Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 52-30 (2019) 56–61

Optimal irrigation management for large-scale arable farming using model predictive control ⋆

L.P.A. Schoonen * A.T.J.R. Cobbenhagen * W.P.M.H. Heemels *

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

Abstract: The productivity and financial success of large-scale arable farming operations depends highly on how effectively resources are distributed among fields. Therefore, it is of interest to develop methods to determine (near optimal) resource inputs. In this paper we focus on computing the optimal irrigation policy for large-scale arable farming operations, where the number of irrigation machinery is much smaller than the number of fields that require irrigation inputs. We propose a model predictive control (MPC) framework that simultaneously computes the optimal division of irrigation over the fields and which irrigation machines should be allocated to which fields, such that the profit at the end of the season is maximized. The fact that the optimization of irrigation and allocation of irrigation machinery is done simultaneously makes our approach vastly different from strategies available in the literature. Another important novelty of our work is that we link short-term effects of crop growth to long-term effects on profit. The proposed framework has reasonable computation times when optimizing on a daily basis over many fields and irrigation machines and guarantees a feasible solution. The main principles of our approach are more widely applicable. Using simulations we demonstrate the robustness of the scheme with respect to changes in weather and benchmark it with respect to a heuristic approach.

Keywords: Irrigation; Resource allocation; Large-scale Arable farming; Model predictive control; Optimization

1. INTRODUCTION

As stated by United Nations (2019), agriculture is responsible for about 69% of annual water withdrawals globally and is therefore the largest water consumer. In particular, irrigation has a large contribution to this water consumption. Proper irrigation can be crucial for crop survival, and could potentially double or triple yields. Therefore, it is crucial to manage this water effectively, such that water use is minimized, while good yields are maintained.

In this paper we focus on optimal irrigation management for large-scale arable farming. For large-scale arable farming the challenge lies in deciding when, where and how much water should be applied to the various (sub)fields. There is often a large variety in crop states and soil conditions between the various (sub)fields and a limited number of irrigation machinery is available. Optimal decision making techniques can improve the effectiveness of the applied irrigation, while satisfying constraints on equipment limitations. These decision making techniques have the potential to increase overall productivity of a farm and/or decrease the total resource usage.

Studies on decreasing the dependency on irrigation inputs focus on two main aspects. On the one hand the focus lies on increasing the water productivity and design of more drought tolerance crops, as presented by Kijne et al. (2003), Spitters and Schapendonk (1990) and El-Wahed and Ali (2013). On the other hand, the focus lies on better managing irrigation inputs; Robert et al. (2016) list an overview of decision making models in agriculture. In particular, Saleem et al. (2013) and Lozoya et al. (2014) show how model predictive control can be used for closed-loop and real-time irrigation control, trying to minimize water deficit and irrigation inputs. McCarthy et al. (2014) shows a model predictive control strategy that takes long-term effects on yield into account when controlling short-term irrigation inputs, although the method remains unclear. Furthermore, studies have been performed to determine crop response to water stress for different growing stages, aiming for better irrigation scheduling, as presented by Zhang and Oweis (1999) and Sun et al. (2006).

In general for irrigation management, 1) there is a lack of a formal and clear mathematical description of the control method, 2) the long-term effects are often not linked to short-term irrigation inputs, or the method is not clearly described, 3) it is often assumed that water levels can be controlled on a daily basis, making these methods unsuitable for day-to-day management of large-scale arable farming operations, where the number of

⋆ The authors acknowledge the support of 'Toeslag voor Topconsortia voor Kennis en Innovatie' (TKI HTSM) from the Ministry of Economic Affairs, the Netherlands.

2405-8963 Δ 2019, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2019.12.489

irrigation systems is often at least an order of magnitude smaller than the number of (sub)fields, and 4) scaling to a large number of fields results in large computation times. In the literature cited above, at least one of the four main issues applies.

As a solution to these issues, we consider optimal irrigation management for large-scale arable farming as a multi-agent resource allocation (MARA) problem. The goal for such problems is to allocate *agents* that deliver resources, to *clients* that produce products from the received resources. In the context of this paper, the agents represent irrigation machinery and clients represent arable farm fields. The main contribution is a framework that can be used to compute the optimal short-term actions for resource distributions over a large number of clients with a limited number of available agents, in order to maximize profit over the season. Although the main framework is applicable to many situations, we focus on allocating irrigation machinery. This framework, based on model predictive control (MPC) ideas (see e.g., Mayne (2014) for a survey), has the following key properties: 1) short-term irrigation actions are computed, while optimizing long-term profit, 2) irrigation machine limitations are taken into account and the framework provides a feasible allocation, 3) the framework can handle heterogeneity between fields, but also heterogeneity in irrigation machines, 4) the framework is robust against errors in weather predictions, and 5) the framework shows reasonable computation times when optimizing over many (sub)fields and agents, such that decision making can occur with a high resolution. The benefits of this framework are demonstrated with extensive simulations.

The remainder of this paper is structured as follows. In Section 2 the LINTUL2 simulation model is introduced, which is used in the control framework. Subsequently, a formal description of the problem.Section 3 presents the proposed framework to solve this problem. The performance of this controller is demonstrated using simulation examples in Section 4. Finally, we provide conclusions and recommendations for future work.

The following mathematical notations will be used. Let $\mathbb{N}$ denote the set of natural numbers (including zero), $\mathbb{B}$ the binary set defined as $\mathbb{B} := \{0, 1\}$, and $\mathbb{R}$ the set of real numbers. If M is a matrix, then $[M]_{ij}$ denotes the entry at row i and column j of M and $[M]_j$ denotes column j of M. If v is a vector, $[v]_j$ denotes the j-th element of v. We use $\mathbf{1}_n$ and $\mathbf{0}_n$ to denote the vectors of size n that contain all ones and all zeros, respectively. Operators '=', '$\leq$' or '$\geq$' in matrix equations denote element-wise comparisons.

2. SYSTEM CONTEXT AND PROBLEM FORMULATION

2.1 Crop growth dynamics in LINTUL2

For the crop growth and water dynamics, we use the LINTUL2 simulation model[1]. LINTUL2 calculates total daily biomass production of crops on field level and takes weather influences and water stress effects into account. This simulation model can handle a variety of crops that grow storage organs, but for illustrative purposes, we use an application to spring wheat. In this section we present the basic assumptions, states and state equations of the LINTUL2 model that are used to compute an optimal irrigation policy.

Total biomass production: Let $\mathcal{T} := \{1, 2, ..., T\}$ denote the set of time steps (days) of the growing season, with T the last day of the growing season and let $t \in \mathcal{T}$ denote a specific day of the growing season. LINTUL2 calculates total daily biomass production $G(t)$ (in $\mathrm{g\,m^{-2}}$) at time $t \in \mathcal{T}$, based on intercepted irradiation in the following way:

$$G(t) = 0.5 \cdot LUE \cdot Q_0(t) \cdot \left(1 - e^{-k \cdot LAI(t)}\right) \cdot T_{red-tran}(t), \tag{1}$$

where LUE is the 'light use efficiency' in $\mathrm{g\,MJ^{-1}}$, used as a way of modelling photosynthesis by assuming that the conversion from light into biomass occurs at a constant ratio (it is crop specific and assumed to be constant in this paper), $Q_0(t)$ is the total irradiation in $\mathrm{MJ\,m^{-2}\,d^{-1}}$, $LAI(t)$ is the leaf area index in $\mathrm{m^2\,m^{-2}}$ and k is a crop specific attenuation coefficient in $\mathrm{m^2\,m^{-2}}$. Lastly, $T_{red-tran}(t) \in [0, 1]$ is used to model the influence of water stress to the biomass growth.

Water dynamics: The transpiration reduction factor $T_{red-tran}(t)$ is modelled as a relation between the actual water level (wa in mm) and some important bounds on the water level. Figure 1 shows this model. Transpiration reduction ranges from 0 (i.e., full transpiration reduction) to 1 (i.e., no transpiration reduction). Between the critical water level for optimal growth (cr) and wet soil (wet), plants can transpire optimally and growth is optimal. Field capacity (fc) lies between cr and wet. LINTUL2 assumes that all water content above field capacity drains within one day and is therefore considered to be lost. Below the critical bound, the transpiration reduction factor goes linearly to zero, which is indicated as the wilting point (wp). The same holds above the wet level, until the soil is completely saturated (sat). When the water level drops below the wilting point, plants can not transpire and it is assumed that growth stops completely.

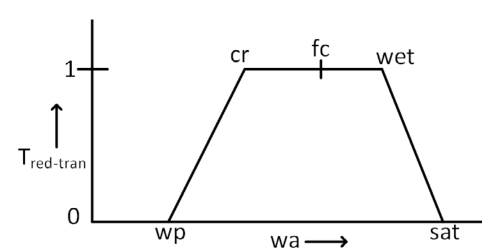

Fig. 1. Transpiration reduction factor ($T_{red-tran}$) as a function of the total soil water amount (wa). Here wp stands for wilting point, cr for critical water level, fc for field capacity, wet for wet soil and sat for completely saturated soil.

The actual water level as computed in the LINTUL2 model is naturally influenced by rainfall, intercepted rain, runoff, evaporation, transpiration, drainage and exploration. Every factor is influenced by the weather, but some also

[1] LINTUL2 is available via Wolf (2011), including historic weather data from 1966-1999, Eelde, The Netherlands. LINTUL2 is based on Spitters and Schapendonk (1990) and a predecessor of LINTUL3, where nitrogen stress effects are added, Shibu et al. (2010).

depend on crop states and the actual soil water level. The modelled dynamics for the soil water amount is

$$wa(t+1) = wa(t) + u(t) + d(t). \tag{2}$$

Hence, the water amount behaves as an integrator, with $u(t)$ the irrigation input and $d(t)$ the natural disturbance to the water level.

Biomass partitioning and final yield: The total daily biomass production as described in (1) is partitioned between the various organs of the crop: the leafs, stems, roots and storage organs. In this work, we measure yield by the mass of the storage organs (note that the control procedure proposed in this paper is applicable whenever the yield is defined as any linear combination of the model states). The final yield, y_T in $\mathrm{g\,m^{-2}}$, is therefore equal to

$$y_T = \sum_{t=1}^{T} p_{so}(t) G(t), \tag{3}$$

where $p_{so}(t) \in [0,1]\,, t \in \mathcal{T}$ is the fraction of growth that is allocated to the storage organs, which changes throughout the various growing stages.

2.2 Multi-agent resource allocation (MARA)

We define $\mathcal{M} := \{1, 2, ..., m\}$ as the label set of clients, with m the total number of (sub)fields and $\mathcal{N} := \{1, 2, ..., n\}$ the label set of agents, with n the total number of agents. Every client has the dynamics as described previously in Section 2.1. For resource allocation, we define matrix $A_t \in \mathbb{B}^{n \times m}$ as the allocation matrix and $D_t \in \mathbb{R}^{n \times m}$ as the resource delivery matrix. If $[A_t]_{pj} = 1$, agent $p \in \mathcal{N}$ delivers resources to client $j \in \mathcal{M}$ at time t, otherwise $[A_t]_{pj} = 0$. Element $[D_t]_{pj} \in \mathbb{R}_{\geq 0}$ denotes the quantity of resources (i.e., water) delivered by agent $p \in \mathcal{N}$ to client $j \in \mathcal{M}$ at time t.

We define $C_{max} \in \mathbb{R}^n$, $Q_a \in \mathbb{N}^n$ and $Q_c \in \mathbb{N}^m$ such that $[C_{max}]_p$ denotes the maximum quantity of resources agent $p \in \mathcal{N}$ can deliver, $[Q_a]_p$ the maximum number of clients agent p can serve per time step and $[Q_c]_j$ the maximum number of agents that may serve client $j \in \mathcal{M}$ per time step. Furthermore, $U_{max,t} \in \mathbb{R}$ is the maximum amount of available resources and $R^j_{max,t} \in \mathbb{R}_{\geq 0}$ the maximum amount of resources client j may receive.

From here on 'clients' will be referred to as '(sub)fields' since they represent arable farm fields in the context of this paper. We will use 'agents' and 'irrigation machinery' interchangeably.

2.3 Profit function

The MARA problem is to find a policy for daily irrigation distributions that maximizes the profit, while considering system limitations. Similarily to Cobbenhagen et al. (2018), we define the following profit function,

$$J = \sum_{j=1}^{m} \pi\, y_T^j - \sum_{t=1}^{T} (\sigma^\top A_t + \rho^\top D_t) \mathbf{1}_m, \tag{4}$$

which describes the profit over the entire growing season. Here $\pi \in \mathbb{R}$, $\sigma \in \mathbb{R}^n$ and $\rho \in \mathbb{R}^n$ are used to apply costs on yield, agent operation and resource deliveries, respectively. Hence, (4) is the difference between the estimated profit from yield summed over all fields j and the agent operation costs and resource delivery costs summed over all time steps t.

3. PROPOSED OPTIMIZATION FRAMEWORK

3.1 Model predictive control (MPC)

Maximizing (4) at the start of the season, such that we obtain an open-loop optimal irrigation policy for the entire growing season, requires some crucial assumptions on weather predictions and leaves the farmer with limited flexibility during the growing season. Solving the problem for the entire growing season for a large number of agents and (sub)fields, would also result in large computational efforts, making the problem possibly infeasible to solve.

In order to guarantee a feasible solution and to take actual crop states and weather predictions into account, we use a model predictive control strategy with a limited prediction horizon of $N \in \mathbb{N}$ days and implement this with receding horizon, such that optimal decisions are made on a daily basis, with the knowledge of system states and weather predictions for that specific day.

In the remainder of this section we will demonstrate how the profit function was adjusted, how we modelled the crop growth dynamics, and which allocation constraints were added. This leads to (18), the formalization of our optimiztion framework.

3.2 Adjusted profit function

For the MPC framework, we adjust profit function (4) to

$$J_t = \sum_{j=1}^{m} \pi\, y_t^j - \sum_{i=t}^{t+N-1} (\sigma^\top A_{i|t} + \rho^\top D_{i|t}) \mathbf{1}_m, \tag{5}$$

such that y_t^j is a prediction at time $t \in \mathcal{T}$ of the final yield of field $j \in \mathcal{M}$, $D_{i|t}$ and $A_{i|t}$ are the predicted resource deliveries and agent allocations at time $t \in \mathcal{T}$ for time i in the prediction horizon $\mathcal{H}_t^N := \{t, t+1, ..., t+N-1\}$. The receding horizon implementation means that at each time $t \in \mathcal{T}$, we compute the optimal allocation with respect to J_t, such that we obtain an optimal irrigation sequences $A_{i|t}$ and $D_{i|t}$ for $i \in \mathcal{H}_t^N$, and implement only the first step (i.e., $A_t = A_{t|t}$ and $D_t = D_{t|t}$).

3.3 Field states and yield prediction

Water dynamics: As mentioned in Section 2.1, two bounds on the water level are considered between which the crop can exhibit optimal growth and no water is lost due to excessive drainage (see Figure 1): the critical water level for optimal growth (cr) and field capacity (fc).

Let $\epsilon^j_{1,i|t}, \epsilon^j_{2,i|t} \in \mathbb{R}_{\geq 0}$ be the violation of the critical bound and field capacity, respectively. Ideally $\epsilon^j_{1,i|t} = \max(0,\ cr^j_{i|t} - wa^j_{i|t})$, indicating how far the water level is below the critical bound. In the optimization framework this is modelled as

$$\epsilon^j_{1,i|t} \geq cr^j_{i|t} - wa^j_{i|t}, \tag{6}$$

such that $\epsilon^j_{1,i|t}$ is not strictly constrained. The value of $\epsilon^j_{1,i|t}$ is therefore not unique. However, the solver will chose the lowest value since $\epsilon^j_{1,i|t}$ has a negative effect on the yield prediction and appears linearly in the profit function via (10) and (11). Similarly,

$$\epsilon^j_{2,i|t} \geq wa^j_{i|t} - fc^j_{i|t} \tag{7}$$

holds for the bound on field capacity. This is used to prevent irrigating above field capacity, which is taken care of via (8) and the cost on resource deliveries in the profit function. The values of $cr^j_{i|t}$, $fc^j_{i|t}$, and the natural disturbances $d^j_{i|t}$ are predicted with the LINTUL2 model, given the current state of the (sub)fields and weather predictions. The water dynamics of (sub)fields $j \in \mathcal{M}$ in (2) are rewritten in terms of optimization parameters as

$$wa^j_{i+1|t} = wa^j_{i|t} + u^j_{i|t} + d^j_{i|t} - \epsilon^j_{2,i|t}. \tag{8}$$

Crop growth dynamics: Crop growth as described in (1) can be split into two parts: maximum potential daily growth and water stress (i.e., transpiration reduction). Maximum potential growth depends on the daily total radiation Q_0 and the leaf area index LAI and we assume that both of these values can be predicted for a relatively short prediction horizon (i.e., $N \leq 10$). It is therefore assumed that irrigation inputs during the prediction horizon have limited influence on the short-term change in LAI. We define $G^j_{max,i|t}$ as the predicted value of maximum daily growth for (sub)fields j (i.e., in absence of water stress):

$$G^j_{max,i|t} = 0.5 \cdot LUE \cdot Q^j_{0,i|t} \cdot \left(1 - e^{-k \cdot LAI^j_{i|t}}\right). \tag{9}$$

The values of $G^j_{max,i|t}$ for $i \in \mathcal{H}^N_t$, $j \in \mathcal{M}$ are predicted at time $t \in \mathcal{T}$, with the assumption that growth is not limited by water stress effects.

We only consider water stress as a result of water shortage, i.e., we assume that $wa(t) \leq wet$, for all $t \in \mathcal{T}$. If the water level drops below the critical bound cr, $T_{red-tran}$ goes linearly to zero (Figure 1). Hence, we can use (6) to express $T^j_{red-tran,i|t}$ as

$$T^j_{red-tran,i|t} = 1 - \frac{\epsilon^j_{1,i|t}}{cr^j_{i|t} - wp^j_{i|t}},$$

where $\epsilon^j_{1,i|t}$ is an optimization parameter and $cr^j_{i|t}$ and $wp^j_{i|t}$ are time-varying parameters determined using LINTUL2. This allows us to simplify (1) to

$$G^j_{i|t} = G^j_{max,i|t}\left(1 - \frac{\epsilon^j_{1,i|t}}{cr^j_{i|t} - wp^j_{i|t}}\right), \tag{10}$$

with optimization parameters $G^j_{i|t}$ and $\epsilon^j_{1,i|t}$. The values of $G^j_{max,i|t}$, $cr^j_{i|t}$ and $wp^j_{i|t}$ are estimated for $i \in \mathcal{H}^N_t$ at time $t \in \mathcal{T}$ with the LINTUL2 model, given the current field states and a weather forecast. For these predictions it is assumed that $T_{red-tran} = 1$ (i.e., no water stress).

Prediction of the final yield: A key characteristic of our optimization scheme is that the water limited growth during the prediction horizon can be used to make a prediction of the yield at the end of the season in a linear fashion. Specifically, we introduce the approximation rule

$$y^j_t = y^j_{min,t} + (y^j_{max,t} - y^j_{min,t})\frac{\sum_{i=t}^{t+N-1} G^j_{i|t}}{\sum_{i=t}^{t+N-1} G^j_{max,i|t}}, \tag{11}$$

where the value of $y^j_{min,t}$ is the predicted yield if there would be no growth during prediction horizon and $y^j_{max,t}$ is the predicted yield if growth would be optimal during the prediction horizon, both estimated with the LINTUL2 model. For both these predictions it is assumed that growth is optimal for the remainder of the growing season (after the prediction horizon). The value of the fraction in (11) is the relative growth during prediction horizon N. The reasoning why we can estimate the final yield based on a linear relation with short-term growth originates from the LINTUL2 model and its validity is further demonstrated in Appendix A.

3.4 Constraints on multi-agent resource allocation

The following allocation constraints are imposed in the framework. Firstly, the total amount of resources delivered to (sub)field j by all agents should equal the resource inputs of (sub)field j:

$$[A_{i|t}]_j^\top [D_{i|t}]_j = u^j_{i|t}. \tag{12}$$

Secondly, the sum of the total resource deliveries should satisfy agent capacity:

$$D_{i|t}\mathbf{1}_m \leq C_{max}, \tag{13}$$

Thirdly, the number of (sub)fields one agent can serve during one time step is assumed to be limited and we constrain the number of agents that may serve a (sub)field as

$$A_{i|t}\mathbf{1}_m \leq Q_a, \tag{14}$$

$$A^\top_{i|t}\mathbf{1}_n \leq Q_c. \tag{15}$$

Lastly, the total amount of resource deliveries of all agents (e.g., in case there is a limited supply) and the total resource inputs per (sub)field (e.g., if required by regulations) is constrained by, respectively,

$$\sum_{i=t}^{t+N-1} \mathbf{1}^\top_n D_{i|t}\mathbf{1}_m \leq U_{max,t}, \tag{16}$$

$$\sum_{i=t}^{t+N-1} u^j_{i|t} \leq R^j_{max,t}. \tag{17}$$

3.5 Optimization framework

The complete optimization problem for MARA becomes a mixed integer bilinear program (MIBP) with a linear cost function:

$$\boxed{\begin{aligned} \max_{\substack{wa^j_{i|t}, G^j_{i|t}, \epsilon^j_{1,i|t},\\ \epsilon^j_{2,i|t}, u^j_{i|t}, y^j_t, A_{i|t}, D_{i|t}}} \quad & \sum_{j=1}^{m} \pi\, y^j_t - \sum_{i=t}^{t+N-1} (\sigma^\top A_{i|t} + \rho^\top D_{i|t})\mathbf{1}_m \\ \text{subject to } & (6) - (17) \end{aligned}} \tag{18}$$

Here $wa^j_{i|t}$, $G^j_{i|t}$, $\epsilon^j_{1,i|t}$, $\epsilon^j_{2,i|t}$, y^j_t, $A_{i|t}$, $D_{i|t}$, and $u^j_{i|t}$ are the decision variables.

The MPC controller proposed in this paper is implemented with a receding horizon, which means that (18) is solved at every time step $t \in \mathcal{T}$ and then $u^j_t = u^{j*}_{t|t}$ is executed for all (sub)fields j, where '$*$' denotes an optimal maximizer to (18). At every iteration the time-varying parameters ($d^j_{i|t}$,

$cr^j_{i|t}$, $fc^j_{i|t}$, $G^j_{max,i|t}$, $y^j_{min,t}$, and $y^j_{max,t}$) are estimated with the LINTUL2 model given a certain weather forecast and current field states.

4. SIMULATION RESULTS

The performance of the proposed controller is demonstrated using simulations in MATLAB where SCIP (see Gleixner et al. (2017)) was used to solve the optimization. For demonstration purposes, we use the LINTUL2 for the field dynamics. For these simulations full state knowledge of the field states was assumed and initial conditions and crop specific parameters are not changed from the standard settings of the provided LINTUL2 model. The performance is compared to a heuristic controller that applies 25 mm of water when the water level drops below the critical bound. This heuristic is based on Janssens (2015) and personal communication with van den Borne (2019). In order to cover a large set of possible scenarios, the performance is evaluated over weather data from 1966 to 1999 in Eelde, The Netherlands.

The performance of the MPC and heuristic strategies are evaluated with respect to (4). We assume that all fields are 1 ha and that there is homogeneity in initial conditions and crop parameters for the various fields. The price of spring wheat is set to €0.175 kg^{-1}, the irrigation delivery cost is set to €2 mm^{-1} ha^{-1} and operation costs are €100 per allocation (Smale (2017); Janssens (2015)).

The following two simulation experiments have been performed: a single field and a single agent (Section 4.1), and 30 fields with only 3 agents (Section 4.2). Each experiment has been evaluated over the 33 years of the weather data set. A prediction horizon of $N = 10$ days was used for the MPC, as we assume that weather predictions during this time frame are fairly accurate. Outside of this prediction horizon, the daily averages over all years are used as a prediction for the weather in order to predict y_{min} and y_{max}. To demonstrate robustness with respect to weather predictions, we apply a random disturbance to all weather variables, taking *lag0 correlations* between daily total radiation, wind speed, vapour pressure, and temperature into account, similar to the method described by Shahin et al. (2013).

4.1 Performance comparison for a single field

For the proposed MPC method, (18) was evaluated daily, with $m = 1$, $n = 1$, $N = 10$, $\pi = 9$, $\sigma = 100$, $\rho = 2$, $Q_a = Q_c = 1$ and $C_{max} = 30$ for all 33 years in the data set. Notice that we use a different value for π as the value we use for evaluating the final profit. This choice is made since the benefit of irrigation past the prediction horizon is not accounted for in the optimization framework. Using $\pi = 9$ shows good performance over a wide range of weather data, when evaluating the profit over the complete season, this was the only instance of parameter tuning necessary in the entire MPC framework.

Table 4.1 shows the results for 1 field (note that the values for the averages and standard deviations are rounded).The result of 1996 was omitted for the computation of relative profit, as the heuristic controller made a negative profit. The table shows that the proposed controller with these specific cost parameters makes a good trade-off between yield and irrigation inputs and could potentially improve profitability by 44% on average. When final yield is more important, a higher value for π can be chosen.

Table 1. Performance comparison for 1 field.

	Mean	SD	Min.	Max.
Profit MPC [€]	838	284	206	1398
Profit Heuristic [€]	657	279	-7	1398
Yield MPC [tons ha^{-1}]	7.5	1.1	5.5	9.3
Yield Heuristic [tons ha^{-1}]	8.1	8.0	6.6	9.4
Irrigation MPC [mm]	85	54	0	180
Irrigation Heuristic [mm]	126	54	0	250
Irrigation moments MPC	3	1.9	0	6
Irrigation moments Heuristic	5	2.1	0	10
Relative profit MPC/Heuristic	1.44	0.78	0.77	5.5

4.2 Performance comparison for 30 fields and 3 agents

We now consider the situation where $m = 30$ fields of 1 ha that need to be managed by $n = 3$ agents. When more irrigation is needed than the overall capacity, the heuristic controller applies irrigation to the fields that had the least water stress during previous days of the growing season (i.e., they have the highest yield potential, when considering homogeneous fields). For the MPC method we use $Q_a = \mathbf{1_n}$, $Q_c = \mathbf{1_m}$, $U_{max,t} = \infty$, and $R^j_{max} = \infty$, and the remaining parameters have the same values as described in Section 4.1.

Table 4.2 shows the result for the 33 years (note that the values for the averages and standard deviations are rounded) and Figure 2 shows the histogram of the relative profit. For the relative profit, the results of 1976 were omitted since the heuristic controller made a negative profit. For the chosen control parameters, the proposed controller uses significantly less irrigation and irrigation moments on average, while the average yield is only 5% less. Also, with the chosen control parameters, it shows a potential average increase in profit of 19%. The average computation time of one optimization step was 105 seconds, with a maximum of 1082 seconds, which is well within an acceptable range when optimizing on a daily basis.

Table 2. Performance comparison for 30 fields with 3 agents

	Mean	SD	Min.	Max.
Profit MPC [€]	26099	8299	5736	42485
Profit Heuristic [€]	22405	8721	-1495	41176
Yield MPC [tons ha^{-1}]	7.0	1.1	3.9	9.0
Yield Heuristic [tons ha^{-1}]	7.4	0.8	5.7	8.8
Irrigation MPC [mm]	1922	1541	0	5679
Irrigation Heuristic [mm]	2779	1278	225	5275
Irrigation moments MPC	65.2	52.5	0	194
Irrigation moments Heuristic	111.1	51.1	9	211
Profit MPC/Heuristic	1.19	0.16	0.94	1.61

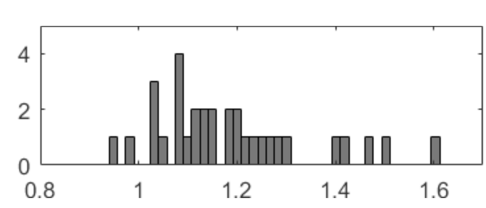

Fig. 2. Histogram of relative profit for 30 fields

5. CONCLUSIONS AND FUTURE WORK

This paper introduced a scalable MPC controller for irrigation scheduling of a large-scale arable farm in order to maximize profit at the end of the growing season. The field dynamics are based on the well known LINTUL2 simulation model and constraints are added for proper agent allocation. With simulations we have demonstrated the potential influence on profit of the proposed controller, when we compare this with a commonly used heuristic policy. Also, the main framework for resource distributions is applicable to many other situations.

Full state knowledge was assumed and hence it is of importance to further research what states can be measured directly and how to design observers for the remaining states of interest. The sensitivity of the optimization procedure with respect to the cost parameters and estimations made with the LINTUL2 model should also be studied. Inclusion of additional costs/constraints such as agent inventory and travel distance are reserved for future work. This will add to an even more complete and realistic framework of a large-scale farm. Adding more constraints could potentially increase the computational effort and thus it may be interesting to research methods that efficiently solve the optimization problem (e.g. a distributed framework).

REFERENCES

Cobbenhagen, A.T.J.R., Antunes, D.J., van de Molengraft, M.J.G., and Heemels, W.P.M.H. (2018). Heterogeneous multi-agent resource allocation through multi-bidding with applications to precision agriculture. *IFAC-PapersOnLine*, 51(23), 194–199.

El-Wahed, M.H.A. and Ali, E.A. (2013). Effect of irrigation systems, amounts of irrigation water and mulching on corn yield, water use efficiency and net profit. *Agricultural Water Management*, 120, 64–71.

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, Takustr.7, 14195 Berlin.

Janssens, P. (2015). Grotere oogst na irrigeren. *Boerenbond, Management&Techniek*, 18.

Kijne, J.W., Barker, R., and Molden, D. (2003). *Water Productivity in Agriculture: Limits and Opportunities for Improvement*. CABI.

Lozoya, C., Mendoza, C., Mejia, L., Quintana, J., Mendoza, G., Bustillos, M., Arras, O., and Solis, L. (2014). Model predictive control for closed-loop irrigation. *IFAC Congress*, 47, 4429–4434.

Mayne, D.Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, 50(12), 2967 – 2986.

McCarthy, A.C., Hancock, N.H., and Raine, S.R. (2014). Simulation of irrigation control strategies for cotton using model predictive control within the VARIwise simulation framework. *Computers and electronics in agriculture*, 101, 135–147.

Robert, M., Thomas, A., and Bergez, J.E. (2016). Processes of adaptation in farm decision-making models. a review. *Agronomy for Sustainable Development*, 36–64.

Saleem, S.K., Delgoda, D.K., Ooi, S.K., Dassanayake, K.B., Liu, L., Halgamuge, M.N., and Malano, H. (2013). Model predictive control for real-time irrigation scheduling. *IFAC Agricontrol*, 46, 299–304.

Shahin, A., Abourizk, S.M., and Mohammed, Y. (2013). A weather generator for use in construction simulation models. *Intenational Journal of Architecture, Engineering and Construction*, 2(2).

Shibu, M.E., Leffelaar, P.A., van Keulen, H., and Aggarwal, P.K. (2010). Lintul3, a simulation model for nitrogen-limited situations: Application to rice. *Europ. J. Agronomy*, 32(255-271).

Smale, A. (2017). Prijs tarwe oude oogst omhoog. URL `https://www.boerderij.nl/Akkerbouw/Achtergrond/2017/7/Prijs-tarwe-oude-oogst-omhoog-154897E/`. Date accessed: 1-7-2019.

Spitters, C.J.T. and Schapendonk, A.H.C.M. (1990). Evaluation of breeding strategies for drought tolerance in potato by means of crop growth simulation. *Plant and Soil*, 123, 103–203.

Sun, H.Y., Liu, C.M., Zhang, X.Y., Shen, Y.J., and Zhang, Y.Q. (2006). Effects of irrigation on water balance, yield and WUE of winter wheat in the North China Plain. *Agricultural water management*, 85, 211–218.

United Nations (2019). The united nations worlds water development report 2019. *UNESCO Publishing*.

van den Borne, J. (2019). Variabel beregenen. URL `https://www.vandenborneaardappelen.com`. Pers. comm.

Wolf, J. (2011). Lintul-2, a simple crop growth model for both potential and water limited growing conditions. URL `https://models.pps.wur.nl/`. Date accessed: 2-7-2019.

Zhang, H. and Oweis, T. (1999). Water-yield relations and optimal irrigation scheduling of wheat in the mediterranean region. *Agricultural Water Management*, 38, 195–211.

Appendix A. RELATION GROWTH AND YIELD

Figure A.1 shows the relation between the fraction of maximum growth (i.e., water stress) during 10 days and the expected yield at the end of the growing season if the growth would be optimal for the remainder of the growing season. The relation between water stress and yield is linear for most periods of the year (except when the crop goes from the growing stage to the anthesis stage, $k = 80$ in the figure), hence we assume that a linear approximation is sufficient for the prediction of yield in (11) for all days of the growing season.

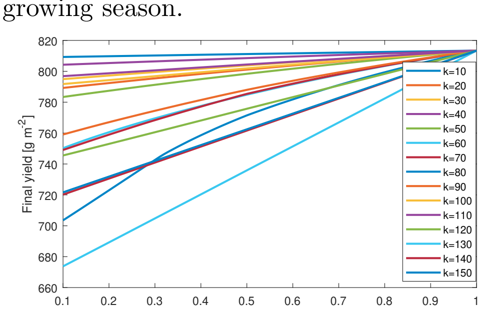

Fig. A.1. k indicates the first day of water stress. This figure is made with weather data of 1976, Eelde, The Netherlands