Optimal Irrigation Allocation for Large-Scale Arable Farming

A. T. J. R. Cobbenhagen, L. P. A. Schoonen, M. J. G. van de Molengraft, and W. P. M. H. Heemels, *Fellow, IEEE*

***Abstract*— In this article, we propose an optimization framework that computes the allocation of irrigation machinery (agents) and water to arable fields by maximization of a profit function in a receding horizon fashion using realistic models for crop growth dynamics. A key advantage of the proposed framework is that it can use many of the existing crop growth models for prediction and hence can be applied to a wide variety of crops, soils, locations, and weather patterns. The output of the framework is a feasible allocation of agents to fields and delivery of water over the growing season, in such a way that it allows for practical implementation by farmers. The allocation is feasible as the framework takes into account relevant real-world constraints such as the water-carrying capacity and application rates of the agents, but also traveling costs and refilling of water at designated locations. A realistic case study using validated crop-growth models by agronomists is used to show the strengths and the generality of the framework.**

***Index Terms*— Large-scale systems, model predictive control (MPC), multi-agent systems, precision agriculture.**

I. Introduction

THIS work is motivated by two global trends. First of all, agriculture is responsible for approximately 69% of annual water withdrawals globally, of which irrigation is a large portion [1]. Second, global food demand is rising [2]. Therefore, a major question in arable farming is when, where, and how much irrigation should be applied in order to produce food of high quality in large quantities while minimizing waste of water. Typically, there are many more farmlands than there are irrigation machines, and farmers, therefore, need to carefully choose which fields to irrigate in addition to how much water should be applied. We use "fields" to denote a connected piece of arable farmland. Fields typically represent a local context characterized by, e.g., location, soil composition, and groundwater levels. Fields can be of any size: from a large scale (i.e., order of hectares) to a near individual crop scale (i.e., square meters). Many uncertainties play a role when making irrigation-related decisions, which the farmer has to do on a daily basis. There are uncertainties due to disturbances such as (future) weather influences (e.g., rain and solar irradiation), but also uncertainty due to lack of accurate knowledge of the state of the crops and limited predictability of future crop states due to model mismatch and/or parameter uncertainty.

In the present article, we adopt a model-based approach to design a framework in order to compute the optimal actions for the water delivering agents. "Optimal" here is in the sense of maximization of a profit function. In this framework, we use crop growth models in order to predict the states of the crops in the future. There is a vast amount of literature concerned with the modeling of crop growth. For survey articles, e.g., [3], [4]. These models are usually simulation models made by agronomists, ecologists, or biologists to better understand how crops grow and to predict crop yield. For a given set of parameters, inputs, and disturbances, such models predict growth during the growing season. Examples include light interception and utilization simulator (LINTUL) [5], WOrld FOod STudies (WOFOST) [6], Simulateur mulTIdiscplinaire pour les Cultures Standard (STICS) [7], and AquaCrop [8]. These models have been validated extensively.

In order to design control algorithms for crop growth, one might be inclined to convert and approximate available simulation models into a set of differential/difference equations or partial differential equations. In some applications, these are good enough as a "first principle" model (e.g., design of state-estimators [9]). However, these models tend to generalize poorly when applied to other types of crops, soils, locations, and weather patterns and they require extensive calibration. This makes them rather unsuitable for larger-scale decision making on crop management such as irrigation as considered in the present work. In fact, crop growth simulation models, as discussed in the previous paragraph, already suffer from the effects of climate change and are continuously adapted [10]. We argue that there is a tremendously important and interesting opportunity for system and control engineers and theorists to apply their knowledge of modeling in cooperation with the domain experts of agriculture in order to obtain crop growth models that: 1) generalize well to different situations (such as climatic effects, different crop types, etc.) and 2) can make more use of the existing methods in control theory to design controllers.

In the present article, we take an alternative approach and we focus on a control design that is not driven by any severe modification of the crop growth simulation models. In fact, our approach does not even require an analytical expression of the crop-growth model, which allows for the

Manuscript received March 1, 2021; revised August 2, 2021; accepted September 5, 2021. Date of publication November 9, 2021; date of current version June 23, 2022. This work was supported by "Toeslag voor Topconsortia voor Kennis en Innovatie" [TKI Top Sector High Tech Systems and Materials (HTSM)] from the Ministry of Economic Affairs, The Netherlands. Recommended by Associate Editor D. Raimondo. *(Corresponding author: A. T. J. R. Cobbenhagen.)*

The authors are with the Control Systems Technology Group, Department of Mechanical Engineering, Eindhoven University of Technology, 5612 AZ Eindhoven, The Netherlands (e-mail: a.t.j.r.cobbenhagen@tue.nl; l.p.a.schoonen@student.tue.nl; m.j.g.v.d.molengraft@tue.nl; m.heemels@tue.nl).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TCST.2021.3118296.

Digital Object Identifier 10.1109/TCST.2021.3118296

use of software models. This results in a modular approach, where one can substitute a new crop growth model into our framework and thereby making it generally applicable. In our proposed optimization framework (defined in Section II-C), we make use of only a very small set of assumptions on the crop growth dynamics that are general for many different crops in different conditions.[1] The optimization framework we present can therefore make use of a large collection of existing crop growth (simulation) models that have been verified and validated in experiments and are continuously updated by domain experts.

A. Problem Description and System Context

One of the key decisions many farmers have to make on a daily basis is which fields should receive irrigation and how much irrigation they should receive. The difficulty in those decisions is due to several factors such as uncertainty in crop states, in how the crops will develop, in what the soil-water content is, and in the future weather. Due to the presence of these unknown disturbances and uncertainties, farmers need to make use of feedback and thus make their decisions in an "online" fashion. Constraints on irrigation agents' capacities and rates also need to be taken into account, as well as the costs of operating these agents. "Optimal" decision making in irrigation can then be defined in various different ways, dependent on the user. Some farmers may want to maximize their economic profit over the season by delivering a large number of high-quality crops. Other farmers may want to deliver the most amount of crops by doing the least amount of actions or using the least amount of water.

There are many forms of applying irrigation such as flood irrigation, drip irrigation, spray irrigation, and nebulizer irrigation [11]. In this work, we consider agents to be of the latter two types. More specifically, we assume that the agents are mobile and can move over a field to deliver the water. These include, for instance, popular machines such as traveling gun irrigation machines [12], [13].

Summarizing, the problem discussed in this work is the following. How to design a framework for optimal online automated decision making in irrigation for large-scale arable farming that takes into account the limitations of the agents and predictions of crop growth using any crop-growth model.

B. Previous Work

Due to the importance and difficulty of optimal irrigation, it has been a topic of interest for many years. Recent developments in the application of optimal control to irrigation can be found in [14] and [15] and in model predictive control (MPC) for irrigation in [16]–[19]. Optimal control for large-scale irrigation channels that do not make use of crop growth models can be found in [20] and in [21] for a human-in-the-loop application. An overview of the usage of MPC in agriculture can be found in [22].

In many of the irrigation scheduling policies available in the literature, it is assumed that every field that requires irrigation can actually be irrigated, e.g., [14], [16], [18], [19]. However, the number of irrigation machines that farmers have to their availability is typically much less than the number of fields (that require irrigation). Hence, there is a need to not only compute how much irrigation a field would optimally need (ignoring the possible unavailability of machinery) but there is also a need to select the allocation of which (sub)fields receive irrigation by which delivery agent.

For the above reasons, we present in this article a framework for optimal irrigation scheduling in arable farming by a group of heterogeneous agents. An important novelty of this framework is that it computes the optimal allocation of irrigation while taking into account constraints on agent allocations such as:

1) The maximum resource carry capacity;
2) The maximum rate of irrigation;
3) Minimization of agent travel distances;
4) Returning to a refill station to replenish water;
5) Time-scale separation between crop growth dynamics and allocation.

Item 5) is of interest as the dynamics of crops have time-scales of several days, whereas allocation can occur much faster and has a time-scale of hours. For an overview of multi-agent resource allocation, see [23] and references therein.

The framework we propose is based on ideas from MPC. Within this context, a key novel insight is that we provide a technique to link short-term water stress effects within the (short) prediction horizon to the long-term effects of the crop growth during the remainder of the season.

C. Contributions

Summarizing, the main contributions of this work are threefold. First, we present a problem definition of irrigation scheduling in large-scale arable farming for optimal control including the (often ignored) allocation constraints of irrigation machinery, and we provide a mathematical formulation of this problem.

Second, the optimal control framework we employ can make use of a wide selection of existing crop growth model. This allows for a large applicability to different types of crops, soils, locations, and weather patterns.

Finally, the proposed optimization framework can handle a wide variety of real-world constraints on the agent allocation such as carrying-capacity, irrigation rate, and replenishing of resources.

This article is a significant extension of our preliminary works [24], [25]. The extensions include (but are not limited to) the minimization of agent travel distances, monitoring of resources of agents and returning to a refill location, and an extensive case study of the proposed framework over 33 growing seasons. Furthermore, several constraints in the framework have been adapted in order to significantly reduce computation time.

D. Structure

The remainder of this article is structured as follows. We end this section with mathematical notations. In Section II,

[1]The reason why we require the small set of assumptions will become clear in Section V. In short, the assumptions allow us to use the simulation models without doing an excessive amount of simulations at each time for all fields on a farm. This results in a scalability of our framework.

we pose a mathematical abstraction of the problem formulation described in Section I-A. Constraints on the irrigation machinery (as mentioned in Section I-B) are given in Section III, followed by a model of soil water and soil water/crop interaction in Section IV. The proposed solution to the formulated problem (the optimization framework) is described in Section V. In Section VI, we demonstrate the effectiveness of the proposed framework with a realistic case study using crop-growth models developed by agronomists.

E. Mathematical Notation

Let $\mathbb{N}_0$ denote the set of natural numbers including zero. Moreover, $\mathbb{N} := \mathbb{N}_0 \setminus \{0\}$, and $\mathbb{B} := \{0, 1\}$. If M is a matrix, then $[M]_{ij}$ is used to denote the element on the ith row and j-th column of M. We use $\mathbf{1}_n$ and $\mathbf{0}_n$ to denote vectors of size n consisting of all ones and all zeros, respectively. Likewise, $\mathbf{1}_{n\times m}$ and $\mathbf{0}_{n\times m}$ denote matrices of size n-by-m, consisting of all ones and all zeros, respectively. The n-by-n identity matrix is denoted by I_n. Furthermore, $\mathrm{Tr}(\cdot)$ denotes the trace of a matrix and $\mathrm{diag}(c) \in \mathbb{R}^{n\times n}$ denotes a diagonal matrix with the elements of vector $c \in \mathbb{R}^n$ on the diagonal. Operators "$=$," "$\leq$" or "$\geq$" in matrix (in)equalities denote element-wise comparisons.

II. System Context

In this section, we propose a mathematical abstraction of the problem formulation described in Section I-A. First, we introduce some notation and properties of the agents. Second, we introduce the objective function to be maximized by our framework.

A. Multi-Agent System Set-Up and Notations

We define $\mathcal{M} := \{1, 2, \ldots, m\}$ as the label set of (sub)fields, with m the total number of fields. Similarly, we define $\mathcal{N} := \{1, 2, \ldots, n\}$ as the label set of agents, with n the total number of irrigation agents. Typically, in real-world problems, $n \ll m$. Let T denote the known length of the growing season and hence any day t is in the set $\mathcal{T} := \{0, 1, \ldots, T-1\}$.

In addition to fields on which crops grow, agents can also be allocated to places for combined storage and refilling. We call these locations "stores" and label each store by an index from the set $\mathcal{V} := \{m+1, \ldots, m+\upsilon\}$, where $\upsilon \in \mathbb{N}$ is the total amount of stores. We use "point of interest" or "POI" to denote a field or a store. POIs are labeled from the set $\mathcal{M}^+ := \mathcal{M} \cup \mathcal{V} = \{1, 2, \ldots, m+\upsilon\}$. At each store $k \in \mathcal{V}$, an agent can refill water at a rate at most $r_{k-m} \in \mathbb{R}_{\geq 0}$.

It is assumed that each day $t \in \mathcal{T}$ is divided into $S \in \mathbb{N}$ time slots (or subtimes) in which agents can be allocated. Each agent is allocated only once per time slot. The time slots are labeled from the set $\mathcal{S} := \{1, 2, \ldots, S\}$ in chronological order. Crop growth dynamics typically has a time scale of days to weeks, whereas agents allocation can have time scales of hours. By introducing the time slots we separate these time scales in our set-up. That is, we can accommodate the fast allocation of agents, whilst leveraging on the fact that the crop states do not vary significantly on that time scale.

We define the allocation matrix $A_t^s \in \mathbb{B}^{n\times(m+\upsilon)}$ such that $[A_t^s]_{ij} = 1$ if and only if agent $i \in \mathcal{N}$ is allocated to POI $j \in \mathcal{M}^+$ in time slot $s \in \mathcal{S}$ of day $t \in \mathcal{T}$, otherwise $[A_t^s]_{ij} = 0$. Likewise, we define the delivery matrix $D_t^s \in \mathbb{R}_{\geq 0}^{n\times m}$ such that $[D_t^s]_{ij}$ denotes the amount of irrigation that agent $i \in \mathcal{N}$ delivers to field $j \in \mathcal{M}$ in time slot $s \in \mathcal{S}$ of day $t \in \mathcal{T}$. Furthermore, we define the daily allocation and delivery matrices as

$$A_t = \sum_{s\in\mathcal{S}} A_t^s \text{ and } D_t = \sum_{s\in\mathcal{S}} D_t^s$$

respectively, for $t \in \mathcal{T}$.

Agent capacity is captured by the vector $\bar{c} \in \mathbb{R}_{\geq 0}^n$, where the i-th element $\bar{c}_i$ of $\bar{c}$ denotes the maximum amount of water agent $i \in \mathcal{N}$ can store. Similarly, the i-th element $\bar{d}_i$ of $\bar{d} \in \mathbb{R}_{\geq 0}^n$ denotes the maximum amount of water that agent $i \in \mathcal{N}$ can deliver per slot. In other words, $\bar{d}$ can be used to represent the rate at which irrigation can be applied.

B. Objective

As mentioned in Section I-A, the objectives of a farmer can be of all sorts, including maximizing economic profit, minimizing number of irrigation agent deployments, or saving water. In our optimization framework, we model the objective function as an economic profit function, i.e., the total income over the season, minus the total costs incurred. Other perspectives can be incorporated by appropriately changing the weights of the income and costs terms in the objective function.

The income is related to the quantity and quality of the crops at the end of the season. The costs include the costs of water for irrigation and costs of agent operation such as fuel/energy and operator costs. We now show how these incomes and costs lead to the objective function of interest in this article.

First, let us look at the incomes due to harvest. Let $x_t^j \in \mathbb{R}^s$ denote the state vector[2] of a field $j \in \mathcal{M}$ at time $t \in \mathcal{T}$. In this work, we assume that the profit is determined by a linear combination of the elements in the state x_T^j at the end of the season (i.e., at time T). Let $\pi \in \mathbb{R}^s$ denote the weights such that

$$\pi^\top x_T^j$$

denotes the profit obtained from field $j \in \mathcal{M}$. Note that we use the state at time T, the end of the season, as we assume that the profit is obtained from the harvest.

Second, we look at the costs. We distinguish three main costs: operation costs of agents, resource costs and travel costs of the agents. For operation costs, we define $\sigma \in \mathbb{R}_{\geq 0}^n$ such that the ith element σ_i of σ denotes the costs of operation for agent $i \in \mathcal{N}$ per slot. This is the cost that needs to be paid when an agent is allocated to a field and can represent costs such as fuel/energy or labor costs (e.g., [26] for the importance of minimizing electricity costs). It is assumed that if an agent is allocated to a store, no operations costs are incurred. For the resource costs, we define $\rho \in \mathbb{R}_{\geq 0}$ to denote the cost per

[2]In the context of arable farming, the states typically include the biomass of different organs of the crop, the leaf size, and water content in the crop.

unit irrigation. Hence, the total amount of expenses over the growing season incurred by operation and resource costs can be written as

$$\sum_{t\in\mathcal{T}} \sigma^\top A_t \begin{bmatrix}\mathbf{1}_m\\ \mathbf{0}_\upsilon\end{bmatrix} + \rho\, \mathbf{1}_n^\top D_t \mathbf{1}_m.$$

For the travel costs, we define the matrix $E \in \mathbb{R}_{\geq 0}^{(m+\upsilon)\times(m+\upsilon)}$ such that $[E]_{jk}$ denotes the cost incurred from moving from field j to field k (with $j, k \in \mathcal{M}^+$). Hence, the total costs due to agent movements can be computed using[3]

$$\sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{M}^+}\sum_{k\in\mathcal{M}^+} \left[A_t^s\right]_{ij}\left[A_t^{s+1}\right]_{ik}[E]_{jk} = \sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}} \mathrm{Tr}\Big(A_t^s E \big(A_t^{s+1}\big)^\top\Big) \quad (1)$$

where we define $A_t^{S+1} = A_{t+1}^1$ for $t \in \mathcal{T}$. Furthermore, we have chosen to not take the very last travel cost into account and therefore define $A_T^{S+1} = \mathbf{0}_{n\times(m+\upsilon)}$.

Note that the matrix E is not necessarily symmetric, i.e., the cost of traveling from $i \in \mathcal{M}$ to $j \in \mathcal{M}$ does not have to be equal to the cost of traveling from j to i. Also, the diagonal of E can be comprised of only zeros since, technically, the agent does not move from one field to another in that case and does not incur costs. However, we can use a non-zero diagonal of E to penalize agents that remain stationary on a field without delivering resources. Note that it is assumed that the agents are able to move from any POI to any other POI within one time slot. Situations with infeasibly large movements are not prevented, but are less likely to occur due to the penalization in the objective function.

Combining these costs, we obtain the objective function J as

$$J = \pi^\top \sum_{j\in\mathcal{M}} x_T^j - \sum_{t\in\mathcal{T}}\Bigg(\sigma^\top A_t \begin{bmatrix}\mathbf{1}_m\\ \mathbf{0}_\upsilon\end{bmatrix} + \rho\, \mathbf{1}_n^\top D_t \mathbf{1}_m + \sum_{s\in\mathcal{S}} \mathrm{Tr}\Big(A_t^s E \big(A_t^{s+1}\big)^\top\Big)\Bigg). \quad (2)$$

C. Lay-Out of Proposed Solution

In essence, the objective is to maximize (2) subject to constraints on agent allocation and growth of the crops, in the presence of uncertainty about the state of our system, uncertainty in weather predictions, and other disturbances.

To tackle this problem, we will formalize the agent constraints in Section III, followed by a model of soil water and soil water/crop interaction in Section IV. In order to deal with the uncertainty in our system, the proposed framework will make use of principles from MPC. The adaptation needed to do so is described in Section V.

III. Agent Allocation Constraints

In this section, we will provide a mathematical formulation of several important constraints on agent allocation. Section III-A introduces essential constraints that are required for any use of the proposed framework. The other subsections contain additional constraints that may be added if desired.

[3]Note that for (1) to hold, it is assumed that an agent can be allocated to only one field at a certain slot. This will be enforced in (3).

A. Essential Allocation Constraints

As already stated in Section II-A, at each $s \in \mathcal{S}$, an agent must be allocated to exactly one POI. We, therefore, impose for all $i \in \mathcal{N}$ and all $s \in \mathcal{S}, t \in \mathcal{T}$

$$\sum_{j\in\mathcal{M}^+} \left[A_t^s\right]_{ij} = 1$$

or stated otherwise, for all $s \in \mathcal{S}, t \in \mathcal{T}$

$$A_t^s \mathbf{1}_{m+\upsilon} = \mathbf{1}_n. \quad (3)$$

A vital constraint in the agent allocation is that if an agent $i \in \mathcal{N}$ delivers resources to field $j \in \mathcal{M}$, it must be allocated. Formally stated

$$\left[D_t^s\right]_{ij} > 0 \Rightarrow \left[A_t^s\right]_{ij} = 1 \quad (4)$$

for all $s \in \mathcal{S}$ and $t \in \mathcal{T}$. Furthermore, each agent $i \in \mathcal{N}$ cannot deliver more than its maximum delivery rate $\bar{d}_i$

$$\left[D_t^s\right]_{ij} \leq \bar{d}_i. \quad (5)$$

We can enforce (4), (5) simultaneously by the following linear inequalities, for all $s \in \mathcal{S}, t \in \mathcal{T}, i \in \mathcal{N}$ and $j \in \mathcal{M}$:

$$\left[D_t^s\right]_{ij} \leq \bar{d}_i \left[A_t^s\right]_{ij}$$

or, equivalently, for all $s \in \mathcal{S}, t \in \mathcal{T}$

$$D_t^s \leq \mathrm{diag}\big(\bar{d}\big) A_t^s \begin{bmatrix} I_m \\ \mathbf{0}_{\upsilon\times m}\end{bmatrix}. \quad (6)$$

B. Refilling Constraints

Let $c_t^s \in \mathbb{R}_{\geq 0}^n$ denote the amount of water stored by each agent at slot $s \in \mathcal{S}$ of day $t \in \mathcal{T}$. In order to respect the carrying capacity of each agent it must be true that, for all $s \in \mathcal{S}, t \in \mathcal{T}$

$$\mathbf{0}_n \leq c_t^s \leq \bar{c}. \quad (7)$$

Furthermore, we assume that the dynamics of c_t^s for all $s \in \mathcal{S}, t \in \mathcal{T}$ equals

$$c_t^{s+1} = c_t^s - D_t^s \mathbf{1}_m + R_t^s \mathbf{1}_\upsilon \quad (8)$$

subject to (7), where we define $c_t^{S+1} := c_{t+1}^1$ and $R_t^s \in \mathbb{R}_{\geq 0}^{n\times\upsilon}$ is a decision variable such that $[R_t^s]_{i(k-m)}$ denotes how much water is refilled to agent $i \in \mathcal{N}$ at store $k \in \mathcal{V}$ in slot $s \in \mathcal{S}$ on day $t \in \mathcal{T}$. Similar to (4), we must have that an agent must be allocated to a store in order for it to receive water and, hence, for all $i \in \mathcal{N}$, $k \in \mathcal{V}$, $s \in \mathcal{S}$ and $t \in \mathcal{T}$

$$\left[R_t^s\right]_{i(k-m)} \leq r_{k-m}\left[A_t^s\right]_{ik} \quad (9)$$

where $r_{k-m} \in \mathbb{R}_{>0}$ is the maximum refill rate per allocation for store $k \in \mathcal{V}$. This can be written in matrix form as

$$R_t^s \leq A_t^s \begin{bmatrix} \mathbf{0}_{m\times\upsilon} \\ \mathrm{diag}(r)\end{bmatrix}.$$

IV. Model of Water Limited Crop Growth

So far, we have dealt with the agent allocation side of the optimization framework. In this section, we will shift our attention to the crop growth side of the framework. First, this section presents the soil water dynamics, and second, the soil water/crop interaction.

A. Soil Water Dynamics

Let w_t^j denote the soil water content of field $j \in \mathcal{M}$ at time $t \in \mathcal{T}$. We assume the (unsaturated) water dynamics occurs according to the following discrete-time dynamics:

$$w_{t+1}^j = w_t^j + d_t^j - \upsilon_t^j + p_t^j \tag{10}$$

where $d_t^j = \sum_{i \in \mathcal{N}} [D_t]_{ij}$ is the amount of water field j receives at time t, υ_t^j is a disturbance term that includes effects such as surface run-off, transpiration, and drainage, and p_t^j is the precipitation. We obtain υ_t^j from the prediction model and weather predictions, and p_t^j from weather predictions/measurements.

We assume that the soil for each field has a maximum carrying capacity for water called field capacity (fc). It is assumed that if w_t^j exceeds fc^j, this amount of water will be lost due to deep drainage within one day and the water level will reduce to $w_t^j = \mathrm{fc}^j$. Therefore, in effect, we adapt the unsaturated dynamics (10) to

$$w_{t+1}^j = \min\Big\{\mathrm{fc^j},\ \mathrm{w_t^j + d_t^j - v_t^j + p_t^j}\Big\}. \tag{11}$$

B. Soil Water/Crop Interactions

We assume that the prediction model G for the crop state x_t^j is of the form

$$x_{t+1}^j = x_t^j + G\Big(x_t^j, w_t^j, d_t^j, \upsilon_t^j, p_t^j, \tau_t^j, i_t^j\Big)$$

for $t \in \mathcal{T}$, where $\tau_t^j \in \mathbb{R}$ is the daily average temperature and $i_t^j \in \mathbb{R}_{\geq 0}$ the total daily solar irradiation. Furthermore, in order to model the soil water/crop interaction, we assume that G can be decomposed as the product of a water-limiting factor $\theta_t^j \in [0, 1]$ and the potential growth G^* as

$$G(\cdot) = \theta_t^j\Big(w_t^j, \mathrm{cr}_t^j\Big) \cdot G^*\Big(x_t^j, \tau_t^j, i_t^j\Big).$$

Here, $\mathrm{cr}_t^j(x_t^j, \upsilon_t^j, p_t^j, \tau_t^j, i_t^j)$ is the critical water level which is dependent on the evapo-transpiration (which is dependent on crop states and weather influences such as precipitation, temperature, and irradiation), and G^* is the maximum amount of growth that is theoretically possible in the case where there is no water shortage, i.e., the potential growth. Note that this implies that G^* is not a function of d_t^j. The dependency of θ_t^j on d_{t-1}^j is through w_t^j by (11).

Remark 1: As an illustrative example of a G^*, let us consider the crop growth model LINTUL2. In this model, the optimal growth is given by $G^* = (1/2)\lambda\, i_t^j (1 - e^{-\kappa L_t})$, where λ is the light-use efficiency in $\mathrm{g\,MJ^{-1}}$, i_t^j is the total daily solar irradiation in MJ, κ is an attenuation coefficient, and L_t is the leaf area index (LAI).

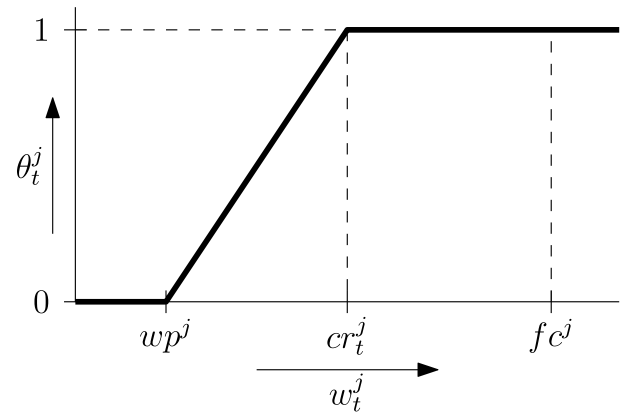

Fig. 1. Dependency of growth-reduction factor θ_t^j on the soil water content w_t^j for any $t \in \mathcal{T}$ and $j \in \mathcal{M}$.

In order to compute the optimal actions more efficiently, we now exploit a key principle of soil water/crop interactions that is common to many crop growth models. This principle is used in many different crop growth models such as (but not limited to) LINTUL2/3 [5], WOFOST [6], STICS [7], AquaCrop [8], and a derivation thereof in [27]. In these models (as well as in our optimization framework), the water limiting factor is dependent on w_t^j in the manner as shown in Fig. 1.

In Fig. 1, wp^j is the wilting point below which growth does not occur. It holds that $\mathrm{wp}^j < \mathrm{cr}_t^j < \mathrm{fc}^j$ for all $j \in \mathcal{M}$, $t \in \mathcal{T}$. The key principle mentioned earlier is that the growth is optimal for water levels above the critical water level and reduces linearly for values below the critical water level until the wilting point.

Any crop growth model that satisfies the following assumption can be used in our framework as a prediction model.

Assumption 1 (Crop growth model for prediction): The crop growth model used in prediction can output the maximum possible growth G^* under optimal irrigation (i.e., potential growth), the critical water level cr_t^j and soil water effects υ_t^j, for all $t \in \mathcal{T}$.

V. Model Predictive Control Scheme

Optimization of (2) subject to the allocation constraints of Section III before the growing season starts, is generally a computationally demanding problem. Furthermore, the growth of crops is highly susceptible to changes in weather and other disturbances. It is mainly for these reasons that the proposed framework incorporates ideas from MPC (e.g., [28]). Instead of an entire growing season, the optimization only considers a finite horizon of $H \in \mathbb{N}$ days. Within these H days, the optimization framework computes the optimal amount of irrigation and which agent delivers water to which field.

In order to take into account a shrinking horizon when $t \in \mathcal{T}$ is close to T, we make use of $H_t := \min\{H, T - t\}$ instead of H in the remainder. Let $\mathcal{H}_t^H := \{t, t+1, \ldots, t+H_t-1\}$ denote the label set of prediction steps in the (receding) horizon for time $t \in \mathcal{T}$. In the formulation of the optimization problem, we use the subscript "$h|t$" (e.g., $D_{h|t}$) to denote the value of a parameter at time $h \in \mathcal{H}_t^H$, when the optimization is performed at $t \in \mathcal{T}$.

In the remainder of this section, we will first adapt the objective function (2) such that it can be used in the MPC setting, followed by the required adaptations to the allocation constraints.

A. Adaptation of the Objective Function

We adapt (2) to the following objective function that is to be maximized at each time $t \in \mathcal{T}$:

$$J_t = \pi^\top \sum_{j\in\mathcal{M}} x^j_{T|t} - \frac{T-t}{H_t} \times \sum_{h\in\mathcal{H}^H_t}\Bigg(\sigma^\top A_{h|t}\begin{bmatrix}\mathbf{1}_m\\ \mathbf{0}_\upsilon\end{bmatrix} + \rho\,\mathbf{1}_n^\top D_{h|t}\mathbf{1}_m + \sum_{s\in\mathcal{S}}\mathrm{Tr}\Big(A^s_{h|t}E\Big(A^{s+1}_{h+1|t}\Big)^\top\Big)\Bigg) \tag{12}$$

where we define $A^{S+1}_{h|t} := A^1_{h+1|t}$ and $A^{S+1}_{t+H|t} := \mathbf{0}_{n\times(m+\upsilon)}$.

Since we are optimizing over a finite horizon, we have made the following two modifications: 1) a prediction $x^j_{T|t}$ at time t for the yield at the end of the season and 2) a multiplication of $(T-t)/H_t$ of the resource costs in the prediction horizon. We now elaborate on each of these choices separately in Sections V-A1 and V-A2.

1) Prediction of the Yield at the End of the Season: If we are optimizing over a finite horizon that does not include the end of the season, then our approach requires an estimate of $x^j_{T|t}$, the state at the end of the season. We will use the following insights in order to obtain such an estimate.

First, for all $j \in \mathcal{M}, h \in \mathcal{H}^H_t$, let $\tilde{x}^j_h$ be such that $\tilde{x}^j_t = x^j_t$ and $\tilde{x}^j_{h+1} = \tilde{x}^j_h + G^*(\tilde{x}^j_h, \tau^j_{h|t}, i^j_{h|t})$. The maximum possible growth (i.e., there is a complete absence of water stress) for field $j \in \mathcal{M}$ during the considered prediction horizon is then equal to

$$G^j_\circ = \sum_{h\in\mathcal{H}^H_t} G^*\Big(\tilde{x}^j_h, \tau^j_{h|t}, i^j_{h|t}\Big).$$

Second, let $x^j_{t,\min}$ denote the prediction at $t \in \mathcal{T}$ of the state at T if there would be no growth during the prediction horizon and optimal growth afterward: $\theta^j_h = 0$, $h \in \mathcal{H}^H_t$ and $\theta^j_\iota = 1$, $\iota \in \{t+H_t, t+H_t+1, \ldots, T\}$. Similarly, let $x^j_{\max}$ denote the maximal possible yield, which can be achieved in complete absence of waterstress during and after the prediction horizon, i.e., $\theta^j_\iota = 1$, $\iota \in \{t, t+1, \ldots, T\}$.

The proposed estimate of $x^j_{T|t}$ is then linear in the $\theta^j_{h|t}$ for $h \in \mathcal{H}^H_t, j \in \mathcal{M}$

$$x^j_{T|t} = x^j_{t,\min} + \Big(x^j_{t,\max} - x^j_{t,\min}\Big)\frac{\sum_{h\in\mathcal{H}^H_t}\theta^j_{h|t}G^*\Big(\tilde{x}^j_h, \tau^t_{h|t}, i^j_{h|t}\Big)}{G^j_\circ}. \tag{13}$$

Remark 2: For a further interpretation of this, let us reconsider the crop growth model LINTUL2 as given in Remark 1. Assuming that the LAI does not change much during the prediction horizon, then

$$x^j_{T|t} \approx x^j_{t,\min} + \Big(x^j_{t,\max} - x^j_{t,\min}\Big)\frac{\sum_{h\in\mathcal{H}^H_t}\theta^j_{h|t}Q_{h|t}}{G^j_\circ}.$$

The fraction can therefore be interpreted as an average of the reduction in growth, weighted according to the solar irradiation. If we furthermore assume that $Q_{h|t}$ does not change much in the prediction horizon, then the fraction is approximately equal to the average $\theta^j_{h|t}$ over the prediction horizon.

2) Scaling of Resource Costs: Since we are interested in the profit over the entire season, we must not only give a prediction of the yield at the end of the season, but also a prediction of the resource costs in the remainder of the season. The assumption we used in the adapted cost (12) is that the resource costs during the prediction horizon are indicative of the resource costs over the entire remainder of the growing season. In fact, we assume that the daily average within the prediction horizon is equal to the daily average over the remainder of the season. Hence, the resource cost during the prediction horizon is multiplied by $(T-t)/H_t$. Note that if $t \in \mathcal{T}$ is such that $T-t \le H$, then the factor $(T-t)/H$ should be equal to one and hence the introduction of H_t in (12).

B. Adaptation of Constraints

1) Allocation and Refilling Constraints: The adaptation of the essential allocation constraints is rather straightforward and we obtain

$$D^s_{h|t} \le \mathrm{diag}\big(\bar{d}\big)A^s_{h|t}\begin{bmatrix}I_m\\ \mathbf{0}_{\upsilon\times m}\end{bmatrix} \tag{14a}$$

$$A^s_{h|t}\mathbf{1}_{m+\upsilon} = \mathbf{1}_n \tag{14b}$$

for all $h \in \mathcal{H}^H_t$ and $s \in \mathcal{S}$.

As with the essential allocation constraints, the adaptation of the refilling constraints (7)–(9) follows readily and we have, for all $s \in \mathcal{S}, h \in \mathcal{H}^H_t$:

$$c^{s+1}_{h|t} = c^s_{h|t} - D^s_{h|t}\mathbf{1}_m + R^s_{h|t}\mathbf{1}_\upsilon \tag{15a}$$

$$\mathbf{0}_n \le c^s_{h|t} \le \bar{c} \tag{15b}$$

$$\mathbf{0}_n \le c^1_{H+1|t} \le \bar{c} \tag{15c}$$

$$R^s_{h|t} \le A^s_{h|t}\begin{bmatrix}\mathbf{0}_{m\times\upsilon}\\ \mathrm{diag}(r)\end{bmatrix} \tag{15d}$$

where $c^1_{t|t} = c_{t-1}$ (i.e., the current inventory) and furthermore

$$c^{S+1}_{h|t} = c^1_{h+1|t} \tag{16a}$$

$$D^{S+1}_{h|t} = D^1_{h+1|t} \tag{16b}$$

$$R^{S+1}_{h|t} = R^1_{h+1|t}. \tag{16c}$$

2) Water Dynamics: In order to model (11), we use the substitution of w^j_t to $w^j_{h|t}$ for $t \in \mathcal{T}, h \in \mathcal{H}^H_t$ and $j \in \mathcal{M}$. However, for ease of computation, we also replace the min-function by a set of linear constraints. For brevity of notation, let $z = w^j_{h|t} + d^j_{h|t} + p^j_{h|t} - \upsilon^j_{h|t}$, then (11) (with the aforementioned substitutions), can be rewritten as

$$w^j_{h+1|t} = \min\big\{z, fc^j\big\} = z - \underbrace{\max\big\{z - fc^j, 0\big\}}_{=:\epsilon^j_{h|t}}. \tag{17}$$

The max-function can be modeled alternatively by introducing the three constraints

$$\epsilon_{h|t}^{j} \geq 0 \tag{18a}$$

$$\epsilon_{h|t}^{j} \geq z - fc^{j} \tag{18b}$$

$$\epsilon_{h|t}^{j}\left(\epsilon_{h|t}^{j} - z + fc^{j}\right) = 0. \tag{18c}$$

The complementarity condition (18c) ensures that either $\epsilon_{h|t}^{j} = 0$ or $\epsilon_{h|t}^{j} = z - \text{fc}^{j}$. Interestingly, (18c) can be removed in the optimization problem, while still satisfying $\epsilon_{h|t}^{j} = \max\{0, z - \text{fc}^{j}\}$, as explained next. The water level $w_{h|t}^{j}$ has a linear effect on the objective function through $\theta_{h|t}^{j}$ and $x_{T|t}^{j}$. A larger $\epsilon_{h|t}^{j}$ would yield a lower water level, which can only have negative effects on the crop growth in the long term (see Fig. 1). Hence, we set $\epsilon_{h|t}^{j}$ to be optimization variables in our framework and remove the complementarity constraint (18c) as the optimal solution will have the smallest value of $\epsilon_{h|t}^{j}$ that satisfies (18a), (18b), thereby automatically satisfying (18c) without having to impose it. This value for $\epsilon_{h|t}^{j}$ will be equal to the true maximum. This effect has been verified for the discussed framework and no violations of these constraints were observed in the simulation results presented in Section VI. In summary, we model the water dynamics as follows:

$$w_{h+1|t}^{j} = w_{h|t}^{j} + d_{h|t}^{j} + p_{h|t}^{j} - v_{h|t}^{j} - \epsilon_{h|t}^{j} \tag{19a}$$

$$\epsilon_{h|t}^{j} \geq w_{h|t}^{j} + d_{h|t}^{j} + p_{h|t}^{j} - v_{h|t}^{j} - fc^{j} \tag{19b}$$

$$\epsilon_{h|t}^{j} \geq 0. \tag{19c}$$

3) Soil Water/Crop Interactions: In order to model $\theta_{h|t}^{j}$ for $w_{h|t}^{j} \geq \text{wp}^{j}$, we use a similar approach as in Section V-B2 above. It is evident from Fig. 1 that, for $w_{h|t}^{j} \geq \text{wp}^{j}$, it holds that

$$\theta_{h|t}^{j} = \min\left\{1, \frac{w_{h|t}^{j} - \text{wp}^{j}}{\text{cr}_{h|t}^{j} - \text{wp}^{j}}\right\} = 1 - \frac{\overbrace{\max\left\{0, \text{cr}_{h|t}^{j} - w_{h|t}^{j}\right\}}^{\eta_{h|t}^{j} :=}}{\text{cr}_{h|t}^{j} - \text{wp}^{j}}$$

for all $t \in \mathcal{T}$, $h \in \mathcal{H}_{t}^{H}$ and $j \in \mathcal{M}$. Similar to how we replaced the max-function in (17) with $\epsilon_{h|t}^{j}$, we introduce the auxiliary variable $\eta_{h|t}^{j}$ to represent $\max\{0, \text{cr}_{h|t}^{j} - w_{h|t}^{j}\}$. This can be written to be identical to the constraints

$$\eta_{h|t}^{j} \geq 0 \tag{20a}$$

$$\eta_{h|t}^{j} \geq \text{cr}_{h|t}^{j} - w_{h|t}^{j}. \tag{20b}$$

and the complementarity constraint

$$\eta_{h|t}^{j}\left(\eta_{h|t}^{j} - \text{cr}_{h|t}^{j} + w_{h|t}^{j}\right) = 0. \tag{21}$$

Note that $\theta_{h|t}^{j}$ appears linearly in the objective function through $x_{T|t}^{j}$. A smaller value of $\theta_{h|t}^{j}$, and thus a larger value of $\eta_{h|t}^{j}$, has a reducing effect on the objective function. Hence, the optimal solution will have the smallest $\eta_{h|t}^{j}$ satisfying (20) and thus satisfy (21) automatically. Hence, we can remove the complementarity constraint without changing the optimal solution. Then, for $w_{t}^{j} \geq \text{wp}^{j}$

$$\theta_{h|t}^{j} = 1 - \frac{\eta_{h|t}^{j}}{\text{cr}_{h|t}^{j} - \text{wp}^{j}}. \tag{22}$$

The critical water level $\text{cr}_{h|t}^{j}$ is dependent on the evapotranspiration of the crop and, hence, it is dependent on several crop states and weather effects such as temperature and solar irradiation. This would imply that (22) is a non-linear constraint. However, we can find a representable value of the critical water level within the prediction horizon by using $\tilde{x}_{j}^{t}$, i.e., by replacing $\text{cr}_{h|t}^{j}$ in (22) by $\text{cr}_{t}^{j}(\tilde{x}_{h}^{j}, v_{t}^{j}, p_{t}^{j}, \tau_{t}^{j}, i_{t}^{j})$. This substitution makes the constraint linear ($\theta_{h|t}^{j}$ and $\eta_{h|t}^{j}$ are the only remaining decision variables). In cases where it is unlikely that the water level reaches levels below the wilting point (i.e., $w_{h|t}^{j} < \text{wp}^{j}$), the constraints (20), (22) suffice to model the dependence of $\theta_{h|t}^{j}$ on $w_{h|t}^{j}$ as shown in Fig. 1.

C. Implementation of Control Framework

The workings of the control framework are summarized in the following algorithm.

Algorithm 1 (Irrigation control framework): For each time $t \in \mathcal{T}$, the following tasks are to be performed in order.

1) Obtain the current states x_{t}^{j} of all fields $j \in \mathcal{M}$.
2) Obtain the weather predictions required for the prediction model (temperature, irradiation, and precipitation).
3) Compute, for all $j \in \mathcal{M}$, $G_{\circ}^{j}$, $x_{t,\min}^{j}$, $x_{t,\max}^{j}$, $v_{h|t}^{j}$, and $\text{cr}_{h|t}^{j}$ (see Section V-A1).
4) Maximize (12), subject to (13)–(16), (19), (20), and (22).
5) Execute according to $A_{t|t}^{s}$, $D_{t|t}^{s}$ and $R_{t|t}^{s}$, $s \in \mathcal{S}$.
6) Repeat for the next day.

The optimization problem in step 4 is a mixed-integer bilinear problem (MIBP). It becomes a mixed-integer linear problem (MILP) if we do not take into account agent travel costs (i.e., E is the zero matrix).

VI. Performance Analysis

A. Case Study Over 33 Years

The framework as described in Section V was implemented into MATLAB [29] and the optimization problem was solved on a notebook with a 2.60-GHz Intel Core i7 processor and 8 GB of RAM, using SCIP [30] and YALMIP [31]. We used the validated and realistic LINTUL2 [5], [32], [33] as the prediction model. We consider the crop spring wheat and the crop parameters used are the standard, validated ones that come with the LINTUL2 model [32]. The "real" dynamics of the field were simulated also using LINTUL2. However, several soil/water parameters (such as fc^{j} and wp^{j}) of the real fields were perturbed (at most 10%) from the values of the prediction model in order to introduce heterogeneity in the field and model mismatch. Further heterogeneity in the fields is introduced by selecting *a priori* a constant growth reducing/improving factor per field as opposed to the ideal setting given by the prediction model. These factors range from 0.8 to 1.2 with an average of 1.0. Both the growth reducing factors and the change in parameters are unknown to the optimization solver. Weather data are obtained from 1966 to 1998 (33 seasons) and the location is Eelde, The Netherlands [32]. The weather used in the prediction step of the framework is perturbed from the weather data. Similar as in [25], this weather perturbation is done using a first-order linear autoregressive model with the same lag 0 and lag 1 correlations between the weather states as the historical data, see [34].

The system of consideration in this case study consists of $N = 5$ agents and $M = 100$ fields of equal area, and $\upsilon = 1$ store. Agents can be allocated $S = 2$ times per day and the prediction horizon is $H = 12$. Furthermore, $T = 160$, $\pi = 500$, $\rho = 1, \sigma = 500\mathbf{1}_n, \bar{d} = 100$. We consider x_t^j to be the mass of the storage organ ("weight storage organ" or "WSO" in LINTUL2) to represent the yield at the end of the season.

The objective function could represent the economic profit and ρ could have, for instance, the unit of \$/liter water (similar for π and σ). However, the objective is not only strictly financial as some parameter tuning can be done in order to balance the estimate of the "true" economic profit and other effects such as penalizing excessive agent movements beyond economic reasons. Introducing currency can therefore be misleading and is therefore omitted here.

We compare our framework in this setting to the following heuristic, which is based on [35]. If the water level is below the critical water level (i.e., $w_t^j \leq \text{cr}$), then 30 mm of water is applied to field $j \in \mathcal{M}$ at day $t \in \mathcal{T}$. If this happens to more than nS fields, then the nS fields with the largest WSO at that time are irrigated. We do not take into account the travel costs (i.e., $E = \mathbf{0}_{m+\upsilon \times m+\upsilon}$) and refilling constraints in this case study in order to have a fair comparison between the proposed framework and the heuristic.

B. Results

Over the 33 considered seasons, the value of the objective function is on average 5.31% higher when using the MPC compared to the heuristic, with a maximum of 18.3% in 1967. The objective function for the MPC was smaller than the heuristic in three seasons, but only by a negligible amount, namely 0.8% in 1972, 0.3% in 1990, and 0.7% in 1995. Note that the heuristic does not take into account all allocation constraints and violations in maximum delivery rate were present. Fig. 2 shows an overview of the objective function increase of the MPC scheme and Table I shows the performance compared over the years.

Fig. 3 shows several more details of the results. On average over the 33 seasons: the WSO is 5.29% larger with the MPC compared to the heuristic, while the MPC uses 5.17% less water. These are significant improvements in the domain of agriculture.

The bottom two graphs in Fig. 3 also show the "water use efficiency" η_w and the "irrigation moment efficiency" η_i. We have defined these to be equal to, respectively

$$\eta_w = \frac{\sum_{j\in\mathcal{M}} x_T^j}{\mathbf{1}_n^\top \sum_{t\in\mathcal{T}} D_t \mathbf{1}_m} \text{ and } \eta_i = \frac{\sum_{j\in\mathcal{M}} x_T^j}{\mathbf{1}_n^\top \sum_{t\in\mathcal{T}} A_t \begin{bmatrix} \mathbf{1}_m \\ \mathbf{0}_\upsilon \end{bmatrix}}$$

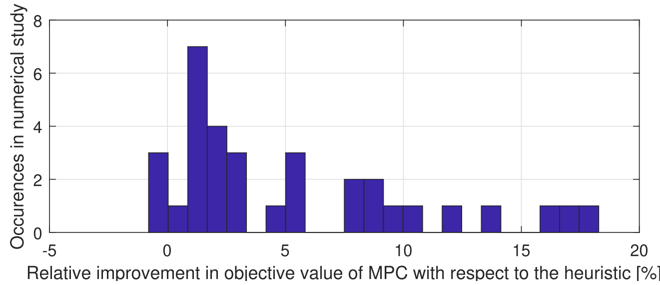

Fig. 2. Histogram of the relative increase of the objective function value of the MPC relative to the heuristic.

TABLE I

ANNUAL AVERAGE PERFORMANCE COMPARISON FOR 100 FIELDS, FIVE AGENTS, ONE STORE

	Mean	SD	Min.	Max.
Objective value MPC [$\times 10^4$]	36.6	4.7	23.5	48.0
Objective value Heuristic [$\times 10^4$]	34.8	4.4	23.4	44.4
Yield MPC [tons ha^{-1}]	7.4	0.95	4.7	9.7
Yield Heuristic [tons ha^{-1}]	7.0	0.88	4.7	8.9
Irrigation MPC [mm]	109.5	54.4	11.8	232.4
Irrigation Heuristic [mm]	110.5	46.0	24	209
Agent allocations MPC	3.8	1.66	0.88	9.1
Agent allocations Heuristic	3.7	1.53	0.80	7.0
Objective value MPC/Heuristic	1.053	0.054	0.992	1.183
Irrigation MPC/Heuristic	0.948	0.162	0.490	1.228

for each year. That is, they are the total WSO divided by the total amount of irrigation and the total WSO divided by the total number of agent allocation. On average over the 33 seasons, the water use efficiency for the MPC is 11.64% higher than for the heuristic and the irrigation moment efficiency is 11.7% higher. Calculation times for the MPC were on average 3 min and at most 20 min for a single day, which allows for practical implementation on a daily basis.

The year with the most irrigation and most agent allocations (irrigation moments) is 1996, which is also the year with the least amount of precipitation. In 1972, 1979, and 1998, the opposite was true: there was a lot of rain and hence only a handful of irrigation moments were required. In those years, the performance of the MPC is only marginally better than the heuristic. In general, the relative benefit of the MPC is low in extremely wet years and high in extremely dry years.

C. Traveling Costs

In order to demonstrate the traveling costs, we have done a numerical study over five seasons (1977–1981). We simulated the case of $N = 2$ agent and $M = 25$ fields. Furthermore, $\pi = 100, \rho = 1, \sigma = 10\,\mathbf{1}_n$, and the remaining parameters are the same as discussed in Section VI-A. We studied the case in which the 25 fields are arranged in a 5-by-5 grid and that the distances are given based on their central locations. For example, field 1 and field 8 are a distance $(1^2 + 2^2)^{1/2}$ apart since field 8 is one row up and two columns further. The distance from the fields to the virtual field is assumed to be equal to 2 and $[E]_{ii} = 0$ for all $i \in \mathcal{M}^+$.

For each season, we used our framework to optimize the objective function both with and without the traveling costs, in parallel. Table II shows several results. It is evident that including the traveling costs induces an additional punishment on the number of agent allocations (i.e., besides the costs

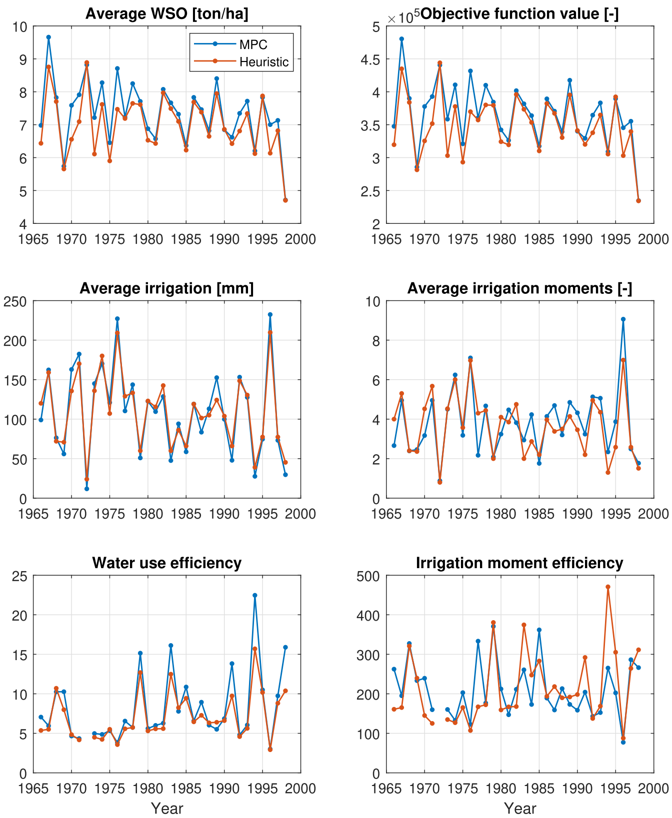

Fig. 3. Simulation results over various years. Here, "average" denotes the average over all fields. There was a lot of rain in 1972 and hence only a handful of irrigation moments were required. For visual clarity, the values for the efficiencies were omitted in this graph for 1972. The values for water use efficiency are 74.94 and 37.04 for the MPC and heuristic, respectively. The values for the irrigation moment efficiency are 1 002 and 1 111 for the MPC and heuristic, respectively.

TABLE II

PERFORMANCE COMPARISON WHEN TAKING INTO ACCOUNT TRAVELING COSTS RELATIVE TO NOT TAKING TRAVELING COSTS INTO ACCOUNT. ANALYSIS FOR 25 FIELDS, TWO AGENTS, ONE STORE

Year	Agent allocations	Travel costs	Objective w/ travel costs
1977	−13.2%	−12.2%	+1.91%
1978	−14.3%	−22.2%	+0.97%
1979	+2.1%	−9.5%	+0.88%
1980	−17.0%	−6.3%	−0.69%
1981	+4.5%	−15.0%	+3.23%
Mean	−7.6%	−13.0%	+1.26%

induced by σ). This also shows in the results: in three out of five seasons there were significantly less agent allocations when taking traveling costs into account. In all five seasons, the traveling costs where indeed significantly reduced in the optimizations with the traveling costs in the objective function. On average, the traveling costs were reduced by 13%. The last column of Table II shows the relative differences in the values of the objective function. For the optimizations that did not take into account the traveling costs, we computed the traveling costs *a posteriori* and corrected the objective value for these costs. The objective value when taking the travel costs into account in the optimization was on average 1.26% larger. Calculation times per day were on average 16.8 s when not taking the travel costs into account (MILP) and 55.1 s when we do take them into account (MIBP). The maximum times were 180 and 236 s respectively.

VII. CONCLUSION AND FUTURE WORK

In this work, we introduced a problem of irrigation scheduling in large-scale arable farming, together with a mathematical formalization of this problem amendable for online optimization. In particular, we introduced a model predictive control framework in order to translate this problem into a real-time feasible formulation. This framework is practical to use as it can make use of a wide variety of crop growth model and it allows many real-world constraints on the agent dynamics and crop growth to be included. Although this article is concerned with the delivery of water, the proposed framework could be generalized to other resources as well.

There are several interesting directions for further research. First, a sensitivity analysis on the effects of the weather disturbances, crop growth model parameters, and MPC framework parameters are of great interest to further assess the robustness of the framework. Second, it is interesting to compare the suggestions of the MPC framework to historical data of farming practices. Both for the validation of our framework as well as to investigate whether we can derive better heuristics for farmers that do not use the framework. Finally, we assume full-state information on the crops. It is interesting to study whether the presented framework can be adapted to be used for the allocation of sensing agents for state estimation.

Furthermore, we presented a realistic case study based on a well-known crop growth model (LINTUL2) and real weather data over 33 growing seasons, in order to demonstrate the effectiveness of the introduced framework. We introduced both parameter uncertainties and weather uncertainties and compared our framework to a state-of-the-art heuristic. For the chosen parameters, we were able to obtain a 5.31% larger value of the objective function on average. Our proposed framework was able to generate 5.29% more yield, while using 5.17% less water. These are significant improvements that help in tackling two global trends: to reduce water use in agriculture whilst providing more food.

REFERENCES

[1] "The United Nations world water development report 2019," UNESCO World Water Assessment Programme, Paris, France, Tech. Rep., 2019.

[2] "The future of food and agriculture: Trends and challenges," Food Agricult. Org. United Nations, Rome, Italy, Tech. Rep., 2017.

[3] M. K. van Ittersum, P. A. Leffelaar, H. van Keulen, M. J. Kropff, L. Bastiaans, and J. Goudriaan, "On approaches and applications of the Wageningen crop models," *Eur. J. Agronomy*, vol. 18, nos. 3–4, pp. 201–234, Jan. 2003.

[4] S. M. Siad, V. Iacobellis, P. Zdruli, A. Gioia, I. Stavi, and G. Hoogenboom, "A review of coupled hydrologic and crop growth models," *Agricult. Water Manage.*, vol. 224, Sep. 2019, Art. no. 1057461.

[5] M. E. Shibu, P. A. Leffelaar, H. van Keulen, and P. K. Aggarwal, "LINTUL3, a simulation model for nitrogen-limited situations: Application to rice," *Eur. J. Agronomy*, vol. 32, no. 4, pp. 255–271, May 2010.

[6] A. de Wit *et al.*, "25 years of the WOFOST cropping systems model," *Agricult. Syst.*, vol. 168, pp. 154–167, Jan. 2019.

[7] N. Brisson *et al.*, "An overview of the crop model STICS," *Eur. J. Agronomy*, vol. 18, nos. 3–4, pp. 309–332, Jan. 2003.

[8] P. Steduto, T. C. Hsiao, D. Raes, and E. Fereres, "AquaCrop—The FAO crop model to simulate yield response to water: I. Concepts and underlying principles," *Agron. J.*, vol. 101, no. 3, p. 426, 2009.

[9] N. B. Rossello, R. F. Carpio, A. Gasparri, and E. Garone, "A novel observer-based architecture for water management in large-scale (Hazelnut) orchards," *IFAC-PapersOnLine*, vol. 52, no. 30, pp. 62–69, 2019.

[10] S. Asseng, Y. Zhu, and E. Wang, "Crop modeling for climate change impact and adaptation," in *Crop Physiology*, V. O. Sadras and D. F. Calderini, Eds., 2nd ed. Amsterdam, The Netherlands: Elsevier, 2015, ch. 20, pp. 505–546.

[11] L. García, L. Parra, J. M. Jimenez, J. Lloret, and P. Lorenz, "IoT-based smart irrigation system: An overview on the recent trends on sensors and IoT systems for irrigation in precision agriculture," *Sensors*, vol. 20, no. 4, p. 1042, 2020.

[12] R. J. Smith, M. H. Gillies, G. Newell, and J. P. Foley, "A decision support model for travelling gun irrigation machines," *Biosyst. Eng.*, vol. 100, no. 1, pp. 126–136, May 2008.

[13] J. Rolim and J. L. Teixeira, "The design and evaluation of travelling gun irrigation systems: Enrolador software," *Engenharia Agrícola*, vol. 36, no. 5, pp. 917–927, Oct. 2016.

[14] N. Kalboussi, S. Roux, K. Boumaza, C. Sinfort, and A. Rapaport, "About modeling and control strategies for scheduling crop irrigation," *IFAC-PapersOnLine*, vol. 52, no. 23, pp. 43–48, 2019.

[15] J. Baayen *et al.*, "An overview of continuation methods for non-linear model predictive control of water systems," *IFAC-PapersOnLine*, vol. 52, no. 23, pp. 73–80, 2019.

[16] S. K. Saleem *et al.*, "Model predictive control for real-time irrigation scheduling," *IFAC Proc. Volumes*, vol. 46, no. 18, pp. 299–304, Aug. 2013.

[17] A. C. McCarthy, N. H. Hancock, and S. R. Raine, "Simulation of irrigation control strategies for cotton using model predictive control within the VARIwise simulation framework," *Comput. Electron. Agricult.*, vol. 101, pp. 135–147, Feb. 2014.

[18] C. Lozoya *et al.*, "Model predictive control for closed-loop irrigation," *IFAC Proc. Volumes*, vol. 47, no. 3, pp. 4429–4434, 2014.

[19] C. Shang, W. H. Chen, A. D. Stroock, and F. You, "Robust model predictive control of irrigation systems with active uncertainty learning and data analytics," *IEEE Trans. Control Syst. Technol.*, vol. 28, no. 4, pp. 1493–1504, Jul. 2019.

[20] R. R. Negenborn, P.-J. van Overloop, T. Keviczky, and B. De Schutter, "Distributed model predictive control of irrigation canals," *Netw. Heterog. Media*, vol. 4, no. 2, pp. 359–380, 2009.

[21] A. Sadowska, P.-J. van Overloop, J. M. Maestre, and B. De Schutter, "Human-in-the-loop control of an irrigation canal using time instant optimization model predictive control," in *Proc. Eur. Control Conf. (ECC)*, Jul. 2015, pp. 3274–3279.

[22] Y. Ding, L. Wang, Y. Li, and D. Li, "Model predictive control and its application in agriculture: A review," *Comput. Electron. Agricult.*, vol. 151, pp. 104–117, Aug. 2018.

[23] Y. Chevaleyre *et al.*, "Issues in multiagent resource allocation," *Informatica*, vol. 30, no. 1, pp. 3–31, Jan. 2006.

[24] A. T. J. R. Cobbenhagen, D. J. Antunes, M. J. G. van de Molengraft, and W. P. M. H. Heemels, "Heterogeneous multi-agent resource allocation through multi-bidding with applications to precision agriculture," *IFAC-PapersOnLine*, vol. 51, no. 23, pp. 194–199, 2018.

[25] L. P. A. Schoonen, A. T. J. R. Cobbenhagen, and W. P. M. H. Heemels, "Optimal irrigation management for large-scale arable farming using model predictive control," *IFAC-PapersOnLine*, vol. 52, no. 30, pp. 56–61, 2019.

[26] D. Kiedanski, D. Kofman, A. Orda, J. Horta, and A. Otero, "Exploiting flexibility in irrigation while maintaining optimal crop productivity," in *Proc. IEEE Int. Conf. Commun., Control, Comput. Technol. Smart Grids*, Oct. 2019, pp. 1–7.

[27] N. Pelak, R. Revelli, and A. Porporato, "A dynamical systems framework for crop models: Toward optimal fertilization and irrigation strategies under climatic variability," *Ecol. Model.*, vol. 365, pp. 80–92, Dec. 2017.

[28] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.

[29] *MATLAB, 9.3.0.713579 (R2017b)*, MathWorks, Natick, MA, USA, 2017.

[30] A. Gleixner *et al.*, "The SCIP optimization suite 5.0," ZIB, Berlin, Germany, Tech. Rep. 17-61, 2017.

[31] J. Löfberg, "YALMIP: A toolbox for modeling and optimization in MATLAB," in *Proc. CACSD Conf.*, Taipei, Taiwan, 2004, pp. 284–289.

[32] P. A. Leffelaar and J. Wolf. (2011). *LINTUL-2, a Simple Crop Growth Model for Both Potential and Water Limited Growing Conditions*. [Online]. Available: http://models.pps.wur.nl/node/944

[33] A. J. Haverkort, A. C. Franke, J. M. Steyn, A. A. Pronk, D. O. Caldiz, and P. L. Kooman, "A robust potato model: LINTUL-POTATO-DSS," *Potato Res.*, vol. 58, no. 4, pp. 313–327, Dec. 2015.

[34] A. Shahin, S. M. AbouRizk, and Y. Mohamed, "A weather generator for use in construction simulation models," *Int. J. Archit., Eng. Construct.*, vol. 2, no. 2, pp. 73–87, Jun. 2013.

[35] P. Janssens, "Grotere oogst na irrigeren," *Boerenbond, Manage. Techniek*, vol. 18, pp. 18–19, 2015.

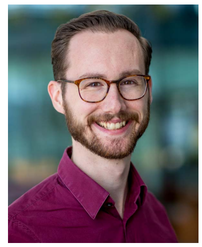

A. T. J. R. (Roy) Cobbenhagen received the B.Sc. degree *(cum laude)* in mechanical engineering, the M.Sc. degree *(cum laude)* in automotive technology, and the Ph.D. degree in control systems from Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands, in 2013, 2016, and 2020, respectively.

Since completing his Ph.D. degree on optimal multi-agent allocation with applications to precision agriculture, he has been a VP Research with Lightyear, Helmond, The Netherlands, a producer of long-range solar electric vehicles, where he leads research on vehicle integrated photovoltaics, efficient electric vehicles, design optimization, and autonomous driving.

Dr. Cobbenhagen received the Award for Best M.Sc. Thesis of the TU/e Automotive Technology Program 2016.

L. P. A. (Luc) Schoonen received the B.Sc. degree in mechanical engineering and the M.Sc. degree (Hons.) in mechanical engineering from Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands, in 2017 and 2019, respectively.

His research interests include seed propagation and breeding for robust open-pollinated varieties, and agroforestry development. His aim is to take a holistic approach and develop sustainable and scalable farming strategies that produce healthy products without the use of agrochemicals and inorganic fertilizers.

M. J. G. (René) van de Molengraft received the M.Sc. degree *(cum laude)* in mechanical engineering and the Ph.D. degree from Eindhoven University of Technology, Eindhoven, The Netherlands, in 1986 and 1990, respectively.

In 1991, he returned to the control group at Mechanical Engineering, after having fulfilled his military service. He is currently an Associate Professor and a Group Lead of the Robotics Group, Control Systems Technology Section, Eindhoven. He founded the Tech United RoboCup Team, and gained five world champion titles in the MSL robot soccer league in 2012, 2014, 2016, 2018, and 2019, and a world champion title in the @Home carerobot league in 2019.

W. P. M. H. (Maurice) Heemels (Fellow, IEEE) received the M.Sc. degree *(summa cum laude)* in mathematics and the Ph.D. degree *(summa cum laude)* in control theory from Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands, in 1995 and 1999, respectively.

From 2000 to 2004, he was with the Electrical Engineering Department, TU/e, as an Assistant Professor, and from 2004 to 2006 with the Embedded Systems Institute (ESI) as a Research Fellow. Since 2006, he has been with the Department of Mechanical Engineering, TU/e, where he is currently a Full Professor. He held as a Visiting Professor position at Swiss Federal Institute of Technology (ETH Zürich), Zürich, Switzerland, in 2001, the University of California at Santa Barbara, Santa Barbara, CA, USA, in 2008, and the University of Lorraine, Nancy, France, in 2020.

Dr. Heemels was a recipient of a Personal VICI Grant awarded by Dutch Research Council (NWO). He has been the Chair of the IFAC Technical Committee on Networked Systems since 2017. He has served/s on the Editorial Board of *Automatica*, *Nonlinear Analysis: Hybrid Systems*, *Annual Reviews in Control*, and IEEE TRANSACTIONS ON AUTOMATIC CONTROL.