Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 51-23 (2018) 194–199

Heterogeneous multi-agent resource allocation through multi-bidding with applications to precision agriculture ⋆

A.T.J.R. Cobbenhagen * D.J. Antunes *
M.J.G. van de Molengraft * W.P.M.H. Heemels *

* *Control Systems Technology group, Dept. of Mechanical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands (e-mail: a.t.j.r.cobbenhagen@tue.nl).*

Abstract: In this paper we consider the problem of allocating multiple resources to a number of clients by a group of heterogeneous agents over time such that the clients can produce products while maximizing a profit function. We propose an approximate optimization framework in which every client provides multiple bids from which the agents choose such that an allocation is feasible and that the profit function is maximized over time. The proposed framework exploits decomposition techniques that can be used for large-scale multi-agent resource allocation problems in which the cost objective is additive, the dynamics of product generation is non-linear and the agents have different capabilities. Interestingly, the decomposition can be solved in a distributed fashion, enabling application to large-scale problems. We apply this decomposition to the management of resources and agents in precision agriculture as an inspirational and important application domain of the obtained results. We show that our framework can be used in order to schedule the time, location and quantity of resources that every agent must provide whilst optimizing the profit of the entire farm over the growing season.

Keywords: Multi-agent resource allocation, Agriculture, Optimal control, Approximate dynamic programming, Distributed optimization

1. INTRODUCTION

In this paper we consider the problem of delivering multiple resources to various clients in order to optimize the production of certain products. The delivery has to be carried out by a group of heterogeneous agents. A prime instance of this problem can be found in precision farming of crops. The resources in precision farming include water, fertilizer and pesticides, which the (sub)fields (the clients in this case) need in order to grow crops effectively. The objective is to devise a control strategy for the allocation of agents and resources such that the profit at the end of the season is maximized. The profit is defined as the difference between the income generated from the amount of harvested crops and the costs incurred by the usage of agents and resources. The (centralized) optimization of resource usage over all fields, whilst taking into account agent constraints, is a large and complex problem as the number of (sub)fields and the number of agents differ by several orders of magnitude and the optimization problem per (sub)field is already rather complex due to the highly nonlinear behaviour of crop growth. The decomposition proposed in this work must therefore significantly improve on the centralized global optimization problem by being computationally tractable and yet close to optimal.

⋆ This work is supported by 'Toeslag voor Topconsortia voor Kennis en Innovatie' (TKI HTSM) from the Ministry of Economic Affairs, the Netherlands.

Precision agriculture concerns the knowledge of when, where and how much resources should be delivered in order to grow crops most effectively. One way of increasing this knowledge is by using optimal decision making techniques. Since there can be a large variance in crop states over a farmland, increasing the temporal and spatial resolution of decision making increases the quality of decision making. Increasing the spatial resolution implies that a field on a farmland can be subdivided into multiple subfields. It is therefore very likely to have in the order of one thousand subfields. The agents in precision agriculture are the irrigation machines, fertilizer machines, pesticide spraying machines, etc. Typically these are in the order of 1-10 per type per farm and thus there is a large asymmetry between the number of agents and clients in this application.

There has been a wide variety of studies done in the field of multi-agent resource allocation (MARA), for recent overviews see Chevaleyre et al. (2006) and Shoham and Leyton-Brown (2009). A common theme in these problems is the analysis of MARA over networks, (e.g., Lesser et al. (2005); Nowzari et al. (2017); Xiao et al. (2004); Obando et al. (2017)), bandwith allocation (e.g., Xiao et al. (2004)) and computation time allocation (e.g., Bredin et al. (2000)). These approaches do not fit the problem of precision agriculture as the dynamics of crop growth are highly nonlinear. The optimization for a single client/subfield is nonconvex, which prevents the usage of techniques such as ADMM (Boyd et al. (2010)) in order to distribute the optimization problem.

Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2018.12.034

Contrarily to these mentioned works, our approach for the decomposition results naturally from an approximate dynamic programming viewpoint of the problem for precision agriculture. As mentioned before, the precision agriculture problems entail approximately thousands of fields and hence dynamic programming becomes sheer impossible due to the *curse of dimensionality* (Bertsekas (2017)). In our approach, we resolve this by an approximation in which the clients solve smaller optimization problems in parallel and a subsequently an optimization is done by the agents as a group.

In MARA problems, task allocation of the agents often happens on a different time scale than the dynamics of the effects of the resources on product generation at the clients. In fact, this is typically the case in the mentioned precision agriculture application. Tasks such as irrigation or applying fertilizer can be done by agents in the time-span of several hours, whereas the effects of these actions can be seen in time-spans of days to weeks. Contrarily to the works mentioned previously, the decomposition that we will propose in this paper, explicitly exploits this time scale separation. In the example of precision farming, the allocation of agents should occur at an hourly to daily basis, whereas the demand of resources that a field needs, should be satisfied on a daily to weekly basis. The reason for this is that the farmer needs to know what the agents are required to do on an hourly basis since that is the order of magnitude of time a task generally takes, whereas the crop growth tends to develop at a much slower rate.

The main contributions of this paper can be summarized as follows. First, a formal introduction of what we call the 'dynamic heterogeneous multi-agent resource allocation' (D-H-MARA) problem relevant to precision agriculture applications is given. The second contribution is the structured description of a decomposition that can be applied to many D-H-MARA problems in which the cost function is nonconvex, the agents are heterogeneous, the product dynamics is nonlinear and the time scales of resource allocation and product dynamics are different. Finally, the third contribution is the application of this novel framework in the context of precision agriculture. This will show that with our new framework we can schedule the actions of all agents with a high temporal resolution whilst optimizing the profit for the entire farm over the growing season. This is a key result in precision farming.

The remainder of the paper is structured as follows. We will first define the distributed heterogeneous MARA we consider and state our assumptions in Section 2. Then we will introduce the multi-bidding decomposition in Section 3 and the allocation constraints in Section 4. The resulting optimization problem of the decomposition is presented in Section 5. We will demonstrate the technique using a simulation example in Section 6 and then provide conclusions and recommendations for future work.

2. PROBLEM FORMULATION

2.1 Mathematical notation and terminology

When using the '$=$', '$\leq$' or '$\geq$' operators in matrix equations, we mean an element-wise comparison. We use $\mathbf{1}_n$ and $\mathbf{0}_n$ to denote the n-dimensional vectors of all ones and all zeros, respectively, and I_n denotes the $n \times n$ identity matrix. If M is a matrix, then $[M]_{ij}$ denotes the element on the i-th row and j-th column of M and $[M]_j$ denotes the j-th column of M. We use $\operatorname{diag}(v)$ to denote a diagonal matrix with the elements of vector v on the diagonal. We use $\mathbb{N}_{[\alpha,\beta]}$ to denote the set of natural numbers in the closed interval $[\alpha, \beta]$ with $\alpha, \beta \in \mathbb{N}$; the binary set $\mathbb{B}$ is defined as $\mathbb{B} = \mathbb{N}_{[0,1]} = \{0, 1\}$ and $\mathbb{R}_{\geq 0}$ denotes the set of non-negative real numbers.

Throughout the rest of this paper we will use the word 'clients' instead of 'fields' as to illustrate that the proposed decomposition has many more applications in which a large, complex and dynamic MARA optimization is performed.

2.2 System parameters and description

Let n denote the number of agents, labelled from the set $\mathcal{N} := \{1, 2, \ldots, n\}$ and let m denote the number of clients, labelled from the set $\mathcal{M} := \{1, 2, \ldots, m\}$. There are a total of p different types of products that clients can produce and there are r different types of resources that the group of agents can deliver to the clients. Each resource is labelled from the set $\mathcal{R} := \{1, 2, \ldots, r\}$. The time horizon of the optimization is denoted by $T \in \mathbb{N}$ and in the example of precision agriculture, this would typically be the duration of the growing season in days.

We define the amount of resources a client $j \in \mathcal{M}$ receives at time $t \in \mathbb{N}_{[0,T]}$ as $u_t^j \in \mathbb{R}^r_{\geq 0}$ and $\rho \in \mathbb{R}^r_{\geq 0}$ denotes a vector where each entry $\iota \in \mathcal{R}$ coincides with the cost per unit for resource ι. Finally, let $\sigma \in \mathbb{R}^n_{\geq 0}$ denote the vector where each entry σ_i, $i \in \mathcal{N}$, is the cost of operation per time unit of agent i. This could also include the cost of the agent's operator.

Similarly, we denote the state of a client $j \in \mathcal{M}$ at time $t \in \mathbb{N}_{[0,T]}$ by $z_t^j \in \mathbb{R}^o$ (o thus equals the state dimension) and assume that the discrete-time dynamics equals

$$z_{t+1}^j = z_t^j + \tilde{f}^j(z_t^j, u_t^j, w_t^j),$$

for some mapping $\tilde{f}^j : \mathbb{R}^o \times \mathbb{R}^r_{\geq 0} \times \mathbb{R}^v \to \mathbb{R}^o$, where w_t^j is a disturbance (assumed to have dimension v) at time t. The probability density function (pdf) of a disturbance w_t^j is equal to $h_t^j : \mathbb{R}^v \to [0,1]^v$ for all $j \in \mathcal{M}$ at time $t \in \mathbb{N}_{[0,T]}$. We assume that these pdfs are known and that they can vary over time.

Not all of the states are interesting for the final profit. Crop growth models, for instance, typically have a large amount of states, including the leaf area size, soil-water content, root depth and more (e.g., see Shibu et al. (2010)). These are relevant for describing the state of the product (e.g., the biomass of storage organ) but not for the final profit. Let $x_t^j \in \mathbb{R}^p_{\geq 0}$ denote the vector with the amounts of each products at client $j \in \mathcal{M}$ at time $t \in \mathbb{N}_{[0,T]}$. It is assumed that this can be computed from the state of client j by using a constant matrix $H^j \in \mathbb{R}^{p \times o}$ such that

$$x_t^j = H^j z_t^j.$$

We can thus state that the dynamics of the products at a client $j \in \mathcal{M}$ equals

$$x_{t+1}^j = x_t^j + f^j(z_t^j, u_t^j, w_t^j), \tag{1}$$

with $f^j = H^j \tilde{f}^j$. Let $\pi \in \mathbb{R}^p_{\geq 0}$ denote the price vector, where the i-th element denotes the price per unit of product type i.

As mentioned before, we consider the problem in which the allocation of agents happens on a faster time scale than the effects of the resources on growth of products at clients. From now on, for convenience, we will call the slow time scale 'time' and the fast time scale 'subtime'. Between two instances of time, we assume that there are $q \in \mathbb{N}$ units of subtime, which means that the agents can be allocated q times to different clients within one time unit. If the amount of subtimes q is chosen such that one unit of subtime corresponds to a typical task duration of an agent (including recovery time), then it is realistic to make the assumption that each agent visits at most one client at an instance of subtime.

We assume that the agents can have different maximum capacities of each resource. Let $C_{\max} \in \mathbb{R}^{r\times n}_{\geq 0}$ denote the capacity matrix such that $[C_{\max}]_{ki}$ denotes the maximum quantity of resource $k \in \mathcal{R}$ that agent $i \in \mathcal{N}$ can deliver at an instance of subtime. Note that $[C_{\max}]_{ki} = 0$ means that agent i cannot deliver resource k.

2.3 Profit function

We aim to optimize the total profit of the system at the end of the time horizon. This means that at a time $\iota \in \mathbb{N}_{[0,T-1]}$, the goal is to maximize

$$J_\iota = \mathbb{E}\left\{\sum_{j\in\mathcal{M}}\left(\pi^\top x_T^j - \sum_{t=\iota}^{T-1}\left(\rho^\top u_t^j + \sigma^\top [A_t]_j\right)\right)\right\}, \qquad (2)$$

where $A_t \in \mathbb{N}^{n\times m}_{[0,q]}$ is the allocation matrix of which $[A_t]_{ij}$ denotes the number of times agent $i \in \mathcal{N}$ serves client $j \in \mathcal{M}$ between time t and time $t+1$. The decision variables are the amount of delivered resources u_t^j, for all $t \in \mathbb{N}_{[\iota,T]}$ and all $j \in \mathcal{M}$, and the allocations of agents and clients $[A_t]_{ij} \in \mathbb{N}_{[0,q]}$, for all $t \in \mathbb{N}_{[\iota,T]}$. We use the expectation operator $\mathbb{E}$ to denote the expected values with respect to the disturbances w_t^j.

2.4 D-H-MARA problem definition

Let us now formally define a dynamic heterogeneous MARA.

Definition 1. (D-H-MARA). A dynamic heterogeneous multi-agent resource allocation (D-H-MARA) is defined as

$$\Sigma = (m, n, \pi, \rho, \sigma, T, q, C_{\max}, \{\tilde{f}^j, H^j\}_{j=1}^m, J_\iota),$$

with the parameters as previously discussed in this section.

We call it heterogeneous as the agents can have different capabilities in terms of delivering types and quantities of resources, but also the clients can produce different types of products with different dynamics.

The D-H-MARA *problem* is then finding a *policy* μ_t for u_t^j and A_t that maximizes (2) given a D-H-MARA and pdfs h_t^j. The policy is therefore of the form

$$(u_t^1, \ldots, u_t^m, A_t) = \mu_t(z_t^1, \ldots, z_t^m).$$

3. DECOMPOSITION THROUGH MULTI-BIDDING

In this section we will use approximate dynamic programming (see Bertsekas (2017)) in our approach of solving the D-H-MARA problem, i.e., finding a policy that maximizes (2).

3.1 Rewriting the cost function

The objective function in (2) is a summation over the clients. Note that there is a coupling between the clients and the agents not only due to the term $[A_t]_{ij}$ but also through the constraints on the term u_t^j which depend on the clients (see Section 4). We will now decompose the problem of computing the required amount of resources and the allocation problem into separate problems.

Let us rewrite (2) using the dynamics (1) as

$$J_\iota = \mathbb{E}\left\{\sum_{j\in\mathcal{M}}\left(\pi^\top x_\iota^j + \sum_{t=\iota}^{T-1}\pi^\top f^j(z_t^j, u_t^j, w_t^j) - \rho^\top u_t^j - \sigma^\top [A_t]_j\right)\right\}$$

Since we are interested in the selection of u_ι^j and A_ι, we observe that the term $\pi^\top x_\iota^j$ is not influenced by either of them and can thus be removed from the cost function of the optimization problem. Hence, we obtain the following adjusted cost function (with the same optimizer)

$$J_\iota' = \sum_{j\in\mathcal{M}}\left(\pi^\top \mathbb{E}\left\{\sum_{t=\iota}^{T-1} f^j(z_t^j, u_t^j, w_t^j)\right\} - \rho^\top \mathbb{E}\left\{\sum_{t=\iota+1}^{T-1} u_t^j\right\}\right.$$
$$\left. -\rho^\top \mathbb{E}\left\{u_\iota^j\right\} - \sigma^\top \mathbb{E}\{[A_\iota]_j\} - \sigma^\top \mathbb{E}\left\{\sum_{t=\iota+1}^{T-1} [A_t]_j\right\}\right),$$

where we also split the summations over u_t^j and $[A_t]_j$. Note that J_ι' is a function of all u_ι^j and this is potentially a highly nonlinear and nonconvex function. We keep the notation of the expectation operator over the u_t^j and A_t since these are policies dependent on the future states and disturbances.

3.2 Sampling of allocation space using bids

We will now take samples in the space of possible allocations and propose the approximate decomposition. We assume that every client computes an amount of *s scenarios* at each time instant, labelled from the set $\mathcal{S} := \{1, 2, \ldots, s\}$. We call these scenarios *bids* and formally define these as follows.

Definition 2. (Bid). A bid $b_k^j(\iota)$, $k \in \mathcal{S}$, of a client $j \in \mathcal{M}$ at a time $\iota \in \mathbb{N}_{[0,T]}$ is defined as the triple

$$b_k^j(\iota) = \left(g_k^j(\iota), \ell_k^j(\iota), \bar{u}_k^j(\iota)\right) \in \mathbb{R}^p \times \mathbb{R}^r_{\geq 0} \times \mathbb{R}^r_{\geq 0},$$

where

$$g_k^j(\iota) = \mathbb{E}\left\{\sum_{t=\iota}^{T-1} f^j(z_t^j, u_t^j, w_t^j)\right\}$$

is the expected total future gain of products given that the client receives at least $\bar{u}_k^j(\iota)$ resources at time ι. Similarly,

$$\ell_k^j(\iota) = \mathbb{E}\left\{\sum_{t=\iota+1}^{T-1} u_t^j\right\}$$

is the expected total future use of resources needed to realize $g_k^j(\iota)$ if $\bar{u}_k^j(\iota)$ resources are delivered at time ι.

Using these bids, we are essentially sampling in the space of allocations and resource deliveries such that the optimization of resources can be done by each client separately and thus in parallel (distributed) and the agents use these results to solve the allocation. The best way to interpret $g_k^j(\iota)$ and $\ell_k^j(\iota)$ is that $\pi^\top g_k^j(\iota) - \rho^\top \ell_k^j(\iota)$ is an approximation of the *cost-to-go* for client j at time ι if the client receives $\bar{u}_k^j$ resources at time ι. Also note that the pdfs h_t^j of the disturbances are taken into account in the computation of the bids.

The decomposition we propose assumes that each client computes $g_k^j(\iota)$ for several values of $\bar{u}_k^j(\iota)$ such that u_t^j follows a certain *base policy* (Bertsekas (2017)) for $t \in \mathbb{N}_{[\iota+1,T-1]}$. One such example of a base policy is that the clients assume that they will not receive any resources after time ι and thus $u_t^j = 0$ for all $t \in \mathbb{N}_{[\iota+1,T-1]}$ and $\ell_k^j(\iota) = \mathbf{0}_r$. A second example is the policy in which the agents compute the amount of resources needed for largest growth of the crop.

The results of these computations are then passed to the agents in the form of bids. The agents will then use these bids to allocate the resources while taking into account the allocation constraints. Let $\delta \in \mathbb{B}^{s\times m}$ denote the bid decision matrix. If the agents choose that the k-th bid is the scenario for client j that they wish to execute, then $\delta_k^j = [\delta]_{kj}$ equals one, otherwise it is zero. We assume that one bid is selected and thus $\delta_k^j = 1 \Rightarrow \delta_{k'}^j = 0$ for $k \neq k'$ for all $j \in \mathcal{M}$. We use this bid decision matrix to provide the following approximation of J'_ι

$$\tilde{J}_\iota = \sum_{j\in\mathcal{M}} \left(\sum_{k\in\mathcal{S}} \delta_k^j \left(\pi^\top g_k^j - \rho^\top \left(\bar{u}_k^j + \ell_k^j \right) \right) - \sigma^\top [A_\iota]_j \right). \qquad (3)$$

Also, we assume that the value of the future number of agents that will be allocated to all clients, i.e., $\sum_{t=\iota+1}^{T-1} \sum_{j\in\mathcal{M}} [A_t]_j$, does not vary significantly with respect to $\bar{u}_\iota$ and therefore leave it out of the objective function.

4. OPTIMIZATION CONSTRAINTS

In this section we will show how the assumptions on bid selection, agent capabilities and allocation constraints translate to mathematical formulations that can be used in the optimization problem.

All decision variables, bid parameters and constraints in this section are evaluated for the same time $\iota \in \mathbb{N}_{[0,T-1]}$. For notational clarity, we omit this dependency on time.

4.1 Bid selection constraints

We assume that the agents always select exactly one bid for each client and hence

$$\sum_{k\in\mathcal{S}} \delta_k^j = 1, \quad j \in \mathcal{M},$$

or otherwise stated,

$$\delta^\top \mathbf{1}_s = \mathbf{1}_m. \qquad (4)$$

4.2 Allocation constraints

Let $A(\tau) \in \mathbb{B}^{n\times m}$ denote the allocation matrix at subtime τ such that if $a^{ij}(\tau) := [A(\tau)]_{ij}$ equals one, a agent i delivers resources to client j at subtime $\tau \in \{1, 2, \dots, q\} =: \mathcal{Q}$. As stated in Section 2, we can assume that at every instant of subtime, an agent can deliver resources to at most one client and every client can receive resources from at most one agent. This results in

$$A(\tau)\mathbf{1}_m \leq \mathbf{1}_n, \qquad (5a)$$

$$A(\tau)^\top \mathbf{1}_n \leq \mathbf{1}_m, \qquad (5b)$$

for all $\tau \in \mathcal{Q}$.

4.3 Agent constraints

Let $D(\tau) \in \mathbb{R}_{\geq 0}^{r\times m}$ denote the receive matrix such that $[D(\tau)]_j = d^j(\tau)$ is the vector of resources delivered to client $j \in \mathcal{M}$ at subtime τ. Furthermore, let $C(\tau) \in \mathbb{R}_{\geq 0}^{r\times n}$ denote the delivery matrix such that $[C(\tau)]_i = c^i(\tau)$ is the vector of resources that agent $i \in \mathcal{N}$ has delivered at subtime τ. Under the assumption that all resources that are delivered by an agent are received by a client, we have that

$$c^i(\tau) = \sum_{j\in\mathcal{M}} d^j(\tau)\, a^{ij}(\tau), \quad \tau \in \mathcal{Q}, i \in \mathcal{N}, \qquad (6)$$

or stated otherwise,

$$C(\tau) = D(\tau)A(\tau)^\top, \quad \tau \in \mathcal{Q}.$$

Since $[C(\tau)]_{ki}$ denotes the k-th resource that agent i delivers, we can incorporate capacity constraints to the agents by the capacity matrix $C_{\max} \in \mathbb{R}_{\geq 0}^{r\times n}$ and enforcing the constraints

$$C(\tau) = D(\tau)A(\tau)^\top \leq C_{\max}, \quad \tau \in \mathcal{Q}. \qquad (7)$$

Note that these inequalities include products of decision variables $d^j(\tau)$ and $a^{ij}(\tau)$ and thus these are *bilinear* constraints.

4.4 Flow constraints

If a client receives resources from an agent, then that agent must be allocated to that client. We must therefore have

$$d^j(\tau) = \sum_{i\in\mathcal{N}} a^{ij}(\tau) c^i(\tau), \quad j \in \mathcal{M}, \tau \in \mathcal{Q}.$$

Substitution of (6) yields

$$d^j(\tau) = \sum_{i\in\mathcal{N}} a^{ij}(\tau) \sum_{h\in\mathcal{M}} a^{ih}(\tau)\, d^h(\tau),$$

and since $a^{ij}(\tau)a^{ih}(\tau) = 0, j \neq h$ due to (5a) and $(a^{ij}(\tau))^2 = a^{ij}(\tau)$ as $a^{ij}(\tau) \in \mathbb{B}$, we can write

$$d^j(\tau) = \sum_{i\in\mathcal{N}} a^{ij}(\tau)a^{ij}(\tau)\, d^j(\tau) = \sum_{i\in\mathcal{N}} a^{ij}(\tau)\, d^j(\tau)$$

and therefore

$$\left(1 - \sum_{i\in\mathcal{N}} a^{ij}(\tau)\right) d^j(\tau) = \mathbf{0}_r, \quad \tau \in \mathcal{Q}, j \in \mathcal{M}. \qquad (8)$$

We observe that if an agent is allocated to client j, then $1 - \sum_{i\in\mathcal{N}} a^{ij}(\tau) = 0$ and $d^j(\tau)$ is not restricted further by this constraint. If no agent is allocated to client j, then $1 - \sum_{i\in\mathcal{N}} a^{ij}(\tau) = 1$ and $d^j(\tau) = \mathbf{0}_r$ is the only solution.

These constraints can be written into one matrix equation as

$$\mathbf{0}_{r\times m} = D(\tau)\big(I_m - \operatorname{diag}\big(A(\tau)^\top \mathbf{1}_n\big)\big), \quad \tau \in \mathcal{Q}.$$

These equations are also bilinear in the decision variables.

4.5 Inter time scale constraints

In accordance with previous definitions, we must have that the allocation in subtime must be congruent with the allocation in time, i.e.,

$$A_t = \sum_{\tau\in\mathcal{Q}} A(\tau). \tag{9}$$

Furthermore, the resources that are delivered to client j in all instances of subtime together should be equal to the amount received in an instant of time $t \in \mathbb{N}_{[0,T]}$

$$u_t^j = \sum_{\tau\in\mathcal{Q}} d^j(\tau), \quad j \in \mathcal{M}, \tag{10}$$

where we must impose that the delivered resources are equal to the amount stated in the accepted bid and thus

$$\sum_{k\in\mathcal{S}} \delta_k^j\, \bar{u}_k^j = u_t^j, \quad j \in \mathcal{M}. \tag{11}$$

4.6 Overview of constraints

Summarizing the effects of the constraints, we have that

i) exactly one bid is accepted for each client through (4);
ii) the amount of resources delivered to a client is equal to the amount stated in the accepted bid, which is guaranteed by (11);
iii) through (10) we have that the delivered amount of resources to a client equals the total amount delivered to that client in subtime;
iv) there are no free deliveries through (8), i.e., if resources are delivered to a client, then an agent must be allocated;
v) it is guaranteed that the agents deliver the resources below their maximum capacity through (7);
vi) through (9) we have that the number of usages of agents in time is equal to the number of usages in subtime.

5. OPTIMIZATION OF THE DECOMPOSITION

Using the constraints presented in the previous section, we can rewrite $\tilde{J}_t$ in (3) and thereby the optimization problem for the agents becomes a *Mixed Integer Bilinear Program* (MIBP) with a linear objective function:

$$\max_{\substack{\delta, A(\tau),\\ D(\tau)}} \sum_{j\in\mathcal{M}}\sum_{k\in\mathcal{S}} \delta_k^j(\pi^\top g_k^j - \rho^\top \ell_k^j) - \sum_{\tau\in\mathcal{Q}} (\rho^\top D(\tau) + \sigma^\top A(\tau))\mathbf{1}_m$$

subject to, for all $\tau \in \mathcal{Q}$,

$$\begin{aligned} \delta^\top \mathbf{1}_s &= \mathbf{1}_m,\\ A(\tau)\mathbf{1}_m &\le \mathbf{1}_n,\\ A(\tau)^\top \mathbf{1}_n &\le \mathbf{1}_m,\\ -\big[\textstyle\sum_{\tau\in\mathcal{Q}} D(\tau)\big]_j + \sum_{k\in\mathcal{S}} \delta_k^j \bar{u}_k^j &= \mathbf{0}_r, \qquad j \in \mathcal{M},\\ D(\tau)\big(I_m - \operatorname{diag}\big(A(\tau)^\top \mathbf{1}_n\big)\big) &= \mathbf{0}_{r\times m},\\ D(\tau)A(\tau)^\top &\le C_{\max}, \end{aligned}$$

$$\delta \in \mathbb{B}^{s\times m},\ A(\tau) \in \mathbb{B}^{n\times m},\ D(\tau) \in \mathbb{R}_{\ge 0}^{r\times m}. \tag{12}$$

This optimization problem has $m(s + q(n + r))$ decision variables and

- $m(1 + qr)$ linear equality constraints;
- $q(n + m)$ linear inequality constraints;
- qrm bilinear equality constraints;
- qrn bilinear inequality constraints.

This has now become an optimization problem of a smaller complexity than a centralized optimization in which the optimal resource computation and the allocation is solved simultaneously.

Now that the optimization problem of the decomposition is clear, let us formally define the decomposition.

Definition 3. (Decomposition). A decomposition with multi-bidding of a D-H-MARA Σ is defined as the pair (Σ, s) with s the number of bids per client per time instance and the optimization problem (12) is solved at every time instance.

If the clients' bids are too high relative to the agents capacities, the optimization problem might not provide a feasible solution. We will give a sufficient condition on the bids to ensure feasibility. To this extend, we use the notion of *zero bids*. A zero bid $b_k^j(\iota)$ is defined as a bid $k \in \mathcal{S}$ of an agent $j \in \mathcal{M}$ at a time $\iota \in \mathbb{N}_{[0,T-1]}$ such that $\bar{u}_k^j(\iota) = \mathbf{0}_r$. The sufficient condition for feasibility is given in the following proposition.

Proposition 1. (Sufficient condition for feasibility). If all clients have a zero bid among their bids, i.e.,

$$\forall j \in \mathcal{M}, \exists k \in \mathcal{S} : b_k^j = (\cdot, \cdot, \mathbf{0}_r),$$

then the optimization (12) has at least one feasible solution.

Proof. If all clients $j \in \mathcal{M}$ place at least one bid where $\bar{u}_k^j(\iota) = \mathbf{0}_r$, then these bids can be selected and through (11) no resources are delivered. This implies that no agents need to be allocated and hence all allocation constraints are satisfied. ■

In the example of precision agriculture, inclusion of a zero bid is a reasonable assumption as not all fields need to be served at each time instant.

6. SIMULATION EXAMPLE

As a simple, yet illustrative example, we present the results of a simulation of farm management of spring wheat. To this extend we use the celebrated and well known LINTUL2 crop growth model, the documentation of which can be found in Spitters and Schapendonk (1990). In LINTUL2 it is assumed that the crop grows under water-limited stress and other limiting factors such as nitrogen stress and pests are assumed to be absent. We use spring wheat as an example since the parameters in LINTUL2 have been extensively validated with regard to this particular crop.

This simulation example considers $m = 30$ subfields of equal area and $n = 3$ irrigation agents. There is one product ($p = 1$) represented by the mass of the storage organ of the crop. Water is the only resource ($r = 1$) that the agents can deliver in this example. The agents' capacities are equal to $C_{\max} = [1\ 1\ 2]$ [mm water/subtime] (arbitrarily chosen). We evaluate the dynamics for each day and assume that agents can be allocated twice a day

($q = 2$). We set $\pi = 100$ [€/(kg/m^2)], $\rho = 1$ [€/mm water] and $\sigma = \mathbf{0}_n$ (i.e., agent allocation is 'free'). We have rather arbitrarily set these costs and reserve an investigation into real costs for future work. At every time instant, the bids for all agents are computed for $\bar{u}_k^j \in \mathbb{N}_{[0,5]}$ [mm water/time], and thus $s = 6$. All clients compute the bids under the assumption that they will not receive any resources in the future and thus $\ell_k^j = 0$ for all $k \in \mathcal{S}$ and $j \in \mathcal{M}$.

The simulation was carried out in MATLAB where SCIP was used to solve the optimization problem (12) (see Gleixner et al. (2017)). The simulation uses historical weather data (supplied with LINTUL2, weather data recorded in 1991 in Wageningen, The Netherlands) of rain, temperature and irradiation to compute the predicted future amount of product. However, in the evaluation of the state updates, a random disturbance from a normal distribution with a mean of -5 and standard deviation of 5, is applied to the amount of rain (excluding negative values of rain).

Under these weather conditions, if we would keep the water at field capacity (i.e., maximum capacity such that there is no water stress for the crop), 11 909 units of water would be needed to achieve the theoretical maximum yield per field of 7.676 ton/ha of dry matter of the storage organ (DM). This may give the largest amount of product, but not the largest amount of profit due to the large amount of irrigation needed. In fact, the agents do not have the capacity to deliver 11 909 units of water in this timespan and thus an evaluation of profit J_0 in (2) is meaningless in this case. If no irrigation would be applied and the fields would only receive water due to rain, then the yield would be 5.215 ton/ha of DM at a profit of $J_0 = 1.56 \cdot 10^6$. Using our proposed framework, a feasible allocation was computed in which the clients were able to generate an average of 6.739 ton/ha of DM using 1 358 units of water at a total profit of $J_0 = 2.02 \cdot 10^6$ [€]. This yield is 88% of the theoretical maximum yield and the profit increased by 29% compared to no irrigation. This is summarized in Table 1.

Table 1. Simulation results

	Resource [mm]	Yield [ton/ha]	Profit J_0 [€]
Optimal yield	11 909	7.676	-
No irrigation	0	5.215	$1.56 \cdot 10^6$
Proposed	1 358	6.739	$\mathbf{2.02 \cdot 10^6}$

7. CONCLUSIONS AND FUTURE WORK

This paper described an approximate dynamic solution to the dynamic heterogeneous MARA (D-H-MARA) such that a large part of the allocation optimization can be distributed over the clients (i.e., the computation of bids). Many relevant problems fit in the class of D-H-MARA problems as the description is general and, as shown, encompasses allocation problems in precision agriculture.

The performance of the decomposition is for a large part based on the quality of the bids. In future work we will analyse the situation in which agents try to improve the quality of their bids by communicating with other clients (either all clients or clients in a certain neighbourhood) before submitting them to the agents. Further exploration of the relation between the problem presented in this paper and the 'knapsack problem' is reserved for future work (Kellerer et al. (2004)). Another direction of research is the study of the effects of other baseline policies for the clients to use in order to compute their bids as mentioned in Section 3.

For the application in precision agriculture, further research is needed into the use of more encompassing crop growth dynamics, i.e., functions $\hat{f}^j$. Further research in temporal ordering constraints of resources is also needed, as well as resource constraints over the entire season.

REFERENCES

Bertsekas, D.P. (2017). *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, Massachusetts, fourth edition.

Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2010). Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends Mach. Learn.*, 3(1), 1–122.

Bredin, J., Maheswaran, R.T., Imer, Ç., Başar, T., Kotz, D., and Rus, D. (2000). A game-theoretic formulation of multi-agent resource allocation. In *Proc. fourth Int. Conf. Auton. agents - AGENTS '00*, 349–356. ACM Press, New York.

Chevaleyre, Y., Dunne, P., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J., and Sousa, P. (2006). Issues in Multiagent Resource Allocation. *Informatica*, 30(1), 3–31.

Gleixner, A., Eifler, L., Gally, T., Gamrath, G., Gemander, P., Gottwald, R.L., Hendel, G., Hojny, C., Koch, T., Miltenberger, M., Müller, B., Pfetsch, M.E., Puchert, C., Rehfeldt, D., Schlösser, F., Serrano, F., Shinano, Y., Viernickel, J.M., Vigerske, S., Weninger, D., Witt, J.T., and Witzig, J. (2017). The SCIP Optimization Suite 5.0. Technical Report 17-61, ZIB, Berlin.

Kellerer, H., Pferschy, U., and Pisinger, D. (2004). *Knapsack Problems*. Springer-Verlag, Berlin, Heidelberg, first edition.

Lesser, V., Ortiz, C.L., and Tambe, M. (eds.) (2005). *Distributed sensor networks*. Springer Science & Business Media, New York.

Nowzari, C., Preciado, V.M., and Pappas, G.J. (2017). Optimal Resource Allocation for Control of Networked Epidemic Models. *IEEE Trans. Control Netw. Syst.*, 4(2), 159–169.

Obando, G., Quijano, N., and Gauthier, A. (2017). Distributed Resource Allocation Among a Subset of Nodes of a Graph. *IFAC-PapersOnLine*, 50(1), 8417–8422.

Shibu, M., Leffelaar, P., van Keulen, H., and Aggarwal, P. (2010). LINTUL3, a simulation model for nitrogen-limited situations: Application to rice. *Eur. J. Agron.*, 32(4), 255–271.

Shoham, Y. and Leyton-Brown, K. (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.

Spitters, C. and Schapendonk, A. (1990). Evaluation of breeding strategies for drought tolerance in potato by means of crop growth simulation. *Plant Soil*, 123(2), 193–203.

Xiao, L., Johansson, M., and Boyd, S. (2004). Simultaneous Routing and Resource Allocation Via Dual Decomposition. *IEEE Trans. Commun.*, 52(7), 1136–1144.