft's TOBT is later than $t_{nextslot}$.

- The delay an agent suffers when it has to decommit will mainly depend on the time it decommits; here we assume a constant value for this delay.
- When an aircraft opts to postpone its decision to reserve a slot until the next round, and it turns out that another agent has reserved the previously earliest slot, then the new $t_{nextslot}$ is simply the old $t_{nextslot}$ plus the de-icing time which we assumed to be equal for all aircraft.

Armed with these simplifications, we can develop a strategy for an aircraft agent.

Strategy 2 (Decommitment Strategy). Reserve the earliest available slot if the expected cost of reserving this slot is less than the expected cost of reserving a slot the next round; otherwise, postpone the reservation decision until the next round.

We will now introduce a number of functions to be able to define the expected cost of reserving the earliest available slot, which takes into account the results of having to decommit. First of all, an agent has to pay the decommitment penalty δ ; second, if t_d denotes the time decommitment occurs, and then the aircraft has wasted $(t_d - t)$ minutes (where t is the time at which the slot was reserved). We assume that this quantity $(t_d - t)$ will in fact delay de-icing by $(t_d - t)$ minutes. For simplification purposes, we will assume (i) that $(t_d - t)$ is a fixed value, and so we define the cost of decommitment for agent a as:

$$dcp = \delta + l(t_d - t, a) \quad (7)$$

The cost specified in Equation 7 is the immediate cost that the aircraft incurs — this is the cost the aircraft will have to pay anyway — without taking into account that there is a chance of having to decommit again in the future. Using the above definitions, an agent can calculate the expected cost of reserving a slot at time t_1 with earliest available slot time t_2 :

$$E_{res}(t_1, t_2) = p_d(t_1, \tau(a)) \cdot dcp + (1 - p_d(t_1, \tau(a))) \cdot l(t_2 - \tau(a), a) \quad (8)$$

Note that a more realistic model for the cost of reserving a slot would be forward recursive: in case an aircraft has to decommit, it will have to try to get a slot again in subsequent rounds, again with the possibility of having to decommit, adding to its cost. Equation 8 effectively cuts off this forward recursion after one step, by taking into account only the immediate cost for decommitment. If we now have potential reservation time t_1 , time t_1^+ of the next reservation decision, starting time of the next slot t_2 , and starting time of the next slot after that t_3 (in our case t_3 equals t_2 plus standard de-icing time), then the expected cost of waiting until the next round is given by the following function:

$$E_{wait}(t_1, t_2, t_3) = p_T(t_1, t_2) \cdot E_{res}(t_1^+, t_3) + (1 - p_T(t_1, t_2)) \cdot E_{res}(t_1^+, t_2) \quad (9)$$

in which $p_T(t_1, t_2)$ stands for the probability of another agent having reserved between time t_1 and t_1^+ the slot starting at t_2 . This probability function is based on the number of aircraft in the system, and the scarcity of the de-icing resources. We assume aircraft take-off times are independent of each other and are uniformly distributed over time, and so we model the probability $p_T(t_1, t_2)$ with a Poisson distribution ($f(k; \lambda) = \frac{e^{-\lambda} \lambda^k}{k!}$) where:

$$p_T(t_1) = 1 - f\left(0, \frac{t_1^+ - t_1}{|D| \cdot T}\right) = 1 - e^{-\frac{|A| \cdot (t_1^+ - t_1)}{|D| \cdot T}} \quad (10)$$

where T is the time in minutes over which these aircraft are distributed (e.g., we could have a simulation run of $T = 360$ minutes in which $|A| = 100$ aircraft have to be deiced using $|D| = 4$ de-icing stations). Equation 9 basically expresses that by not reserving a slot this round, there is a chance that another agent reserves the previously earliest available slot, and you consequently have to schedule a later slot t_3 (which will result in more delay); on the other hand, if no agent has reserved the slot starting from t_2 , then this possibility is still open to you at time t_1^+ . By this time, the probability of decommitment will have lowered (i.e. $p_d(t_1^+, t_2) < p_d(t_1, t_2)$), and thus reserving this slot at time t_1^+ will have a lower expected cost. The agent strategy we propose in this section is simple: in case $E_{res} < E_{wait}$, the agent will reserve at time t_1 the slot starting at $t_{nextslot}$, otherwise it will wait until the next round. In the next section, we will investigate whether reasoning about decommitment in this way results in improved performance.

IMPROVISATION SUPPORT HYPOTHESES

In this section, we will compare the two coordination mechanisms of previous section with each other, and also with a naive, baseline scheduling strategy:

Strategy 3 (Naive Scheduling Strategy). The Naive Scheduling Strategy randomly selects the next aircraft to schedule, and it assigns to this aircraft the first available slot after its target off-block time.

We judge the algorithms on three criteria: the first one is the total delay in minutes generated by both mechanisms on identical instances. The second criterion is similar: now we measure the total delay cost of all aircraft, given by the sum of the delay costs of all agents. Recall that the cost of one agent a is given by $l(d, a)$, where d is the agents' delay —this means that we do not take auction fees and decommitment penalties into account when calculating the global cost. Hence, the first two criteria measure the efficiency of the coordination mechanisms. As a third criterion, we also record the standard deviation of delay in minutes, summed over all agents. The standard deviation can be interpreted as a measure of fairness: if it is low, then all agents suffer a comparable amount of delay.

As the focus of the auction mechanism is on the agents' preferences, and the decommitment penalty mechanism focuses on timeliness, we can formulate the following hypothesis:

Hypothesis 1. The auction mechanism will outscore the decommitment mechanism on total cost; the decommitment penalty mechanism will outscore the auction mechanism on total delay.

Because the auction mechanism assigns aircraft a high priority, we can expect that this mechanism will not score as high on 'fairness'.

Hypothesis 2. As the number of aircraft competing for a limited number of resources increases, the decommitment penalty mechanism will score increasingly better on standard deviation compared to the auction mechanism.

Finally, we would expect that both coordination mechanism described in the previous section outperform the uninformed random-order scheduling of aircraft.

Hypothesis 3. Both the auction and the decommitment penalty mechanism will outperform the Naive Scheduling Strategy in terms of efficiency and fairness.

We conducted these experiments using only a single de-icing station with a single de-icing bay, and a de-icing time of 5 minutes. Target Off-Block Times (τ) are randomly distributed over six simulation hours. De-icing slots may be allocated after the initial six hours; in fact, the simulation continues until all aircraft have received a de-icing slot. For these parameters, the number of aircraft n that can maximally be serviced without any delay equals

$$n = \frac{6 \times 60}{5} = 72, \text{ assuming a maximally convenient distribution of TOBTs. This means that with a random}$$

distribution of τ , we can expect some delays regardless of the scheduling strategy in case we have more than 72 aircraft. Some further parameter values include: α_a (coefficient of the delay cost function l for agent a) is randomly distributed over $[0.5, 1.0]$; the fixed decommitment penalty $\delta = 5$; the maximum decommitment probability $c = 1.0$; In the auction setting, slots are auctioned half an hour in advance; and the time in between two rounds in decommitment penalties is set to 5 minutes. The number of aircraft in the experiment ranges from 10 to 90. The results of the experiments are displayed in Table 1 and Figure 1.

Number of aircraft	NSS			DC			Auction		
	Delay	Cost	STDV	Delay	Cost	STDV	Delay	Cost	STDV
10	0	0	0	0	0	0	20	13.2	1.33
15	7	5.875	1.81	3	2.42	0.875	25	18.3	1.45
20	57	43	5.15	36.2	26.4	2.36	82	52.6	5.28
25	58	45.5	5.8	33.5	25.3	2.89	72	52.8	2.45
30	187	144.5	11	57.8	41	2.85	112	81	4.35
35	157	126	9.4	67.6	57.8	3.4	136	96	5
40	315	230	14.5	148	106	6.85	226	152	5.6
45	284	201	10.4	107	74.5	2.98	194	132.5	3.75
50	530	418	17.3	320	235	7.8	371	254	8.8
55	540	387	18.1	380	255	8.6	375	231	6.55
60	730	513	17.9	360	260	6.9	435	289	7.2
65	1980	1425	68	690	525	8.4	628	438	8.5
70	2190	1500	52.1	1230	895	13	1205	790	16
75	3500	2550	74	2150	1610	20.8	2010	1350	36
80	5400	4060	101.5	2560	1895	22.9	2550	1615	43
85	6700	5200	118	4100	3100	24	3850	2520	60.5
90	9230	6300	127	5770	4300	39.8	5900	3620	87.5

Table 1 Total Delay and standard deviation NSS, DC, and Auction strategies




Figure 2 Total Delay cost and standard deviation in NSS, DC, and Auction

The first thing that catches the eye in Table 1 is that the NSS strategy is outperformed by the two other mechanisms on all counts, except for runs having a very small number of aircraft, in which case the auction setting does not perform very well. The reason for this is that in the auction setting we sell slots starting from specific times, such as

10:00, 10:05, etc. In case there is a mismatch with aircraft TOBTs, for example if $\tau(a_i) = 10:03$ for some aircraft a_i , then small delays will be incurred by the aircraft. As competition for the de-icing resources increases, these small delays become less significant. We therefore conclude that Hypothesis 3 has been confirmed.

With regard to Hypothesis 1, the results are less clear. We can see a marked difference between the total cost of the decommitment mechanism and the auction mechanism, especially as the congestion at the airport increases. For the total delay, however, we can't conclude that the decommitment mechanism produces schedules with shorter delay; not as the number of aircraft increases, anyway. For situations having less than fifty aircraft, the decommitment mechanism does produce lower-delay schedules, although in these cases total delay for the airport is not so high anyway. Hypothesis 2 is confirmed by the results. As soon as the airport starts getting congested — from around 70 aircraft — the standard deviation for the auction mechanism shoots up, leaving the decommitment mechanism 'behind'. Note that Figure 1 shows that for the less efficient NSS strategy, airport congestion starts from around 60 aircraft. As a final conclusion, we can remark that the auction mechanism is the most efficient choice for congested airports.

However, when there are relatively few aircraft that need to be de-iced, the auction mechanism (at least in its current implementation) is not as efficient. The increased efficiency of the auction mechanism does come with a price, however, namely that delay is distributed more unevenly over the aircraft.

CONCLUDING REMARKS

In this paper we presented two coordination mechanisms for scheduling agents' tasks on scarce resources. Both of them managed to improve on a naive scheduling approach, without the need for any cooperation between the agents. Hence, our coordination mechanisms can be used as a decision support tool for human operators, who do not have the ability to evaluate all different scheduling alternatives. We can ask ourselves, however, how we can use these research results in the setting of emergency response and incident management. Obviously, these mechanisms can not be applied directly, as they have been developed with the airport de-icing problem in mind. However, at a more abstract level we find that situational context aware but self-interested agents, tasks, and resources to perform these tasks are also relevant concepts in the area of emergency response. Consequently, coordination (or multi-agent scheduling) mechanisms very similar to the ones presented here may prove applicable and more robust in incident management domains.

Of course, in emergency response the concept of self-interested of an agent is not as important as in the ultra-competitive airline industry. As such, it may prove beneficial to look for coordination mechanisms in which the agents more intensively cooperate. In general, we would expect that the more information we have, and the more agents are consulted, the better the (coordination) solution we are able to find. However, the validity of one approach does not rule out the validity of another. It might be that a cooperative coordination mechanism is used when a high rate of inter-agent communication poses no problem. Then, if the situation deteriorates and communication becomes more difficult (or if physical and cognitive load of a human operator passes the acceptable threshold, causing further incidents), a coordination mechanism requiring only limited but situated agent cooperation can provide a useful fallback option for support in human improvisation.

ACKNOWLEDGMENTS

This research is supported by NWO (Netherlands Organization for Scientific Research), Grant No.CSI4006, which is gratefully acknowledged. We would also like to thank Peet van Tooren, Jeroen Valk, Tamás Máhr, Gerald de Jong, and Jan Peter Larsen (Almende B.V.) as well as Jaap van den Herik (Maastricht University) for their help in this research.

REFERENCES

1. J. Collins, M. Tsvetovas, R. Sundaeswara, J. van Tonder, M. Gini, and B. Mobasher. (1999) Evaluating risk: flexibility and feasibility in multi-agent contracting. In O. Etzioni, J. P. Müller, and J. M. Bradshaw, editors, *Proceedings of the Third International Conference on Autonomous Agents (Agents'99)*, pages 350–351, Seattle, WA, USA. ACM Press.

2. R. R. Dynes. (1983) Problems in emergency planning. *Energy*, 8, 1983.
3. R. R. Dynes. (1993) Disaster reduction: The importance of adequate assumptions about social organization. *Sociological spectrum*, 13, 1993.
4. R. R. Dynes. (2000) *Governmental systems for disaster management*.
5. J. Hale. (1997) A layered communication architecture for the support of crisis response. *Journal of Management Information Systems*, 14, 1997.
6. W. Herroelen, B. D. Reyck, and E. Demeulemeester. (1998) *Resource-constrained project scheduling: a survey of recent developments*. *Comput. Oper. Res.*, 25(4):279–302.
7. P. J. 't Hoen and J. A. L. Poutré. (2003) A decommitment strategy in a competitive multi-agent transportation setting. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 1010–1011, New York, NY, USA. ACM Press.
8. R. Kolisch and S. Hartmann. (1999) Heuristic algorithms for solving the resource-constrained project scheduling problem: Classification and computational analysis. In J. Weglarz, editor, *Project scheduling: Recent models, algorithms and applications*, pages 147 – 178, Kluwer, Amsterdam, the Netherlands. Kluwer academic.
9. R. Kolisch and S. Hartmann. (2005) Experimental investigation of heuristics for resource-constrained project scheduling: An update. *European Journal of Operational Research*, 23, May 2005.
10. J.-S. Liu and K. P. Sycara. (1996) *Multiagent coordination in tightly coupled task scheduling*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pages 164-171, May 23 1996.
11. T. O. Paulussen, N. R. Jennings, K. S. Decker, and A. Heinzl. (2003) *Distributed patient scheduling in hospitals*. In G. Gottlob and T. Walsh, editors, *IJCAI*, pages 1224–1232. Morgan Kaufmann.
12. E. L. Quarantelli. (2003) *A half century of social science disaster research: Selected major findings and their applicability*.
13. T. Sandholm and V. Lesser. (2001) Leveled commitment contracts and strategic breach. *Games and Economic Behaviour*, 25:212–270.
14. D. E. Smith, J. Frank, and A. K. Jónsson. (2000) Bridging the gap between planning and scheduling. *Knowl. Eng. Rev.*, 15(1):47–83.
15. V. Teague and L. Sonenberg. (2002) Investigating commitment flexibility in multi-agent contracts. *Game Theory and Decision Theory in Agent-Based Systems*, pages 267–292.
16. J. E. Trainor. (2004) *Search for a system: Multi-organizational coordination in the September 11th world trade center search and rescue response*.
17. J. Valk. (2005) *Coordination among Autonomous Planners*. PhD thesis, Delft University of Technology, December.
18. I. Vermeulen, S. Bohte, D. Somefun, and J. L. Poutré. (2006) Improving patient schedules by multi-agent Pareto appointment exchanging. In *Proceedings of 2006 IEEE International Conference on E-Commerce Technology (CEC/EEE 2006)*, San Francisco, California, June 26-29.
19. M. de Weerd, A. ter Mors, and C. Witteveen. (2005) *Multi-agent planning: An introduction to planning and coordination*. Technical report, Dept. of Software Technology, Delft University of Technology.
20. B. Yang, J. Geunes, and W. J. O'Brien. (2001) *Resource-constrained project scheduling: Past work and new directions*. Technical report, Department of Industrial and Systems Engineering, University of Florida.