

Combinatorial optimisation of a large, constrained simulation model: an application of compressed annealing

Graeme J. Doole^{1,2,3} and David J. Pannell^{1,2}

¹ School of Agricultural and Resource Economics, Faculty of Natural and Agricultural Sciences, University of Western Australia, 35 Stirling Highway, Crawley, Western Australia 6009.

² Cooperative Research Centre (CRC) for Plant-Based Management of Dryland Salinity, University of Western Australia, 35 Stirling Highway, Crawley, Western Australia 6009.

³ CRC for Australian Weed Management, Waite Road, Urrbrae, PMB 1, Waite Campus, Glen Osmond, South Australia 5064.

Abstract. Simulation models are valuable tools in the analysis of complex, highly-constrained economic systems unsuitable for solution by mathematical programming. However, model size may hamper the efforts of practitioners to efficiently identify the most-valuable configurations. This paper investigates the efficacy of a new metaheuristic procedure, compressed annealing, for the solution of large, constrained systems. This algorithm is used to investigate the value of incorporating a sown annual pasture, French serradella (*Ornithopus sativa* Brot. cv. *Cadiz*), between extended cropping sequences in the central wheat belt of Western Australia. Compressed annealing is shown to be a reliable means of considering constraints in complex optimisation problems in agricultural economics. It is also highlighted that the value of serradella to dryland crop rotations increases with the initial weed burden and the profitability of livestock production.

Keywords. combinatorial optimisation, crop rotation, simulated annealing.

JEL classification codes. C63, Q15.

Paper presented at the 51st Australian Agricultural and Resource Economics Society Conference 2007.

1. Introduction

Simulation involves experimentation with the control parameters in a given model to gain an understanding of a system's management (Law and Kelton, 2000). This approach is valuable in economic analysis since the complexity of many problems precludes their definition in a formal, analytically-tractable framework suitable for solution by mathematical programming. This includes problems in which it is important to represent numerous discontinuities, strong nonlinearities, a high number of integer variables, and multidimensionality. However, model size may severely constrain meaningful experimentation, thereby complicating the identification of the most beneficial management strategies.

Moreover, the addition of optimisation algorithms to simulation models, particularly those incorporating constraints, has traditionally been a non-trivial exercise. A simulation model may be sampled repetitively to approximate continuous functions for solution by mathematical programming (Plambeck *et al.*, 1996; van Kooten *et al.*, 1997). Alternatively, an approximate value function may be determined through the solution of a simplified version of the model using dynamic programming (Woodward *et al.*, 2005). However, many large simulation models, such as that considered below, require major simplification if they are to be transcribed into a form amenable to solution by these techniques. This increases the probability that critical elements of the decision problem will be lost if the configuration space is pared to allow solution by traditional methods.

A metaheuristic is a general computational algorithm that uses stochastic-search mechanisms to identify high-quality solutions in large and/or complex search spaces. Such procedures are valuable for the optimisation of simulation models in which analytical gradients cannot be efficiently computed. Global optimisation of simulation models is often impractical because of their size, particularly as the problems posed by some frameworks will be nondeterministic

polynomial-time hard (NP-hard)¹ (Du and Ko, 2000). Nonetheless, those near-optimal policies identified by metaheuristics are computed directly for the original problem posed by the simulation framework and not an approximation. Metaheuristics have previously been employed to identify profitable solutions in simulation models in agricultural economics (e.g., Mayer *et al.*, 1996). However, these procedures are limited in their capacity to consistently incorporate resource restraints. This is a significant inadequacy as agricultural systems are typically highly-constrained. A typical example is the definition of fixed factor supplies in classical farm-planning models solved by linear programming.

This paper demonstrates the efficacy of a new metaheuristic algorithm, compressed annealing (Ohlmann *et al.*, 2004; Ohlmann and Thomas, 2007), for the identification of profitable solutions in large, constrained solution landscapes. The portability of this procedure is aided by the free availability of code and its ease of integration with models implemented in Microsoft Excel®. Its flexibility as a computational algorithm for combinatorial systems is displayed in an application exploring the value of an annual legume pasture to dryland cropping systems in the central wheat belt of Western Australia, where price relativities and technological development have promoted the use of extended crop sequences.

The paper is structured as follows. Section 2 describes the method of compressed annealing (CA). Section 3 presents a background to the empirical application and describes the simulation model optimised through CA. Section 4 examines the value of a sown annual pasture for extended cropping sequences in the wheat belt of Western Australia using this framework. Section 5 presents a summary of results and discusses further research.

¹ The common understanding of an NP-hard problem is one that cannot be solved exactly in polynomial time.

2. Simulated-annealing (SA) algorithms

This section describes the SA metaheuristic and its extension to deal with constrained problems. Other popular metaheuristics are genetic algorithms (Coello Coello and Montes, 2002) and tabu methods (Glover, 1990). The relative superiority of each procedure is problem-dependent since, according to the no-free-lunch theorem, all search algorithms will perform equally well when averaged over all feasible cost functions (Ho and Pepyne, 2002). Metaheuristics are popular algorithms for general optimisation given the unavailability of tailored algorithms for many important combinatorial problems, including that studied here. The focus on simulated annealing in this paper is motivated by (a) the recent development of a simulated-annealing procedure for the solution of constrained problems and an associated convergence proof, (b) the lack of an equivalent framework for other metaheuristics, and (c) the high computational effort and complex mappings typically required when constraints are incorporated in a genetic algorithm (Coello Coello and Montes, 2002).

2.1 Simulated annealing

Simulated annealing (Aarts and Korst, 1989) is currently one of the most popular optimisation algorithms for the solution of discrete problems. This method is based on an analogy between the solution of a combinatorial optimisation problem and the physical process of annealing metals in condensed-matter physics. Cooling (i.e., annealing) a metal sufficiently slowly from a high heat will permit its particles to be frozen in their global-minimum energy state; here, the substance will form a strong solid. High temperatures introduce some randomness into the activity of these particles, allowing them to escape from locally-minimal configurations early in the annealing process. However, this capacity to evade local minima decreases over time as the system is cooled. The substance will reach its state of lowest energy if cooled sufficiently slowly, but will form a weak metastable glass if annealed too quickly (i.e., quenched).

Simulated annealing interprets the performance index of a combinatorial optimisation problem as the energy state of a solid. (The following description is from a minimisation perspective since physical annealing is, by nature, a minimising process.) An initial solution is generated randomly to represent the heating of a metal. The configuration of atoms in this substance (i.e., the current decision vector in an optimisation problem) is then randomly perturbed. Decreases in the energy state are always accepted (this corresponds to a decrease in cost in a minimisation problem). In contrast, increases (i.e., suboptimal movements) are also accepted, but at a probability contingent on the temperature. The configuration of atoms will converge towards its state of minimum energy if allowed to reach thermal equilibrium at a high number of temperatures along a monotonically-decreasing cooling schedule. Equivalently, if a sufficient number of iterations is performed at each level of an artificial temperature variable, and this variable (and therefore the probability of taking sub-optimal movements) is decreased sufficiently slowly over time, the SA algorithm should theoretically converge towards the global optimum of the combinatorial problem (Hajek, 1988).

Though conceptually interesting, the general applicability of annealing in agricultural economics is restricted by its inability to efficiently consider resource constraints, which are of significant importance in the accurate definition of farming systems. The simplest method of incorporating restraints in an annealing model is through only permitting legal transitions; nonetheless, this requires a prohibitive amount of coding in large and complex systems (Aarts and Korst, 1989). Alternatively, the feasibility of a candidate solution may be determined before it is enumerated, with infeasible solutions discarded before evaluation. However, this method is expensive where the configuration space is highly constrained. Another option is to permit infeasible primary transitions only if feasibility is restored by the time that a given number of secondary transitions have been performed. This practice is extremely slow, and

may even fail, in the solution of large combinatorial problems with multiple hard constraints (Abramson *et al.*, 1996).

An alternative is to penalise infeasible solutions through the use of a penalty term in the objective function. (This is equivalent to adjoining a constraint to an objective function using a Lagrangian multiplier in a standard calculus problem.) Penalty terms are important in simulated annealing since infeasible transitions provide additional paths out of local minima (Johnson *et al.*, 1989). Infeasible solutions may contain individual components that are part of the optimal solution; thus, these configurations help to connect feasible regions during the search process. This may improve convergence relative to those annealing applications that solely consider the feasible space (Angel and Zissimopoulos, 2000). In addition, the definition of a penalty function allows the faster generation of new solutions, as feasibility need not be tested before a given configuration is updated. A dynamic penalty factor is generally more flexible and efficient than a static parameter (Ohlmann *et al.*, 2004; Ohlmann and Thomas, 2007). This penalty function should be increased over time, so the probability of accepting infeasible solutions decreases as the procedure converges (Ohlmann *et al.*, 2004). This approach has been termed compressed annealing (Ohlmann *et al.*, 2004), in analogy with the effect of compression on the energy of a metal in physical annealing.

2.2 Compressed annealing

There are a number of statements necessary to correctly define a compressed-annealing problem. This description is from a maximisation perspective since the subsequent application incorporates an objective of profit maximisation.

Definition 1. *An instance of a constrained combinatorial optimisation problem is defined as the triple (Θ, ξ, π) , where $\Theta \subset R^\mu$ is the finite set of all configurations over a μ -*

dimensional space; Θ^f is a subspace of Θ incorporating all feasible solutions; Θ^{if} is a subspace of Θ incorporating all infeasible solutions; ξ is a single-valued mapping (violation function), defined on R^+ , that states the degree of infeasibility accruing to each configuration in Θ^{if} (i.e., $\xi: \Theta^{if} \rightarrow R^+$); and π is a mapping (profit function), defined on R , that assigns a single value to each configuration in Θ (i.e., $\pi: \Theta \rightarrow R$).

Definition 2. A generic constrained combinatorial optimisation problem (**OP**) may be defined $\max \pi(\theta)$ subject to $\theta \in \Theta^f$.

Definition 3. A relaxed constrained combinatorial optimisation problem (**RP**) may be defined $\max J'(\theta, \lambda_i) = \pi(\theta) - \lambda_i \xi(\theta)$, where J' is an auxiliary function and λ_i is an exogenously-defined penalty function at iteration i , where $\lambda_i \in R^+$. The optimal value of the augmented function is denoted J'_{opt} and the associated parameter vector is denoted θ'_{opt} .

Definition 4. The algorithm explores the configuration space by searching neighbourhoods $N(\theta) \subset \Theta$ for all $\theta \in \Theta$. These structures denote the set of all solutions in the neighbourhood of θ . The structure $N(\theta)$ is strongly connected, so for $\theta_i, \theta_j \in \Theta$ and $\theta_i \neq \theta_j$, there exists a path $\theta_0 = \theta_i$, $\theta_l = \theta_j$, and $\theta_{l+1} \in N(\theta_l)$ for $j = [0, 1, 2, \dots, l-1]$ (where l is finite). Moreover, $N(\theta)$ is symmetric (i.e., $\theta_j \in N(\theta_i) \Leftrightarrow \theta_i \in N(\theta_j)$).

The neighbourhood of a given configuration θ describes those points reachable from the current solution. All pathways are required to be symmetric and connected, so that all configurations in the solution space are accessible from one another. This ensures that random-search methods are appropriate for driving the optimisation process. The probability of accessing vector θ_j from θ_i given $N(\theta_i)$ is determined randomly from a generation rule.

The compressed-annealing algorithm is presented in Algorithm 1. The control parameter (i.e., temperature) is defined K . In addition, α is a uniformly-distributed random variable over the semi-open interval $[0,1)$. The general SA procedure enters Algorithm 1 as a special case where $\hat{\lambda} = 0$. Moreover, standard SA with a fixed penalty factor (of $\hat{\lambda}$ magnitude) is recovered from Algorithm 1 if γ is defined as an arbitrarily-high number.

Algorithm 1. A compressed-annealing algorithm.

Purpose: Identify an optimal decision vector θ'_{opt} .

1. Let NT be the number of trials and NA be the number of acceptances. Define an initial control parameter (K_0), the maximum number of configurations to be evaluated for each control parameter ($NTMAX$), a terminal control parameter (ε), a maximum penalty value ($\hat{\lambda}$), a maximum acceptance ratio (Γ), a control-decrement ratio (σ), and a compression rate (γ). Iterations are denoted i .
2. Set $i = 0$, $K = K_0$, and $\lambda_0 = 0$.
3. Generate a random initial configuration θ_0 and evaluate $J'(\theta_0)$.
4. Do until $K < \varepsilon$ (outer loop);
5. Do (inner loop);
 - a. Obtain θ_j from θ_i given $N(\theta_i)$;
 - b. Evaluate $J'(\theta_j)$ (obtain from simulation model);
 - c. Calculate $\Delta J' = (J'(\theta_j) - J'(\theta_i))$;
 - d. If $\Delta J' \geq 0$, then $\theta_i = \theta_j$ and $NA = NA + 1$;
 - e. If $\Delta J' < 0$ and $\exp(\Delta J' / K) > \alpha$, then $\theta_i = \theta_j$ and $NA = NA + 1$;
 - f. Else, if $\Delta J' < 0$ and $\exp(\Delta J' / K) \leq \alpha$, then $\theta_i = \theta_i$ (i.e., the current configuration remains unchanged);
6. If $NA = \Gamma \cdot NTMAX$, then Exit Do;
7. If $NT = NTMAX$, then Exit Do;
8. $NT = NT + 1$;
9. Loop.
10. $K_{i+1} = \sigma K_i$ and $\lambda_{i+1} = \hat{\lambda}(1 - e^{-\gamma})$.
11. If the terminal configuration has not changed for ten consecutive Markov chains (i.e., ten consecutive levels of the control parameter) then Exit Do.
12. $i = i + 1$.
13. Loop.
14. End.

The inner loop (steps 5–9 in Algorithm 1) is terminated at a given level of K if a maximum number of trials is performed or the number of acceptances reaches a pre-defined proportion of this maximum. The latter improves efficiency through minimising computational effort at high temperatures. This inner loop is the Metropolis algorithm (Metropolis *et al.*, 1953) that uses Monte-Carlo sampling to attain quasi-equilibrium at each value of the control parameter. This requires the evaluation of a large number of configurations at each level of K using the Boltzmann criterion. This states that solutions of higher or equal profitability are always accepted, but worse (i.e., less profitable) configurations are accepted with a probability dependent on the Boltzmann acceptance criterion, $P(\Delta J') = \exp(\Delta J' / K)$. Inferior solutions are only accepted if $\exp(\Delta J' / K) > \alpha$. The acceptance of less-profitable configurations allows escapement from the attraction basins surrounding individual local maxima.

The Boltzmann criterion states that the probability of accepting a suboptimal solution increases for (a) “less-suboptimal” configurations (i.e., smaller $-\Delta J'$), and (b) higher values of the control parameter K . This control parameter is decreased incrementally from K_0 at each iteration (according to the rule $K_{i+1} = \sigma K_i$), so the probability of accepting less-profitable solutions is reduced over time. Only solutions with equivalent or greater profitability are accepted near the usual termination of the algorithm (i.e., as $K \rightarrow 0$); here, the procedure is analogous to a random direct-search algorithm. Also, the penalty multiplier is increased according to the limited-exponential penalty function (Ohlmann and Thomas, 2007) at each iteration. This is defined as $\lambda_{i+1} = \hat{\lambda}(1 - e^{-\gamma})$, where λ_{i+1} is the penalty function at iteration $i+1$, $\hat{\lambda}$ is the maximum penalty (maximum pressure), and γ is the penalty rate (rate of compression). (The penalty function increases from zero to $\hat{\lambda}$ at a velocity determined by γ .) The procedure terminates if the control parameter is smaller than ε or the terminal configuration has stabilised for ten consecutive Markov chains. The flexibility of the

algorithm is evident in the need to only define a suitable decision vector and associated performance index, the state space does not enter the algorithm at any stage as all transitions are embedded in the simulation model.

Simulated annealing is statistically guaranteed to converge to a global optimum, but this requires either (a) the evaluation of an infinite number of configurations at each level of the control parameter, or (b) a cooling schedule that cools no faster than $T_i = T_{max} (\ln i)^{-1}$, where i is a temperature stage and T_{max} is an initial temperature of sufficient magnitude (Hajek, 1988; Ohlmann *et al.*, 2004). Both requirements are impractical for the optimisation of large combinatorial problems in finite-time (Ingber, 1993); thus, in practice, most practitioners instead adopt a “quenching” schedule. This exploits the capacity of the simulated-annealing algorithm to escape from local optima, even though the system is cooled too fast to ensure an exhaustive search of the solution space. This approach is valuable for the identification of near-optimal solutions in large combinatorial problems, particularly those that are NP-hard.

3. Herbicide resistance in the Western Australian wheat belt

3.1 Background

Traditional crop rotations in the central wheat belt of Western Australia incorporated regular phases of annual legume pasture, primarily subterranean clover (*Trifolium subterraneum* L.) or *Medicago* species, that were sown once and then regenerated after surviving a short crop phase as dormant (hard) seed. Only one or two years of consecutive crop could be grown without harming the ability of the legume to regenerate with sufficient vigour (Puckridge and French, 1983). In addition, the high relative profitability of grains encouraged short pasture phases. This rotation produced high cereal yields, improved soil structure, increased livestock

production, and permitted regular weed control through grazing and cultivation (Norman et al., 2000).

However, price relativities and technological development have promoted reductions in the proportion of pasture included in land-use rotations in this region since the early 1970s. Primary drivers of prolonged periods of cropping have been: (a) low wool prices, (b) low nitrogen fertiliser prices, (c) the availability of grain legumes, (d) the introduction of selective herbicides (those that specifically target weeds with only a minor effect on crop yield) for the in-crop control of weeds, and (e) the wide-scale adoption of reduced-cultivation methods that decrease sowing costs. These incentives have prompted the adoption of continuous- (or near-continuous-) cropping rotations in a large number of farming systems throughout this region (Reeves and Ewing, 1993).

However, though appearing economic in the short-term, extended periods of cropping detrimentally impact soil physical properties, decrease soil nitrogen reserves, limit land use diversity for the control of cereal diseases, and place heavy reliance on herbicides for weed control (Reeves and Ewing, 1993). Accordingly, significant research effort has been directed towards the development of pasture legumes for sowing between long crop sequences to overcome these potential constraints to crop production (Revell and Thomas, 2004). In contrast to ley-farming systems, legume pastures must be sown in each phase in these rotations as long cropping sequences prevent the regeneration of productive swards. Perhaps the most-promising pasture species recently introduced for these systems is the aerial-seeded annual legume, French serradella (Norman et al., 2000). This pasture is affordable to sow frequently as aerial seeding allows serradella seed to be gathered with the machinery typically used for cereal harvesting (Revell and Thomas, 2004).

Monjardino *et al.* (2004) determined the value of a three-year serradella phase in an extended crop rotation, but did not consider one- or two-year phases, which are potentially more attractive given the high relative profitability of cereal production in the study region. Moreover, these authors employed simulation to identify the most-profitable combination of weed treatments in each rotational system. This empirical application focuses on determining the value of serradella in a variety of alternative contexts through using compressed annealing to optimise a complex simulation model.

3.2 Description of the simulation model

The Ryegrass and Integrated Management (RIM) framework is used by producers and researchers to analyse the agronomic and economic implications of alternative approaches to managing the weed annual ryegrass (*Lolium rigidum* Gaud.) in the central wheat belt of Western Australia (Pannell *et al.*, 2004). RIM is a deterministic simulation model, implemented in Microsoft Excel®, that describes the multiple-cohort dynamics of ryegrass plants and seeds across a twenty-year horizon. Though focussed on weed management, this model also incorporates realistic production dynamics for the primary land uses planted in this region. These are barley (*Hordeum vulgare* L.), canola (*Brassica napus* L.), lucerne (*Medicago sativa* L.), lupins (*Lupinus* spp.), self-regenerating subterranean clover (*Trifolium subterraneum* L.), French serradella, unsown pasture, and wheat (*Triticum aestivum* L.). Accordingly, this framework permits the economics of incorporating serradella pasture in extended crop sequences in the central wheat belt to be determined with reasonable accuracy.

Weed dynamics are influenced by a number of different forms of weed control, including (a) crop rotation, (b) non-selective herbicides, (c) selective herbicides, (d) biological methods (e.g., grazing), and (e) physical control (e.g., burning). The model represents forty biological,

chemical, and physical weed treatments available at different times of the year. Not all methods may be used in each enterprise; for example, a grain crop cannot be grazed.

The RIM model used here incorporates production parameters for a typical deep sandplain soil in this region. Around five hundred parameters are incorporated in the model to describe the underlying agronomic, biological, and economic relationships (Pannell *et al.*, 2004). The practitioner specifies the number of herbicide applications that may be used from a given herbicide group before this mode-of-action is rendered ineffective because of herbicide resistance (Llewellyn and Powles, 2001).

Two state variables describe the dynamics of the ryegrass population: (a) the number of annual ryegrass seeds per square metre present in the top 5 cm of the soil in period sx in year t ($x_{sx,t}$), and (b) the number of ryegrass plants per square metre in period sw in year t ($w_{sw,t}$). Time periods (sx and sw) denote the occasions during the growing season that the state variables are calculated. There are two sets of periods, one for each state variable (Table 1).

Table 1. Index number and the description of the time during the growing season that annual ryegrass seed ($x_{sx,t}$) and plant ($w_{sw,t}$) populations are calculated.

Ryegrass seeds	Ryegrass plants
1. Immediately prior to the break of season	-
2. First chance to seed	1. First chance to seed
3. 10 days after the break of season	2. 10 days after the break of season
4. 20 days after the break of season	3. 20 days after the break of season
5. Time of post-emergent herbicide application	4. Time of post-emergent herbicide application
6. Spring	5. Early spring
7. Prior to harvest	6. Weed seed production/Prior to harvest
8. After summer	-

The number of land uses that may be grown in rotation is denoted n ; thus, $n = 8$ in this application. The producer is assumed to maximise net present value (NPV) per hectare (π),

$$\pi = \sum_{t=1}^{20} (1+r)^{-t} \left[\Theta^L(w_{5,t}, u_t^L, L_{t-1}, L_{t-2}, L_{t-3}, t) - c^{n,L} - c^{u,L} u_t^L \right] L_t, \quad (1)$$

where r is the discount rate; L_t is a $n \times 1$ indicator vector denoting the land use at time t ; L superscripts denote values specific to a given land use; Θ^L is a $1 \times n$ vector denoting enterprise revenue as a function of the weed population in spring ($w_{5,t}$), the $j_{max} \times 1$ vector of weed treatments (u_t^L) (where j_{max} denotes the maximum number of weed treatments available to a producer in each year), and prior land uses in the sequence ($L_{t-1}, L_{t-2}, L_{t-3}$); $c^{n,L}$ is a $1 \times n$ vector that denotes those costs not directly attributable to weed treatments in a given enterprise; and $c^{u,L}$ is a $1 \times j_{max}$ vector of weed-treatment costs.

The vectors L_t and u_t^L are the control variables selected exogenously by the user. Their elements are specified as unity when the associated enterprise/treatment is employed and zero when it is not. The model automatically calculates the net present value (NPV) accruing to each defined rotation and corresponding set of weed treatments.

The revenue vector ($\Theta^L(w_{5,t}, u_t^L, L_{t-1}, L_{t-2}, L_{t-3}, t)$) varies for grain and pasture enterprises. Profit for each year of pasture is calculated through the multiplication of the relevant stocking rate (measured in dry stock equivalents (DSE)) by the gross margin received for a single stock unit in a standard sheep enterprise in this region. The stocking rate depends on the type of pasture and the duration of the pasture phase. The revenue received for a grain crop in year t depends on the annual ryegrass population in that year, phytotoxic damage from selective-herbicide application, and soil fertility as determined by prior enterprises, particularly legumes.

The objective function is calculated subject to two interdependent state-transition equations that describe the population dynamics of ryegrass seeds and plants.

The initial ryegrass seed population is set exogenously, i.e., $x_{1,1} = x_0$ where x_0 is predefined.

The standard RIM model incorporates a default initial ryegrass seed burden of $x_0 = 500$ seeds m^{-2} , which is based on the average number of annual ryegrass seeds encountered in fields throughout the central wheat belt (Pannell *et al.*, 2004). In addition, $x_{1,t} = x_{8,t-1}$ for $t = [2,3,\dots,20]$. The motion equations governing the dynamics of the seed population across a single year are,

$$x_{sx+1,t} = \begin{cases} x_{sx,t} (1 - M_{seed,sx}) (1 - g_{sx}), & \text{for } sx = [1,2,\dots,5], \\ x_{sx,t} (1 - M_{seed,sx}) + \frac{sw_{5,t}^{adj} R}{w_{5,t} (\varpi + w_{5,t}^{adj} + \psi d_o)}, & \text{for } sx = 6, \\ x_{sx,t} (1 - M_{seed,sx}) (1 - f_x^L u_t), & \text{for } sx = 7, \end{cases} \quad (2)$$

where $M_{seed,sx}$ is the rate of natural mortality for ryegrass seeds in period sx , g_{sx} is the proportion of seeds germinating in period sx , s denotes the sub-lethal effect of selective herbicides, $w_{5,t}^{adj}$ is the weed population in early spring adjusted downward to represent the lower seed production of younger (later-germinating) plants, R denotes the maximum seed production of ryegrass (in seeds $m^{-2} yr^{-1}$), ϖ represents the effect of intra-specific competition on seed production, ψ represents the strength of the relationship between grain crop density (d_o) and seed production, and f_x^L is a $1 \times j_{max}$ vector with each element describing the proportion of the seed population killed by the associated weed treatment in land use L .

In contrast to the seed bank, the initial ryegrass plant population is zero in each year (i.e., $w_{1,t} = 0$ for $t = [1,2,\dots,20]$) given the annual life cycle of ryegrass.

The motion equations governing the growth of the weed population across a single year are,

$$w_{sw+1,t} = \begin{cases} w_{sw,t} (1 - M_{plant,sw}) (1 - f_w^{sw,L} u_t) + g_{sx} x_{sx,t}, & \text{for } sw = [1,2,\dots,4] \text{ and } sx = sw + 1, \\ w_{sw,t} (1 - M_{plant,sw}) (1 - f_w^{sw,L} u_t), & \text{for } sw = 5, \end{cases} \quad (3)$$

where $M_{plant,sw}$ is the rate of natural mortality rate for ryegrass plants in period sw and $f_w^{sw,L}$ is a $1 \times j_{max}$ vector with each element describing the proportion of the weed population killed by the associated treatment in period sw in land use L .

The set of control vectors must also satisfy a set of feasibility constraints. These may be denoted collectively as,

$$\Xi(u_{j,t}) \in \Theta^f. \quad (4)$$

This constraint is violated if incompatible treatments are present in the solution. For example, this may involve the production of both hay and silage in a single year or surpassing the maximum number of permissible applications available for a selective-herbicide group. In addition, a set of terminal constraints is enforced. An average seed burden is considered as 500 seeds m^{-2} . The terminal constraints are therefore (1) $x_{1,1} \geq x_{8,20}$ if $x_{1,1} \leq 500$ seeds m^{-2} , and (2) $500 \geq x_{8,20}$ if $x_{1,1} > 500$ seeds m^{-2} . These provide a strong incentive for sustained weed control and discourage mismanagement towards the end of the horizon.

The relaxed optimisation problem (**RP**) may subsequently be stated as,

$$\max_{u_{j,t}} J'(u_{j,t}, \bar{L}_t, \lambda_t) = \pi(u_{j,t}, \bar{L}_t) - \lambda_t \xi(u_{j,t}), \quad (5)$$

subject to the motion equations (2) and (3) and the binary constraint on the weed-treatment vector (i.e., $u_{j,t}^L \in \{0,1\}$). The land-use sequence is fixed (denoted \bar{L}_t) to permit the comparison of different rotational systems. This approach retains the flexibility of simulation for land-use sequencing, while employing optimisation to remove the complexity associated with the identification of the most-profitable combination(s) of weed treatments. The parameter vector is stated $\sum_{t=1}^{20} \sum_{j=1}^{j_{max}} u_{j,t}^L \in \theta$. The violation function $\xi(u_{j,t})$ is the total number of violations of $\Xi(u_{j,t}) \in \Theta^f$ in a given configuration. Equations (1)-(3) and the violation function are all calculated in the simulation model.

The RIM model is too large to guarantee the identification of a global maximum using compressed annealing. Optimisation of the continuous-cropping rotation in the following application involves the consideration of $2^{395} \approx 10^{119}$ possibilities. If it took one CPU millisecond to evaluate one solution, it would take approximately 10^{108} years to completely enumerate the configuration space. This highlights the difficulty facing a user of such a large simulation model and the value of compressed annealing for the identification of near-optimal solutions. Failure to identify a global optima reduces the precision of model output. However, the enormous size of the model suggests that this loss in precision is likely to be much smaller than that associated with transcription of the model to allow solution through mathematical programming.

The CA procedure converges to different near-optimal solutions in each run of the model for a given set of parameters given its stochastic-search mechanism and the large size of the simulation model in which it is embedded. Therefore, ten runs are performed for each problem instance. The profitability of each rotation is described by the most-valuable configuration identified over ten runs of the algorithm for each problem instance. The mean

and standard deviation over ten runs indicates the variability inherent in these computations, so are examined in Section 4.3 where the convergence of the CA procedure is compared with that of standard SA with a fixed penalty factor.

3.3 Simulated rotations

This analysis involves the evaluation of six rotations (Table 2). These are practically relevant and based on discussion with agronomists with broad experience in research and extension in the study region. The continuous-cropping rotation is based on that in Monjardino *et al.* (2004). Sowing strategies (date of sowing and seeding rates) are determined exogenously to simplify coding given the variety of enterprises considered. Seeding strategies are selected through the comparison of the profitability of alternative combinations. The subterranean clover stand is sown in the first year of the horizon and regenerates from hard seed in subsequent years. The second and third years of a multi-year serradella phase regenerate from the soil seed bank. The grazing rate of serradella is reduced in the first year of a multi-year phase to permit sufficient seed-set for subsequent regeneration.

Table 2. Simulated rotations and their associated reference terms.

Rotation	Reference term
lupin-wheat-wheat-barley	C
subterranean clover-wheat	U+W
serradella-wheat-wheat-barley	S+3C
serradella-wheat-wheat-barley-lupin-wheat-wheat-barley	S+7C
serradella-serradella-wheat-wheat-barley-lupin-wheat-wheat-barley	2S+7C
serradella-serradella-serradella-wheat-wheat-barley-lupin-wheat-wheat-barley	3S+7C

Lupin crops and pasture are the first land uses in each rotation since they permit the most efficient weed control. This increases the accuracy of model output because otherwise profit is severely constrained by an artificially high weed population, especially at high initial seed

densities, until effective treatments available within a lupin crop or pasture phase can be used. The cost of an elevated weed population is high early in the horizon, compared with later periods, due to discounting.

3.4 Parameters for the compressed-annealing algorithm

This section provides a short description of the parameters used to implement the compressed-annealing algorithm. More specific information can be obtained from www.are.uwa.edu.au/home/csa_parameters.

An initial control parameter of $K_0 = 903$ is identified using standard annealing theory (see, for example, Aarts and Korst, 1989, p. 32). The maximum penalty function value ($\hat{\lambda} = 413$) is determined by a process adapted from Ohlmann and Thomas (2007). The parameters $\sigma = 0.925$, $\Gamma = 0.5$, and $\gamma = 0.04$ are identified through extensive experimentation. This is necessary in annealing models since the generality of this algorithm complicates the development of appropriate estimation procedures for these parameters. In addition, $\varepsilon = 0.0001$. However, the compressed-annealing procedure applied here always converges before the control parameter falls beneath this level.

The neighbourhood structure treats each admissible weed-management decision as a member of a string of length $\mu = 20 \cdot j_{\max}$ as there are twenty years in the planning horizon. Here, $u_i = \{0,1\}$ (where $i = [1,2,\dots,\mu]$), with $u_i = 1$ representing the use of a treatment and $u_i = 0$ representing no use of this treatment. Each trial consists of generating a random number $\beta = \mu\eta$ (where η is a uniformly distributed random variable on $[0,1]$) and performing the update $u_{\beta} = 1 - u_{\beta}$.

The maximum number of trials ($NTMAX = 480$) is the largest string length encountered over the evaluated rotations and thus ensures, at least theoretically, that each treatment is visited at each level of the control parameter. The binary control variables (i.e., $u_i = \{0,1\}$) prevent the use of alternative generation mechanisms (those processes used to update the decision vector) to speed the convergence of the algorithm. For example, the Ingber distribution may be used in problems with continuous parameters to guarantee the identification of global optima using the standard quenching schedule $T_i = T_{max} \exp((c - 1)i)$, where $0 < c < 1$ (Ingber, 1993).

The compressed-annealing procedure is coded in Visual Basic for Applications (VBA) (Walkenbach, 2003) since the simulation model it is used to optimise is implemented in Microsoft Excel®. Longer annealing times than those described by the compressed-annealing parameters stated in this section consistently fail to identify superior configurations. This suggests that the adopted parameters are sufficient. The application of Algorithm 1 in the following empirical analysis required two pages of code, which is available from the primary author on request. The ease with which this algorithm is integrated into a spreadsheet environment increases its attractiveness as a tool for applied economic analysis.

4. Results and discussion

This section presents an empirical application of the algorithm and model presented in the previous sections.

4.1 Standard model results

This section presents model output for the standard (i.e., base) set of parameter values and discusses its implications for the management of phase-farming systems in the study region. Profitability is defined in terms of NPV ha⁻¹ according to equation (1).

The profitability of each rotation is presented for a range of initial ryegrass seed densities in Table 3. These densities differ greatly from a small, but significant, 100 seeds m⁻² to an immense 10,000 seeds m⁻². The incorporation of different densities is necessary to generalise the results of this analysis and ascertain the association between the initial seed population and the relative profitability of the simulated rotations.

Table 3. Value of each rotation (NPV ha⁻¹) for a range of initial ryegrass seed densities.

Rotation	Initial ryegrass seed density (seeds m ⁻²)						
	100	250	500	1,000	2,500	5,000	10,000
C	693 ¹	698	695	689	672	670	643
S+7C	685	691	681	676	683	673	667
S+3C	673	664	659	654	659	653	649
3S+7C	671	667	661	662	669	668	665
2S+7C	624	631	635	621	540	533	556
U+W	503	518	509	523	513	519	514

¹ A figure in bold for a rotation represents that this is the most-valuable sequence at that initial seed density.

There is a weak relationship between the initial seed density and the profitability of those rotations incorporating pasture. Annual ryegrass provides valuable early-winter feed, but can also poison stock under certain circumstances. Consequently, its value as a grazing plant is not included in this analysis. In addition, intensive weed management in the initial serradella phases, particularly through grazing and the removal of this pasture with a non-selective herbicide (brown-manuring), greatly reduces the seed population, minimising the effect of weed competition in the subsequent cropping phase. The size of the model also prevents the identification of a global optimum for each problem instance; thus, the need to “quench” the annealing algorithm to discover near-optimal solutions in a practicable period of time introduces some inherent noise into the results reported in Table 3. Additional runs identify that such mild inconsistencies remain for alternative compression and cooling schedules (data not shown).

Continuous cropping is more profitable than those rotations incorporating pasture at initial seed densities up to and including 1,000 seeds m^{-2} (Table 3). All available applications of Group A and B selective herbicides are used in each optimal configuration identified for this sequence. These applications are supplemented by the use of Group C and D chemicals, though less than four applications of each of these herbicide groups are ever used across the twenty-year horizon. The heavy reliance placed on selective herbicides in the continuous-cropping rotation is consistent with evidence that herbicide resistance constrains weed management in extended cropping sequences throughout this region (Llewellyn and Powles, 2001).

Selective herbicides are an efficient form of weed control, but the incorporation of a pasture phase becomes more profitable as the initial seed density increases. This is evident in the superior profitability of the S+7C rotation at initial seed burdens exceeding 1,000 seeds m^{-2} . Moreover, the continuous-cropping rotation is less profitable than three sequences incorporating a pasture phase at the highest initial weed density considered (i.e., 10,000 seeds m^{-2}) (Table 3). This follows directly from the weak relationship between weed density and the profitability of those rotations including pasture highlighted above. For example, in the continuous-cropping rotation, profit decreases by 8.3 percent when moving from an initial seed density of 100 seeds m^{-2} to 10,000 seeds m^{-2} , while in the S+7C rotation it decreases by only 2.6 percent. (A higher initial seed population represents a situation where a producer struggles to adequately control seed burdens, for example because of the variable efficacy of non-chemical weed treatments.) However, the S+7C sequence is more profitable than continuous cropping at an initial seed density of 500 seeds m^{-2} when the yield of those crops following the legume pasture is promoted by the magnitude inferred by field experiments implemented on an acidic sandy soil in Revell and Thomas (2004). This highlights the (a) usefulness of serradella as an economic break pasture in phase-farming systems in the central

wheat belt, and (b) the need for more accurate information regarding the extent to which the yield of subsequent crops is enhanced by serradella.

The high comparative profitability of cereal production in these farming systems suggests an inverse association between the value of a rotation and the proportion of pasture it incorporates. However, this relationship is not strongly evident in Table 3. The least-profitable sequence is the traditional ley-farming sequence (i.e., the U+W rotation) that consists half of pasture. Yet, the S+7C rotation is the most-profitable sequence at a high initial weed burden since it incorporates a high proportion of crop and the single year of pasture is valuable for weed management. Moreover, the S+3C and 3S+7C sequences are both more valuable than the 2S+7C rotation (Table 3), despite each of them incorporating more years of pasture over the planning horizon.

An isolated year of serradella in the S+3C sequence obtains high weed control, has a low opportunity cost by virtue of its short duration, and supports heavier grazing rates than admissible in the subterranean-clover phase as the pasture is not required to regenerate from hard seed. In contrast, establishment costs and the need for light grazing to permit sufficient pasture regeneration decrease the profitability of the first year of a multiple-year serradella phase. The high production of a three-year serradella stand and the benefits of an extended pasture phase for improving soil fertility help to compensate for the high costs incurred in its first year. In comparison, the magnitude of these benefits is not sufficient to offset the first-year loss in the 2S+7C system because of its shorter length. This occurs despite the higher proportion of crop incorporated in this rotation. These results highlight that the value of pasture for farming systems in the study region is multifaceted. The primary benefits considered here are (1) improved weed control, (2) nitrogen fixation and increases in soil organic matter, (3) reduced disease incidence, and (4) higher grazing income.

Table 4 presents the specific strategies identified in the first three runs of the CA algorithm (these runs are arranged in order of increasing net present value) for the C rotation at the standard set of parameter values. (Table 4 is reported for an initial ryegrass density of 500 seeds m⁻².) This provides more detailed insight into the nature of the solutions identified by the CA procedure.

Table 4. Treatments used in the optimal configurations identified in the first three runs of the CA procedure for the C rotation.

Run 1 (NPV=\$663)	Run 2 (NPV=\$674)	Run 3 (NPV=\$685)
Treatments	Treatments	Treatments
15 applications of knockdown herbicide	15 applications of knockdown herbicide	15 applications of knockdown herbicide
2 applications of Hoegrass® (Group A fop)	2 applications of Hoegrass®	2 applications of Hoegrass®
2 applications of Select® (Group A dim)	2 applications of Select®	2 applications of Select®
1 post-emergent application of Glean® (Group B)	2 pre-emergent applications of Glean®	2 pre-emergent applications of Glean®
1 post-emergent application of simazine (Group C)	1 post-emergent application of simazine	2 post-emergent applications of simazine
3 applications of trifluralin (Group D)	2 applications of trifluralin	1 application of trifluralin
Swathe barley twice	Swathe barley twice	Swathe barley twice
Use seed catching 7 times and windrowing 7 times	Use seed catching 4 times and windrowing 9 times	Use seed catching 5 times and windrowing 8 times

The alternative solutions reported in Table 4 are typified by common (i.e., core) forms of weed management and a number of heterogenous treatments. For example, non-selective herbicides are applied prior to sowing in every cereal crop and harvest treatments, such as seed-catching, are used in nine out of the first ten years in all solutions for this scenario. However, the management of selective herbicides is somewhat different, each solution employing a different number of applications of Group C (i.e., simazine) and Group D (i.e., trifluralin) chemicals. The timing of these applications also varies between model runs. The

existence of a common core set of management strategies gives confidence that these elements should be adopted. In addition, the non-core strategies (those that are selected in some model runs but not others, or with differing frequencies) provide an opportunity for producers to select from them according to personal preference, without significantly penalising the profitability of their farming system.

4.2 Implications of higher livestock profitability

Stocking rate is a key determinant of the profitability of pasture enterprises in Australian farming systems. Although the results discussed earlier are realistic for the assumed level of livestock profitability, this does vary significantly from time to time. The value of each rotation is presented in Table 5 for a range of alternative sheep gross margins that could be sustained over a twenty-year period. (Table 5 is calculated for an initial ryegrass density of 500 seeds m⁻².) For comparison, estimated gross margins for 2002 and 2005 are \$21.50 per dry sheep equivalent (DSE) (Agriculture Western Australia, 2002) and \$15.69 DSE⁻¹ (Agriculture Western Australia, 2005) respectively.

Table 5. Value (NPV ha⁻¹) of each rotation for a range of sheep gross margins.

Rotation	Sheep gross margin (\$ DSE ⁻¹)						
	10	12.50	15 [base value]	17.50	20	22.50	25
C	695¹	695	695	695	695	695	695
S+7C	668	671	681	683	702	712	733
S+3C	617	642	659	671	709	720	747
2S+7C	591	594	635	637	672	674	717
3S+7C	570	629	661	669	716	736	764
U+W	411	453	506	556	598	669	705

¹ A figure in bold represents the most-valuable rotation at that sheep gross margin.

The profitability of those rotations containing pasture is strongly related to the sheep gross margin (Table 5). The U+W rotation remains the least-profitable configuration, despite the

inflated worth of livestock production. This reflects the low stocking rate of regenerating subterranean clover and the low proportion of crop in this sequence. The continuous-cropping rotation is the most-profitable sequence up to a sheep gross margin of \$17.50. However, the S+3C and 3S+7C rotations are both more profitable than the S+7C and C rotations at sheep gross margins above \$17.50 DSE⁻¹ as the high profitability of livestock husbandry offsets the loss of additional years of crop over the planning horizon. The profitability of the S+3C sequence is also favoured by the value of regular pasture phases for annual ryegrass management. The most-valuable rotation for sheep gross margins above \$17.50 DSE⁻¹ is the 3S+7C sequence. The extended serradella phase has a high stocking capacity, but sheep production is required to be highly lucrative before the displacement of a high number of cropping years is justified. However, the importance of grazing to the profitability of an extended serradella phase yields it sensitive to large declines in the sheep gross margin. For example, profit decreases by nearly 15 percent for the 3C+7C rotation as the gross margin is reduced from \$15 DSE⁻¹ to \$10 DSE⁻¹ (Table 5). In addition, the profitability of the U+W sequence falls by around 19 percent as it contains a high proportion of pasture.

The mean sheep gross margin over the twenty-year period must be around the \$21.50 DSE⁻¹ estimate (Agriculture Western Australia, 2002) formulated for the study area in a year of highly-favourable sheep meat prices in order for those rotations containing serradella to dominate the continuous-cropping sequence (Table 5). The mean sheep gross margin is unlikely to remain so high over a twenty-year period, particularly since anthelmintic resistance threatens to reduce profit by around \$6.60 DSE⁻¹ over the next decade (Besier and Love, 2003) and wool prices, on average, are predicted to halve in real terms by 2029 (Sackett, 2004). In contrast, cereal production is forecast to remain highly profitable in Australia due to increased demand in developing countries (Food and Agriculture Organisation, 2002). Therefore, improvements in the profitability of livestock husbandry

obtained through enhanced genetic selection, management practices, marketing strategies, and product quality must indeed be substantial to encourage the adoption of regular pasture phases.

4.3 Convergence of the CA algorithm

Mean profit and its standard deviation are computed for each rotation (at an initial ryegrass density of 500 seeds m⁻²) using the CA algorithm and a standard SA procedure incorporating different levels of a fixed penalty factor.² (This approach is similar in principle to the comparative method implemented by Ohlmann and Thomas (2007).) These summary statistics are more appropriate than the value of the most-profitable configuration as they provide some indication of the convergence performance of these metaheuristics over a given number of runs. These results are presented in Table 6.

Table 6. Value of each rotation identified by compressed annealing and a standard simulated-annealing algorithm incorporating alternative levels of a fixed penalty factor.

Rotation	Algorithm			
	CA	SA ($\lambda = 137$)	SA ($\lambda = 275$)	SA ($\lambda = 413$)
C	674 ¹ (14.7) ²	-1659 (1454.1)	-7395 (3737.2)	-9856 (4096.7)
S+7C	649 (10)	655 (6.46)	647 (16.2)	-3625 (6358.9)
S+3C	621 (12.6)	629 (2.71)	-7252 (1293.8)	-8753 (3829.8)
3S+7C	620 (10.5)	75 (846)	606 (18.9)	-8578 (3831.3)
2S+7C	605 (16.3)	595 (16.8)	581 (40.8)	-9017 (5640.1)
U+W	496 (12.5)	-2582 (1671.1)	-6255 (3335.2)	-9085 (4739.7)

¹ Mean profit calculated to no decimal places.

² Standard deviation of profit calculated to one decimal place.

A higher and more stable level of mean profit is identified for the S+3C and S+7C rotations using standard simulated annealing and a fixed penalty value of $\lambda = 137$ (Table 6). However,

² The adopted set of fixed penalty values ($\lambda=\{137,275,413\}$) represents different proportions of the estimated maximum compression value. Lower values are inappropriate since errors are not penalised sufficiently to prevent them from recurrently entering optimal configurations. Higher penalty values are not required since all rotations failed to converge at $\lambda=413$; here, infeasible trials are penalised too greatly for sufficient exploration of the solution space.

the differences are marginal. For example, mean profit for the S+7C rotation increases by less than 1 percent. In addition, not all rotations converge to a feasible configuration with a fixed penalty factor. This is evident in the presence of negative values, which signify convergence to a heavily-penalised solution, for some rotations in all columns for the standard SA results. (Losses increase with higher levels of the fixed penalty factor in Table 6 because infeasibilities in the final configuration are obviously penalised by a greater amount.) In contrast, the CA approach has the important advantage of flexibility as it reliably converges for each rotation in Table 6. This flexibility seems to be the primary benefit of using compressed annealing in the place of standard annealing with a fixed penalty factor.

The standard deviation of profit is reasonably low for those solutions identified by CA (Table 6). The capacity of this algorithm to consistently identify solutions of comparable quality infers that this procedure is identifying areas of the configuration space that are in the vicinity of the true optimum. This motivates the hypothesis that the pay-off hypersurface is reasonably flat near its peak (Pannell, 2006) for the defined sequences and that multiple runs of a compressed-annealing algorithm may be used to effectively delineate this region.

Ohlmann and Thomas (2007) found the use of CA markedly decreased solution time (up to 41 percent in one case) in a number of constrained Travelling Salesman Problems. This outcome is not evident in this study where all runs converge in a similar number of iterations (around 110). This may reflect the size of the configuration space that constrains the capacity for either algorithm to conduct a faster search than the other. Stochastic-search procedures are generally much less efficient than gradient-based algorithms as, by definition, metaheuristics use no curvature properties to guide exploration. Consequently, high computational times are perhaps the greatest constraint accruing to the application of annealing algorithms, though a

single run of the standard model (i.e., those in Section 4.1) converges in less than 10 minutes on a Pentium IV 3.2 GHZ desktop computer with 2 GB of memory.

5. Summary and conclusions

The empirical application presented in this paper demonstrates the value of compressed annealing for the identification of near-optimal configurations in large, constrained simulation models. Metaheuristic algorithms are valuable tools for the optimisation of such models since the processes incorporated therein are left intact and not approximated via econometric techniques. Conceptually, the consideration of infeasible solutions in simulated annealing helps to improve convergence as these provide additional paths out of local optima (Johnson *et al.*, 1989). However, the primary benefit of a monotonically-increasing penalty function is shown to be its flexibility relative to a static parameter. The variability present in the near-optimal configurations identified by the annealing procedure arises naturally from the need to accelerate the annealing process in order to obtain results in a practicable period of time. This inherent variation may be exploited in an applied sense in that producers may select from a number of possible strategies of similar profitability. The presence of common characteristics in many solutions also highlights some capacity for the formulation of integrated weed management strategies for extension and further evaluation.

The empirical analysis presented in this paper focuses on determining the value of serradella pasture for improving the profitability of dryland cropping systems in the central wheat belt of Western Australia. Incorporation of a single year of serradella pasture in a crop rotation is more profitable than a representative continuous-cropping system at high initial seed densities. However, inclusion of this pasture is justified at lower weed burdens if the yield of subsequent crops is enhanced by the magnitude inferred by field experiments. A three-year phase of serradella is also competitive at high initial seed burdens, particularly if livestock

production is highly profitable. Serradella is costly to establish and must be lightly grazed if required to regenerate. The high production in the second and third years of a three-year serradella stand and its significant investment in soil organic matter help to offset these costs. In contrast, the shorter duration of a two-year phase does not provide sufficient compensation, despite displacing fewer years of crop in the rotation.

Notwithstanding these insights, there are two important extensions to this analysis of phase-farming systems worthy of examination. First, the value of serradella for weed management in the presence of established herbicide resistance has not been determined. The existence of herbicide resistance may be expected to enhance the value of pasture phases since these allow the use of cost-effective forms of non-selective control, such as brown-manuring and grazing. Second, the comparative value of perennial pastures, such as lucerne (*Medicago sativa* L.), has not been considered. This is particularly relevant since the shorter growing season and shallow rooting depth of annual crops and pastures in the central wheat belt reduces their capacity to intercept rainfall, contributing to the development of soil salinisation.

The compressed-annealing algorithm may also be extended in a number of directions. First, the land-use sequence may be selected optimally through including its determination in the sampling procedure. However, this would result in a much larger solution time and also complicate the construction of an appropriate search algorithm given the relationship between a land-use rotation and the set of weed treatments employed therein. Second, extension of the CA procedure to sample the parameter space more efficiently through the adoption of new generation mechanisms will increase the speed and the efficacy of the algorithm. The binary nature of the control variables incorporated in the RIM model hampers the use of such practices here. Third, integration of stochastic simulation models with metaheuristic

procedures developed for the optimisation of variable systems (e.g., Andradottir, 1995) would aid the identification of profitable configurations in these frameworks.

References

- Aarts, E.H.L. and Korst, J. (1989), *Simulated annealing and Boltzmann machines: a stochastic approach to combinatorial optimisation and neural computing*, Wiley, New York.
- Abramson, D., Dang, H. and Krishnamoorthy, M. (1996), 'A comparison of two methods for solving 0-1 integer programs using a general purpose simulated annealing algorithm', *Annals of Operations Research* 63, pp. 129-50.
- Agriculture Western Australia (2002), *Gross margins guide 2003: Western Australia*, Agriculture Western Australia Miscellaneous Publication 18/2002, Perth.
- Agriculture Western Australia (2005), *Gross margins guide 2005: Western Australia*, Agriculture Western Australia, Perth.
- Andradottir, S. (1995), 'A method for discrete stochastic optimisation', *Management Science* 41, pp. 1946-61.
- Angel, E. and Zissimopoulos, V. (2000), 'On the classification of NP-complete problems in terms of their correlation coefficient', *Discrete Applied Mathematics* 99, pp. 261-77.
- Besier, R.B. and Love, S.C.J. (2003), 'Anthelmintic resistance in sheep nematodes in Australia: the need for new approaches', *Australian Journal of Experimental Agriculture* 43, pp. 1383-91.
- Coello Coello, C.A. and Montes, E.M. (2002), 'Constraint-handling in genetic algorithms through the use of dominance-based tournament selection', *Advanced Engineering Informatics* 16, pp. 193-203.
- Du, D. and Ko, K. (2000), *Theory of computational complexity*, Wiley, New York.
- Food and Agriculture Organisation (2002), *World agriculture: towards 2015/2030*, Food and Agriculture Organisation, Rome.
- Glover, F. (1990), 'Tabu search: a tutorial', *Interfaces* 20, pp. 74-94.
- Hajek, B. (1988), 'Cooling schedules for optimal annealing', *Mathematics of Operations Research* 13, pp. 311-29.
- Ho, Y. and Pepyne, D.L. (2002), 'Simple explanation of the no-free-lunch theorem and its implications', *Journal of Optimisation Theory and Applications* 115, pp. 549-70.
- Ingber, L. (1993), 'Simulated annealing: practice versus theory', *Mathematical Computing and Modelling* 18, pp. 29-57.
- Johnson, D.S., Aragon, C.R., McGeoch, L.A. and Schevon, C. (1989), 'Optimisation by simulated annealing: an experimental evaluation. Part I: graph partitioning', *Operations Research* 37, pp. 865-92.

- Law, A.M. and Kelton, W.D. (2000), *Simulation modelling and analysis*, McGraw-Hill, New York.
- Llewellyn, R.S. and Powles, S.B. (2001), 'High levels of herbicide resistance in rigid ryegrass (*Lolium rigidum*) in the wheat belt of Western Australia', *Weed Technology* 15, pp. 242-48.
- Mayer, D.G., Belward, J.A. and Burrage, K. (1996), 'Use of advanced techniques to optimise a multi-dimensional dairy model', *Agricultural Systems* 50, pp. 239-53.
- Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A. and Teller, E. (1953), 'Equation of state calculations by fast computing machines', *Journal of Chemical Physics* 21, pp. 1087-92.
- Monjardino, M., Pannell, D.J. and Powles, S.B. (2004), 'The economic value of pasture phases in the integrated management of annual ryegrass and wild radish in a Western Australian farming system', *Australian Journal of Experimental Agriculture* 44, pp. 265-71.
- Norman, H.C., Ewing, M.A., Loi, A., Nutt, B.J. and Sandral, G.A. (2000), 'The pasture and forage industry in the Mediterranean bioclimates of Australia', *Options mediterraneennes* 45, pp. 469-76.
- Ohlmann, J.W., Bean, J.C. and Henderson, S.G. (2004), 'Convergence in probability of compressed annealing', *Mathematics of Operations Research* 29, pp. 837-60.
- Ohlmann, J.W. and Thomas, B. (2007), 'A compressed annealing approach to the travelling salesman problem with time windows', *INFORMS Journal on Computing* in press.
- Pannell, D.J. (2006), 'Flat-earth economics: the far-reaching consequences of flat payoff functions in economic decision making', *Review of Agricultural Economics* 28, pp. 553-66.
- Pannell, D.J., Stewart, V., Bennett, A., Monjardino, M., Schmidt, C. and Powles, S.B. (2004), 'RIM: a bioeconomic model for integrated weed management of *Lolium rigidum* in Western Australia', *Agricultural Systems* 79, pp. 305-25.
- Plambeck, E.L., Fu, B.R., Robinson, S.M. and Suri, R. (1996), 'Sample-path optimisation of convex stochastic performance functions', *Mathematical Programming* 75, pp. 137-76.
- Puckridge, D.W. and French, R.J. (1983), 'The annual legume pasture in cereal-ley farming systems of Southern Australia: a review', *Agriculture, ecosystems and environment* 9, pp. 229-67.
- Reeves, T.G. and Ewing, M.A. (1993), 'Is ley farming in Mediterranean zones just a passing phase?' in *Proceedings of the XVII International Grasslands Congress, Hobart, Tasmania*, International Grasslands Society, Quebec, pp. 2169-77.
- Revell, C.K. and Thomas, D.K. (2004), 'Management of crop weeds through the strategic use of annual pasture', in Sindel, B.M. and Johnson, S.B. (Eds), *Proceedings of the 14th Australian Weeds Conference*, Weed Society of New South Wales, Wagga Wagga, pp. 145-52.
- Sackett, D. (2004), *Will wool growing be a viable business in 2029?*, Holmes, Sackett and Associates, Sydney.
- van Kooten, G.C., Young, D.L. and Krautkraemer, J.A. (1997), 'A safety-first approach to dynamic cropping decisions', *European Review of Agricultural Economics* 24, pp. 47-63.
- Walkenbach, J. (2004), *Excel 2003: power programming with VBA*, Wiley, Indianapolis.

Woodward, R.T., Wui, Y.-S. and Griffin, W.L. (2005), 'Living with the curse of dimensionality: closed-loop optimisation in a large-scale fisheries simulation model', *American Journal of Agricultural Economics* 87, pp. 48-60.