

Exploring Profit – Sustainability Trade-offs in Cropping Systems Using Evolutionary Algorithms

P. deVoil^a, W. A. H. Rossing^b and G. L. Hammer^{a,c}

^aAgricultural Production Systems Research Unit, Department of Primary Industries, Qld, Australia.
Peter.deVoil@dpi.qld.gov.au

^bBiological Farming Systems Group, Wageningen University, The Netherlands.

^cSchool of Land and Food Sciences, The University of Queensland, Brisbane, Australia.

Abstract: Models that implement the bio-physical components of agro-ecosystems are ideally suited for exploring sustainability issues in cropping systems. Sustainability may be represented as a number of objectives to be maximized or minimized. Models provide a means to evaluate system performance for alternative strategic, tactical and operational decisions. The full decision space is usually very large and simplifications based on insight into agronomic relations and farming practice are necessary to safeguard computational feasibility. Different optimisation approaches have been proposed in the literature, usually based on mathematical programming techniques. The disadvantage of these techniques is that they can only deal with greatly simplified system descriptions, thus excluding the direct use of bio-physical system models. Here, we present a global search approach based on an Evolutionary Algorithm (EA). We introduce a multi-objective evaluation technique within this EA framework, linking the optimisation procedure to the APSIM cropping systems model, and perform simulations in parallel over an existing computer network.

A case study addressing crop choice and sowing rules in north-east Australian cropping systems is used to illustrate the EA methodology. Sustainability of these systems is evaluated in terms of economic performance (gross margin, financial risk) and resource use (erosion). Because of the limited size of the problem, the quality of the EA optimisation can be assessed by comparison to a large sample of the full problem domain. Results demonstrate that the EA procedure, parameterized with generic parameters from the literature converges to a useable solution set within a reasonable amount of time. Frontier “peels” or pareto-optimal solutions as described by the multiobjective evaluation procedure provide useful information for discussion on trade-offs between conflicting objectives. Application to a more realistic decision problem is recommended to evaluate the usefulness of the approach in informing discussions on agro-ecosystem design.

Keywords: *Evolutionary Algorithms; Agricultural Systems Modelling; Cropping System Design*

1. INTRODUCTION

We seek to optimise a cropping systems design problem involving crop choice and sowing rules in North-East Australian cropping systems. The systems are evaluated in terms of economic performance (gross margin, financial risk) and resource use (erosion). We investigate the use of Evolutionary Algorithms (EA), a class of stochastic optimisation techniques based on concepts of natural selection and genetic inheritance that improve by selection, recombination and mutation. While such optimisation of agricultural systems is not unique (Mayer, 2002; El-Nazer, 1986; Rossing et al., 1997) to our knowledge there is no existing work that demonstrates the use of EAs in cropping systems design, which allow the concept of pareto – optimality.

While in single objective optimisation the optimal (highest ranking) solution is usually clearly defined, this does not hold for multiobjective

problems. Instead of a single optimum, there is a set of compromise solutions, known as the Pareto-optimal solution after the Italian economist Vilfredo Pareto (1848-1923). These solutions represent the best tradeoff across all criteria as no other solutions are superior to them when all objectives are considered. Figure 1 shows these sets as a series of “peels”. To arbitrate between members of the Pareto-optimal set, a human decision maker is needed. These two aspects of multiobjective optimisation, search and decision making (Zitzler, 1999), may be combined in three ways in optimisation algorithms (Horn, 1997): (i) decision making before search where the objectives are aggregated into a single objective, by implicitly taking into account the decision maker’s preferences; (ii) search before decision making, where the decision maker selects an acceptable compromise solution from the set of solutions resulting from the search process; (iii) decision making during search, where preferential information is used to guide the search process. Methods that aggregate

multiple objectives into a single objective function are known to be sensitive to the shape of the Pareto-optimal front (Zitzler, 1999). In this paper, no preference information is used, consistent with the aim of enhancing learning associated with design of cropping systems.

Since the pioneering work on evolutionary algorithms for multiobjective optimisation in the mid-1980s (Schaffer, 1984), many variations have been published. In general terms, a multiobjective EA should meet the following requirements (Zitzler, 1999):

- The distance of the non-dominated front found to the true Pareto-optimal front is minimized.
- Solutions are adequately (usually uniformly) distributed along the front.
- The spread of the obtained non-dominated front is maximised, i.e. for each objective a wide range of values is covered by the non-dominated solutions.




Figure 1: The top 3 Pareto-Optimal solution frontiers joined by solid lines. Objectives were to maximise Gross Margin and minimise Erosion.

Inherent to the heuristic nature of EAs, convergence to the true Pareto-optimal front is not guaranteed, nor are there hard rules that describe ‘optimal’ configuration: the method of fitness assignment, selection of parents, generation of offspring and replacement of parents affects performance in a problem-specific manner. Since few general guidelines are available, the empirical configuration of an EA is part and parcel of its application.

The fitness of an individual is based on the rank determined by a non-dominated sorting procedure. The procedure starts by finding the set of individuals that are Pareto-optimal in the current population, to which the highest rank is assigned. This set is called the first Pareto stratum. The individuals that are Pareto optimal in the

remaining population are assigned the next highest rank. These individuals comprise the second Pareto stratum. This process continues until all the individuals are ranked. Figuratively speaking, the subsequent non-dominated fronts are peeled off step by step as illustrated in figure 1.

In this paper we construct an EA to investigate applicability and usefulness to cropping systems design. Its performance is compared to a method in which the solution space is sampled using a factorial design.

2. METHODS

2.1. Cropping Systems

Dryland cropping in the grain region of northern Australia is characterised by the opportunity to produce a range of cereal, pulse, oilseed, forage and fibre crops. Both summer and winter crops are grown, with yields largely determined by water supply from either in-season rainfall or water stored in the soil at planting. High variability in seasonal rainfall implies that prospects for any one crop are often risky (Hammer et al., 1996), although major differences in water requirement exist among crops. Fallowing the soil between crops is a recommended activity to reduce dependence on in-season rainfall by building up moisture levels in the soil. Disadvantages of fallowing are the lack of financial return, and possible resource degradation through increased soil erosion and solute leaching. Thus, crop choice and fallowing are critical elements in the balance between economic and ecological objectives for northern Australian cropping systems.

Although performance of cropping systems is primarily assessed in economic terms, farmers and legislators increasingly consider efficiency of resource use as a performance criterion. The farmer can optimise a cropping systems performance through appropriate choice and management of crops and fallow. Cropping systems management is a complex optimisation problem. Performance indicators can be assigned as measures of the objectives to be maximised or minimised, with management actions or rules as the variables to be optimised.

The number of combinations of decisions on land cultivation, crop choice and crop management is extremely large, and restrictions have to be made to define a relevant yet feasible optimisation problem: some decisions are set at agronomically plausible levels, while others, considered less obvious, are candidates for optimisation.

This experiment presents the EAs operation in a known environment. The variables optimised are a discrete choice of land use in four seasons, and a threshold level of plant available water to consider planting (one of four) crops. The cropping system is designed as a two-year rotation, with cotton and sorghum as potential summer crops, and wheat and chickpea as winter crops. Fallowing may take place in either season. Decisions to be optimised comprise which crop to plant and when, and the timing of planting in relation to current soil moisture conditions.

2.2. Problem Representation

APSIM (McCown, et al, 1995) version 2.00 was configured to simulate crop rotations grown on a Brigalow soil type of 5% slope at Dalby (27S, 153E) over a 40 year historical weather record. At the start of each cycle within the rotation, soil water level was set to 260mm available soil water, representing a 100% full profile to 1800mm depth. Sowing rules for each crop consisted of a fixed sowing window, i.e. a period during which sowing was possible, the requirement of 30 mm rainfall over a 3 day period prior to sowing, and a minimum necessary level of plant available soil water. Sowing windows were 1 October to 30 November for cotton, 15 September to 15 January for sorghum, and 15 May to 1 August for both chickpea and wheat. At sowing, cotton was fertilized with 100 kg N/ha, wheat and sorghum with 50 kg N/ha and chickpea was left unfertilized. Prices of inputs and outputs (Table 1) are assumed fixed over the period of the simulation.

Each individual within the EA encodes land use type and requirements of plant available soil water at sowing for each crop into 8 ‘genes’, as described in Table 2. Within the GENIAL framework, a gene is represented by a real number, a useful extension of the binary encoding methods of the more traditional binary EA implementations.

The simulations expose each system to 40 years of climate, evaluated in similar terms as in Carberry et al. (1999): 1) Average yearly gross margin (\$/ha); 2) Risk of economic loss, defined as the percentage of years when actual gross margin was less than \$250 (an estimate of the fixed costs required to maintain a typical farm at Dalby); and 3) Soil loss, defined as the simulated average annual soil erosion (t/ha).

Table 1: Costs and prices used in the optimization studies. Values reflect levels prevalent in 1999.

Crop	Market Price	Fixed costs	Variable costs
Cotton	\$450/bale	\$811/ha+ \$90/ha N	\$80/bale
Wheat	\$150/t	\$180/ha+ \$45/ha N	
Sorghum	\$180/t	\$114/ha+ \$45/ha N	
Chickpea	\$300/t	\$180/ha	

As each APSIM simulation (one individual of the EA population) takes several minutes to complete, we were keen to exploit the implicit parallelism of the EA paradigm by evaluating simulations in parallel, potentially reducing the time taken for one generation of the EA from several hours to a matter of minutes. Given that our hardware resources were limited, we developed a networked, fault tolerant system that: marshalled simulation tasks at a central server; employed workstation clients that connect to this server, collected simulation tasks, performed them, and returned results to this server. The EA process creates the tasks, uploads them to the server, waits until they are finished, retrieves the results and proceeds to the next generation.

Table 2: Functions and values for ‘genes’ used in the optimisation procedure.

Esw: Extractable Soil Water

Gene	Range	Function
Summer Crop 1,2	1,2,3 for Cotton, Sorghum, or Fallow	Land use in 1 st and 2 nd summer
Winter Crop 1,2	1,2,3 for Wheat, Chickpea, or Fallow	Land use in 1 st and 2 nd winter
Cotton Esw	100, 160, 220mm	mm extractable soil water required for a cotton planting
Sorghum Esw	20, 120, 220mm	mm extractable soil water required for a sorghum planting
Wheat Esw	20, 120, 220mm	mm extractable soil water required for a wheat planting
Chickpea Esw	20, 120, 220mm	mm extractable soil water required for a chickpea planting




Figure 2. The problem domain. Simulation results for cropping system over 4 seasons with choice of 3 crops, and 3 levels of starting water for planting each crop.

2.3. Numerical Experiments

We created 6 numerical experiments to firstly ensure that the EA was capable of finding its way to an optimal solution set within a known domain of 2800 points (Figure 2), and secondly to examine how the choice of EA parameters affected efficiency.

In an effort to understand how genetic operators affected the behavior of the EA, 2 selection pressures were trialled: i) Small population ($N = 50$), high replacement rate ($Q = 50\%$) (high selection pressure), and ii) large population ($N = 200$), low replacement rate ($Q = 25\%$) (low selection pressure)

This choice of population size and replacement rate has operational considerations: for peak efficiency, the number of new individuals introduced to the population should be a multiple of the number of client workstations available.

For both selection pressure routes, the weight of two genetic operators, probability of uniform crossover p_{uni} , and probability of mutation p_{mut} , were varied. The three cases were ($p_{uni} = 10\%$, $p_{mut} = 10\%$), ($p_{uni} = 10\%$, $p_{mut} = 30\%$), and ($p_{uni} = 30\%$, $p_{mut} = 10\%$), giving 6 experiments in total.

3. RESULTS

When the feasible solution space was explored using simulations with all possible combinations of crops, fallows, and planting conditions, a wide range of outcomes resulted (Fig. 2). The most desirable outcomes combined high gross margins with low erosion and low financial risk. At gross margins up to about \$750/ha, there were systems with erosion as low as 1-2t/ha among the combinations examined. To reach higher levels of gross margin a trade-off with increased erosion was required. A similar outcome was evident with financial risk. It was only at high gross margin

levels that the trade-off between profit and risk came into play. These envelopes represented the Pareto-optimal frontiers for this problem. There were numerous redundant combinations behind these frontiers, where for any given level of gross margin, high erosion or financial risk resulted.

Figures 4 & 5 show the results of searching this solution space with the two configurations of the evolutionary algorithm. The boxes indicate the position of the pareto optimum frontier for the trade-off between gross margin and erosion as the execution of the EA advanced through generations. Financial risk was closely related to GM for this cost/price regime (data not shown). In both cases, the Pareto-optimum frontier closely reflected that found when the full solution space was explored (compare Fig. 2). This indicated the efficacy of the EA in searching the feasible space using the multi-objective criterion and Pareto-optimum methodology. It was clear, however, that the Pareto-optimal frontier had more uniform coverage with the higher population size (Fig. 4). Five of the six experiments behaved in a similar manner (data not shown), arriving at a stable population close to the optimal frontier within 20 generations. The combination of small population and high random mutation rate (population = 50, $p_{uni} = 10\%$, $p_{mut} = 30\%$) failed to converge to a stable frontier.

An adaptation landscape of crop ‘genes’ is shown in Figure 3. The figure shows the change in frequency of ‘alleles’ retained in the population as it progressed through generations of selection. At the outset, all crop types (including fallow) were given equal representation. The population moved rapidly to an increased proportion of winter crop, preference for wheat over chickpea as the winter crop, a reduced proportion of cotton as summer crop, and more winter fallow. These shifts reflected the price and cost differentials defined for these simulations (Table 1) and the

greater consequences on erosion of including cotton or chickpea in the crop rotation. The frequencies associated with this landscape include all combinations remaining in the population, even though most will be sub-optimal. Hence, it is instructive to examine the Pareto-optimal set individually.

The attributes of the Pareto-optimal set from Figure 3 showed that increased intensity of cropping, and the inclusion of summer cropping of sorghum, was needed to achieve high gross margins (Table 3). However, this also increased financial risk. Further, increased winter cropping of wheat was required to minimise erosion. While these trade-offs were specific to the conditions specified for this analysis, they provided an ideal basis to inform discussion about design of cropping systems with practitioners. Our experience with discussion support for decisions associated with single crop management (Nelson et al., 2002) has shown this to be an effective means to interact with decision makers and their key advisers.

Substantial time investment was needed to create the EA system, largely due to the parallel processing capabilities that were needed to overcome computational time constraints. However, the advantage over cropping system optimisation approaches based on mathematical programming is that the full explanatory power of the agro-ecosystem model can be used. With the system now in place, insights on approaches to cropping system management can be enhanced via direct and on-going interaction with managers. Their ideas can be incorporated and discussed. It would be appropriate to address a more realistic decision problem in this way to evaluate the usefulness of the approach in informing discussions on agro-ecosystem design.

4. CONCLUSIONS

We have shown that an EA can be successfully employed to search the feasible solution space for a complex cropping systems design problem that involves multiple criteria. The concept of Pareto-optimal frontier has proved useful for exploring trade-offs among important conflicting criteria (such as economic and environmental consequences) associated with the system design problem. The quantification of attributes of cropping systems that represent Pareto-optimal combinations of economic and environmental objectives provides highly interesting material for informed discussions on strategic decision making with decision makers.




Figure 3. Relative proportion of crop type and fallow in the cropping system versus generation for the numerical optimisation experiment with $N=200$, $p_{uni}=10\%$, $p_{mut}=10\%$.

5. ACKNOWLEDGEMENTS

We thank the CSIRO Division of Plant Industry, Cotton Research Unit, Narrabri, for the use of cotton module OZCOT within the APSIM framework for these simulations. Financial support to Dr. Rossing by OECD and GRDC is gratefully acknowledged.




Figure 4. 60 iterations of the EA in the numerical experiment with $N=50$, $Q=50\%$, $p_{uni}=10\%$, $p_{mut}=10\%$. 1st, 2nd and 3rd Pareto rankings are joined by lines.

Table 3: The Pareto-optimal sets of parameter values from generation 60 of the numerical experiment with $N=200$, $Q=25\%$, $p_{uni}=10\%$, $p_{mut}=10\%$. The first 4 columns describe the crop sequence, the next 3 the soil water thresholds for planting each crop, and the last three show the outcomes for each management system. Note that cotton does not appear in this set.

First Summer	First Winter	Second Summer	Second Winter	Sorghum (mm)	Wheat (mm)	Chickpea (mm)	Gross Margin (\$)	Risk (%)	Erosion (t/ha)
Sorghum	Wheat	Fallow	Wheat	20	120		329	8	1.07
Sorghum	Wheat	Fallow	Wheat	120	120		332	5	1.08
Sorghum	Fallow	Sorghum	Wheat	120	20		828	15	2.04
Sorghum	Wheat	Sorghum	Chickpea	20	120	20	937	19	2.65
Sorghum	Wheat	Sorghum	Wheat	120	120		761	15	1.12




Figure 5. 60 iterations of the EA in the numerical experiment with $N=200$, $Q=25\%$, $p_{uni}=10\%$, $p_{mut}=10\%$.

6. REFERENCES

- Carberry, P.S., G. L. Hammer, H. Meinke, and M. Bange, The potential value of seasonal climate forecasting in managing cropping systems. in *Applications of Seasonal Climate Forecasting in Agricultural and Natural Ecosystems*. Kluwer Academic. pp.167-181, 2000
- El-Nazer, T., and B.A. McCarl, The choice of crop rotation: a modelling approach and case study. *American Journal of Agricultural Economics* 68, 127-136, 1986
- Hammer, G.L., D.P. Holzworth, and R. Stone, The value of skill in seasonal forecasting to wheat crop management in a region with high climatic variability. *Australian Journal of Agricultural Research*, 47: 717-737, 1996.
- Horn, J., The Nature of Niching: Genetic Algorithms and the Evolution of Optimal, Cooperative Populations. PhD thesis, University of Illinois at Urbana, Champaign, Urbana, Illinois, 1997.
- Mayer, D.G., *Evolutionary Algorithms and Agricultural Systems*, Kluwer Academic, Dordrecht, Netherlands. 107pp, 2002.
- McCown et al., APSIM: an agricultural production system simulation model for operational research. *Mathematics and Computers in Simulation*. 39(3): 225:234, 1995.
- Nelson, R.A., D.P. Holzworth, G.L. Hammer, and P.T. Hayman, Infusing the use of seasonal climate forecasting into crop management practice in North East Australia using discussion support software. *Agricultural Systems*, 74:393-414, 2002.
- Rossing, W.A.H., J.E. Jansma, F.J. de Ruijter, and J. Schans, Operationalizing sustainability: exploring options for environmentally friendly flower bulb production systems. *European Journal of Plant Pathology* 103, 217-234, 1997.
- Schaffer, J.D., *Multiple Objective Optimisation with Vector Evaluated Genetic Algorithms*. PhD Thesis, Vanderbilt University, 1984.
- Zitzler, E., *Evolutionary Algorithms for Multiobjective Optimisation: Methods and Applications*. PhD Thesis, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, 1999.