



A hierarchical partitioning method for optimizing irrigation strategies

J.-E. Bergez ^{a,*}, F. Garcia ^b, L. Lapasse ^a

^a INRA, Unité d'agronomie, BP27, 31326 Castanet-Tolosan Cedex, France

^b INRA, Unité de Biométrie et Intelligence, Artificielle, BP27,
31326 Castanet-Tolosan Cedex, France

Received 7 October 2002; received in revised form 26 June 2003; accepted 25 July 2003

Abstract

Improving water efficiency in irrigated agriculture is a priority for better environmental and economic performance. It can be achieved by reducing the amount of water used and optimizing the timing of application. Among numerous methods, hybrid biophysical/decisional simulation models are effective tools for evaluating and comparing irrigation strategies under different weather conditions. However, optimizing irrigation strategies for some specified agro-environmental expected criteria is a computationally hard problem. In this paper we propose a new approach for optimizing the parameters (quantities, dates, thresholds, etc.) of maize irrigation strategies represented by structured decision rules and simulated with the management-oriented model MODERATO. We introduce a stochastic simulation optimisation method, called Partition-2P (P2P). P2P is designed to completely explore by sampling the domain of the strategy parameters. This exploration is based on a hierarchical decomposition of the domain, heuristically guided toward optimal regions, and allows to consider high-dimensional optimisation problems. We first evaluate P2P by comparing it to a systematic grid search on a simple two-parameter problem, in order to optimise the expectation of the direct margin over 49 years of weather records. We obtain very similar results that confirm the sound behaviour of P2P. We then apply our approach to a more complex optimisation problem that involves an eight-parameters strategy, for which the systematic grid search method is not effective. The best strategy we obtain shows a 28 € ha⁻¹ increase in margin compared to a basic strategy proposed by irrigation advisors. The P2P algorithm is also used in three specific hydraulic contexts concerning the available flow rate. The different optimal parameter obtained

* Corresponding author. Tel.: +33-5-61-28-50-37; fax: +33-5-61-73-55-37.
E-mail address: jbergezt@toulouse.inra.fr (J.-E. Bergez).

for each context agree well with the expert knowledge of irrigation advisors. We finally conclude by discussing some limitations and possible improvements of the presented work.
© 2003 Elsevier Ltd. All rights reserved.

Keywords: Irrigation strategy; Optimisation; Decision support system; Maize

1. Introduction

Managing water properly is certainly one of the major challenges of the 21st century. Agriculture is an important consumer of water and irrigation is essential in some areas to allow acceptable crop yield and quality (Hook, 1994). Improved water use efficiency in irrigated agriculture is therefore a priority for better environmental and economic performance (Howell, 2001).

Irrigation scheduling has been an important topic in agricultural research for several decades and continues to be so today (King et al., 2001). There is a vast amount of literature on irrigation scheduling methods (e.g., Pleban et al., 1984; Plant et al., 1992; Epperson et al., 1993; Botes et al., 1996; Raghuwanshi and Wallender, 1998; Bergez et al., 2002). Biophysical models linked with decisional models can be a helpful tool to solve irrigation scheduling problems. MODERATO (Bergez et al., 2001a) is such a hybrid decision model that includes the main constraints for irrigation, but requires fewer detailed data than classical whole farm models. Given a farmer's irrigation strategy, it simulates the plant–soil system with a dynamic biophysical model and takes account of some within-field variability due to the time required to irrigate the whole field. MODERATO is an effective tool for evaluating and comparing irrigation strategies under different weather conditions (Bergez et al., 2002).

In MODERATO, irrigation strategies are represented as sets of decision rules with lots of parameters (thresholds, quantities, dates, etc.). The design of innovative strategies that perform well for some specified agro-environmental criteria thus leads to large optimisation problems. In this paper, we propose a new simulation-based optimisation approach that handles this difficulty. This approach combines a hierarchical partitioning algorithm called P2P with MODERATO. The use of the P2P algorithm is a further step in optimising strategies and therefore provides some important improvements on the work carried out by Bergez et al. (2002). After an introduction to P2P as a simulation optimisation method, we present MODERATO and the irrigation problem we considered. Experimental results we obtained are then presented and discussed.

2. Materials and methods

2.1. A hierarchical partitioning method for optimising agricultural systems

The problem of optimising the management of the irrigation season can be defined as finding the best sequence of decision rules that maximises the expected value

of an objective function, J , which commonly is defined as the net margin of the crop. This optimisation problem can be seen as a control problem, for which dynamic programming (DP) methods exist (Kennedy, 1986). However, DP methods are often inappropriate when faced with large state and decision spaces, despite some promising improvements of DP methods that have been shown recently with the reinforcement learning approach (Sutton and Barto, 1998).

A more flexible and realistic formulation of the optimisation problem consists of searching by means of simulation for the best values for the parameters of a pre-defined irrigation strategy. In this case the simulation model is considered as a black box function where the vector of strategy parameters $\theta = (\theta_1, \dots, \theta_p)$ is taken as the input variable and where the output is the objective function J . In most of the simulation models of agricultural production systems, the objective function J also depends on the weather, which has to be considered as an unknown and uncontrollable random variable ξ . Optimising a strategy thus consists of searching for the set of parameters θ^* that maximises the expected value of the objective function according to the random weather series:

$$\theta^* = \operatorname{argmax}_{\theta \in \Theta} E(J(\theta, \xi)), \quad \text{with } \theta \in \Theta, \quad (1)$$

where Θ is the value domain of θ .

Several efficient methods (stochastic approximation, random search, stochastic branch-and-bound, etc.) have been recently developed for solving this stochastic simulation optimisation problem, which is one of the most difficult problems of mathematical programming (Azadivar, 1999; Swisher et al., 2000; Fu, 2001). Stochastic iterative methods for deterministic optimisation problems like genetic algorithms, simulated annealing or tabu search (Reeves, 1995) have also been adapted for the stochastic environment associated with simulation. In these approaches, the criterion $E(J(\theta))$ is estimated by averaging the objective function $J(\theta, \xi)$ over a large number of sampled variables ξ_i :

$$E[J(\theta, \xi)] \approx \frac{1}{N} \sum_i J(\theta, \xi_i) \quad (2)$$

When the ξ_i are fixed, the objective function J thus becomes a deterministic function of the input variables θ , and a range of classical optimisation algorithms can be used.

Such evolutionary algorithms and other standard deterministic optimisation methods have all been applied to agricultural simulation models in recent years (Mayer et al., 1998a) and have performed well. However, most of these optimisation problems consist of optimising management decisions on existing (past) climatic series (e.g., Li and Yost, 2000; Mayer et al., 1996, 1998b, 2001; Parsons, 1998; Reddy et al., 1995). Very few optimisation methods have been applied so far to agricultural systems for optimising the expected value of an objective function, as described above. Cros et al. (2001) used the Kiefer–Wolfowitz stochastic gradient method (1952) to derive the best values of some parameters involved in a grazing management strategy. It appears that this approach is reliable, but the algorithm requires a difficult parameterisation that is specific to the application. Botes et al. (1996) optimised irrigation

strategies by using the Nelder–Mead simplex algorithm (1965). Expected values of strategies are estimated by Eq. (2), using $N = 20$ years' weather data.

In this paper we develop a new method for stochastic simulation optimisation problems with continuous variable space $\Theta \subset R^p$, where we assume that parameters θ_j are bound-constrained, which means that the domain Θ is a hyper-rectangle on R^p . This algorithm is designed to completely explore (by sampling) the domain Θ , in order to find local and global optimal values for $\theta \in \Theta$. For small values of p , the simplest method for such problems is a systematic grid-search (SGS), which consists of estimating the objective function on all points located on finer and finer grids over Θ , until a maximum precision is reached for each variable. For large-dimension problems ($p > 2$) this method is not efficient since the number of grid points grows exponentially with p . A more useful approach is therefore to prioritise the evaluation of points θ within promising regions of Θ that are supposed to contain optimal solutions. Its main advantage is to maximise, for a given budget of time or simulation runs, the chance of finding a good solution.

The method we propose falls within this approach and is inspired by the DIRECT (Jones et al., 1993) or the MCS (Huyer and Neumaier, 1999) algorithms dedicated to deterministic problems. It belongs to the family of stochastic branching methods, like stochastic branch and bound or nested partitions methods. It is based on a hierarchical decomposition of Θ into 2^p -trees, and we therefore call it Partition- 2^p , or P2P. The principle of the algorithm is simple. At each iteration of the search, a promising region is selected from a list of “pending regions” (see below). This selected region, a p -dimensional rectangle, is then broken down into 2^p smaller ones, whose sides are half the length of the original rectangle. Each of these 2^p new pending regions is then randomly sampled, and the sample points are used for estimating by simulation some indices of the region, on the basis of which the region is ranked in the pending list. From the initial region Θ , this algorithm generates a sequence of rooted trees of rectangles, whose leaves form a partition of Θ (see Fig. 1). The maximum depth of the principal tree is achieved when pending regions cannot be broken down any further (the maximum precision is reached).

Indices are used as heuristics for selecting promising regions to explore. The efficiency of the algorithm thus relies mainly on the choice of these indices. For instance, selecting the pending regions with the lowest depth leads to a systematic grid-search method. Conversely, selecting regions with maximal average values on sample points leads to an optimistic best-first search that favors a fast convergence to local optima. A useful trade-off can also be achieved by considering both the average value and the variance of sample points in the region. Indeed, for two regions with a similar mean, a better potential increase is to be expected in the one having the larger variance.

In the optimisation runs with P2P presented in this paper, the average value f of a region was the only index used for ranking the pending cells. For each of the sampled points θ_j within a region, the simulation model is run and the value $f(\theta_j) = E(J(\theta_j))$ is estimated by averaging the objective function $J(\theta_j, \xi)$ for N sampled variables ξ_i

$$f(\theta_j) \approx \frac{1}{N} \sum_i f(\theta_j, \xi_i) \quad (3)$$

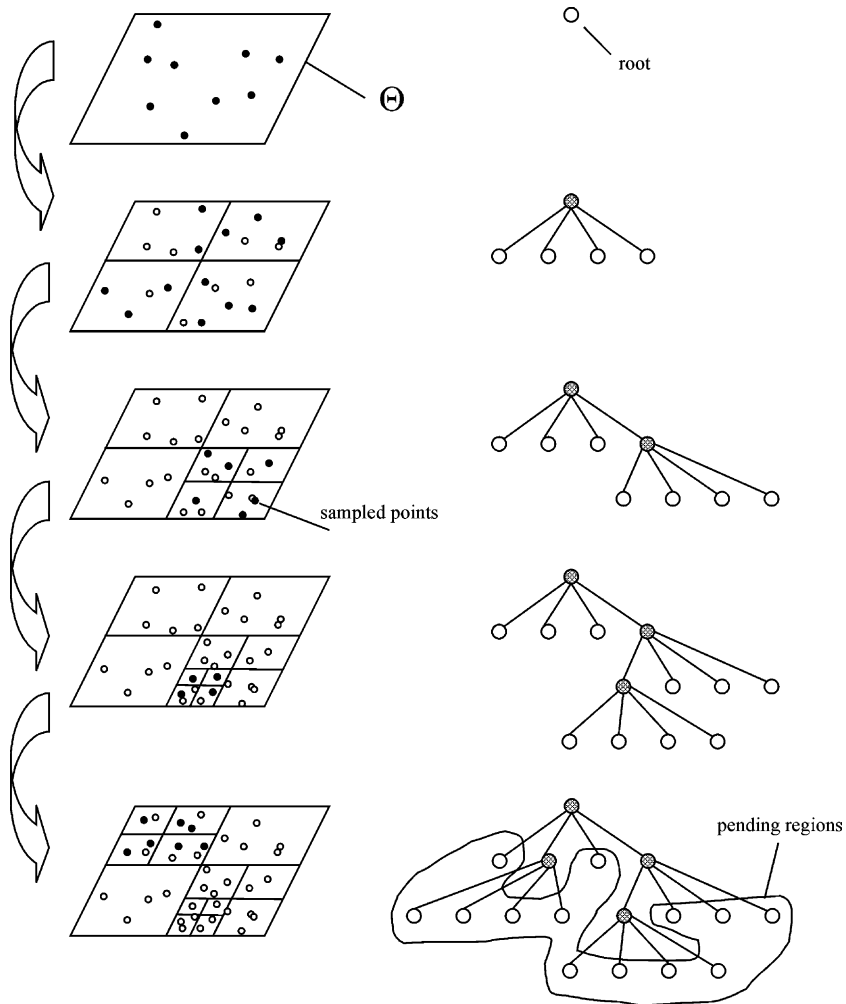


Fig. 1. Description of the P2P algorithm for a two-parameter optimisation problem. From the initial region Θ describing the space of the two parameters, P2P generates a sequence of trees of rectangular regions that form a partition of Θ . Each new created region is evaluated by simulation on sampled points. Pending regions are ranked and the best one is developed first. P2P stops when a minimal discretisation has been reached for all pending regions.

The number M of sampled points θ_j drawn randomly on each new region can be either fixed (e.g., one point in the centre of the region), or based on the depth of the cell (e.g., proportional to the volume of the region). The average value f of the region is then estimated by

$$f \approx \frac{1}{M} \sum_j f(\theta_j) \quad (4)$$

As f is just an empirical indicator, its exact value does not need to be precisely estimated, and small numbers N and M can be sufficient in practice (Ólafsson and Shi, 1999).

The algorithm P2P stops when the maximal depth or minimal discretisation has been reached for each pending region, or more generally when a given number of iterations has been performed. If the objective function $E(J(\theta))$ is continuous, the algorithm is guaranteed to converge to the globally optimal parameter θ^* when the minimal discretisation tends to 0 and the number N of sampled variables ζ_i tends to infinity. During the search, a list managing the best regions is maintained, sorted only on the index f . This list constitutes the output of the algorithm.

2.2. *The irrigation management simulation model*

MODERATO is a management-oriented model aimed at improving the irrigation strategy for an irrigated corn crop. Details of the decision model may be found in Bergez et al. (2001a). In this paper, we will deal only with decision management for a case in which irrigation is to be optimised. MODERATO includes in the irrigation decision-making process the main constraints for irrigation (flow, volume and working rate) and simulates the plant–soil system with a dynamic biophysical model (a detailed description of the equations of the biophysical model is given by Wallach et al. (2001)). We specified the system boundaries in a way that reflects the constraints and opportunities faced by farmers (Cox, 1996). The constraints for irrigation included in MODERATO are divided in three categories:

- *Equipment constraints*: maximum flow available due to pumping capacity per hectare; maximum and minimum amount of water allowed per irrigation.
- *Resource constraints*: a limited total amount of water for irrigation.
- *Regulatory or human constraints*: irrigation bans on some days of the week, farmer's reluctance to work for all or part of certain days, decreased flow rate availability during the irrigation period due to administrative restrictions, etc.

A major part of the decision-making process followed by farmers consists of deciding when to start a new irrigation round, the quantity of water to apply and when to stop irrigation. These decisions are taken on the basis of almost daily observations of the physiological stage of the development of the crop and the soil water content, given some equipment constraints (flow rate and available water). In surveys carried out in south-western France (Darracq, 1996; Clavé, 2001), it was found that farmers use different rules to irrigate and different indicators to trigger the rules. For example, to decide when to start irrigation, different indicators were used: when leaves start to roll, when the soil changes colour, when the soil thrown up by moles becomes drier, when tensiometer values reach a given threshold, when the neighbour starts to irrigate, etc. All these indicators are difficult or even impossible to translate into a decision tool. Unlike Shaffer and Brodahl (1998), who decided to introduce a lot of complexity into their decision rule editor, we chose to simplify this complex decision strategy and to base the elementary rules on a two-step condition algorithm:

Fig. 2. Template describing the irrigation rule for use at plant emergence. Editing fields allow the strategy to be described. See the text for a description of the actual rule.

1. a condition describing the crop development [Cond1];
2. a condition describing the water status in the system [Cond2].

The Boolean process to decide on any rule is then very simple: if Cond1 is true, then Cond2 is checked; if Cond2 is true, then an action occurs.

Two choices are offered to describe each of the two conditions: either an indicator directly related to weather data and easily obtained by the user (calendar date for Cond1, sum or average rainfall or potential evapotranspiration for Cond2) or a simulated indicator, calculated by the biophysical model (stage of the crop described by accumulated thermal units (Muchow et al., 1990) for Cond1, fraction of available water for Cond2). For example, Fig. 2 shows the rule for irrigation for plant emergence, which reads:

If, 15 days after sowing (Cond1, Choice 1) or 70 °C day after sowing (Cond1, Choice 2), the cumulative rainfall has been less than 20 mm (Cond2, Choice 1), then 20 mm of irrigation is applied. The irrigation round will stop if the cumulative rainfall since the beginning of the round is more than 10 mm. The applied amount of irrigation has to be subtracted from the total amount of water available for irrigation.

Once it is decided to irrigate, the amount must be decided. Two possibilities exist: either to apply a fixed amount, or one based on the soil water depletion. The latter is described either by the ratio of actual available water to the available water capacity, or by the amount of water necessary to recharge the soil. Both possibilities allow a partial refill of the soil. A complete irrigation strategy is therefore defined by a set of five elementary rules: (i) a rule to determine if irrigation is to be used to facilitate plant emergence (sowing rule), (ii) a rule to decide when to start the main irrigation period (starting rule), (iii) a rule to determine when to start a new irrigation (returning rule), (iv) a rule to decide when to stop irrigation (ending rule) and (v) a rule to delay irrigation due to weather conditions (delaying rule).

2.3. Testing the optimisation method P2P

A three-step approach will be used to demonstrate the usefulness of the developed algorithm:

Table 1
Description of the basic strategy (BAS) as put forward by irrigation advisors

Operation	Rules
Sowing	Sowing is between 20 April and 30 May as soon as the cumulative rainfall during the previous 3 days is less than 15 mm. Variety Cécilia ^a is sown at 80,000 plants ha ⁻¹ .
Fertilisation	A single application of 200 kg N ha ⁻¹ is made at sowing.
Harvest	The crop is harvested when grain moisture content reaches 20% or accumulated thermal units from sowing reach 2100 ATU ^b and if the cumulative rainfall during the previous 3 days is less than 15 mm. In any case, the crop must be harvested before 15 October.
Irrigation	<p><i>Sowing rule:</i> Irrigation to facilitate plant emergence (either because it is too dry or because a crust has been created by heavy rainfall on silty soil) is not taken into account. Irrigation to dissolve fertiliser is not taken into account.</p> <p><i>Starting rule:</i> The main irrigation period starts from 650 ATU as soon as the soil water deficit reaches 60 mm. 20 mm is applied.</p> <p><i>Returning rule:</i> 25 mm is applied every 8 days if no rainfall occurs.</p> <p><i>Ending rule:</i> For the irrigation cycle following 1st September, if the soil water deficit is greater than 90 mm before the irrigation cycle starts, a last irrigation cycle is performed; otherwise irrigation is stopped. 20 mm is applied.</p> <p><i>Delaying rule:</i> Precipitation delays irrigation. When the cumulative rainfall over the 5 previous days is more than 10 mm, a delay of one day is applied for every 4 mm. The delay cannot exceed 7 consecutive days.</p>

^a Cécilia is a late growing variety requiring 1045 ATU from sowing to flowering and 1990 ATU from sowing to maturity (35% grain humidity).

^b Accumulated thermal units (see text for explanation).

1. We will simulate a basic strategy (BAS) (see Bergez et al., 2002). The general strategy is given in Table 1. This strategy is the result of discussions with the irrigation advisors and is therefore expected to be a reasonable approximation to current practices. The irrigation equipment used for the study allows a 3.5 mm day⁻¹ flow rate. No limit to the amount of available water and no flow rate restriction during summer (except those due to the equipment) are imposed.
2. A comparison between the systematic grid-search method (SGS) and the P2P algorithm will be carried out on two parameters of the basic strategy: the soil water deficit ($D1$) and the time (expressed as accumulated thermal units) to start the main irrigation campaign ($T1$). As before, the irrigation equipment used for the study allows a 3.5 mm day⁻¹ flow rate. No limit to the amount of available water and no flow rate restriction during summer (except those due to the equipment) are imposed. $D1$ will be searched in [0; 150] and $T1$ in [0; 1990].
3. The basic irrigation strategy is extended as follows:
 “The main irrigation period starts from $T1$ as soon as the soil water deficit reaches $D1$. An amount $I1$ is applied. Once an irrigation cycle ends, a new cycle starts when the soil water deficit reaches $D2$. An amount $I2$ is applied. For the irrigation cycle following $T3$, if the soil water deficit is greater than $D3$ before this irrigation cycle starts, a last irrigation cycle is performed; otherwise the irrigation campaign ends. An amount $I3$ is applied”.

P2P will be used to optimise the set of the eight parameters (Table 2) in a first step using the same hydraulic context as before (a 3.5 mm day⁻¹ flow rate and no water limitation) and in a second step when varying the available flow rate: 2.5, 3.5 and 4.5 mm day⁻¹ and limiting the amount of water to 200 mm.

All the simulations will be performed using the same soil and weather data. The soil for this study is a medium clay-silt, 0.8 m deep, with clay accumulation at depth, locally called “Boulbènes moyennes” (fluvisol). This type of soil is representative of a large area of the Midi-Pyrénées and has a 150 mm cumulative available water capacity.

The climate used is from the weather records of Toulouse-Blagnac from 1949 to 1997. On average, July and August receive a total of 92 mm rainfall and the cumulative potential evapotranspiration (ET₀) is 290 mm. The average climatic moisture deficit (ET₀ minus rainfall) for this 2-month period is around 200 mm. However, there is a large variation in rainfall during the two summer months as it ranges from 30 to 240 mm, underlining the unpredictable nature of rainfall in the area. Cumulative ET₀ is less variable, ranging from 235 to 372 mm.

The objective function to be maximised is the expectation of the direct margin (i.e., the gross margin minus specific costs for a given activity, here irrigation) over all the years of weather recording. The direct margin regarding irrigation may be written as:

$$J(\theta_i, \xi_j) = Y(\theta_i, \xi_j) \cdot P - [F + q(\theta_i, \xi_j) \cdot p + n(\theta_i, \xi_j) \cdot c], \quad (5)$$

where $J(\theta_i, \xi_j)$ is the direct margin for climate ξ_j and the strategy θ_i , $Y(\theta_i, \xi_j)$ is the grain yield obtained under climate ξ_j and using the strategy θ_i , P is the selling price for maize, F is the operational costs for maize production, $q(\theta_i, \xi_j)$ is the amount of

Table 2
The eight parameters of the irrigation strategy to be optimised

Name	Meaning	Unit	Min	Max	Δmin
<i>T1</i>	Accumulated thermal unit to start the irrigation campaign	°C day	400	1000	5
<i>D1</i>	Soil water deficit to start the irrigation	mm	50	130	3
<i>I1</i>	Irrigation applied at the first irrigation	mm	5	50	2
<i>D2</i>	Soil water deficit to start a new irrigation cycle	mm	50	130	3
<i>I2</i>	Irrigation depth applied after the first irrigation round	mm	5	50	2
<i>T3</i>	Accumulated thermal units to stop the irrigation	°C day	1400	1800	5
<i>D3</i>	Soil water deficit to stop irrigation	mm	50	130	3
<i>I3</i>	Irrigation applied at the last irrigation round	mm	5	50	2

Min and max show the range for each parameter within which the optimum will be sought. Δ min is the criterion used to stop the division of the parameter.

water used under climate ξ_j and using the strategy θ_i , p is the cost of irrigation water, $n(\theta_i, \xi_j)$ is the number of irrigation cycles performed under climate ξ_j and using the strategy θ_i and c is the cost of carrying out a new irrigation cycle.

The average selling price for maize (grain) is assumed to be 106.71 € Mg⁻¹ in the Toulouse area. Operational costs (seed, weeding, fertiliser, insurance) are assumed to be 327.77 € ha⁻¹. The cost of irrigation water is assumed to be 0.76 € mm⁻¹ and that of setting up a new irrigation cycle is assumed to be 7.62 € (labour).

In this study, the parameters used to tune the P2P algorithm are as follows: (i) the simulations are based on the same $N = 49$ climatic years as described previously; (ii) the pending cells are ranked depending on the expectation f of the objective function; (iii) Only the best 500 cells are kept; (iv) the number of points per cell is computed in order to reach a single point at the lowest depth of cell division; (v) pending cells are excluded (i.e., these cells will not be divided) if the expectation of the objective function for this cell is 10% less than that of the first cell in the best-cell sorted list.

2.4. Programming aspects

The optimisation algorithm P2P was programmed as an independent library in C++ and linked directly to the biophysical and decisional model MODERATO. The algorithm was developed in such a way that it can be used for other models and is not specifically linked to the optimisation problem developed in this study. Two modifications to the main program were needed, being quite straightforward: (i) the first allows for a modification of the parameters of the strategy by sending them as an array to the main program by the optimisation module; (ii) the second is to calculate the objective function for each run and then to send it as an array to the optimisation module from the main program.

3. Results

3.1. Results from the basic strategy

The main results from the “basic” (BAS) strategy are given in Table 3. With this strategy, a 9.54 Mg ha⁻¹ average grain yield is obtained, using, on average, 170 mm of irrigation water. The average margin is 506 €. Drainage is low, 3 mm on average. In 29 cases out of 49, the soil water deficit is at least 60 mm (the soil water deficit threshold to start the irrigation campaign) when the accumulated thermal units reach 650 °C days (the accumulated thermal units to start the irrigation campaign). Therefore irrigation starts immediately at 650 °C day.

3.2. Comparing the SGS and P2P methods

For a two-parameter optimisation problem, the optimal solution obtained with SGS is $T1 = 520$ °C day and $D1 = 40$ mm. For this set of parameters, the margin is

Table 3

Results from the simulation using the basic strategy as expressed by the irrigation advisors (“BAS”), the optimum two-parameter strategy (“2P”) and the optimum eight-parameter strategy (“8P”) calculated by P2P

	BAS	2P	8P
Yield (Mg ha ⁻¹)	9.54 [1.02] (7.63; 11.07)	9.87 [0.9] (7.88; 11.07)	9.52 [0.92] (7.55; 11.15)
Drainage (mm)	3 [12] (0; 77.71)	3 [14] (0; 80)	3 [15] (0; 101)
Irrigation (mm)	170 [38] (70; 240)	196 [47] (95; 265)	166 [45] (46; 224)
Margin (/ha)	506 [134] (254; 737)	515 [125] (227; 714)	534 [118] (312; 710)
Efficiency (Mg mm ⁻¹)	0.06 [0.02] (0.04; 0.15)	0.05 [0.02] (0.03; 0.11)	0.07 [0.04] (0.04; 0.22)

Tabulated values are: average [standard deviation] (minimum; maximum).

515 €. Fig. 3A shows isomargin contours in the parameter space. It can be seen that a large area of the parameter space gives roughly the same margin ($J(\theta) > 500$ €).

With P2P applied to the same two-parameter problem (Fig. 3B), the optimal set of parameters is obtained with $T1$ ranging between 443 and 451 and $D1$ ranging between 37 and 42. For this set of parameters, the margin is 515 €. The second best set for P2P is $T1$ in the range 521–529 and $D1$ in the range 37–42 but the margin is only 0.10 € less. Actually, for the 500 best sets of parameters from P2P, there is a difference in margin of less than 1% (from 515 to 511 € ha⁻¹). This is in good agreement with the large area of similar results obtained using SGS and shown in Fig. 3A.

P2P provides a range for each optimised parameter. We call strategy 2P the strategy obtained by using the middle of this range for each parameter. The simulation of 2P leads to an average grain yield of 9.87 Mg ha⁻¹, i.e., 0.33 Mg ha⁻¹ more than with the basic strategy (see Table 3, strategy 2P). On average 196 mm of water is used, which is 26 mm more than with the basic strategy. This is due to the fact that the main irrigation campaign starts earlier and ends at the same time. In half of the 49 years, irrigation starts before 29/06 for BAS and before 14/06 for 2P. The average margin is 515 €. The same average drainage is obtained for both strategies.

3.3. Eight-parameter strategy

Table 4 gives the ranges of the eight parameters that lead to the maximum expectation of the objective function. We call strategy 8P the strategy obtained by using the middle of the optimal range for each parameter. Following 8P, irrigation has to start a little later than with 2P, at 480 °C days and a 52 mm soil water deficit. The amount applied is then 49 mm. The next irrigation is due when the soil water deficit reaches 57 mm; 46 mm of irrigation is then applied. Finally, the last irrigation cycle is performed if at 1459 °C day the soil water deficit is more than 102 mm; 49 mm irrigation is then applied.

As for the two-parameter-optimised strategy, there is a difference in margin of less than 1% between the 500 best sets of parameters. Within the 100 best sets,

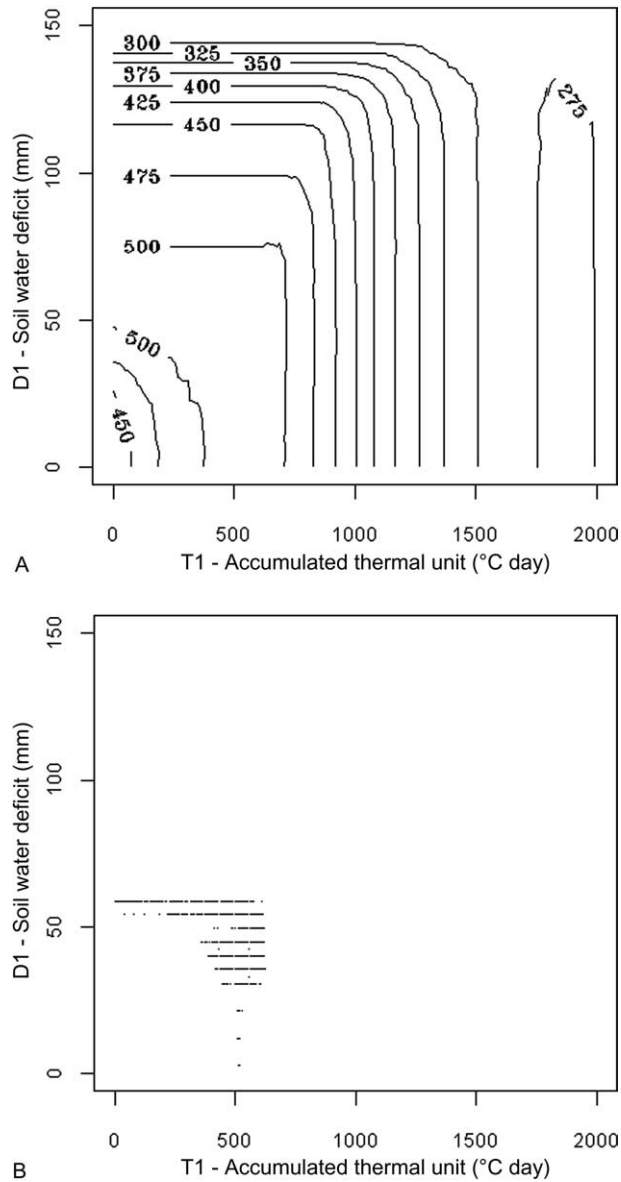


Fig. 3. (A) Isomargin graph when using SGS for each point of the grid when varying $T1$ and $D1$; (B) the 500 best strategies when using P2P. Margin difference for these points is less than 1%.

among all parameters, $T1$ is the most variable as it ranges from 405 to 545 °C days. The others vary much less. This is analogous to the results for the two-parameter strategy. As shown in Fig. 3A, the margin is roughly the same for a large range of $T1$ values.

Table 4
Range of the eight parameters maximizing the expectation of the objective function calculated with P2P

Parameter	Unit	Min	Max
<i>T1</i>	°C day	475	484
<i>T3</i>	°C day	1456	1462
<i>D1</i>	mm	50	55
<i>D2</i>	mm	55	60
<i>D3</i>	mm	100	105
<i>I1</i>	mm	47	50
<i>I2</i>	mm	44	47
<i>I3</i>	mm	47	50

See Table 2 for explanation of the parameters.

The simulation of the strategy 8P leads to a 9.52 Mg ha⁻¹ grain yield (Table 3, strategy 8P). This is slightly less than with 2P (−0.35 Mg ha⁻¹), but the margin is increased: +19 € ha⁻¹. Less water is used (166 vs 196 mm). A large amount of irrigation is applied at each run (46 mm). As a consequence, each irrigation cycle lasts longer and fewer cycles are needed (3.7 vs 7.1 for BAS). Water use efficiency (mg of grain yield produced per mm of irrigation water) has increased.

Fig. 4 shows the cumulative frequencies for grain yield, margin, starting and ending dates obtained for the 49-year climatic series and using the BAS, 2P and 8P strategies. There is a 50% chance to exceed 9.56 Mg ha⁻¹ with BAS, 10.08 Mg ha⁻¹ with 2P and 9.67 Mg ha⁻¹ with 8P (Fig. 4A). Regarding the margin, and keeping the same 50% probability, the figures are 498, 528 and 538 € ha⁻¹, respectively, for BAS, 2P and 8P (Fig. 4B). Actually, 2P gives a better chance to reach a higher grain yield but leads to an earlier start (Fig. 4C) and a later end of irrigation (Fig. 4D) than 8P. As a result, margin is decreased due to a higher amount of irrigation.

P2P performed similarly to SGS for the two-parameter problem in terms of results and run time. However, when the number of parameters to be optimised increases, SGS takes much longer to run. This is part of the well-known curse of dimensionality. P2P allows a quicker convergence to the optimal set of parameters. In our example, the 8P strategy was obtained after roughly 1.6×10^6 simulations (a simulation being a call to the bio-decisional model, i.e., a run for one climate and one set of parameters). On our PC configuration, we run around 25 simulations per second. So, an optimisation required around 14 h. With SGS, it would have been necessary to complete the total map of the parameter space, to run more than 6×10^{11} simulations (actually more than 700 years of computation).

The use of P2P to analyse the change in flow rate gives also interesting results (Fig. 5). These figures are radar plots showing the variability of the eight estimated parameters. Each axis represents one of the eight parameters. All the axes have been normalised to have a [0; 1] range. The values in square brackets to the right of the parameter name indicate the range for this axis. Regarding the rule to begin irrigating, if the flow rate is low (Fig. 5A), irrigation has to start slightly earlier

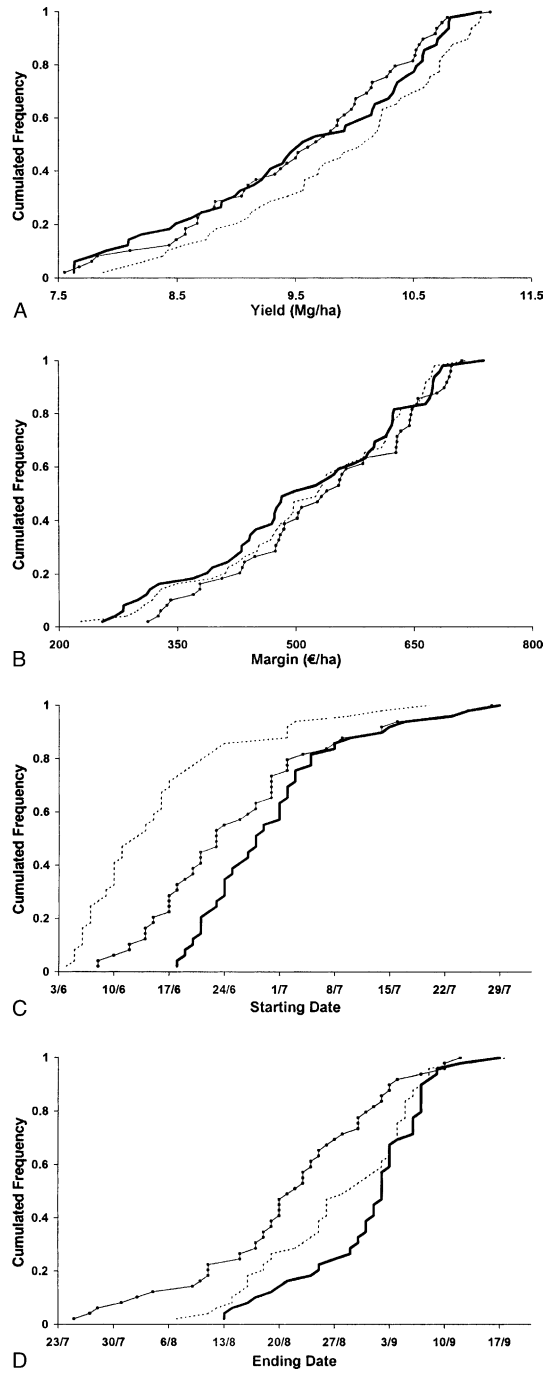


Fig. 4. Cumulative frequency for the basic strategy (bold line), two-parameter optimised strategy (dotted line) and eight-parameter optimised strategy (circles). (A) Grain yield; (B) margin; (C) starting date of irrigation and (D) ending date of irrigation.




Fig. 5. Radar plots showing the eight optimised parameters for three different flow rates: (A) 2.5 mm day⁻¹, (B) 3.5 mm day⁻¹ and (C) 4.5 mm day⁻¹. For each parameter, the range of the axis is given in square brackets. The meaning of the different parameters is given in Table 2.

(T1), at a lower soil water deficit (D1). For higher flow rates (Figs. 5B and C), irrigation can wait as it is easy to rapidly apply the amount required. A larger amount (I1) is applied using a higher flow rate. Regarding the returning rule, the higher flow rate allows waiting until a higher soil water deficit is reached (D2), but the same amount is applied (I2). The decision to stop irrigation (T3) is the latest at the lowest flow rate.

4. Discussion

Starting from a basic strategy as proposed by advisors in the field, the use of P2P to parameterise the decision rules improved the direct margin. From an average 506 € ha⁻¹ using strategy BAS, we obtained 534 € ha⁻¹ with strategy 8P. Irrigation has to start and finish earlier and more water must be applied. Although the grain yield was slightly decreased, the margin was increased because of fewer irrigation cycles. The cost of an irrigation cycle (labour, energy, etc.) is thus important in the search for an optimal strategy: it pays to sacrifice some yield if the number of irrigations can be reduced. This has also been shown by other authors and the idea has been incorporated into new irrigation schedules (Kirda and Kanber, 1999). For example, Norwood and Dumier (2002) found that reducing the quantity of water, even if there is a slight fall in the grain yield, gives more profit and better water use efficiency.

Strategy 8P shows that irrigation should start in mid-June. However, from discussions with irrigation advisors and farmers (Clavé, 2001; Rouffaud, 2001), it seems quite difficult to start that early. Other constraints than those specified in the decisional model have to be taken into account (fertilisation, setting up pipes, etc.). Due to technical problems with equipment, farmers usually start in early to mid-July. Using a model (biophysical and decisional) thus allows constraints and shortcomings of the model or our knowledge to be revealed.

The P2P algorithm was used in three specific hydraulic contexts for one set of soil and climatic conditions. The results agree well with what was already known by irrigation experts: when the flow rate is small, it is worth starting earlier and applying less water. There is much interest in testing different constraints and soil and weather scenarios to provide advice. However, one of the difficulties is to represent the results in an understandable way. In this paper we used a radar plot to represent the strategy. Other forms of representation should be developed in consultation with advisors if we want these results from research be used by advisors and farmers (Cox, 1996; Walker, 2002).

As mentioned by Bergez et al. (2002), a calculated optimal strategy depends on four factors: (i) the biophysical model used to describe the system behaviour (i.e., equation and parameters), (ii) the decisional model and the variables used to trigger the decision, (iii) soil and weather and (iv) the constraints. Adaptation to varying soil and constraints is part of the work of the advisors. Helping them to provide proper strategies is part of the research work and as stated before, finding ways of presenting the results is part of this work.

The first two points are much more difficult to work with. By definition a model is an imperfect representation of reality (Whisler et al., 1986; Boote et al., 1996). When searching the parameter space to find the optimal set of parameters, one may be led to simulate rather extreme soil and plant conditions. It is therefore important that the biophysical model be robust (fairly accurate outside the proper limits of the model). It must also be sensitive (able to predict a small change in the functioning processes when a technical action is modified). For example, regarding maize irrigation, experiments show that a single irrigation at tasselling can increase yield by 29%. Additional

irrigations during the vegetative and grain filling stages increased yield an additional 11% and 13%, respectively (Norwood, 2000). Even if the exact values are not obtained when simulating irrigation strategies, it is important that an increase of this order of magnitude be reached when using the model. It is therefore important to compare knowledge and model estimates.

In order to use P2P, it is necessary to describe the parameterised strategy as a set of decision rules. This task is far from obvious and can only be carried out through a collaborative work with experts and advisors. Nevertheless, optimising strategies without this initial knowledge is even more complex (Bergez et al., 2001b).

In this paper, the P2P algorithm was used to calculate optimal strategies for the direct margin criterion. Such computed optimised strategies must be understood as a guide to possible improvement rather than an optimal strategy to be used as is. For example, it is quite rare that a farmer has only maize. Other productions (animal) and crops may interfere in terms of equipment or working schedule. Environmental aspects must also be considered.

Acknowledgements

The authors thank the irrigation advisors for their help in describing the “basic” strategy and for their comments on MODERATO.

References

- Azadivar, F., 1999. Simulation optimization methodologies. In: Proceedings of the 1999 Winter Simulation Conference, December 5–8, 1999, Squaw Peak, Phoenix, AZ, USA, pp. 93–100.
- Bergez, J.E., Debaeke, P., Deumier, J.M., Lacroix, B., Leenhardt, D., Leroy, P., Wallach, D., 2001a. MODERATO: an object-oriented decision tool for designing maize irrigation schedules. *Ecol. Model.* 137, 43–60.
- Bergez, J.E., Eigenraam, M., Garcia, F., 2001b. Comparison between dynamic programming and reinforcement learning: a case study on maize irrigation management. In: Proceedings of the 3rd European Conference on Information Technology in Agriculture (EFITA'01), Montpellier (FR), pp. 343–348.
- Bergez, J.E., Deumier, J.M., Lacroix, B., Leroy, P., Wallach, D., 2002. Improving irrigation schedules by using a biophysical and a decisional model. *Eur. J. Agron.* 16, 123–135.
- Boote, K.J., Jones, J.W., Pickering, N.B., 1996. Potential uses and limitations of crop models. *Agron. J.* 88, 704–716.
- Botes, J.H.F., Bosch, D.J., Oosthuizen, L.K., 1996. A simulation and optimization approach for evaluating irrigation information. *Agric. Syst.* 51, 165–183.
- Clavé, J., 2001. Représentation des règles de décision des irrigants de la vallée du Tarn et de l'Agout dans l'outil informatique MODERATO. ENITA-Clermont (FR), p. 47.
- Cox, P.G., 1996. Some issues in the design of agricultural decision support systems. *Agric. Syst.* 52, 355–381.
- Cros, M.J., Garcia, F., Martin-Clouaire, R., Duru, M., 2001. Simulation optimization of grazing management strategies. In: Proceedings of the 3rd European Conference on Information Technology in Agriculture (EFITA'01), Montpellier (FR).
- Darracq, S., 1996. Identification et modélisation des règles de décision utilisées par les maïsiculteurs en Aquitaine. INA-Paris (FRA), p. 65.

- Epperson, J.E., Hook, J.E., Mustafa, Y.R., 1993. Dynamic programming for improving irrigation scheduling strategies of maize. *Agric. Syst.* 42, 85–101.
- Fu, M.C., 2001. Simulation optimization. In: *Proceedings of the 2001 Winter Simulation Conference*, December 9–12, 2001, Crystal Gateway Marriott, Arlington, VA, USA, pp. 53–61.
- Hook, J.E., 1994. Using crop models to plan water withdrawals for irrigation in drought years. *Agric. Syst.* 45, 271–289.
- Howell, T.A., 2001. Enhancing water use efficiency in irrigated agriculture. *Agron. J.* 93, 281–289.
- Huyer, W., Neumaier, A., 1999. Global optimization by multilevel coordinate search. *J. Global Optim.*, 331–355.
- Jones, D.R., Perttunen, C.D., Stuckman, B.E., 1993. Lipschitzian optimization without the Lipschitz constant. *J. Optim. Theory Appl.* 79, 157–181.
- Kennedy, J.O. 1986. *Dynamic Programming: Application to Agricultural and Natural Resources*. Elsevier Applied Science, London.
- Kiefer, J., Wolfowitz, J., 1952. Stochastic estimation of the maximum of a regression function. *Ann. Math. Stat.* 23., 462–466.
- King, B.A., Wall, R.W., Taberna Jr., J.P., 2001. Visual soil water status indicator for improved irrigation management. *Comput. Electron. Agric.* 32, 31–43.
- Kirda, C., Kanber, R., 1999. Water, no longer a plentiful resource, should be used sparingly in irrigated agriculture. In: Kirda, C., Moutonnet, P., Hera, C., Nielsen, D.R., (Eds.), *Crop Yield Responses to Deficit Irrigation*. Kluwer Academic Publishers, Dordrecht, pp. 1–20.
- Li, M., Yost, R.S., 2000. Management-oriented modeling: optimizing nitrogen management with artificial intelligence. *Agric. Syst.* 65, 1–27.
- Mayer, D.G., Belward, J.A., Burrage, K., 1996. Use of advanced techniques to optimize a multi-dimensional dairy model. *Agric. Syst.* 50., 239–253.
- Mayer, D.G., Belward, J.A., Burrage, K., 1998a. Optimizing simulation models of agricultural systems. *Ann. Oper. Res.* 82, 219–231.
- Mayer, D.G., Belward, J.A., Burrage, K., 1998b. Tabu search not an optimal choice for models of agricultural systems. *Agric. Syst.* 58, 243–251.
- Mayer, D.G., Belward, J.A., Burrage, K., 2001. Robust parameter settings of evolutionary algorithms for the optimisation of agricultural systems models. *Agric. Syst.* 69, 199–213.
- Muchow, R.C., Sinclair, T.R., Bennett, J.M., 1990. Temperature and solar effects on potential maize yield across locations. *Agron. J.* 82, 338–343.
- Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. *Comput. J.* 7, 308–313.
- Norwood, C.A., 2000. Water use and yield of limited-irrigated and dryland corn. *Soil Sci. Soc. Am. J.* 64, 365–370.
- Norwood, C.A., Dumier, T.J., 2002. Transition to dryland agriculture: limited irrigated vs. dryland corn. *Agron. J.* 94, 310–320.
- Ólafsson, S., Shi, L., 1999. Optimization via adaptive sampling and regenerative simulation. In: *Proceedings of the 1999 Winter Simulation Conference*, December 5–8, 1999, Squaw Peak, Phoenix, AZ, USA, pp. 666–672.
- Parsons, D.J., 1998. Optimizing silage harvesting plans in a grass and grazing simulation using the revised simplex method and a genetic algorithm. *Agric. Syst.* 56, 29–44.
- Plant, R.E., Horrocks, R.D., Grimes, D.W., Zelinski, L.J., 1992. CALEX/Cotton: an integrated expert system application for irrigation scheduling. *Am. Soc. Agric. Engrg.* 35, 1833–1838.
- Pleban, S., Heermann, D.F., Labadie, J.W., Duke, H.R., 1984. Real time irrigation scheduling via “Reaching” dynamic programming. *Water Resour. Res.* 20, 887–895.
- Raghuwanshi, N.S., Wallender, W.W., 1998. Optimal furrow irrigation scheduling under heterogeneous conditions. *Agric. Syst.* 58, 39–55.
- Reddy, V.R., Acock, B., Whisler, F.D., 1995. Crop management and input optimization with GLYCIM: differing cultivars. *Comput. Electron. Agric.* 13, 37–50.
- Reeves, C.R., 1995. *Modern Heuristic Techniques for Combinatorial Problems*. McGraw-Hill, New York (originally published in 1993 by Blackwell Scientific Press, Oxford).
- Rouffaud, H., 2001. Contextes, règles de décision et stratégies de conduite de l’irrigation du maïs dans les vallées du Tarn et de l’Agout. ENESAD-Dijon (FR) ITA Productions végétales, 28 p.

- Shaffer, M.J., Brodahl, M.K., 1998. Rule-base management for simulation in agricultural decision support system. *Comput. Electron. Agric.* 21, 135–152.
- Sutton, R.S., Barto, A.G., 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge.
- Swisher, J.R., Hyden, P.D., Jacobson, S.H., Schruben, L.W., 2000. A survey of simulation optimization techniques and procedures. In: *Proceedings of the 2000 Winter Simulation Conference*, December 10–13, 2000, Wyndham Palace Resort & Spa, Orlando, FL, USA, pp. 119–128.
- Walker, D., 2002. Decision support, learning and rural resource management. *Agric. Syst.* 73, 113–127.
- Wallach, D., Goffinet, B., Bergez, J.E., Debaeke, P., Leenhardt, D., Aubertot, J.N., 2001. Parameter estimation for crop models: a new approach and application to a corn model. *Agron. J.* 93, 757–766.
- Whisler, F.D., Acock, B., Baker, D.N., Fye, R.E., Hodges, H.F., Lambert, J.R., Lemmon, H.E., McKinion, J.M., Reddy, V.R., 1986. Crop simulation models in agronomic systems. *Adv. Agron.* 40, 141–208.