# Report Cover Page

| ACERA Project | | |
|---|---|---|
| 0902 | | |
| **Title** | | |
| Strategies for managing invasive species in space: deciding whether to eradicate, contain or control | | |
| **Author(s) / Address (es)** | | |
| Iadine Chades – University of Queensland. i.chades@uq.edu.au | | |
| **Material Type and Status (Internal draft, Final Technical or Project report, Manuscript, Manual, Software)** | | |
| Final Project report | | |

## Summary


Invasive species are a major threat to ecosystems worldwide. Once invasive species have established, even with a determined commitment to control, contain or eradicate an invasive species, it is often difficult to decide on the most efficient and effective management strategy due to the complex interaction of factors such as the extent of the invasion, the ecology of the species, the dynamics of the system, and how the species responds to different management actions. Invasive species by their very nature disperse through a system at a rapid rate, indeed their spread and persistence is often driven by a spatial network of infestations in the landscape. If spatial processes are ignored in our decision-making process we may allocate time and money managing individual infestations while key areas driving the spread of the species are inefficiently controlled. Under such a scenario attaining our goals of eradicating or controlling spread of an invasive species may not be attained, having dire consequences for environments we are trying to protect.

Prioritising the management of invasive species across the landscape to maximise our chance of eradicating an invasive species can be modelled as a Markov decision processes (MDP) and solved using stochastic dynamic programming methods. Managing invasive species spatially is a challenging theoretical problem due to Bellman's curse of dimensionality: the computational complexity grows exponentially with the number of parcels in the system. Another practical issue arises with classic SDP method; the solution proposed only present one optimal strategy amongst a set of possible alternatives constraining decision managers to one option. Using networks as a spatial representation of our optimisation problem we overcome these drawbacks using structured MDP and algebraic decision diagrams (ADD) representations. By doing so we take advantage of the structured colonisation processes between parcels and provide for the first time the set of all optimal strategies to manage invasive species on a network. Under general assumptions we derive rules of thumb for some key network structures.

Using MDP assumes we can observe the system perfectly and determine which parcel is infected or not. In reality we often do not have access to this information. Partially observable Markov decision process (POMDP) model tackles the problem of deciding the most appropriate management action given the difficulty of detecting invasive species. Work by Regan et al (in review) has looked into the optimal management of an invasive weed branched broomrape (Orobanche ramosa) under imperfect detection addressing the key issue of management at a single parcel level. Here, we account for the spatial components of the system and how spatial considerations may affect optimal management of a cryptic invasive species. Using structured POMDP and network representation, we investigate the role space plays in optimally managing invasive species and produce the first spatially explicit decision support tool for managing cryptic invasive species. We provide general rules of thumb that can be applied to a range of species in diverse situations.


| | | |
|---|---|---|
| **ACERA Use only** | Received By: | Date: |
| | ACERA / AMSI SAC Approval: | Date: |
| | DAFF Endorsement: ( ) Yes ( ) No | Date: |

# Strategies for managing invasive species in space: deciding whether to eradicate, contain or control
## ACERA Project No. 0902

Iadine Chades; University of Queensland

Final Report

1st of December 2009

# Final VERSION

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

# Acknowledgements

This report is a product of the Australian Centre of Excellence for Risk Analysis (ACERA). In preparing this report, the authors acknowledge the financial and other support provided by the Department of Agriculture, Fisheries and Forestry (DAFF), the University of Queensland, the Centre of Excellence for Mathematics and Statistics of Complex Systems (MASCOS), the Applied Environmental Decision Analysis (AEDA) research hub, the University of Melbourne, Australian Mathematical Sciences Institute (AMSI), Australian Research Centre for Urban Ecology (ARCUE) and CSIRO sustainable Ecosystems. Iadine Chades would like to thank Regis Sabbadin, Olivier Buffet, Samuel Nicol, Mike Runge, Tara Martin, Eve McDonald-Madden, Peter Baxter, Cindy Hauser, Ross McVinish, Hugh Possingham, Phil Pollett, Yvonne Buckley and Tracey Regan for insightful discussions.

# Disclaimer

# Table of contents

# List of Tables

# List of Figures

# 1. Executive Summary

Our work makes two major contributions to the field of optimal management of invasive species across space and over time i) when we assume perfect detection of the invaders ii) under imperfect detection of the invaders. For both cases we answer the question where should we manage given a state of infection to maximise our chance of eradicating an invasive species in the most cost-efficient way. We derived rules of thumb for the management of motifs and small structured networks (line, source-sink, island and star networks). We developed a general decision tool to solve any invasive network management problem of 15 nodes or less when detection is perfect and 6 parcels or less when detection is imperfect. We solved this challenging problem by taking advantage of the structure of networks. Our decision tool is general and can be adapted to any kind of invasive species management problem.

When detection of invasive species is perfect, we recommend using structured Markov decision process with algebraic decision diagrams to guide the prioritisation of management of invasive species on a network. We found the following rules of thumb for the management of invasive species across space when all parcels are infected:

- Source-sink or directed networks: manage the source first and the sink last.
- Undirected line networks: start from an extremity of the line and keep managing parcels following the same direction.
- Star networks: manage the infected satellite parcels until the number of satellite parcels empty is equal to or greater than the number of satellite parcels infected. Then manage the central parcel and the remaining infected satellite parcels.
- Island networks: start managing anyone parcel and then manage the closest parcels in any direction.

- Island-line networks: start managing the extreme parcel of the network, manage the line, manage the connecting node and manage successfully the parcels that are the closest to the connecting node. If an island network is connected to several lines start managing from the longest line.
- Cluster network: start from the smaller cluster, if the clusters are identical start from any cluster, if a cluster is less connected start from the least connected cluster.

When detection of invasive species is imperfect, we recommend using structured partially observable Markov decision process with algebraic decision diagrams to guide the prioritisation of management and surveillance of cryptic invasive species on a network. The POMDP solution outperforms the MDP solution.

A general rule of thumb when invaders are present across the landscape is to follow the management order derived from the MDP solutions. The POMDP solution usually recommends repeating the management of the most connected parcels even when invaders remain unobserved. This tendency increases as the number of parcels increases and the detection probability or/and the management efficiency decreases.

When we do not have access to prior information on the invasion and the invaders are unobserved, the POMDP solution recommends managing the most connected parcels in priority. Similarly if we were unable to manage a parcel when invaders are unobserved we would survey the most connected parcels in priority.

# 2. Introduction

Invasive species are a major threat to ecosystems worldwide. Once invasive species have established, even with a determined commitment to control, contain or eradicate an invasive species, it is often difficult to decide on the most efficient and effective management strategy due to the complex interaction of factors such as the extent of the invasion, the ecology of the species, the dynamics of the system, and how the species responds to different management actions. Invasive species by their very nature disperse through a system at a rapid rate, indeed their spread and persistence is often driven by a spatial network of infestations in the landscape. If spatial processes are ignored in our decision-making process we may allocate time and money managing individual infestations while key areas driving the spread of the species are inefficiently controlled. Under such a scenario attaining our goal of eradicating or controlling the spread of an invasive species may not be attained, having dire consequences for environments we are trying to protect.

Prioritising the management of invasive species across the landscape to maximise our chance of eradicating an invasive species can be modelled as a Markov decision processes (MDP) and solved using stochastic dynamic programming (SDP) methods. Managing invasive species spatially is a challenging theoretical problem due to Bellman's curse of dimensionality: the computational complexity grows exponentially with the numbers of parcels in the system. Another practical issue arises with classic SDP methods; a single optimal strategy is proposed amongst a set of possible alternatives constraining decision managers to one option. Using networks as a spatial representation of our optimisation problem we overcome these drawbacks using structured MDP (Boutilier et al. 1999) and algebraic decision diagram (ADD) representations (Bahar et al. 1997, Hoey et al. 1999). By doing so we take advantage of the structured colonisation processes between parcels and provide for the first time the set of all optimal strategies to manage invasive species across a network. Under general assumptions we derive rules of thumb for a number of key network structures.

Determining the prioritisation of management of an invasive species as an MDP assumes we can observe our system perfectly and determine which parcel is occupied by an invasive species. In reality this is rarely the case. Partially observable Markov decision process (POMDP) model tackles the problem of deciding the most appropriate management action given the difficulty of detecting invasive species. Work by Regan et al (in review) has looked into the optimal management of an invasive weed, branched broomrape (Orobanche ramosa), under imperfect detection addressing the key issue of management at a single parcel level. Here, we account for the spatial components of the system and how spatial considerations may affect optimal management of a cryptic invasive species. Our objective is to provide a general tool and rules of thumb that can be applied to a range of species in diverse situations. Finding an exact solution to a POMDP is intractable for most problems. However, approximate methods have recently been successfully applied to very complex problems (Boger et al. 2005) prompting an increasing interest in using POMDP. Here we tackle an even more difficult problem, solving a spatial POMDP i.e. a POMDP defined over a spatially explicit problem. Using structured POMDP and network representation, we investigate the role space plays in optimally managing invasive species and produce the first spatially explicit decision support tool for managing cryptic invasive species. We provide general rules of thumb that can be applied to a range of species in diverse situations.

Our report is organised in two parts. We first model and solve our decision problem in the completely observable case using Markov decision process (MDP). In doing so, we describe rules of thumb for managing small size networks when we are perfectly able to detect the invasive species. Given a state of infection, we answer the question which parcel should we manage first? The second part introduces the challenging problem of imperfect detection of an invasive species. On a single parcel we have shown that even under optimal management the probability of eradication of branched broomrape drops to unsatisfactory levels when colonisation is present (Regan et al. in review). Building on these two previous works, we tackle the problem of imperfect detection and spatial management. In solving spatial POMDP on small size networks, we provide guidance to decision managers on where they should manage parcels on an infected network, how long and where they should manage if the invaders are not observed and finally where should they survey in priority for a maximum cost efficiency strategy.

Iadine's ACERA funding has also contributed to a book chapter and four collaborative projects tackling optimisation problems for the management of invasive or endangered species. We provide a list of peer reviewed publications in press, in review or about to be submitted (available upon request):

1. **Chadès, I.** (2010) Markov decision processes in Artificial Intelligence. Chapter 12: Conservation of biodiversity. Hermes Science Publishing, London, United Kingdom.
2. Grechi, I., **Chadès, I**., Buckley, Y., Friedel, M., Grice, T., Possingham, H.P., van Klinken, R. and T.G. Martin Optimal management of commercial invasive species: the case of Buffel Grass Cenchrus ciliaris. *In prep.* Journal of Applied Ecology.
3. Regan*, T.J., **Chadès*, I.**, and H.P. Possingham. Optimal strategies for managing invasive plants in partially observable systems. *In review*. Journal of Applied Ecology (*contributed equally)
4. McDonald-Madden, E., **Chadès, I.**, McCarthy, M.A., Linkie, M. and H.P. Possingham. Allocating conservation resources between areas where persistence of a species is uncertain. *In review*. Ecological applications.
5. Nicol, S., **I. Chadès**, S. Linke and H. P. Possingham. (2009). Conservation decision-making in large state spaces. MODSIM 2009, International Congress on Modelling and Simulation. Also accepted for publication in Ecological Modelling.

# 3. Managing an invasive species in space with perfect detection

## 3.1 Introduction

The role of space in the management of invasive species is critical at the early stage of an invasion process; fast and efficient management will determine the success of eradication or the establishment of invasive species. The spatial spread of invasive species has been the subject of much attention (Hastings et al. 2005) but there are few studies to help inform the management of invasive species spatially (Taylor and Hastings 2004a, Travis and Park 2004). Taylor and Hastings (2004) addressed the basic question of whether it is more efficient to prioritise the removal of outliers or core populations for an invasive grass, *Spartina alterniflora*. Using a structured population model, they found that the optimal strategy alternates between the removal of low and high density plants depending on the annual budget available. Given the uncertainty in future budgets allocated, the authors recommended to prioritise the removal of low density subpopulations. Travis and Park (2004) address the problem of the best way of dividing resources for the management of a source-sink model. Their results indicate that the allocation of resources solely to the source population does not always result in the most effective control strategy. The authors highlight that the most efficient control measure is determined by the nature of dispersal between the source and the sink.

Formal decision theory tools provide a useful avenue to investigate complex interactions by systematically incorporating them in a transparent and consistent manner, allowing the determination of optimal management actions given a specific objective and any constraints imposed on the system (Possingham 2001). Methods exist for determining optimal strategies and several have been applied to non-spatial ecological systems and specifically in invasive species management (Shea and Possingham 2000, Taylor and Hastings 2004a, Travis and Park 2004, Regan et al. 2006). Here, we build on these studies and propose a general approach to tackle the challenging problem of deciding where should we start managing parcels connected in the landscape by the risk of colonisation. Using decision theory and network models where nodes are susceptible or invaded areas (parcels) and links represent the risk of colonisation between parcels, we ask what is the optimal management strategy to eradicate an invasive species? If we were to manage the most connected parcels we may slow down the invasion flow and reduce the spread of the invaders but we are also taking the risk of providing a disturbed area for a reinvasion and therefore wasting time and resources that could have been better allocated (Firn et al. 2008). On the other hand, if we first manage parcels on the edge of the invasion we might control the spread but miss out on the important source parcels which contribute the most to dispersal events.

Using Markov decision processes (MDP) and an advanced stochastic dynamic programming method, from the field of Artificial intelligence (Boutilier et al. 1995, Boutilier et al. 1999, Hoey et al. 1999), we exactly solve the problem of optimally managing invasive species on a network over time. This problem is known for being very difficult to solve as the computational complexity increases exponentially with the number of nodes. Furthermore when solved it is not guaranteed that we will be able to understand the logic of the optimal solution. Traditional methods only provide one optimal strategy amongst several optimal alternatives making the understanding of the solution challenging for non-specialists. It has been therefore difficult to derive with confidence general rules of thumb. Our approach attempts to overcome this difficulty by providing the set of all optimal solutions in a clear and transparent manner. The main benefits of our approach come from its ability to extract and make smart use of structural dependencies.

## 3.2 Problem formulation: States, decisions, transition matrix and costs

We assume a set of finite parcels across space which are either invaded or at risk of invasion. We represent the spatial interaction between parcels using a network representation. A network or graph *G* is defined as a pair (*V*, *E*), where *V* is the set of vertices (or nodes), and *E* the set of edges (or links). Edges describe links between vertices. A graph can be used to describe the connectivity structure in heterogeneous landscapes (Urban and Keitt 2001) or more generally the geometry of the interactions within a complex biological system (Strogatz 2001).

A parcel can be vulnerable to an invasive species or already infected. We model the dynamics of the invasive species with a Susceptible-Infected-Susceptible model (SIS) as a discrete Markov chain. We assume that a susceptible parcel is potentially under threat of an invader. We also assume that the invasion mechanism follows a contact process (Harris 1974), in other words a parcel has a probability of being infected if one of its direct neighbours is infected (see table 1). A distance factor could be included without loss of generality. Contact process has been shown to be a relevant tool for modelling the spread of diseases (Harris 1974), and modelling meta-population dynamics (Snyder and Nisbet 2000, Franc 2004).

We first defined our model as a completely observable Markov decision process (MDP) (Puterman 1994). A MDP consists of four elements: states (*S*), actions (*A*), transition probabilities (*P*) and a reward function (*R*). We recall that using MDP assumes we can consistently detect the presence of the invasive species.

Let *s* in *S* be the state of the system at any given time t, *s* represents the number of parcel occupied by the invasive species: $s = (s_1, s_2, \ldots, s_n)$ where the components are zero-one variables, $s_i \in$ {*infected, susceptible*} . The system has $2^n$ possible states.

Let *a* in *A* be the decision the manager can make at any given time t, *a* represents the parcels managed by the decision-makers: $a=(a_1,a_2,\ldots,a_n)$ where the components are zero-one variables, $a_i \in A_i$={*manage, do nothing*} the set of decisions available in parcel *i*. In an ideal scenario resources to manage an invasive species would be unlimited and one hundred percent efficient. Unfortunately resources are scarce and must be invested in the best location to ensure the best management overall. Here, we assume that due to a fixed budget only *one* parcel can be managed at each time step. Therefore the size of *A* is $|A|= n+1$.

The dynamics of the system is captured by a matrix *P* which contains the transition probabilities of moving from any state to any other state of the system under different actions. Day and Possingham (1995) propose a stochastic meta-population model accounting for the specific characteristics of each patch (variation in patch size and position). A similar approach could be followed here to define the transition matrix at each time step:

- An infected parcel only recovers when managed with a probability of success.
- A contact process is applied to neighbours of contaminated parcels and defines a probability of colonisation given the number of surrounding infected parcels.

We define the reward function (*R*) as the number of parcels susceptible while the management of a parcel incurs a fixed cost (C). The parameters we used are formally defined table 1.

When solving the corresponding optimisation problem we seek to determine an optimal strategy $\pi : S \rightarrow A$ that minimises the expected number of parcels infected over an infinite time horizon. The optimal strategy matches an optimal action to each possible state of the system.

## 3.3 Accounting for the structure of the problem

Markov decision processes have become the model of choice for many resources allocation problems in ecology and conservation biology (Clark and Mangel 2000, Possingham 2001). While classical computational methods for solving MDP, such as value iteration and policy iteration (Puterman 1994) are often effective for small problems when dealing with network optimisation problems we face Bellman's curse of dimensionality: the size of the state space grows exponentially with the number of nodes. The second drawback to classical approaches is the inability to represent optimal solutions in an informative manner: the optimal policy is usually represented as a function that matches an action to each state of the system. One can imagine that extracting general rules of thumb becomes a challenging process when dealing with more than one hundred states.

To overcome the curse of dimensionality and the obscure representation of optimal solutions we introduce the use of structured Markov decision process also known as factored MDP (Boutilier et al. 1995, Boutilier et al. 1999, Koller and Parr 1999, Guestrin et al. 2003). Factored MDP are a way of exploiting and representing the structure of MDP. In Boutilier et al (1999), dynamic Bayesian network (DBN) representations of actions with decision trees are used to represent conditional probability tables and the reward function, the authors proposed a new method called structured policy iteration (SPI). SPI constructs value functions that preserve much of the DBN structure of the problem. Decision trees allow us to deal with bigger size problems but suffer from inefficient representations of certain types of value functions involving redundant use of identical sub-trees. Hoey et al (1999) improved this approach using a more compact and efficient representation called algebraic decision diagrams (ADD, see Figure 1) and adapted the SPI algorithm into SPUDD (Hoey et al. 1999). For the first time in ecology we take into account the structure of the optimisation problem using algebraic decision diagrams (ADD) and dynamic Bayesian networks (DBN).

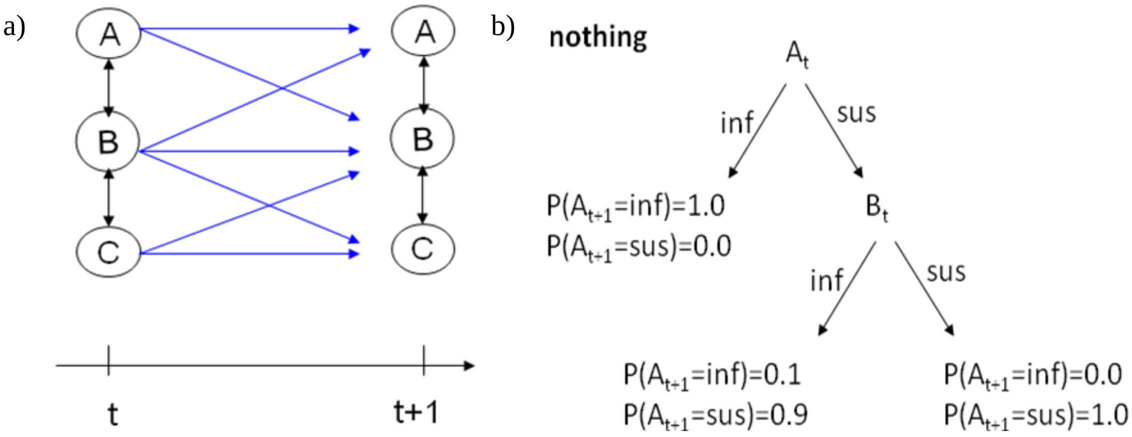


Figure 1 – Structured representation. a) Structure of the network. Parcels A, B, C are connected via a line network (bi-directional black arrows). The dynamics of parcels A and C are independent from one time step to the next time step (blue arrows). b) The ADD representation of the conditional probability table of action “Do nothing” on parcel A. Parcel A and B are connected. The probability that A at time t+1 becomes infected depends on the status of A and B at time t: if A is infected then under the do nothing action A remains infected; if A is susceptible and B is infected then invaders from B might colonise A with probability 0.1. Information about parcel C is not needed. Using ADD compactly represents the transition probabilities.

For the purpose of this study we arbitrarily defined MDP parameters for a hypothetical invasive species as presented in table 1. We have assumed that a linear relationship defines the probability of colonisation related to the number of infected neighbours.

Table 1. MDP parameters assumed for our general case study

| MDP parameters for a hypothetical invasive species | Value |
|---|---|
| Probability of a parcel being infected if one neighbour is infected | 0.01 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j1,t}=\text{infected})$ | |
| Probability of a parcel being infected if two neighbours are infected | 0.02 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j2,t}=\text{infected})$ | |
| Probability of a parcel being infected if three neighbours are infected | 0.03 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j3,t}=\text{infected})$ | |
| Probability of a parcel being infected if four neighbours are infected | 0.04 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j4,t}=\text{infected})$ | |
| Probability of a parcel being infected if five neighbours are infected | 0.05 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j5,t}=\text{infected})$ | |
| Probability of a parcel being susceptible if it has been managed | 0.7 |
| $P(s_{i,t+1}=\text{susceptible} \mid a_{i,t}=\text{manage}, s_{i,t}=\text{infected})$ | |
| Benefits of having a susceptible parcel (Reward) | 100 |
| $R(s_{i,t}=\text{susceptible})$ | |
| Cost of managing one parcel | 100 |
| $C(a_{i,t}=\text{manage})$ | |

## 3.4 Graphical user interface and Youtube video

We developed a research decision tool that first takes as an input an adjacency matrix (see Figure 2), the probability of successful management of a parcel, the probabilities of a parcel being colonised when zero to five neighbours are infected, a reward function and a cost function (not shown). We automatically generate the corresponding Markov decision problems represented using algebraic decision diagrams. Thirdly we solve our optimisation problem using the SPUDD algorithm and provide the set of all optimal solutions represented as an ADD. Finally we can explore our solution space using a graphical user interface (Figure 4). Interested readers are referred to screencast showing live demonstrations of our decision tool:

- On a 3-node network http://www.youtube.com/watch?v=wuOvbCu_nJc
- On a 9-node network http://www.youtube.com/watch?v=UMsKMd-X8QE
- With a user friendly interface http://www.youtube.com/watch?v=csg6eIwH-04 (recommended)

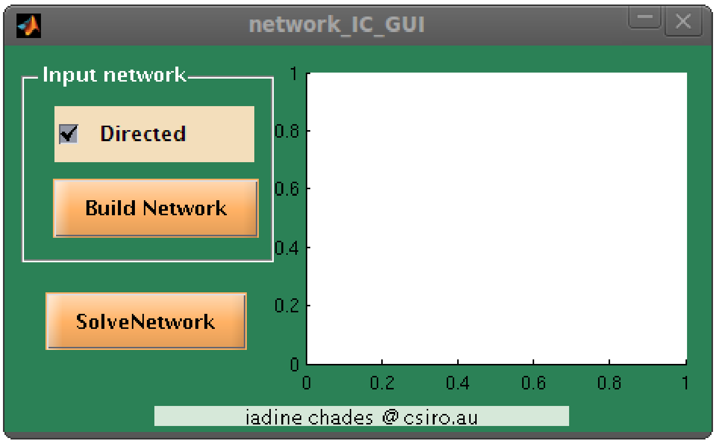


Figure 2 - Graphical user interface of our MDP network. The user first designs the input network (undirected or directed) by selecting “Directed” and pushing the button “Build Network”. A new window appears (Figure 3) where the user designs a network of invasive species. Once the network is designed, the user can call the optimisation program by pushing the button “SolveNetwork”. At the end of the optimisation procedure, a new window appears allowing the user to explore the set of optimal solutions (Figure 4).

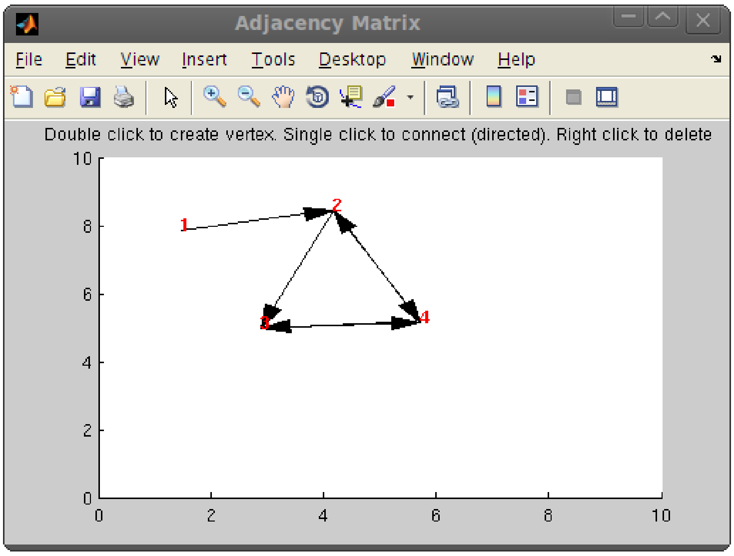


Figure 3– Graphical user interface to design the network of invasive species.

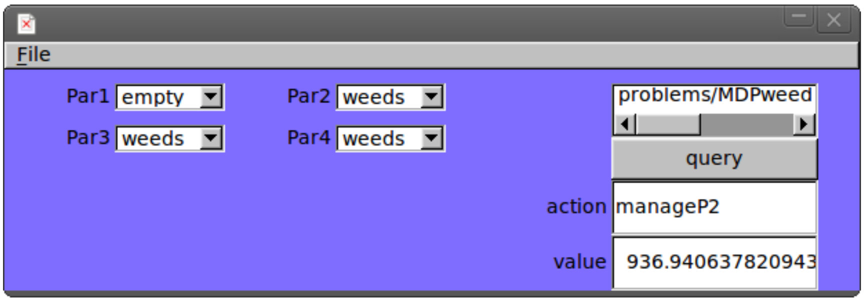


Figure 4– Graphical user interface to explore the set of optimal strategies. The user specifies the state of each parcel and requests the optimal set of actions when pushing the button “query”.

## 3.4 Results

We first solved optimisation problems on small networks that we call motifs. A motif is an entity that may occur several times in larger networks. We then derive rules of thumb to manage larger structured networks.

**Motif network**

The results when managing two patches of similar characteristics are intuitive. If two infected patches have the same probability of colonising each other, the optimal strategy is to manage either one (Figure 5a). If one patch has a higher probability of contaminating its neighbour then it is optimal to start managing the source (Figure 5b). The optimal solution represented as an ADD clearly show that if parcel 1 is infected, state "weeds", the optimal decision is independent from the state of parcel 2: we should manage parcel 1.

Similarly when three patches are linked by a directed probability of dispersal, it is optimal to manage the source before managing the satellites parcels (Figure 6 ab). This kind of network might represent dispersal by water or wind.

In the case of network 7a (Figure 7a), the three parcels are sources but the parcel in the middle can reinfect both of the remaining parcels. Parcel two and three are at the extremity of the colonisation process and have a lower probability of being recolonised once managed than parcel one. The optimal strategy is to start from an extremity of the network then manage parcel one and finally manage the remaining extremity. The rules of thumb for this kind of line network is to start managing an extremity then manage the neighbouring parcel keeping the same direction, until all parcels are managed.

In the case of an undirected triangle (3-island network), all parcels have an equal probability of being reinfected once managed and the managing order does not make a difference (Figure 7b). Prioritising the management becomes optimal when directed probability of colonisation can occur and a source parcel can be identified.

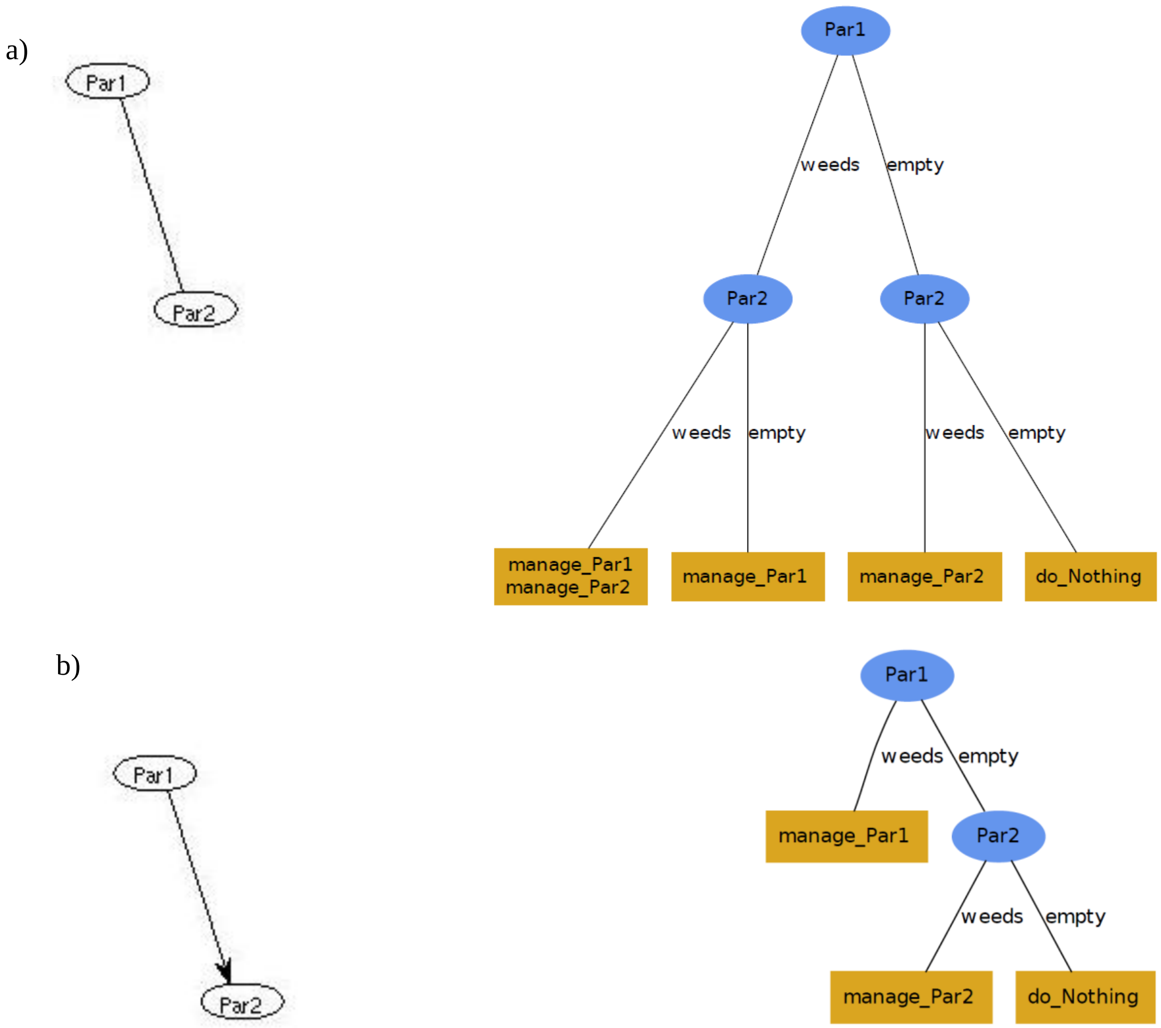


Figure 5 – Optimal management of 2-Parcel networks. Graphs on the left represent the network considered. Graphs on the right represent the set of optimal strategies using an ADD representation. a) The probability of dispersal between both parcels is bidirectional, the optimal strategy is symmetrical. To determine the optimal set of actions, we first check the state of parcel 1 (weeds or empty) and the state of parcel 2. The optimal action is written inside the yellow rectangle. For example if parcel 1 is in state “weeds” and parcel 2 is in state “weeds” we have two optimal actions: manage parcel 1 or manage parcel 2. b) The probability of dispersal goes from parcel 1 (source) to parcel 2 (sink), the corresponding optimal strategy is to manage the source (parcel 1) first.

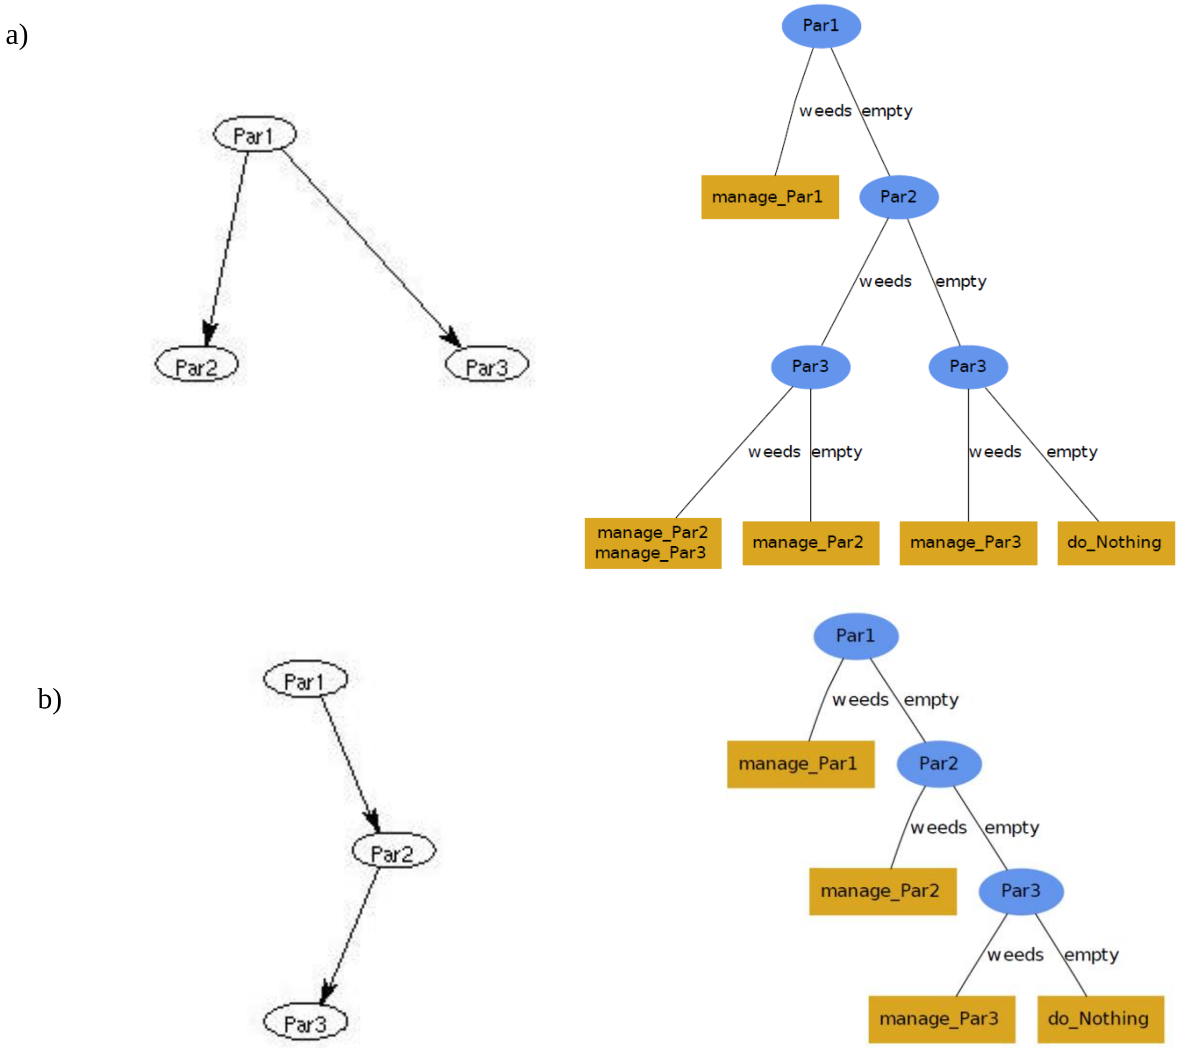


Figure 6 – Optimal management of directed 3-Parcel networks. The optimal strategies start managing the source then manage the sinks.

a)

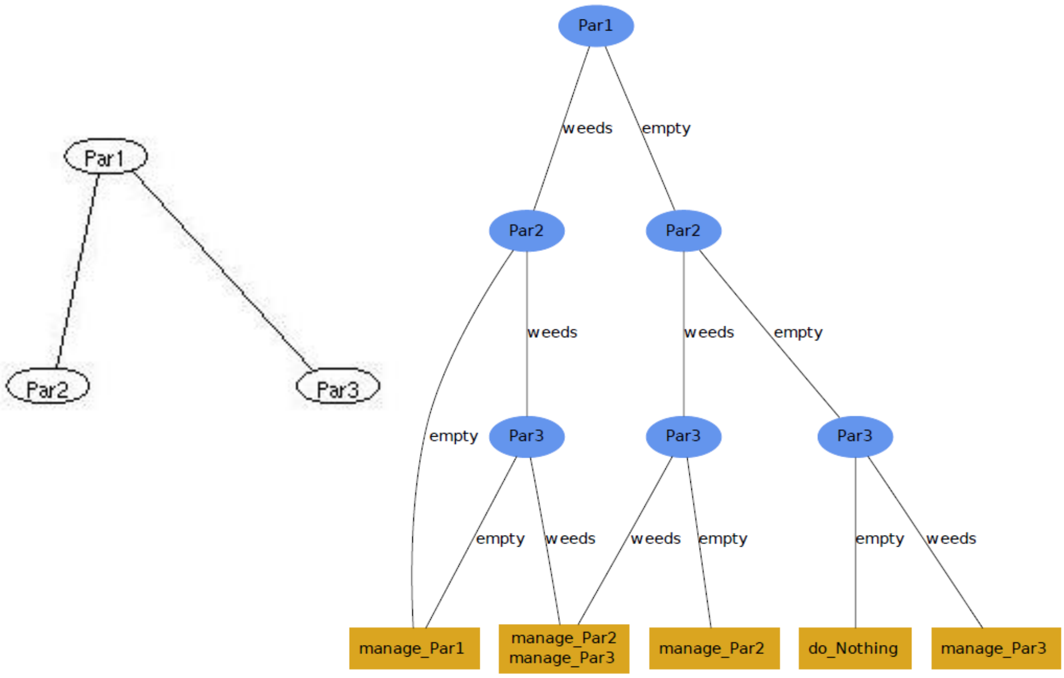


b)

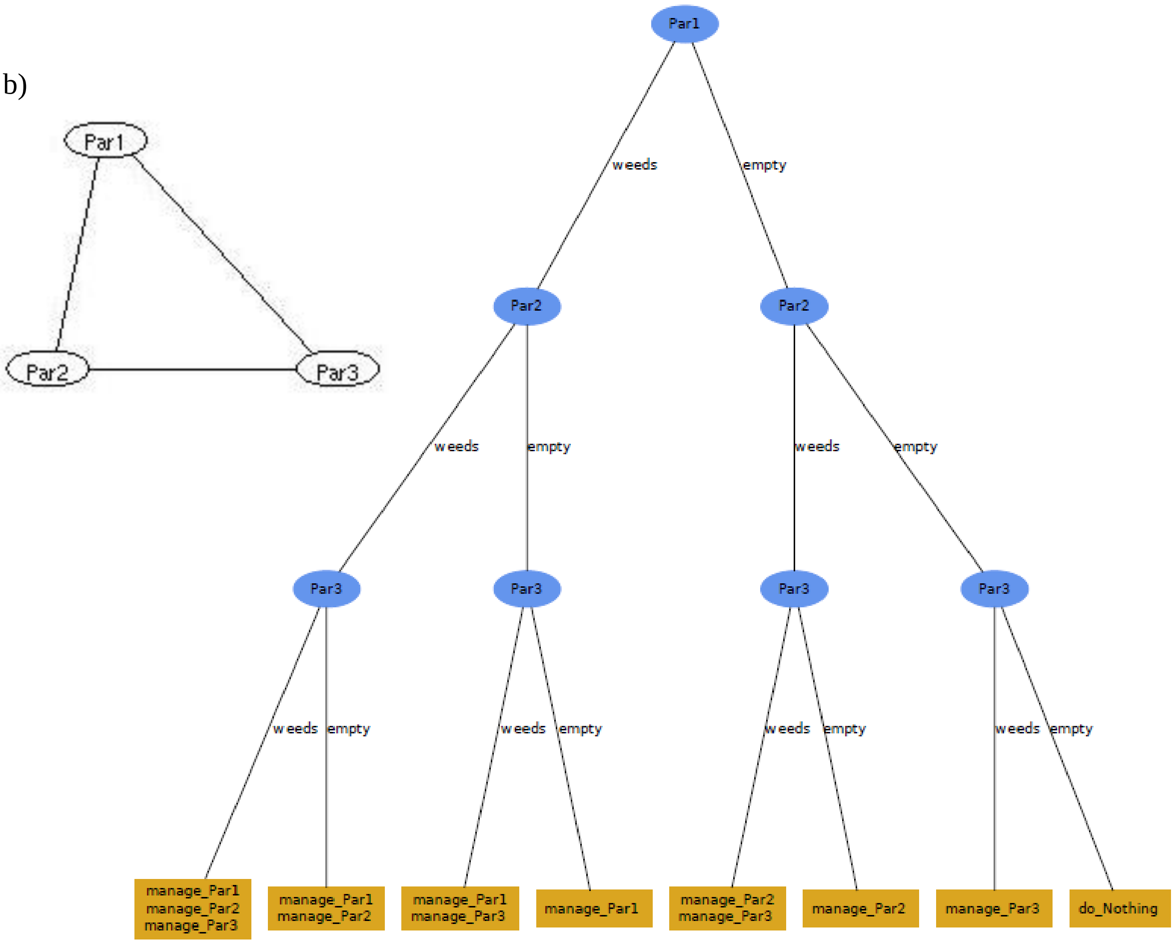


Figure 7– Optimal management of undirected 3-Parcel networks. The optimal strategies start managing from an extremity of the network then manage the closest parcels.

An interesting case where the solution might be difficult to determine is the case of Figure 8. Here it is optimal to start managing from either parcel 1 or 2, then to manage the remaining source before tackling the sink (parcel 3). This case might occur with permanent waterholes where connection to a sink parcel is a consequence of flooding events. Using our knowledge about previously studied 2-parcel motif (Figure 5), we can derive an optimal strategy for this “source-source-sink” network. Network 5 can be seen as a 2-parcel network “source-sink” where parcel one and parcel 2 are aggregated in one source parcel. The optimal strategy remains to first manage the source parcels.

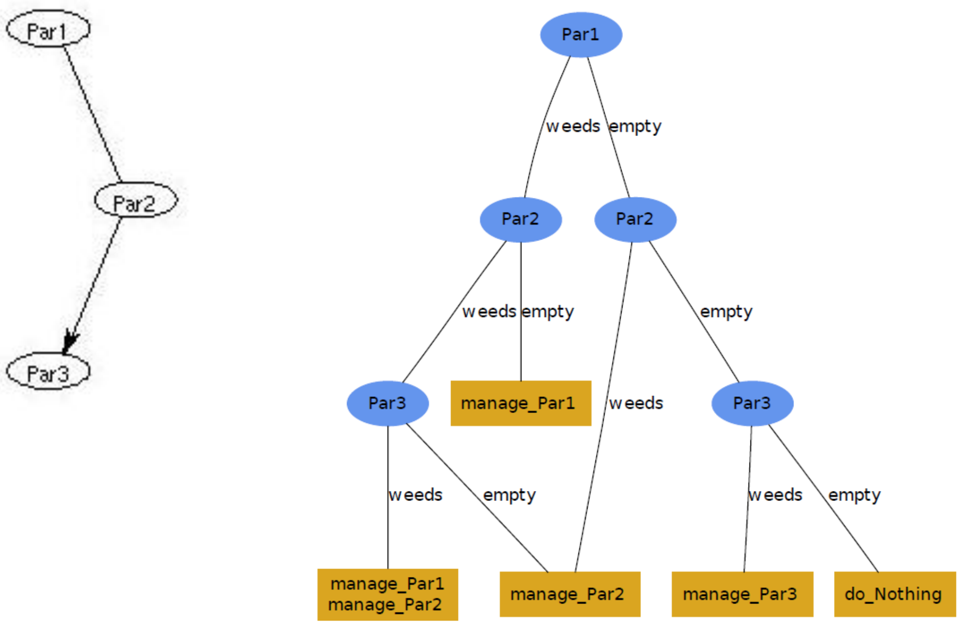


Figure 8 – Optimal management on a 3-Parcel network of type “source-source-sink”. The optimal strategies start managing from one of the source of the network then manage the second source and finally the sink parcel.

### Star network

When dealing with four parcels, we can reuse what we learnt from the previous 3-node and 2-node network management strategies: i) Always start from source nodes when colonisation is directed thus reducing the size of the network as a directed link assumes a source parcel cannot be reinfected by its sink; ii) if the probability of dispersal is bidirectional (or undirected) the optimal strategy is to start from an extremity of the network where the probability of re-invasion is the lowest.

Star structures are made up of a central parcel connected to satellite parcels. The optimal strategy for managing a 4-node star network (Figure 9a) is to first tackle two of the extreme parcels (e.g. parcel one and two) and then tackle the most connected parcel (parcel 4). The remaining extreme parcel (parcel 3) is managed last. When dealing with star networks of any size a general pattern appears: manage the satellite parcels until the number of satellite parcels free of invaders is equal to (or greater than) the number remaining satellites infected, it is then optimal to manage the central node before managing the remaining infected satellite parcels.

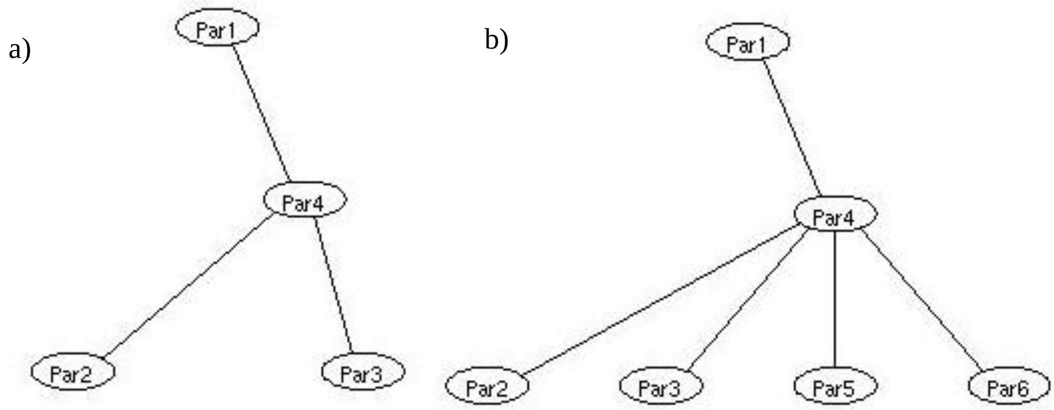


Figure 9 – Star networks of 4 and 6 parcels. The optimal management strategies first tackle the satellite parcels until the number of satellite parcels without invaders is equal to or greater than the number of satellite parcels invaded. Managing the central parcel is then optimal before managing the remaining satellite parcels invaded.

**Island network**

We were able to derive rules of thumb for N-island and N-island-k type networks where N represents the number of nodes involved in the island network and k the number of nodes involved in a line network (Figure 10 and 11). Island networks are a common structure in ecology. If we needed to prioritise our management strategy we would first start from any node and then manage one of the nearest invaded parcels, here the direction of management does not matter. For example the sequential management strategy manage parcel 1,2,3,4,8,5,7 and 6 is one of the optimal solutions for the 8-island network depicted figure 10b. The black arrows figure 10b represents an alternative solution. In the case of island type networks managing parcels successively in the same direction is one of the optimal solutions.

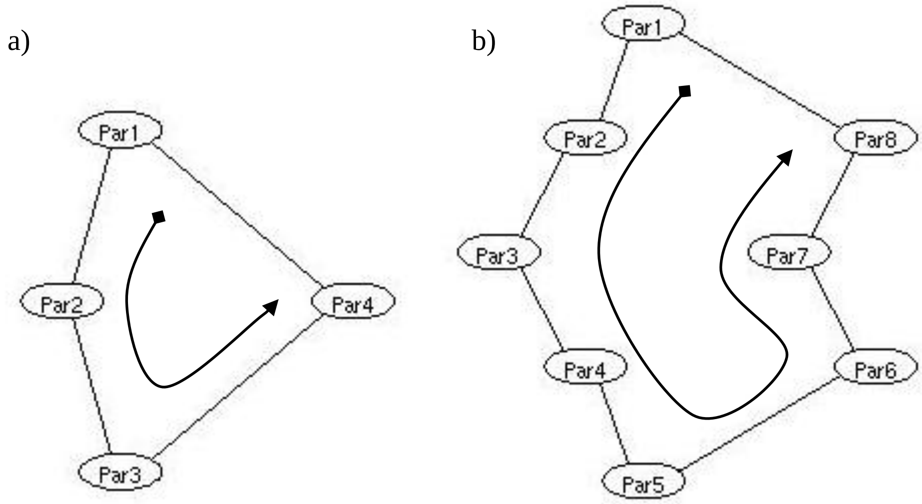


Figure 10 – Island networks of 4 and 8 parcels. The optimal management strategies start from any nodes and then manage one of the nearest invaded parcels. Black arrows represent one of the optimal strategies.

A N-island-k type network is a N-island network where one parcel is connected to a k-line network. When managing N-island-k type networks (see Figure 11), the optimal solution is to first tackle the extreme parcel of the network (here parcel 1) then proceed to the management of the k-line as described previously (Figure 7). Once the connecting parcel is managed (parcel 4) it is optimal to manage successively the parcels that are the closest to the connecting node. The black arrow represents one of the optimal strategies. Parcels in the priority 5 rectangle must be managed before nodes in the priority 6 rectangle. We found the same rule of thumb for the management for N-island-k type with 15 nodes.

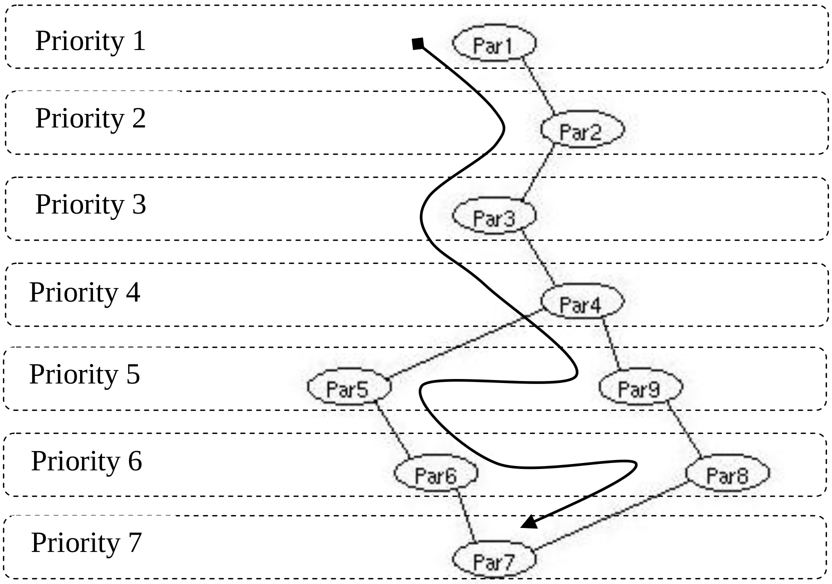


Figure 11 – N-Island-k networks of 9 parcels. The optimal management strategies start from parcel one and manage the nearest invaded parcels successively until parcel 4. From parcel 4 – the connecting node – the optimal strategies manage the nearest parcels from the connecting node (parcel 4). The black arrow represents one of the optimal strategies. The rectangles represent the management priority order. Here, parcels 9 and 5 must be managed before parcels 6 and 8.

In the case where an island has several line type networks connected (Figure 12), the optimal strategy is to start managing the line which has the most extreme parcel (here parcel 7) before managing the island as in the N-island-k case.

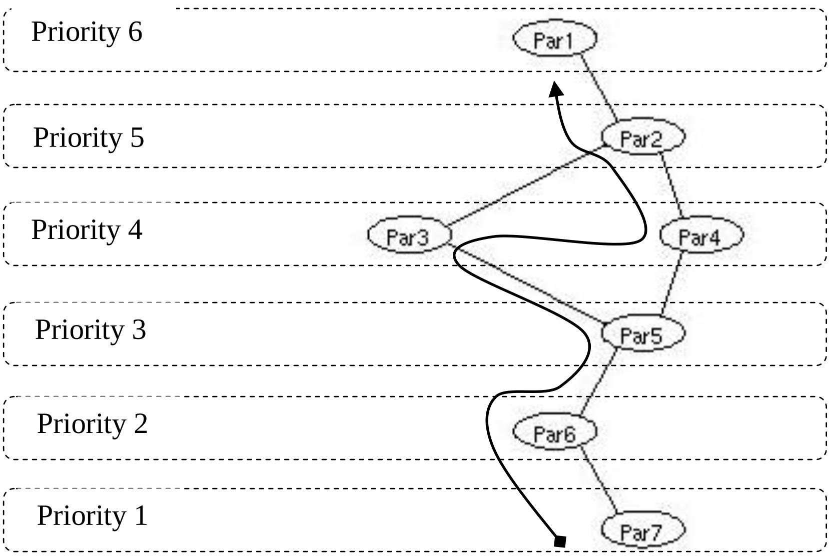


Figure 12 – N-Island-k networks of 7 parcels. The optimal management strategy is to start from parcel seven and manage the nearest invaded parcels successively until parcel 5. From parcel 5 – the connecting node – the optimal strategy is to manage the nearest parcels from 5. The black arrow represents one of the optimal strategies. The rectangles represent the management priority order.

### Cluster network

When managing a cluster type network (Figure 13), it is optimal to start from i) the smaller cluster (Figure 13a) ii) if the network is symetrical (Figure 13b) start from one of the clusters, iii) if one cluster is less connected start from the least connected cluster (Figure 14).

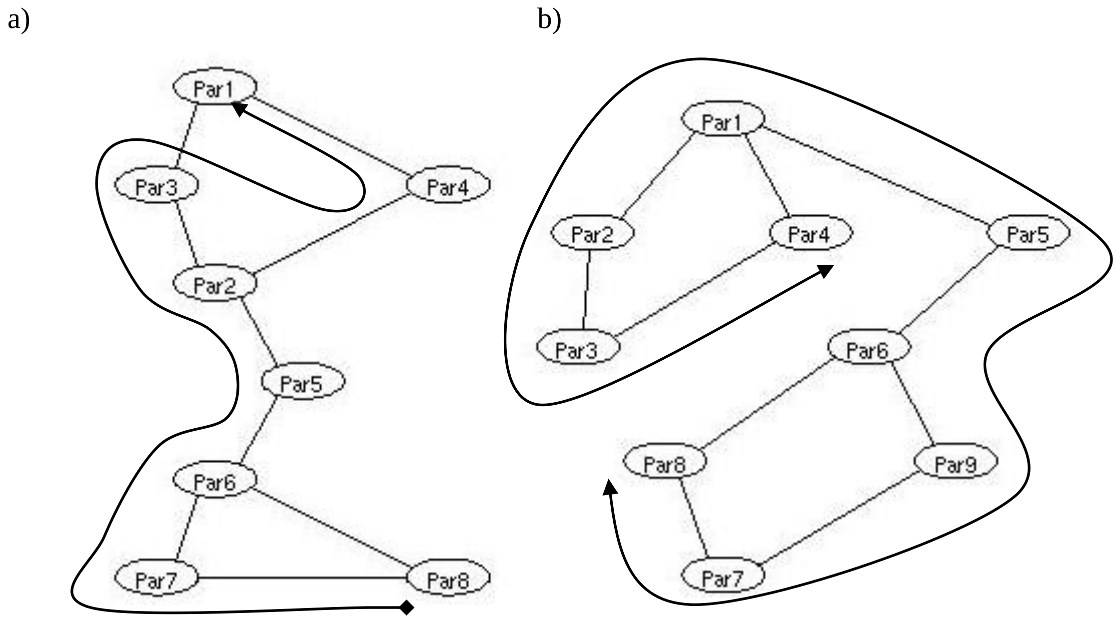


Figure 13 - Cluster network 1. a) The optimal strategies start from the smaller cluster, b) both clusters are symmetrical. The black arrows represent one of the optimal strategies.

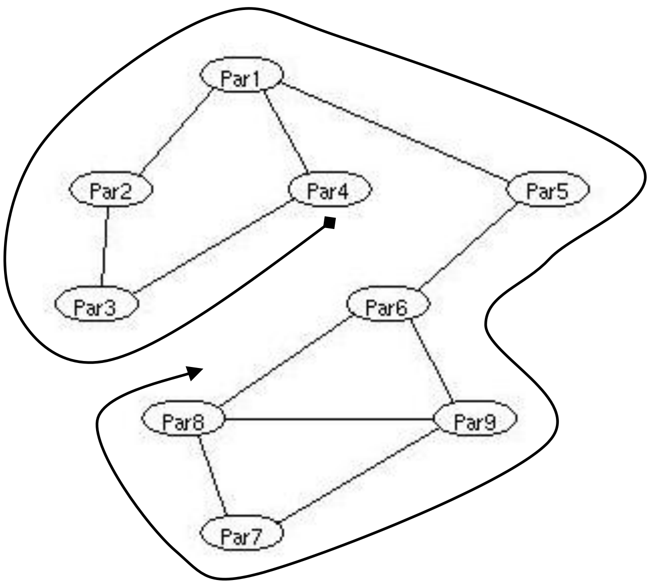


Figure 14 – Cluster network 2. The optimal strategy is to start from the least connected cluster. The black arrow represents one of the optimal strategies.

When we increase the number of parcels between identical clusters the optimal strategy remains the same: first tackle one of the clusters then the connecting nodes and finally the remaining cluster (Figure 15a). However if one of the connecting parcels has a directed link it becomes optimal to start managing from that parcel in priority (Figure 15b).

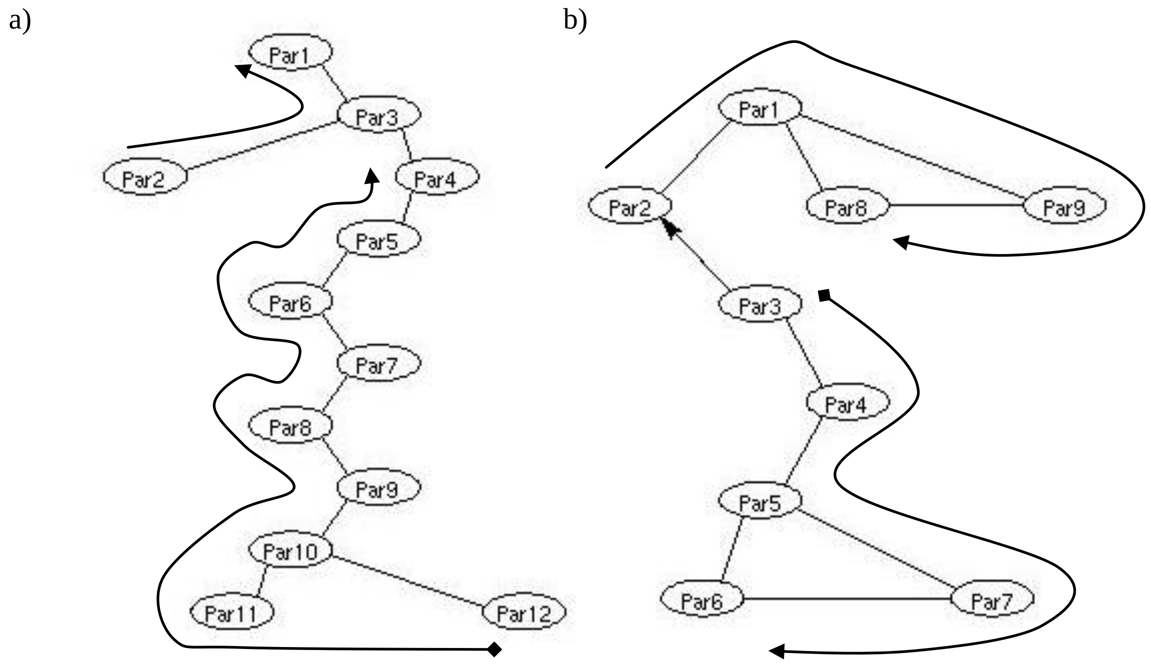


Figure 15 – Cluster network 3. Managing a connecting node is optimal when the probability of dispersal is directional. The black arrows represent one of the optimal strategies. a) undirected cluster network b) directed cluster network

## 3.5 Discussion

Using networks as a model for the spatial management of invasive species we presented a new way of solving sequential decision problems under uncertainty. Using factored MDP and Algebraic decision diagrams (ADD) allowed us to take advantage of the structure of the problem. Our method analyses the dependence and independence of each parcel, generates the corresponding Markov decision problem and performs the optimisation process while keeping the structure of the problem. While traditional methods would give one optimal solution our approach computes the set of all optimal strategies allowing us to choose amongst a set of alternative optimal decisions. In doing so we were able to solve and derive rules of thumb for the management of small networks (motifs) and structured networks under perfect knowledge of probability of colonisation between parcels.

The optimal strategies computed are invariant to management success and to a linear colonisation process. However the colonisation strength drives the optimal management in very different ways (figures 5, 6, 7 and 15). Most invasive processes are driven by directed natural factors such as wind or water dispersal but some invasion processes (bird dispersal, flooding events) are unlikely to follow clear patterns and may only be approximated by distance factors and assumed to be bidirectional around an area. Such information requires intensive and expensive data gathering processes. In the absence of such data we recommend using expert knowledge to elicit the most likely interaction strength or assume that the dispersal rate is bidirectional.

A natural progression of our approach is to ask whether our rules of thumb will hold when we consider an increase in the management budget. The optimal strategies derived with one management action per time step are likely to be robust to an increase in management activities per time step. However we are not able to provide any formal proof of this insight at this stage. An algorithmic solution could provide some answers for small size networks. Indeed, increasing the number of parcels we could manage at each time step can be incorporated in our model at the cost of a significant increase in the computational complexity. The number of management actions will be close to $(n+1)!/(n+1-k))!$ where $k$ represents the number of parcels we could manage every year (e.g. $n=10$ and $k=2$, $|A|\sim110$). Unfortunately large action space remains a challenging problem in the optimisation community. Increasing the size of the action space is likely to rule out our ability to derive rules of thumb for general network structures.

Most of the literature in network theory has focused energy and time on understanding the complexity of network structure to make predictions about various large complex systems such as food webs, epidemics, gene regulation and so on (Dunne et al. 2002, Milo et al. 2002, Proulx et al. 2005, May 2006, Montoya et al. 2006, Jordán et al. 2008). Closer to conservation biology, new and past studies in the management of meta-populations have tackled similar problems (Day and Possingham 1995, Nicol et al. 2009). But none of these works have provided general rules of thumb to prioritise the management over time on structured networks. Our future direction of work will focus on applying structured MDP to the optimal management of the introduced fish Gambusia which is invading endangered mound springs in central Queensland.

In this section, we assumed we were able to detect perfectly the presence of the invaders. Unfortunately some invasive species cannot be detected perfectly. In section 4 we explore solutions for invaded networks when detection is incomplete using partially observable Markov decision processes (POMDP).

# 4. Optimal management of plant invasions on a network when dealing with imperfect detection

## 4.1 Introduction

Invasive plants are a major threat to natural and managed systems. They are notoriously difficult to control or eradicate and require large amounts of effort and resources to manage effectively (Pimental 2002, Panetta and Timmins 2004). With the annual economic impact of weeds in Australia estimated at $3.4 billion (Sinden et al. 2005), it is crucial that resources are distributed optimally and directed towards the most effective management activities and not wasted on management actions that are ultimately unsuccessful. Yet deciding the best course of action for an invasive plant can be excruciatingly difficult due to the complex interaction of factors such as the extent of the invasion, the ecology of the species, the dynamics of the system, and how the species responds to di fferent management actions (Taylor and Hastings 2004b). The role of space in the management of invasive species is critical at the early stage of an invasion process; fast and efficient management will determine the success of eradication or the establishment of invasive species. The spatial spread of invasive species has been the subject of much attention (Hastings et al. 2005) but there are few studies to help inform the management of invasive species spatially (Taylor and Hastings 2004a, Travis and Park 2004). The decision process is exacerbated further by our inability to observe the system perfectly. Imprecise survey techniques, the cryptic nature of some species, and persistent seed banks make it difficult to verify whether the imposed management actions are successful or not (Chadès et al. 2008). Recently, Hauser and McCarthy (2009) have optimised the surveillance effort that should be allocated to detect and destroy an invasive species across space but not over time.

In Regan et al (in review), we developed a POMDP model to determine optimal management strategies for an invasive plant species where the states of the system are not known perfectly. We investigate how the optimal solution changes depending on the value of eradicating the species, the importance of colonisation, and the importance of detection on the optimal solution. We illustrated this method through a case study, Branched broomrape, *Orabanche ramosa*, a parasitic crop weed at the center of a national eradication program in South Australia (Jupp et al. 2002). The efficiency of the optimal management strategy we provided is highly sensitive to the probability of outside colonisation (see Figure 16). When managing a single parcel optimally the probability of eradication of the invasive species drops from 1 to less than 0.6 under a random colonisation probability.

Here we build on this previous work and tackle the problem of managing a cryptic invasive species across space. Using cutting edge methods from Artificial Intelligence we model and solve spatial POMDPs and answer the question of where and how long should we manage or survey in order to maximise the chance of eradicating an invasive species at the landscape level. We provide rules of thumb to help manage small structured networks and a general decision tool that could be applicable to a wide range of agricultural and environmental weeds that are cryptic in nature and/or have a persistent soil seed bank.

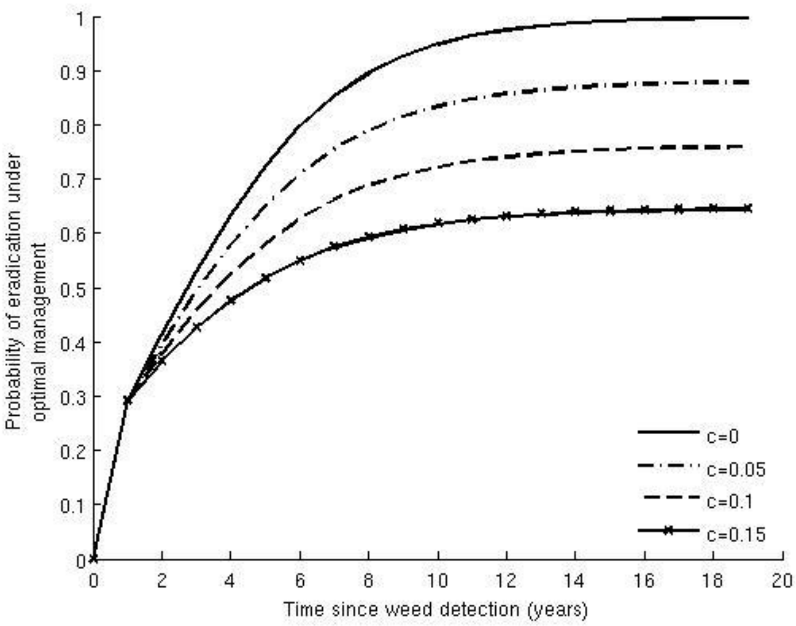


Figure 16 - The probability of eradication given the optimal management strategy for different levels of colonisation. Reward cost ratio is 1:100 and detection probability is 0.7. Thick black line indicates when the management alternative changes (Regan et al, in review).

## 4.2 Observations, observation matrix and belief states

A partially observable Markov decision process (POMDP) is a convenient model for solving sequential decision-making optimisation problems under uncertainty where the decision-maker does not have complete access to the current state of the system. In other words, the manager is not sure whether the invasive species is present or not in a particular parcel. First studied in Operations Research literature, POMDP provides an interesting way of reasoning about trade-offs between actions to gain rewards and actions to gain information.

To take into account the imperfect detection of the invaders we define the finite set of local observations for each parcel at time t, $z_i$=*{Absent, Present}* and their corresponding observation function $o_i$ that maps to each state-action pair a probability distribution over $z_i$. In other words, the probability of detection of the invasive species given that the parcel is infected and that the previous decision is to '*Do nothing*' is defined by $o_i$*(Present|Infected,Do nothing)* .

Let $z$ in $Z$ be an observation of the whole system $z=\{z_1,\ldots,z_n\}$ . We define the observation function of the system as the joint probability of local observation function:

$O(z|s,a)=o_1(z_1|s_1,a_1) \times \ldots \times o_n(z_n|s_n,a_n)$.

Where $s$ in $S$ represents the state of the system at any given time t. $s$ represents the number of parcel occupied by the invasive species: $s = (s_1,s_2,\ldots,s_n)$ where the components are zero-one variables, $s_i \in$ {*infected, susceptible*} . And $a$ in $A$ is the decision the manager can make at any given time t, $a$ represents the parcels managed by the decision-makers:

$a=(a_1,a_2,\ldots,a_n)$ where the components are zero-one variables, $a_i \in A_i$={*manage, survey, do nothing*} the set of decisions available in parcel $i$. In an ideal scenario resources to manage an invasive species would be unlimited and one hundred percent efficient. Unfortunately resources are scarce and must be invested in the best location to ensure the best management overall. Here, we assume that due to a fixed budget only *one* parcel can be managed at each time step. Therefore the size of $A$ is $|A|= n+1$.

As it is neither practical nor tractable to use the history of the action-observation trajectory to compute or represent an optimal solution, belief states are preferred to summarise and overcome the difficulties of incomplete detection. Indeed Aström (1965) has shown that belief states are sufficient statistical tools to summarise all the observable history of a POMDP without loss of optimality. A POMDP can be cast into a framework of a fully observable Markov decision process where belief states represent the continuous but fully observable state space. Here, a belief state *b* is defined as a distribution probability over states in *S*.

In our case, solving a POMDP is finding a strategy $\pi : B \times \tau \mapsto \mathrm{A}$ mapping an allocation of resources $a \in A$ given a current belief state $b \in B$ and a time-step $t \in \tau$. An optimal strategy minimises the expected sum of costs or rewards (*R*) over a finite time horizon, *T*. This expected summation is also referred to as the value function (Cassandra et al. 1995).

A value function essentially ranks strategies by assigning a real value to each *b*. Using the Bellman principle of optimality and the previously-defined POMDP parameters, we can calculate the optimal *t*-step value function from the (*t*–1)-step value function:

$$V_0^*(b) = \min_{a \in \mathrm{A}} \left[ \sum_{s \in S} R(s,a) \Pr(s|b) \right], \quad (1)$$

$$V_t^*(b) = \min_{a \in \mathrm{A}} \left[ \sum_{s \in S} R(s,a) \Pr(s|b) + \sum_{s \in S} \sum_{s' \in S} \sum_{z \in Z} \Pr(s|b) P(s'|s,a) O(z|s',a) V_{t-1}^*(b_z^a) \right], \quad (2)$$

where $Pr(s|b)$ represents the probability of being in state *s* given a belief state *b*, and $b^a{}_z$ is the belief state assuming action *a* and observation *z* and *P* is the state-action transition matrix. Equation (1) minimises the expected sum of instantaneous costs when there is no time left to manage for the species. Similarly when there are *t* steps to go, equation (2) minimises the instantaneous costs and the future expected costs for the remaining *t*–1 steps. Interested readers are referred to Cassandra et al (1995) for further explanations of the dynamic programming equations.

The optimal solution $\pi$ can be represented in two different ways. We can either apply directly the strategy function for each belief state we are in or we can represent the optimal strategy as a policy graph. The policy graph automatically generates all the possible transitions over time given the performed action and the new observation whereas the use of strategy functions requires updating the belief state using Bayes' rule given the performed action and the new observation. After performing action *a* and observing *z*, the updated belief $b_z^a$ can be calculated from the previous belief *b*:

$$b_z^a(s') = \Pr(s'|b,a,z) \ , \tag{3}$$

$$b_z^a(s') = \frac{O(z|a,s')\sum_{s'\in S} P(s'|s,a)b(s)}{\Pr(z|a,b)} \ , \tag{4}$$

With

$$\Pr(z|a,b) = \sum_{s\in S}\sum_{s''\in S} O(z|s'',a)P(s''|s,a)b(s) \ . \tag{5}$$

While various algorithms from the Operations Research and Artificial Intelligence literatures have been developed over the past years, the computational complexity of exact algorithms remains intractable for most problems: finite horizon POMDPs are PSPACE-complete (Papadimitriou and Tsitsiklis 1987) and infinite-horizon POMDPs are undecidable (Madani et al. 2003).

Smallwood and Sondik (1973) have shown that the optimal value function for a finite-horizon POMDP can be represented by hyperplanes, and is therefore convex and piecewise linear. It means that the value function $V_t$ at any horizon *t* can be represented by a set of |S|-dimensional hyperplanes $\Gamma_t=\{\alpha_0,\alpha_1,...,\alpha_n\}$. These hyperplanes are also called α-vectors. A number of exact value function algorithms look at determining the optimal set of α-vectors (Sondik 1971, Monahan 1982, Littman 1996, Cassandra et al. 1997).

Unfortunately, the size of the set $\Gamma_t$ is in $\mathbf{O}(|A|\ |\Gamma_{t-1}|^{|Z|})$ e.g. exponential in the number of observations (table 2). Since each new α-vector requires computation time in $\mathbf{O}(|Z||S|^2)$, the resulting complexity of iteration *t* for exact approaches is in $\mathbf{O}(|A||Z||S|^2|\ \Gamma_t\ -1|^{|Z|})$ (Ross et al. 2008). As a result of this complexity in both memory space and computational time, most of the work on exact approaches has focused on finding efficient ways of pruning the set $\Gamma_t$ (Zhang and Zhang 2002).

Table 2. The complexity of solving a spatial POMDP.

| Number of nodes | States space size \|S\| | Observations space size \|Z\| | Actions space size \|A\| |
|---|---|---|---|
| 2 | 4 | 4 | 5 |
| 3 | 8 | 8 | 7 |
| 4 | 16 | 16 | 9 |
| 5 | 32 | 32 | 11 |
| 6 | 64 | 64 | 13 |
| 7 | 128 | 128 | 15 |

In the last few years off-line approximate methods have been developed to solve POMDPs (Hauskrecht 2000, Pineau et al. 2003, Braziunas and Boutilier 2004, Smith and Simmons 2004, Poupart 2005, Spaan and Vlassis 2005). These methods specify the best action to perform in any state of the system prior to the execution (e.g. off-line procedure). We identified Perseus (Spaan and Vlassis 2005) as the fastest algorithm and the overall best performer. Perseus uses a point-based approach. These approaches approximate the value function by updating it only for some selected belief states. The point-based methods sample belief states by simulating some random interactions in the environment and then update the value function and its gradient over those sampled belief states. Each iteration has polynomial time complexity in $\mathbf{O}(|A||Z||S||B|(|S|+|B|))$ which is an improvement to tackle our problem (e.g. no longer exponential). Unfortunately Perseus does not take advantage of structure in the network. Symbolic Perseus (Poupart 2005) makes the best of both worlds: the factored representations of POMDP using ADD and a fast solving method using Perseus.

For the purpose of this study we developed a research decision tool that takes as input:

- The structure of the spatial interactions using an adjacency matrix;

- The dynamic of the system under no management specifying probabilities of colonisation when one to five neighbours are infected;
- The efficiency of a management action (probability of success)
- The detection probability of invaders on a infected parcels under different decisions (survey, manage, do nothing)
- The cost of managing one parcel
- The cost of surveying one parcel
- The benefits of having one parcel free from invaders.

From this data, our decision tool generates a structured POMDP problem and solves the optimisation problem of maximising the expected sum of benefits over a 100-year time horizon using Symbolic Perseus. The user can then explore the solution state space using a graphical interface. Interested readers could refer to a snapshot of the graphical user interface (Figure 17) and a youtube video showing live demonstrations of our decision tool: http://www.youtube.com/watch?v=Ro-jKitb_-w

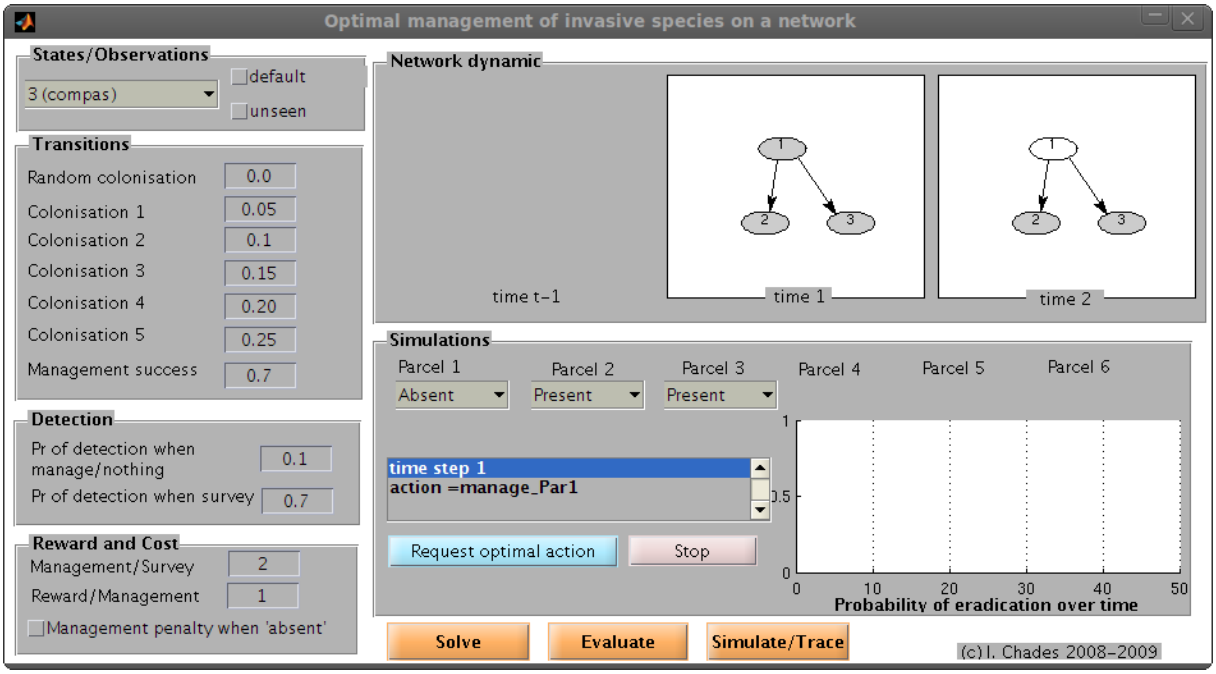


Figure 17– Graphical user interface to manage invasive species under imperfect detection.

So far we have been able to solve optimisation problems on networks of up to 6 nodes. For the purpose of this study we assumed hypothetical parameters representing the spatial problem of managing a cryptic invasive species on a network (table 3). As in the MDP case, we have assumed a linear relationship between the probability of colonisation and the number of infected neighbours.

Table 3. POMDP parameters used.

| POMDP parameters for a hypothetical invasive species | Value |
|---|---|
| Probability of a parcel being infected if one neighbour is infected | 0.05 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j1,t}=\text{infected})$ | |
| Probability of a parcel being infected if two neighbours are infected | 0.1 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j2,t}=\text{infected})$ | |
| Probability of a parcel being infected if three neighbours are infected | 0.15 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j3,t}=\text{infected})$ | |
| Probability of a parcel being infected if four neighbours are infected | 0.20 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j4,t}=\text{infected})$ | |
| Probability of a parcel being infected if five neighbours are infected | 0.25 |
| $P(s_{i,t+1}=\text{infected} \mid a_{i,t}=\{\text{do nothing, survey}\}, s_{j5,t}=\text{infected})$ | |
| Probability of a parcel being susceptible if it has been managed | 0.7 |
| $P(s_{i,t+1}=\text{susceptible} \mid a_{i,t}=\text{manage}, s_{i,t}=\text{infected})$ | |
| Probability of detection of an infected parcel when manage/do nothing | 0.1 |
| $O(z_{i,t+1}=\text{present} \mid a_{i,t}=\{\text{manage, do nothing}\}, s_{i,t+1}=\text{infected})$ | |
| Probability of detection of an infected parcel when survey | 0.9 |
| $O(z_{i,t+1}=\text{present} \mid a_{i,t}=\{\text{survey}\}, s_{i,t+1}=\text{infected})$ | |
| Benefits of having a susceptible parcel | 100 |
| $R(s_{i,t}=\text{susceptible})$ | |
| Cost of managing one parcel | 100 |
| $C(a_{i,t}=\text{manage})$ | |
| Cost of surveying one parcel | 50 |
| $C(a_{i,t}=\text{survey})$ | |

## 4.3 Results

We present the management strategies for a hypothetical invasive species on motifs and small structured networks.

We explore two different scenarios:

i) All the parcels are infected, once managed, invaders remain unobserved from a parcel. In this case our starting belief of eradication of the invasive species is zero.

ii) We have no information about the status of the parcels and invaders are not observed. In this case, we assume the system has the same chance of being in any state. Therefore, if we have $2^n$ states, our starting belief that the invasive species is eradicated on the whole network is $1/2^n$ where n is the number of parcels in the network.

**Line networks**

In the case of two infected parcels connected with a bidirectional link, the POMDP solution recommends managing population 1, then managing population 2 is recommended and in absence of sighting, the solution recommends managing population 1 and 2 before doing nothing (Figure 18). The red line Figure 18a represents our belief that all the parcels are empty, e.g. the probability of eradication of the invasive species. As we manage our network following the recommended POMDP solution, our belief that the invasive species is eradicated increases. The black arrow represents the probability of eradication we would have reached with the completely observable solution (MDP). Indeed, if we had considered the MDP solution we would have stopped managing at the third time step and risked reinvasion with a high probability of failing to eradicate. With the POMDP solution, we account for the low detection probability of the invaders, thus after managing both parcels our belief that the invasive species is eradicated is just above 0.5 (black arrow) which is an unsatisfactory level to stop managing the invasive species. Managing both parcels again increases our belief that the species is eradicated in absence of sighting (0.86). For a 2-line network the POMDP does not recommend surveying. On 3 and 5-line networks, we observed a similar strategy, if we know that all the parcels are infected, we should start managing from an extremity of the network and follow the same direction managing the remaining infected parcels. Once every parcel has been managed once, we then manage the most connected nodes (parcel 2, 3 and 4 Figure 18b), followed by managing the parcels at the extremity of the network. For the 5 line network, the POMDP solution also recommends surveying one of the most connected parcels before doing nothing. This rule of thumb also holds for a 6-population line network. Interestingly while the POMDP solution maintains a good performance and reaches a high probability of eradication before doing nothing, the performance of the MDP solution drops down to 0.2 (black arrow, Figure 18b) reinforcing the importance of using POMDP against MDP under imperfect detection.

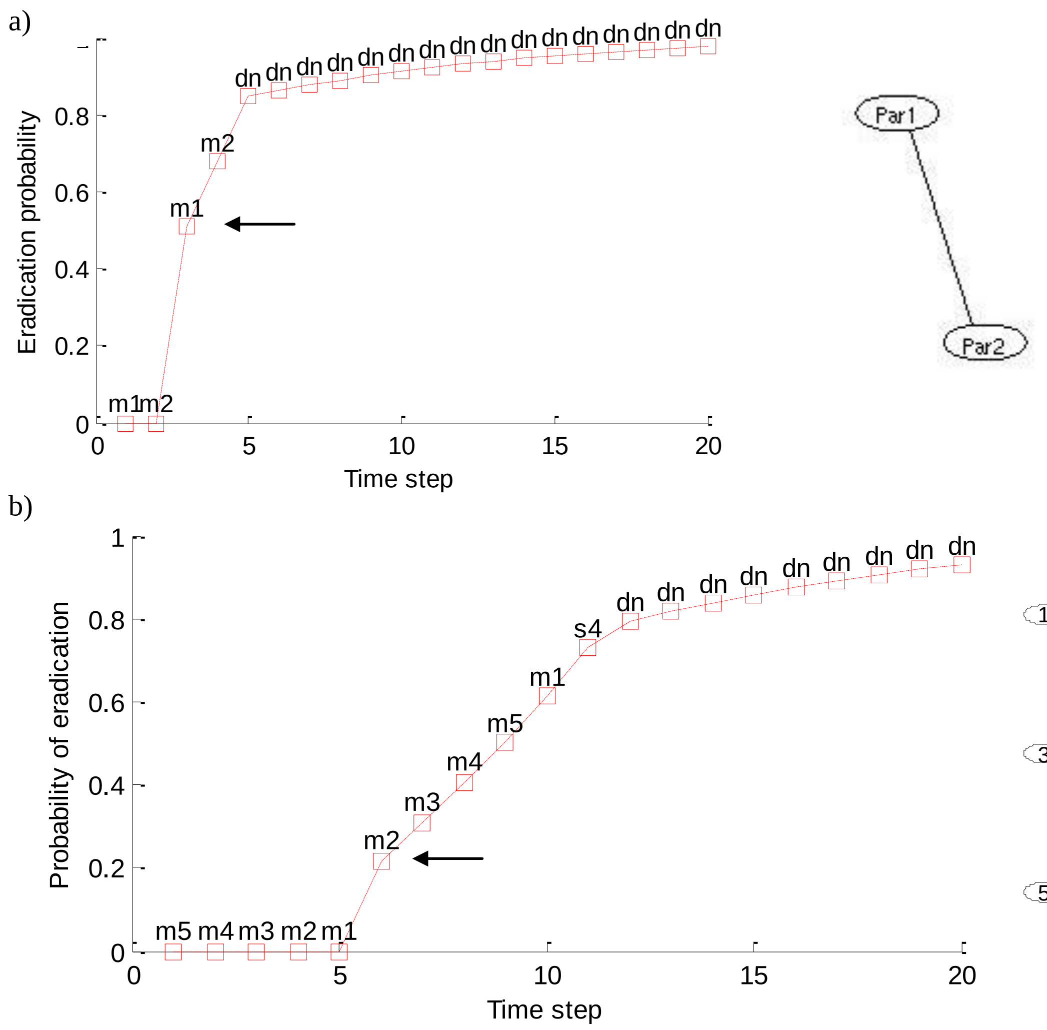


Figure 18 – 2-5-line networks invaded. Strategies discovered when each parcel of the network is invaded and invaders are detected (observation present) at time step 1 and remain unobserved once managed. mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action. The arrow represents the probability of eradication reached by a completely observable strategy (MDP). a) 2-line network b) 5-line network.

If we do not have any information regarding the state of each population and no access to previous history of observation and actions, the strategy when invasive species are unobserved (Figure 19) recommends managing population 1 and population 2 before surveying population 1 (e.g. the first population we managed), then the do nothing action is recommended. Our method in the partially observable case does not produce a set of equivalent solutions, but it is clear that a symmetrical solution starting with managing population 2 is an alternative option. If we were to observe an invader in one or both of the populations our belief that the population is eradicated will drop to zero.

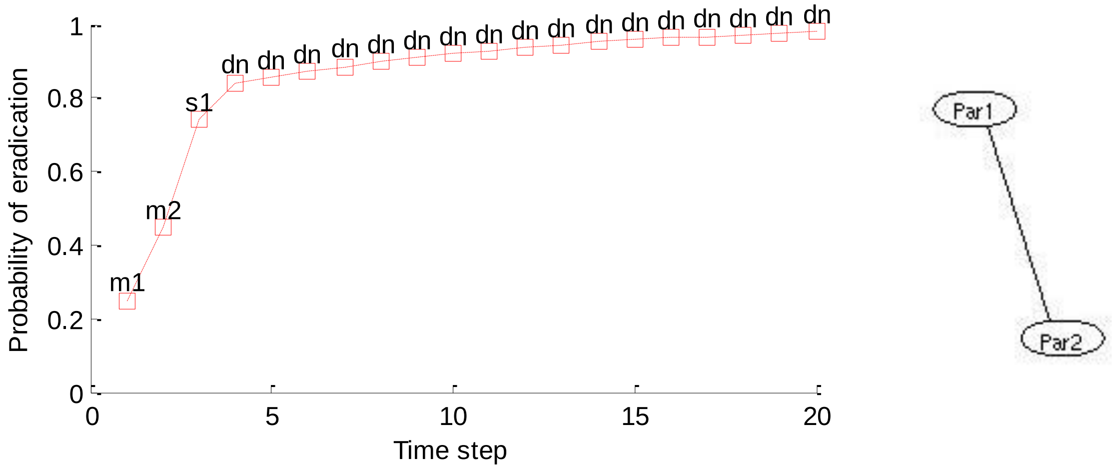


Figure 19 – Line network no prior information. Strategies discovered when we have no prior information about the history of the invasion and the invaders remain unobserved over times (mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action).

If we do not have access to prior information regarding a 3-line network (Figure 20), we would first manage the parcel the most connected before tackling the least connected parcel. When our belief of eradication is high the POMDP solution recommends surveying the most connected parcels. Similar strategies have been observed on 5 and 6 line networks.

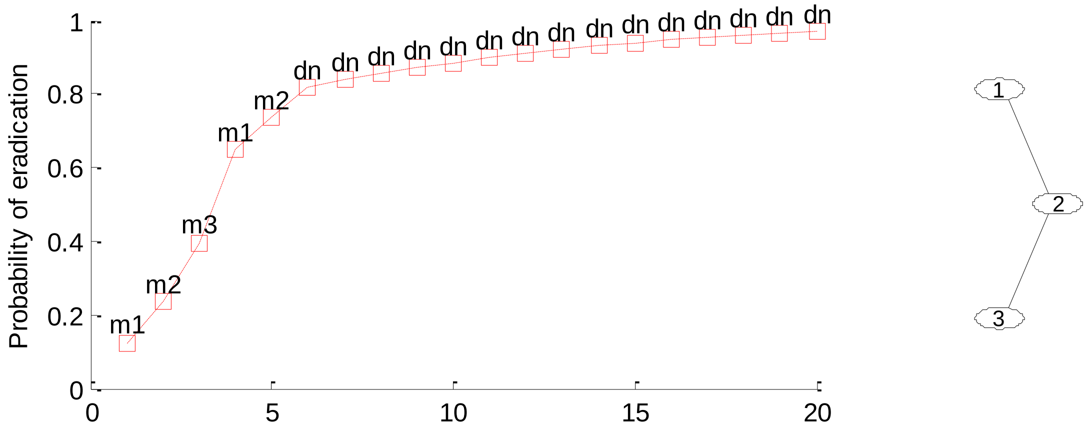


Figure 20 – 3-line network no prior information. Simulation of the POMDP strategies on a 3-line network (mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action).

In the case of a directed interaction between both populations, the POMDP solution recommends managing the source population first. If we had to survey we would also survey the source population first (Figure 21). This rule holds for any directed line network (Figure 22).

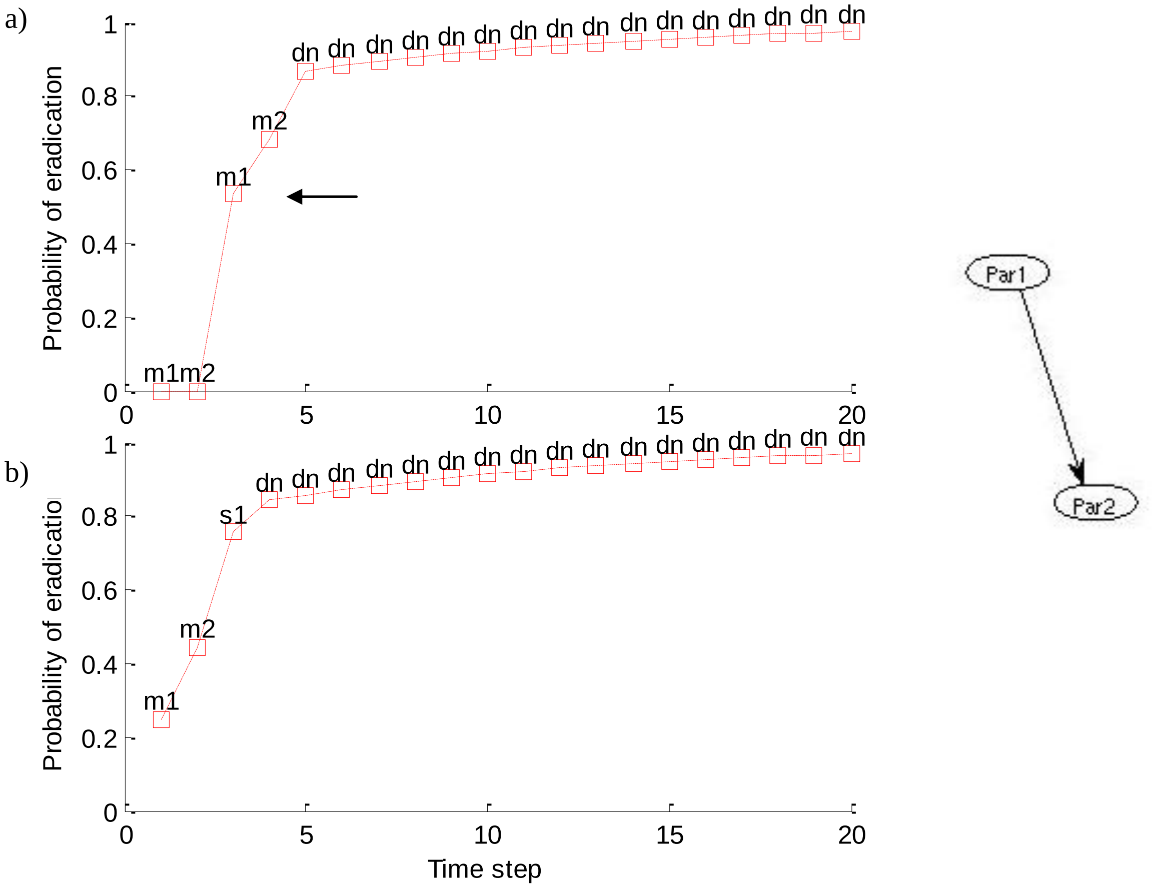


Figure 21 – Simulation of the POMDP strategies on a source-sink scenario. a) invaders are present and remain unobserved once managed b) Invaders are not observed and remain absent (mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action). The arrow represents the probability of eradication reached by a completely observable strategy (MDP).

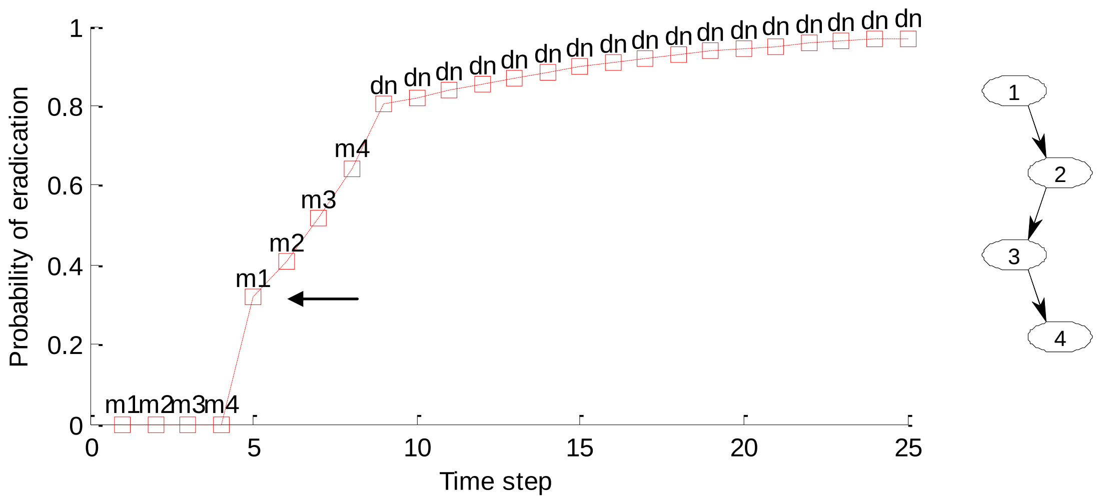


Figure 22 – Simulation of the POMDP strategies on a directed line network of 4 nodes. Invaders are present and remain unobserved once managed (mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action). The arrow represents the probability of eradication reached by a completely observable strategy (MDP).

### Island networks

On an island network of 3 nodes, the starting management node does not matter. We can start managing or surveying from any node and as in the completely observable case, then manage every node which is the nearest from a previously managed node. Intermittent survey action can randomly sample the state of the parcels at high belief of eradication (Figure 23a). We obtain similar results for 4, 5 and 6 island network.

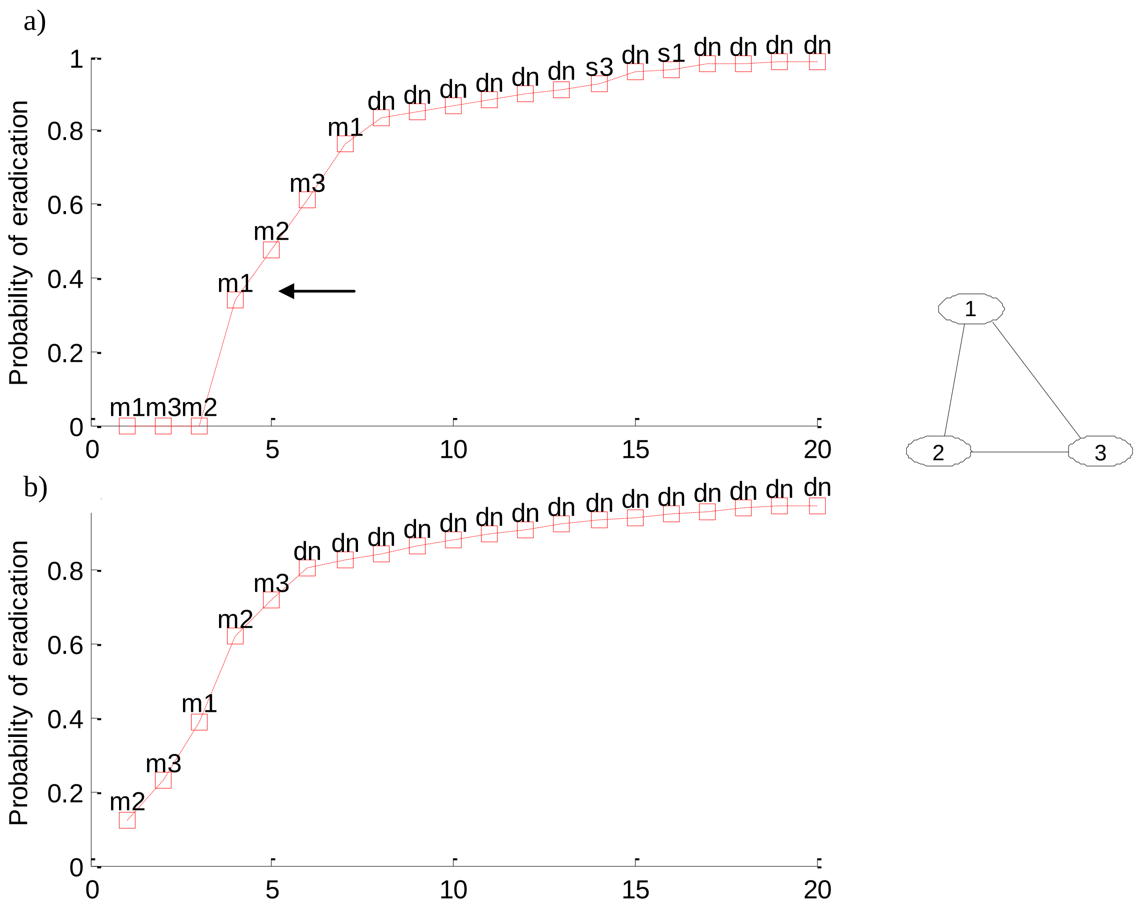


Figure 23 – Simulation of the POMDP strategies on a 3-island network. a) Invaders are present and remain unobserved once managed b) Invaders are not observed and remain absent (mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action). The arrow represents the probability of eradication reached by a completely observable strategy (MDP).

### Star networks

When dealing with star network configurations, the detection probability changes the rule of thumb we have previously discovered when detection was perfect. Instead of managing the most connected parcel when the number of satellites empty is equal to or greater than the remaining infected satellites, we discover that the number of satellites observed as empty must be strictly greater than the number of satellites observed infected. Surveying the central parcel is recommended even when our belief that the invasive species is high. We have found similar results on star networks of 4 and 5 nodes. For 6-node star networks the POMDP strategy recommends to manage one of the satellites twice before tackling the central node (Figure 24). Furthermore as the number of nodes and interactions grow in the system, our confidence in the eradication of the invaders decreases even more. The POMDP solution recommends returning to manage previously managed parcels for a longer time

(Figure 24) and to reach satisfactory eradication probability before doing nothing. A MDP solution would have proposed to wait for the presence of invaders before managing and to stop at a very low level of eradication probability (less than 0.2).

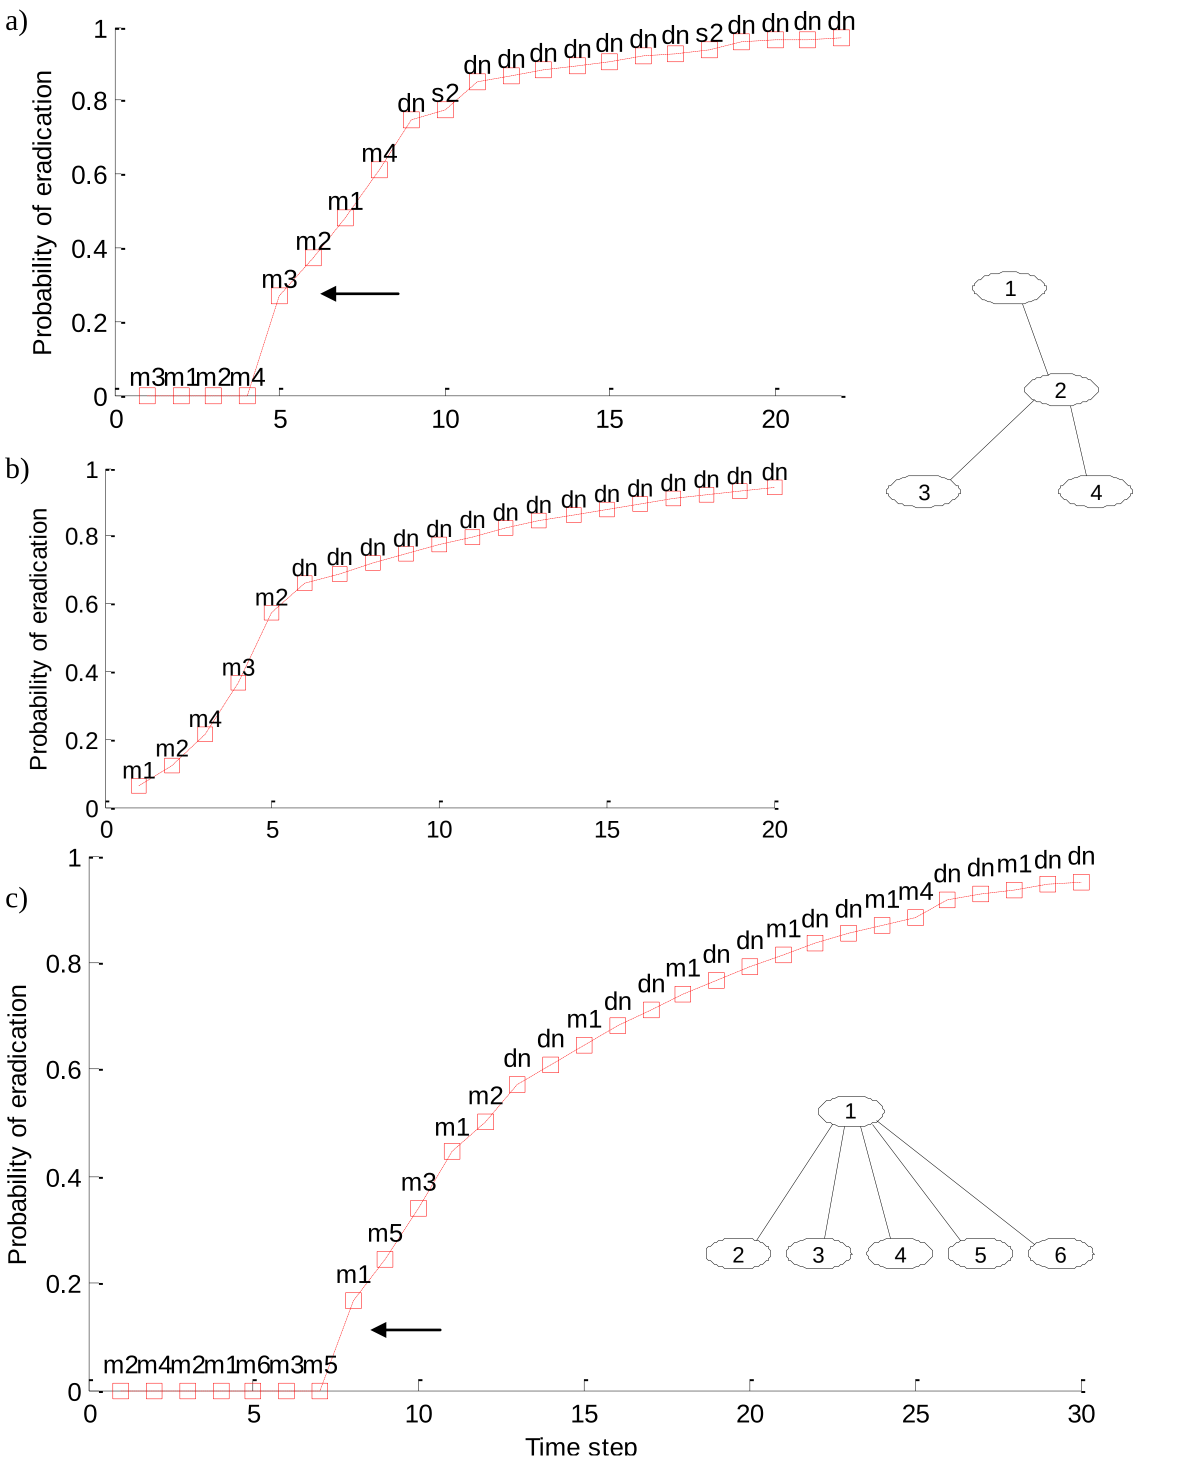


Figure 24 – Simulation of the POMDP strategies on a 4-star network (a, b), 6-star network c). (mi stands for manage parcel i; si stands for survey parcel i; dn is the do nothing action). The arrows represent the probability of eradication reached by a completely observable strategy (MDP).

## 4.4 Discussion

Finding strategies for managing invasive species at a landscape level is a challenging problem that can be modelled as a partially observable Markov decision process. We have shown on small networks that taking into account our ability to detect changes the optimal strategy. When dealing with imperfect detection we must account for the risk of remaining unseen invaders in the landscape. Strategies are now defined over belief states instead of states of the system. Tracking the belief state that all the parcels are empty (probability of eradication) provides guidance about how far we are from our eradication objective. Our results suggest that by using a completely observable strategy we would only manage invasive species when they are observed as present and we would give up on management too soon failing to achieve our eradication objective. The larger and more complex our network is the more relevant the use of our POMDP model becomes.

We discovered rules of thumb for line, source-sink, island and star type networks to help managers prioritise their actions. The managing order from the completely observable case (MDP) is robust and can be reused for the first steps of a partially observable problem, but the difference is, we should manage again previously managed parcels even when invaders are not detected. When the invasive species is unobserved managers should manage in priority parcels that are highly connected. The length of time we should manage each parcel is a difficult general question that depends on the set of parameters we have defined: colonisation process, management efficiency, ability to detect and also cost and benefits. When solving a POMDP we are looking for a strategy over a finite time horizon which will find the most effective cost-benefit strategy to achieve our eradication goal. If the time horizon is too short and the cost of managing is too high, the do nothing action will be recommended by the POMDP solution.

Overall the pattern of the recommended strategies for managing invasive species is in accordance with previous studies on management of cryptic endangered species: It is optimal to first manage, then survey and finally surrender (Chades et al. 2008). Because survey has a relative high cost and does not contribute to the removal of the invasive species, it has a low priority in the management strategy. If managing was significantly more expensive than survey, survey would appear sooner in the optimal strategy. Interested readers are referred to Chades et al (2008) and McDonald-Madden et al (in review) for further results on the sensitivity of POMDP solutions.

Using a hypothetical species, we have demonstrated the interest of POMDP models to optimally manage invasive species across a landscape. The benefits of using POMDP over MDP when species are difficult to detect is clear. However, solving a spatial POMDP on a real case study is yet to be done. Using POMDP comes with a set of modelling and complexity constraints. First, while we can legitimately say that the success of managing invasive species is highly dependent on our ability to detect invaders, having access to detection probability information is rare. Second, solving POMDP remains a challenging exercise for large state problems. To overcome the computational challenge we used a structured representation to model and solve our problem. However we have experienced difficulties solving networks with 7 nodes or more, reducing the range of applications we can tackle at this time.

We have presented exciting results for managing invasive species at a spatial level under imperfect detection. Managers are also interested in rules of thumb to prioritise the surveillance of invaders when detection is imperfect and parcels are at risk of invasion. Under resource constraints surveying a network in the most efficient way can also be modelled in a POMDP framework. Indeed we could also derive rules of thumb to optimally survey across networks. A general result is to survey the most connected nodes in sequence. Given a network structure, transition and detection probabilities, cost, benefit and time horizon, our decision tool is able to produce optimal strategies to survey invasive species on a network in the most cost-efficient way. However being able to specify management constraints into the POMDP model would help its application in ecology. For example, management actions might only be feasible when we observe the invaders (e.g., hand removal, fumigation). Although it is possible to overcome this difficulty temporarily using a high cost function for undesirable state action transitions, improving the definition of POMDP to include such constraints should improve their applicability as well as speeding up the computation of the POMDP. We envisage tackling these topics in future research.

We are currently applying our spatial POMDP on the management of an introduced fish Gambusia which threatens endemic fish and snail species of endangered mound springs in Central Queensland.

# 5. Recommendations

When detection of invasive species is perfect, we recommend using structured Markov decision process with algebraic decision diagrams to guide the prioritisation of management of invasive species on a network. We obtained fast optimal solutions with networks of 15 nodes (in less than an hour). We found the following rules of thumb for the management of invasive species across space when all parcels are infected:

- Source-sink or directed networks: manage the source first and the sink last.
- N-line networks: start from an extremity of the line and keep managing parcels following the same direction.
- N-Star networks: manage the infected satellite parcels until the number of satellite parcels empty is equal to or greater than the number of satellite parcels infected. Then manage the central parcel and the remaining infected satellite parcels.
- N-Island networks: start managing one parcel and then manage the closest parcels in any direction.
- N-Island-k networks: start managing the extreme parcel of the network, manage the k-line, manage the connecting node and manage successfully the parcels that are the closest to the connecting node. If an island network is connected to several lines start managing from the longest line.
- Cluster network: start from the smaller cluster, if the clusters are identical start from any cluster, if a cluster is less connected start from the least connected cluster.

When detection of invasive species is imperfect, we recommend using structured partially observable Markov decision process with algebraic decision diagrams to guide the prioritisation of management and surveillance of cryptic invasive species on a network. The POMDP solution outperforms the MDP solution. We obtained near optimal solutions with networks of 6 nodes (in less than an hour).

A general rule of thumb when invaders are present across the landscape is to follow the management order derived from the MDP solutions. The POMDP solution usually recommends repeating the management of the most connected parcels even when invaders remain unobserved. This tendency increases as the number of parcels increases and the detection probability or/and the management efficiency decrease.

When we do not have access to prior information on the invasion and the invaders are unobserved, the POMDP solution recommends managing the most connected parcels in priority. Similarly if we were unable to manage a parcel when invaders are unobserved we would survey the most connected parcels in priority.

# 6. References


Bahar, R. I., E. A. Frohm, C. M. Gaona, G. D. Hachtel, E. Macii, A. Pardo, and F. Somenzi. 1997. Algebraic decision diagrams and their applications. Formal methods in system design **10**:171-206.

Boger, J., P. Poupart, J. Hoey, C. Boutilier, G. Fernie, and A. Mihailidis. 2005. A Decision-Theoretic Approach to Task Assistance for Persons with Dementia. Page 1293. LAWRENCE ERLBAUM ASSOCIATES LTD.

Boutilier, C., T. Dean, and S. Hanks. 1999. Decision theoretic planning: Structural assumptions and computational leverage. Journal of Artificial Intelligence Research **11**:94.

Boutilier, C., R. Dearden, and M. Goldszmidt. 1995. Exploiting structure in policy construction. Pages 1104-1113. LAWRENCE ERLBAUM ASSOCIATES LTD.

Braziunas, D., and C. Boutilier. 2004. Stochastic Local Search for POMDP Controllers. Pages 690-696. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

Cassandra, A., M. L. Littman, and N. L. Zhang. 1997. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. Pages 54-61 *in* Uncertainty in Artificial Intelligence.

Cassandra, A. R., L. P. Kaelbling, and M. L. Littman. 1995. Acting Optimally in Partially Observable Stochastic Domains. Pages 1023-1023. JOHN WILEY & SONS LTD.

Chadès, I., E. McDonald-Madden, M. A. McCarthy, B. Wintle, M. Linkie, and H. P. Possingham. 2008. When to stop managing or surveying cryptic threatened species. Proceedings of the National Academy of Sciences **105**:13936.

Clark, C. W., and M. Mangel. 2000. Dynamic state variable models in ecology: methods and applications. Oxford University Press, USA.

Day, J. R., and H. P. Possingham. 1995. A Stochastic Metapopulation Model with Variability in Patch Size and Position. Theoretical Population Biology **48**:333-360.

Dunne, J. A., R. J. Williams, and N. D. Martinez. 2002. Network structure and biodiversity loss in food webs: robustness increases with connectance. Ecology Letters **5**:558-567.

Firn, J., T. Rout, H. Possingham, and Y. M. Buckley. 2008. Managing beyond the invader: manipulating disturbance of natives simplifies control efforts. Journal of Applied Ecology **45**:1143-1151.

Franc, A. 2004. Metapopulation dynamics as a contact process on a graph. Ecological Complexity **1**:49-63.

Guestrin, C., D. Koller, R. Parr, and S. Venkataraman. 2003. Efficient solution algorithms for factored MDPs. Journal of Artificial Intelligence Research **19**:399-468.

Harris, T. E. 1974. Contact interactions on a lattice. The Annals of Probability **2**:969-988.

Hastings, A., K. Cuddington, K. F. Davies, C. J. Dugaw, S. Elmendorf, A. Freestone, S. Harrison, M. Holland, J. Lambrinos, and U. Malvadkar. 2005. The spatial spread of invasions: new developments in theory and evidence. Ecology Letters **8**:91-101.

Hauser, C. E., and M. A. McCarthy. 2009. Streamlining 'search and destroy': cost-effective surveillance for invasive species management. Ecology Letters **12**:683-692.

Hauskrecht, M. 2000. Value-function approximations for partially observable Markov decision processes. Journal of Artificial Intelligence Research **13**:33-94.

Hoey, J., R. St-Aubin, A. Hu, and C. Boutilier. 1999. SPUDD: Stochastic planning using decision diagrams. Pages 279–288 *in* Fifteenth Conference on Uncertainty in Artificial Intelligence.

Jordán, F., T. A. Okey, B. Bauer, and S. Libralato. 2008. Identifying important species: Linking structure and function in ecological networks. Ecological Modelling **216**:75-80.

Jupp, P., P. Warren, and N. Secomb. 2002. The branched broomrape eradication program: methodologies, problems encountered and lessons learnt. . Pages 270-273 *in* Proceedings of the 13th Australian Weeds Conference. Plant Protection Society of Western Australia, Perth.

Koller, D., and R. Parr. 1999. Computing factored value functions for policies in structured MDPs. Pages 1332-1339.

Littman, M. L. 1996. Algorithms for sequential decision making. Brown University, Providence, RI.

Madani, O., S. Hanks, and A. Condon. 2003. On the undecidability of probabilistic planning and related stochastic optimization problems. Artificial Intelligence **147**:5-34.

May, R. M. 2006. Network structure and the biology of populations. Trends in Ecology & Evolution **21**:394-399.

Milo, R., S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. 2002. Network motifs: simple building blocks of complex networks. Science **298**:824.

Monahan, G. E. 1982. Survey of Partially Observable Markov Decision Processes: Theory, Models, and Algorithms. MGMT. SCI. **28**:1-16.

Montoya, J. M., S. L. Pimm, and R. V. Solé. 2006. Ecological networks and their fragility. Nature **442**:259-264.

Nicol, S., I. Chades, and H. P. Possingham. 2009. Conservation decision-making in large state spaces.*in* MODSIM.

Panetta, F. D., and S. M. Timmins. 2004. Evaluating the feasibility of eradication for terrestrial weed incursions. Plant Protection Quarterly **19**:5-11.

Papadimitriou, C. H., and J. N. Tsitsiklis. 1987. The Complexity of Markov Decision Processes. MATHEMATICS OF OPERATIONS RESEARCH **12**:441-450.

Pimental, D., editor. 2002. Biological Invasions: Economic and environmental cost of alien plant, animal, and microbe species. CRC Press, Boca Raton, FL.

Pineau, J., G. Gordon, and S. Thrun. 2003. Point-based value iteration: An anytime algorithm for POMDPs. Pages 1025-1032 *in* International Joint Conference on Artificial Intelligence. LAWRENCE ERLBAUM ASSOCIATES LTD.

Possingham, H. P. 2001. The business of biodiversity: applying decision theory principles to nature conservation. Tela–Environment, Economy and Society **9**.

Poupart, P. 2005. Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes. University of Toronto.

Proulx, S. R., D. E. L. Promislow, and P. C. Phillips. 2005. Network thinking in ecology and evolution. Trends in Ecology & Evolution **20**:345-353.

Puterman, M. L. 1994. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, Inc. New York, NY, USA.

Regan, T. J., I. Chadès, and H. P. Possingham. in review. Optimal strategies for managing invasive plants in partially observable systems.

Regan, T. J., M. A. McCarthy, P. W. J. Baxter, F. Dane Panetta, and H. P. Possingham. 2006. Optimal eradication: when to stop looking for an invasive plant. Ecology Letters **9**:759-766.

Ross, S., J. Pineau, S. Paquet, and B. Chaib-draa. 2008. Online Planning Algorithms for POMDPs. Journal of Artificial Intelligence Research **32**:663-704.

Shea, K., and H. P. Possingham. 2000. Optimal release strategies for biological control agents: an application of stochastic dynamic programming to population management. Journal of Applied Ecology:77-86.

Sinden, J., R. Jones, S. M. Hester, D. I. S. Odom, C. Kalisch, R. James, O. J. Cacho, and G. Griffith. 2005. The economic impact of weeds in Australia. Plant Protection Quarterly **20**:25-32.

Smallwood, R. D., and E. J. Sondik. 1973. The optimal control of partially observable Markov processes over a finite horizon. Operations Research **21**:1071-1088.

Smith, T., and R. Simmons. 2004. Heuristic search value iteration for POMDPs. Pages 520-527 *in* Uncertainty in Artificial Intelligence. AUAI Press Arlington, Virginia, United States.

Snyder, R. E., and R. M. Nisbet. 2000. Spatial structure and fluctuations in the contact process and related models. Bulletin of Mathematical Biology **62**:959-975.

Sondik, E. J. 1971. The Optimal Control of Partially Observable Markov Processes.

Spaan, M. T. J., and N. Vlassis. 2005. Perseus: Randomized point-based value iteration for POMDPs. Journal of Artificial Intelligence Research **24**:195-220.

Strogatz, S. H. 2001. Exploring complex networks. Nature **410**:268-276.

Taylor, C., and A. Hastings. 2004a. Finding optimal control strategies for invasive species: a density-structured model for Spartina alterniflora. Ecology **41**:1049-1057.

Travis, J. M. J., and K. J. Park. 2004. Spatial structure and the control of invasive alien species. Animal Conservation **7**:321-330.

Urban, D., and T. Keitt. 2001. Landscape connectivity: A graph-theoretic perspective. Ecology **82**:1205-1218.

Zhang, N. L., and W. Zhang. 2002. Speeding up the convergence of value iteration in partially observable Markov decision processes. Journal of Artificial Intelligence Research **14**:29-51.