# Social networks save energy: optimising energy consumption in an eco-village via agent-based simulation:

Majd Hawasly[1], School of Informatics, University of Edinburgh
David Corne, School of Mathematical and Computer Sciences, Heriot-Watt University
Sue Roaf, School of the Built Environment, Heriot-Watt University

*Abstract*

Energy-conscious communities are continually challenged to optimise the usage of electricity, maximising the benefits obtainable from generating systems (such as photovoltaics and wind turbines) and minimising the reliance on the national supply. Achieving an ideal balance is complicated by many factors, including the fluctuating availability of supplies that depend on solar and wind generation, and the varying patterns of domestic usage. It seems clear that optimising the net energy balance in a community depends on the degree to which householders can be persuaded to modify their usage patterns. We consider two questions that arise in this area. First, given a collection of realistic preferences and constraints on usage patterns for individual households, what degree of energy saving is possible by optimising 'within' these preferences and constraints? Second, what amounts of energy saving are possible when the community exploits its social network by sharing the usage of electrical appliances? These questions are investigated in the context of the planned experimental Riccarton Ecovillage (20 homes). A model of the Riccarton Ecovillage was implemented using the Repast.Net agent-based modelling (ABM) toolkit, and the model was simulated under a range of scenarios. In particular, optimisation methods were wrapped around the simulation model, exploring the space of usage patterns (within given constraints and preferences) to find effective combinations of electricity usage schedules that minimised dependence on the national supply. Our findings are: evolutionary algorithms perform particularly well at this difficult optimisation task; modest savings of 5—10% are achievable under standard assumptions, but savings of 35—40% are achievable in communities that exploit their underlying social network.

## 1. Introduction

Any community of interacting individuals is a *complex system*. In general, such a system of interconnected and interdependent structured components exhibits behaviour that is difficult or impossible to predict when we use simplified models that ignore the temporal and spatial heterogeneities inherent in the system. This is especially the case when the components themselves exhibit nondeterministic stochastic behaviour. Understanding the behaviour of a complex system is a major problem. Even if the structure and behaviour of individual components of a complex system can be captured and described clearly, the behaviour of the complex system depends also on the complicated interaction map of its components, making the whole more than the sum of the parts.

The difficulties in understanding and controlling complex systems are particularly salient in the context of many contemporary challenges, including, in particular, the global need to reduce our

[1] This work was done while the first author was at the School of Mathematical and Computer Sciences, Heriot-Watt University

reliance on fossil fuels. In addressing these challenges, we often need to predict and/or control the behaviour of systems that we don't yet understand. This can be particularly damaging when it goes wrong. However, the relatively new field of science called *complexity science* is attempting to develop appropriate methods to handle and understand system complexity. One prominent tool that emerges from complexity science is the use of agent-based modelling (ABM) – also known as Individual Based Modelling (IBM) (Hood, 1998; Gimblett, 2002; Edmonds, 2005). For example, if we wish to predict the energy consumption of a community of 1,000 people over a year, one way to do this would be to measure the energy consumption of a community of 100 people over 6 months, and multiply that by 20. This is obviously immensely simplified, and makes several strong assumptions about the relationship between the number of people in a community and that community's energy consumption. Perhaps the crudest simplification in this simplified model is the absence of any factor other than the number of people and the time period. At the other extreme, one could imagine a sophisticated model of the community whose energy consumption we are aiming to predict. Rather than use simplifying assumptions, such a model encodes distinct behaviours for each of the individuals in the community, realistic time-varying schedules for factors such as sunlight and weather, and also encodes reasonable and stochastic rules for the interactions between individuals and their responses to external factors. Obtaining predictions from such a model then arises from simulating the model – i.e.: running the agent-based simulation – and observing and collating the results.

Such a simulation based model also allows us to explore 'what-if' questions, and can be used as a tool to explore all variety of questions concerning how the components of a complex system – in this case, behavioural policies of households and communities with regard to energy usage – might be designed to achieve specific outcomes. In particular, by using the agent based simulation within an optimisation procedure, we can automatically search through the space of ideal policies in order to optimise the resulting behaviour of the system.

In the context of an eco-village which has its own means to generate electricity (its own microgrid), one question that can be explored is how the eco-village community might maximise the use of its own supply, and minimise its reliance on the national supply – i.e. how the village can reduce its net use of the national supply. What makes this a complex challenge is the fluctuating and unpredictable nature of the microgrid, which will be dependent on sunlight and wind, and the difficulty of electricity storage. Meanwhile, individual households tend to have patterns of electricity consumption that vary within predictable bounds (e.g. TV mostly in the evening, washing machine usually in the morning on a Monday, etc…), and preferences and 'comfort zones' around those bounds. For example, a household that tends to use its dishwasher between 8 and 10pm might be easily persuaded to schedule this consumption between 7pm and 8:30pm for the benefit of the community, but would not be amenable to shifting this usage to 11:00am.

The first question that we explore is: given a collection of realistic preferences and constraints on usage patterns for individual households, what degree of energy saving is possible by optimising 'within' these preferences and constraints? This question is explored using an agent based simulation model of the planned Riccarton eco-village in which differing but realistic preferences and constraints are assumed for each household, and realistic data is used to inform the time-varying level of electricity supply available from the village's own microgrid.

Second, what amounts of energy saving are possible when the community exploits its social network by sharing the usage of electrical appliances? These questions are investigated in the context of the planned experimental Riccarton Ecovillage (20 homes). A model of the Riccarton Ecovillage was implemented using the Repast.Net agent-based modelling (ABM) toolkit, and the model was simulated under a range of conditions. In particular, optimisation methods were wrapped around the simulation model, exploring the space of usage patterns (within given constraints and preferences) to find effective combinations of electricity usage schedules that minimised dependence on the national supply. Our findings are: evolutionary algorithms perform particularly well at this difficult optimisation task; modest savings of 5—10% are achievable under standard assumptions, but savings of 35—40% are achievable in communities that exploit their underlying social network.

The remainder of the paper is set out as follows. In section 2, we briefly review the concept of an eco-village, and describe some details of the planned eco-village in Riccarton, Edinburgh. In section 3, we briefly introduce the Repast system – the agent based modelling toolkit used in this study, and then broadly describe how it was used to model the Riccarton eco-village. Section 4 then explains how we use optimisation methods wrapped around the agent-based model in order to find sets of usage patterns that combine to maximal overall benefit. This section discusses the two main optimisation experiments. Section 5 presents the results, partly from the viewpoint of competing approaches to the optimisation task, but largely from the viewpoint of the solutions achieved for the task at hand, comparing the relative benefits available via organising consumption in the context of a social network. We present a concluding discussion in section 6.

## 2. Eco-villages

The concept of ecovillages has emerged in response to natural eco-system deterioration and climate change. It is regarded as one kind of the broader concept of intentional communities, where people live together in communities that share common beliefs and intentions (Wikipedia, 2009). Ecovillages are self-sustained, ecologically-sustainable intentional communities, where a small group of people (50-500) can live and develop naturally in a full-featured environment which is also ecologically-, economically- and socially-healthy [Gilman,1991; Kasper, 2008) People living in an ecovillage share the responsibilities of the community, while enjoying the warmth of close relationships and a dense social network.

Ecovillages are characterized usually by:

- harmless integration with nature (Gilman, 1991) through the promotion of recycling behaviour, the avoidance of toxic substances and the exploitation of environment-friendly power sources, using wind turbines, photovoltaic solar systems, geothermal and biomass plants, rather than fossil fuel. Ecovillages have been established to prove that full-featured communities can live and develop with minimum carbon and waste footprint.

- urban design strategies that lower the environmental impact of the community: buildings in an ecovillage are characterized by a higher degree of insulation and proper architectural form to achieve the best energy use efficiency. Buildings are usually oriented to be exposed to the highest level of solar radiation. Also, alternative water and sewage processing systems are utilized.

- 'green' trends in lifestyles, exhibiting more cooperation and less consumerism.

Ecovillages are also ideal examples of real-world complex systems that can be studied and modelled, with the overall aim of understanding how the *interactions* between individuals' behaviour, community polices, and the natural environment impact on key issues such as energy usage and carbon footprints.

## 2.1 Riccarton Ecovillage and Living Laboratory (REALL)

The Riccarton Ecovillage and Living Laboratory is a 20-household village that will be built at the main campus of Heriot-Watt University in Riccarton, Edinburgh. The project has emerged as a multidisciplinary initiative by the Energy Academy in the University, to be a whole-system approach to study the "performance of energy use and supply in buildings" (Energy Academy Heriot Watt, 2009). The project aims to design, build and then thoroughly monitor an inhabited ecovillage. The project is expected to help reach better understanding of the effects of behaviour adaptation on energy demand, the energy performance of different materials, forms and designs, and to thoroughly investigate various power generation methods on site, renewable and non-renewable. The project aims to provide a unique insight into the reality of how occupied buildings actually work in practice to provide data to underpin emerging standards and recommendations designed to meet the low carbon targets for the economy and greenhouse gas reduction, and to inform decision makers, stakeholders, researchers and people with real findings on the new climate-adapting built environment (Roaf et. al., 2009).

The village will comprise five blocks of four houses. All the houses will have the same size, form and orientation, but will differ in their methods of construction and materials. The houses will be occupied by postgraduate students who will participate in researching and monitoring energy performance aspects including "thermal, carbon, acoustic, water, waste, climatic and environmental performance, costs and impacts of occupied homes" (John Gilbert, 2008).

The village will have a private microgrid, containing a wind turbine and photovoltaic solar panels for each house. The microgrid should satisfy the energy needs of the village, but it would be linked to the general grid for backup purposes.

# 3. Modelling the REALL

## 3.1 On Modelling complex systems

A complex 'system' is a collection of interconnected and interdependent structured components (wikipedi, 2009), often exhibiting some characteristic 'emergent' behaviour (Boccara, 2004). In a complex system, the number of features and components is usually moderate or large, and the interactions between the components are usually non-deterministic. From one viewpoint, the 'essence' of a complex system is simply that the outcome of these complex interactions (in other words, the medium and long term behaviour of the system itself), tends to defy understanding and prediction by currently known analytic methods.

A classic example of a complex system is the weather. We are all familiar with the attempts to predict weather by meteorological centres around the world, but perhaps we are less familiar with the fact that the accuracy of weather predictions tends to be rather poor, and degrades sharply the further in advance is the prediction (Stern, 2008). The basis of weather prediction is a model of the atmosphere that treats it as a collection of adjacent air masses, each with dimensions on the order of (typically) 4km. A weather model attempts to simulate the interaction between these masses using physical laws which, given initial estimates of temperature, pressure, wind speed, wind direction and various other measurements for each mass, determine the states of these parameters after a given length of time (e.g. 5 minutes later). This process is iterated in order to develop predictions over longer timescales. When it comes to weather prediction, the physical laws tend to be well understood, but the quality of the prediction relies crucially on our knowledge of the initial conditions.

This example helps introduce some key aspects of ABMs. It is clear to most of us that, as we move ahead in computer processing power, and are able to model the atmosphere with finer-grained detail (larger meshes of smaller air masses), and are able to more accurately measure, and hence provide, initial conditions to the models, then weather prediction will gradually improve. Few have any argument with this, because it is clear that the weather we see is a result of physical laws, and weather simulation simply models and simulates those laws. The key to appreciating the power of ABMs – where we typically model interacting intentional agents, rather than air masses and such – is that much the same can be said of the rules than underpin (e.g.) human interaction, or some other aspects of human behaviour that we may wish to model. If we can characterise (e.g. in terms of ‘what’, and ‘how often’) the typical interactions that occur between the agents that we wish to model, then an ABM simulation of the system is, in the case of most complex systems, a highly effective, and perhaps the only, way to generate predictions of overall behaviour for that system. The increasing popularity and success of ABMs (Hartmann, 2005), now being explored for many applications (some examples are listed in section 3.3), is partly down to the fact, when well designed, they produce plausible and actionable results – essentially their predictions fit data, and do so better than any available non-ABM models. For example, one broad class of particularly simple ABM models are so-called cellular automata (CA), and CA based models of the spread of HIV infected cells within sufferers (in which an ‘agent’ is a cell) seem to reproduce the complex dynamics of healthy vs infected cell concentration over time that is revealed by observations of patients (dos Santos and Coutinho, 2001; Corne and Frisco, 2006). This match with observed data has to date eluded a variety of mathematical modelling approaches, but is achieved by an ABM with a set of simple but plausible rules for the interactions over time between spatially neighbouring agents.

ABMs are a key tool arising from the relatively new field of science called Complexity Science, which is attempting to develop appropriate methods to handle and understand system complexity (Edmonds, 2005). That involves discovering the patterns that recur in complex systems, developing the mathematical foundations of complexity, and creating tool sets to facilitate the study; e.g. to understand complex data (e.g. using machine learning techniques), to create and evaluate models of complex systems (e.g. using cellular automata and agent-based modelling), and to measure the complexity in systems (Shalizi, 2006). Research in Complex Systems is interdisciplinary, spanning physical, ecological, economical, social and political sciences, in addition to computing.

### 3.2 Simulation and Optimization

By enabling simulations of a complex system, an agent-based model (or similar type of model) is excellent for exploring basic `what-if’ s that explore the consequences of the assumptions and starting conditions from which it is developed, but this in itself has limitations for the management of such a complex system. When we have run the model several times and have therefore characterized ‘what will happen’, an entirely typical is that we would rather something different happened, and our goal is to find out how to set up the initial conditions in such a way that our ideal outcome matches the emergent behaviour of the system. For example, our model of the placement of fire exits may, according to our agent-based model, mean that the building takes 20 minutes to evacuate, however we want this to happen in 15 minutes. How do we position the fire exits to achieve this goal? The agent-based model cannot by itself answer our question. Meanwhile a naïve (but sometimes the only viable) approach is to make inspired guesses of suitable designs and then run the model to see the outcome of these guesses – e.g. experiment with different fire exit placements until we happen upon a winning design. But this approach tends to be untenable for two main reasons. First, the number of possible configurations is usually enormous, far too many to explore by repeated simulation; second, and recalling the essence of complex systems, the better designs may well be counter-intuitive, and unlikely to be among those that would be chosen for our simulation experiments.

Agent-based models thereby provide a partial solution to the need to identify ways to *design* aspects of complex systems so that they will meet desired requirements for behaviour, but they cannot do the entire job. Increasingly, this gap is being filled by *optimisation* algorithms. An optimisation is simply an iterative (usually) and fully automated process for finding the set of parameters that is 'best' according to a particular measure of quality. In the context of agent-based models of complex systems, the measure of quality may typically be the distance between the desired outcome (e.g. 95% of the population of a village become regular users of the recycling bins within 2 years) and the actual outcome of the simulation, perhaps averaged over a number of trials (e.g. for a certain configuration of incentives and policies, the simulation may yield 90% of the population as recyclers within 2 years). An optimisation algorithm attempts efficiently to search the space of possible system configurations, aiming to find good solutions reasonably quickly. Optimisation of complex problems and/or simulations is gaining increasing interest (Laguna, 1997; Fu, 2002; Olafsson & Kim, 2002; Fu *et al*, 2005). A number of different optimization algorithms are available, arising from various subfields of mathematics and of computer science. For problems with large discrete solution spaces, metaheuristics and random search provide a variety of methods and techniques that are simple to implement and with good convergence attributes. Genetic algorithms are one of the most promising tools, that are efficient, robust against local optima, and which can scale flexibly according to the problem settings.

### 3.3 The REALL Model: General Considerations

An agent-based model (ABM) of the Riccarton Ecovillage and Living Laboratory (hereafter: the REALL model) was conceived in order to explore aspects of how communities might best work together to achieve energy savings. In particular, envisioning the REALL community as operating its own microgrid, one aspect of interest is the energy balance, i.e. the hour to hour difference between the micro-grid supply and residents' demand, and how this varies with people's behaviour, how they adapt to new circumstances, and the microgrid energy-pricing policy. In the REALL model, this microgrid is assumed to comprise photovoltaic solar panels installed in each house, and one village-wide wind turbine. The electricity supply in the ecovillage is assumed to depend on these renewable sources, reverting to the national supply when necessary. Clearly, the village's microgrid supply will vary considerably according to the changing weather, seasons and time of the day. Energy demand, on the other hand, is linked to the behaviour of the residents and their ability to adapt and control their consumption habits.

The REALL model therefore requires simulation, at some level, of individuals (the inhabitants of the ecovillage) and natural phenomena (weather conditions). This implies the following requirements:

- Realistic time-varying values for the natural phenomena should be modelled. Particularly: wind velocity, solar radiation and temperature, which are all major factors that affect microgrid production.
- The influences of different weather conditions on the power generation devices, solar panels and wind turbines, should be identified, taking into consideration the limitations of these devices and their performance characteristics.
- Finally, simulation of each individual household's energy demand through time is needed. To investigate adaptation of behaviour, the consumption habits of households should be modelled realistically.

The simulation of energy demand in an individual household is a key aspect of the REALL model, and works in the following way. We assume that each household contains a typical array of devices that consume energy, and consider only those devices that are typically power-hungry, such as dishwashers, televisions (which may be switched on for several hours), and washing machines. We assign to each household a realistic but different pattern of usage of these appliances. These patterns are assigned at random, but constrained to within plausible ranges. For example, in one household we may assume that the occupants prefer to use the washing machine twice a week, usually on a Wednesday and a Saturday, for one hour between 8 and 9am on a Wednesday, and between 1pm and 5pm on a Saturday; in the same household we may assume a mean amount of TV viewing of 5 hrs per day, starting and ending within a 6 pm and 1am window. When the REALL model operates, whether or not a particular device is in use in a particular household at a particular time is decided randomly (by the houseSchedule agent, see below), but the random decision is biased by the preferences and constraints that have been set for that household. Finally, a main feature of our experiments with the REALL model is the exploitation of social networks, in which households will from time to time share the use of certain resources with other households. To underpin this, in some experiments a random network of links exists between households, which defines a social network in the village (whose density can be varied). Households' behaviour for such s includes, for each resource, probabilities that govern sharing that resource from time to time with other households that are linked directly with that household in the social network.

In summary, the REALL model is an energy-focused agent-based model of the Riccarton Ecovillage, in which the agents (in more detail below) are the individual households. The model is relatively small scale (since REAL comprises only 20 households), and the interactions that will be modelled are relatively simple. The major source of complexity and interaction will be the interplay between energy supply, weather, and individual behaviour.

## 3.4 Detailed Design of the REALL Model

There are now several open-source software libraries in the ABM research community for developing agent based models (Railsback et al, 2006). For the present work, we used Repast (Recursive Porous Agent Simulation Toolkit), developed at the Argonne National Laboratory – a highly regarded toolkit for this purpose, that has been used for a range of applications ranging through the exploration of business strategies (López-Sánchez *et al.*, 2005), the effects of charging for road use (Takama, 2005), simulation of digital markets (López-Sánchez et al, 2005), battlefield simulation (Baker *et al.*, 2005), the growth of hydrogen transportation infrastructure (Stephan and Sullivan, 2004), the evolution of house prices (Bossomaier et al, 2007), and many more.

We next describe the REALL model in a moderate level of detail, sufficient for readers to understand how such a model is designed, as well as to clarify the underlying assumptions included in our REALL model, and to indicate what is simulated within the model and what isn't. The description is presented in the style of one of the more well-established methods, called the `gaia methodology' (Zambonelli et al, 03), for the design and development of ABMs, and simply amounts to indicating the *roles* of the agents in the model. In the gaia method (of which only a part was used in the design of REALL), roles are capabilities of agents, and the early design of an ABM includes outlining the required roles; this recipe of roles is then the basis for implementing the ABM (in which each agent may have one or more of the distinct roles). In REALL, a broader understanding of the ABM model itself will emerge from observing this set of roles, listed as follows, falling naturally into themed groups:

- Roles concerned with climate conditions:
    - Sun: A SunSimulator agent plays the role of indicating the expected solar radiation levels for given time and date values. Naturally this agent takes into consideration hour to hour and season to season differences.
    - Wind: The role of the WindSimulator agent is to provide realistic values for wind velocity at a fixed elevation, again given specific time and date inputs.
    - Temperature: A TemperatureSimulator agent gives the expected ambient temperature associated with given time and date values.
- Roles concerned with power generation:
    - Solar panels: The role of a PVSimulator agent is to compute an amount of generated electricity when provided with given solar radiation, power and ambient temperature levels. The inner workings of this agent are informed by characteristic performance curves and associated loss equations appropriate for the specific photovoltaic module being simulated, and of course takes account of the size of the simulated solar panel.
    - Wind turbine: The role of the WindTurbineSimulator agent is to compute an amount of generated electricity when provided with wind speed. The inner workings of this agent are informed by characteristic performance plots

associated with the specific turbine device being modelled.

- Roles associated with energy demand
  - Household: A House agent is responsible for simulating the energy production and consumption of an individual household. It interacts with a PV Simulator agent (each household has its own solar panel) in connection with energy production, and it interacts with a House Schedule agent (see below) in connection with the energy consumption of the residents of the house.
  - Household preferences: A HouseSchedule agent has the role of indicating the preferences and habits of the residents of an individual house, in connection with their usage patterns for the household's electrical appliances. In particular, the HouseSchedule agent calculates a level of power demand for its house for any given time period.
- Other roles:
  - Observer: The Observer agent's role is to co-ordinate interaction between the other agents.
  - A Timer agent is responsible for time-keeping – essentially, a common reference frame that can be queried by any agent.

In short, the further aspects of designing an appropriate ABM model amount to designing an appropriate topology of interactions among the various roles, and an appropriate overall control regime (i.e. what happens when in the simulation). The interactions between the roles in the REALL model are illustrated in Figure 1.

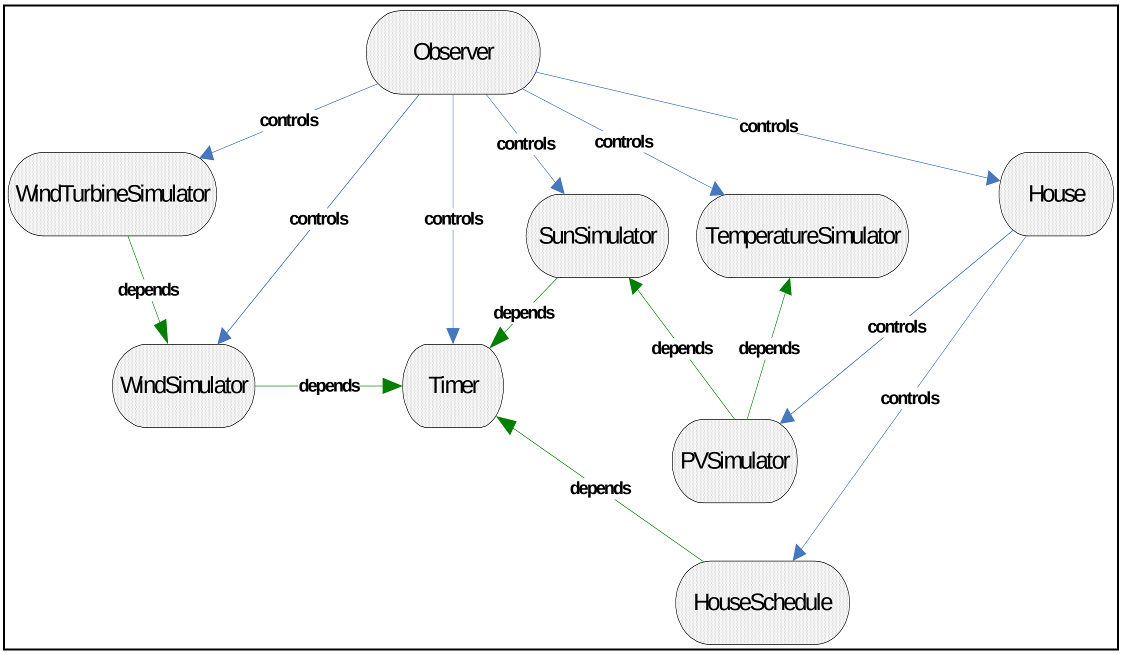


Figure 1. Roles in the REALL model, their interactions and dependencies.

With reference to Figure 1, the Observer controls the other major agents (such as House agents), which in turn control others. These 'control' links typically indicate straightforward operations – e.g. a house agent will ask its PVSimulator agent to provide the amount of energy that it can supply in the next time step. To do this, the PVSimulator agent needs to make queries to the Sun and Temperature agents, and so on.

A broad sketch of the operation of the REALL model is therefore as follows. For convenience we express this mainly in terms of what is being modelled, rather than in 'ABM-speak'. First, the Observer initialises the timer, the wind turbine, the weather (sunlight, wind, temperature), and each household in the village. Each household also initialises its solar panel, and its 'Houseschedule' agent. A single time step in the REALL model corresponds to 3hrs of simulated time. In each time unit, the climate agents calculate realistic values for the mean temperature, solar radiation levels, and wind speed for that time unit. This is done non-deterministically, but within plausible bounds. For example, based on available historical data, the temperature agent knows the normal temperature range for a given time on a given day, and will choose a random temperature within these bounds. Next, the wind turbine and solar panel agents will calculate realistic levels of energy generation given the details supplied by the climate agents and given the characteristics of the devices they are simulating. Independently, each household agent chooses randomly, but within predefined realistic bounds for the individual household, an energy demand for the current time period for each modelled electrical appliance in the household. When a 'social network' is active in the simulation, households also choose, again stochastically, whether or not a device will be shared with another household with which it is directly linked. The entire process repeats over many timesteps; in the experiments reported here, experiments typically covered a month of simulated time, with each timestep representing one hour.

## 4. Using the REALL Model is a Tool for Complexity-Informed Management

### 4.1 Simulation and Stochasticity

Management and planning, in any context, depend on the ability to predict the consequences of actions, designs and policies. In restricted scenarios, and especially when we are predicting something that is only a short step into the future, such predictions can be made with confidence and can validly inform management and planning. For example, when a soccer match has been scheduled between the two major soccer teams in London, we can confidently predict a lower bound on the number of spectators, and the police can make sensible provision for their presence; in the run-up to Christmas, postal services in many countries can predict a massive increase in demand, and arrange accordingly for suitable staffing.

This is well-known and understood. However, for current purposes it is useful to think of these examples in terms of *models*, and consider what we are doing when we make a prediction. In the case of the soccer match, we may have historical data that indicates the attendance for this match in previous years. Our predictive model in this case might simply be to fit a straight line to this attendance/year curve, and read off the prediction for the current year from the line. In the case of the postal service, there may be a precisely analogous situation – perhaps this line will show a downward trend, considering the competition for postal services in many countries over recent years, and the growth of email, but nevertheless it will show a trend, and this will lead to a prediction.

However, major difficulties arise as soon as 'complex interactions' are apparent, which undermine the feeling that historical data, or any other such simplified model, will be valid. For example, the soccer match may occur on the same day as an international cricket match elsewhere in London, perhaps both these events are being televised live for the first time, and both local travel and ticket costs may be double last years' levels, at the same time as unprecedented levels of unemployment.  In the case of the Christmas post, recent strike action by the postal services, in combination with further increase in competition for postal service provision, whereby alternative providers are able to price their services at unprecedentedly low levels, make us far less confident in simply extrapolating from past data.

In essence, the behaviour of a system becomes difficult or impossible to predict as soon as the system includes interactions between varieties of competing 'forces', especially where we have no basis in prior experience in which *the same* or *a very similar system* has operated under the same conditions. The business of complexity science is all about finding ways to arrive at predictions of behaviour in such scenarios, and the way this is done is, in one sense obvious and as follows: 'if you want to see what will happen – try it and see what happens!'. That is easily said, but it seems to imply we cannot predict it in advance – instead we wait for the soccer match to happen, or we await the Christmas postal rush, and record the data we are interested in, too late for any management actions. However, the trick in complexity science that gets around this is to see what happens by *simulating* the system. We build a complex model of the system at the levels of agents, interactions and behaviours (for each of which we can at least make inspired guesses) – rather than a mathematical abstraction based on data that we do not have – and then run the model, thereby simulating the system whose behaviour we are trying to predict.

It is useful to emphasise, at this point, that multiple runs of REALL (or any other ABM with stochastic elements – essentially all useful ABMs) will produce different results. This is entirely analogous to the multiple potential outcomes that arise from, for example, many runs of climate models that attempt to predict future climate parameters. Essentially, any such model has a characteristic probability distribution over the potential outcomes – deploying an ABM helps us gain an understanding of the relative likelihoods of different outcomes, rather than enabling us to confidently predict the future in fine detail. However, what we might generally hope for is the ability to confidently predict that certain outcomes are far more likely than others. For example, if we build into the model certain household polices and preferences for sharing resources, and after ten runs of the model we see energy savings for the community varying from 10% (in the 'worst' run) to 20% (in the 'best'), we can have some confidence in predicting that, to the extent that other aspects of the model are suitably realistic, the given policies and procedures would generally lead to energy savings of 10% or more if implemented by a real-world example of the type of community being modelled.

As a potential tool for managing a real-world complex system, an ABM can be looked upon purely as a prediction tool, in which the model is configured to reflect the current design of the real-world system, and it is then run (several times) to characterise expectations for how the real-world system is likely to behave in the future. However, development and use of an ABM is more involved, and more useful than this implies. In particular, if and when we find that the real-world system has not developed within the bounds of expectation that our model predicted, we have learned something. This could be as simple as finding a 'bug' in the model; however more likely, and more usefully, investigation will find that the disparity is due to aspects of the real-world system that were not thought important enough to include in the model, but which in fact do have a salient effect on behaviour. Perhaps an extreme, but instructive example is the case of the London Millennium Bridge, which swayed violently when pedestrians were first able to walk across it in numbers. What was left out of the 'model' in this case was the fact that pedestrians effect unconscious local interactions whereby they gradually synchronise their steps, and in the real-world case this led to setting up a damaging resonance.

## 4.2 Some Background on Optimization

Earlier in section 3.2, we discussed the distinctions between *simulation* – using an ABM model to see what happens under given starting conditions – and *optimization* – in which we use an automated method that is 'wrapped around' the simulation, and which attempts to find the starting conditions and configurations for the model that lead to ideal outcomes when the simulation is run. Optimization is in fact the main theme of the experiments we report in this article, so we will say a little more here about optimization, providing some brief background material on optimization in general, so as to place into context the optimization methods that are used in the experiments discussed later.

A typical optimization problem can be formulated as follows. Suppose we are given a system, of any kind, which is expressed in the form of a set of *parameters* (e.g. a set of numbers, each allowed to be between 0 and 1). Any particular realisation of these parameters (that is, specific values for each one), is referred to as a specific *configuration* of the system. Suppose further that we are given a quality function which, when applied to a configuration, gives us a number which we can take to be the 'quality' of that configuration. 'Optimization' simply refers to a process of trying to find the 'best' configuration – that is, the configuration that yields the best result from the quality function. Further, an optimisation algorithm is an automatic method for searching the space of configurations with a view to finding good configurations quickly.

In context, a configuration might be a specific subset of parameters for an ABM model, such as, the set of probabilities associated with using specific electrical appliances at certain times of day. The quality function might work as follows: run the ABM model, and return the net use of the national electricity supply over the simulated period. Clearly we would like this to be minimised, so that the 'best' configuration of parameters is one that yields the lowest demand on the national supply (indeed this may be negative, indicating a surfeit of generation in the ecovillage).

Optimization is a large subfield of each of mathematics and computer science, and there are a wide variety of different families of optimisation algorithms (Gray *et al*, 1997; Garcia *et al*, 2006; Griva *et al*, 2009; Weise, 2009) , each applicable to different styles of optimization problem. In navigating this space for current purposes, and using a very broad, but nevertheless valid brush, two salient generalisations can be made: first, it turns out that almost all *interesting* optimization problems (i.e. those of some practical importance to solve) are 'difficult' in a specific technical sense – this amounts to the fact that no algorithm is known that is guaranteed to find the best answer in reasonable time. The consequence of this is that optimisation research is replete with so-called approximate algorithms – these are algorithms which try to find good configurations reasonably quickly, but can never guarantee that they will find the true best result. The second of our salient generalisations is that, when it comes to complex quality functions (essentially, any quality function that is not in itself easily subject to mathematical analysis), the choice of appropriate optimization algorithms narrows down to a single select family known as *black box* or *stochastic search* methods. These are, in essence, trial and error methods, but in which the choice of the next configuration to test is guided (in ways that differ between algorithms) by the quality values that have been calculated for previously tried configurations.

Without getting into too much detail, when we do optimization in the next sections we test three exemplars of stochastic search algorithms. These are 'hillclimbing' (HC), simulated annealing (SA), and an 'evolutionary algorithm' (EA). In general, when a complex optimisation problem has to be solved, stochastic optimisation algorithms tend to have different speed/quality tradeoffs. HC can typically find good solutions quickly, but then be unable to find further improvement; meanwhile a well designed EA can typically find better quality solutions than HC, but takes a relatively long time to do so. SA tends to have a speed/quality trade-off somewhere in between these two.

Given any new complex optimization problem, it is always wise to experiment with such a range of optimisation algorithms, since the aforementioned rules of the thumb can often be violated in practice, depending (in ways which currently defeat the current state of theory in black box optimisation) on the details of the quality function. For example, in some cases, but not in most, HC can provide fast and high quality solutions that are not bettered by either SA or an EA. In our optimisation tests, we therefore try each of these three methods, in order to characterise their performance on the energy-focussed REALL model, and partly to inform the choice of optimisation method in further work on variants of this REALL model.

# 5. Simulation and Optimization Experiments for REALL

In this section we describe the experiments that have been performed with the REALL model, leaving presentation of results for a later section.

Two sets of experiments were performed. In the first set of experiments, the idea was to examine to what extent energy savings can made in the ecovillage within the constraints set by the preferences in individual households. In the second, and perhaps most interesting, set of experiments, the idea was to see what levels of savings were possible if an ecovillage exploited the social network of its occupant households (e.g. watching TV with friends), and these experiments explore how potential savings in demand vary with different densities of social linkage.

Before more detailed explanation, it is useful to note some points common to the two sets of experiments. Each household has a set of preferences which dictate its energy demand behaviour. A household's list of preferences amounts to, for each of up to 10 electrical appliances, an average amount of hours per day using that appliance, and preferred time windows during which that appliance is used. Different, random, but plausible preferences were generated for each household in the REALL model. In all cases, the time step in the simulation, regarding a households' energy consumption, was one hour. That is, the 'HouseSchedule' agent repeatedly chooses, given that household's preferences and constraints, the appliances that will be used (and therefore sets that household's energy demands) for the next hour. Finally, simulations in our experiments are restricted to one summer month. This was a pragmatic choice considering the time demands of the simulations, especially when comparing several optimization methods, each of which needs to run the simulation many times.

## 5.1 Optimizing energy balance in the Riccarton Ecovillage

The microgrid of Riccarton Ecovillage (as modelled) comprises the communal wind turbine and each household's solar panel. The ecovillage could presumably reduce its energy carbon footprint to virtually zero if the generated green electricity is consumed efficiently and no general grid power is imported. The problem, however, is the fluctuations in availability of solar and wind power. An obvious approach to achieving minimal use of grid power would be for residents to adapt their energy consumption behaviour to align with the availability of energy from the microgrid, performing energy-hungry tasks when electricity is normally available and refraining from consumption at other times. In general this may be rather too much to ask, since individual households may have demand patterns that make such alignment difficult. For example, if work commitments mean that a particular household must be empty during the daylight hours in weekdays, it is inevitable that much of that household's energy demand cannot be supplied by solar power. In this set of experiments, however, we adopt and explore the view that households can be easily persuaded to perturb their natural schedules for using appliances *within* the boundaries set by their own preferences. That is, if a household prefers to use its washing machine on Saturdays between 8am and noon, then that household would be amenable to any suggestion within those constraints (e.g. 8am—10am, or 10am—noon) which might emerge from a management process that attempts to optimise energy usage for the village as a whole. The first set of experiments explores this notion by discovering, via optimization, what level of energy savings may be possible in realistic scenarios in which householders shift their schedules within their own constraints and preferences. Notice that this is far from straightforward or predictable: we cannot predict, for example, that everyone should shift their usage towards the sunniest part of the day – this would simply lead to too much demand during the sunniest hours, and under-utilisation (wasted energy or costly storage) of the village microgrid at other times.

### 5.1.1 Optimizing energy balance – a more precise statement

The optimization problem for the first set of experiments can be described slightly more formally as follows: given, for each household in the ecovillage:

- P: a list of residents' preferred times and durations for using specific electrical appliances in their household;
- C: a list of constraints on times and durations of appliance usage that define what is acceptable or possible for the residents and what is not;

find a daily scheduling of consumption events, within the constraints for each household, which minimizes the community's overall need for externally supplied power. In particular, what is optimised are each household's *start times and durations* for use of their electrical appliances, but ensuring that these times remain within the fixed constraints for that household. In this way a daily schedule is optimised for each household.

More formally, $P = \{P_1, P_2, ..., P_{20}\}$ is the set of preferences for each of the 20 households in the Riccarton Ecovillage, where $P_i$ is the preference list for household *i*. $P_i$ itself is composed of a set of consumption events; that is, $P_i = \{e_1, e_2, \ldots, e_m\}$, in which the consumption event $e_i$ expresses the household's preferred *timeslot* and *duration* for using appliance *i*. In detail, $e_i$ comprises four parts:

- Start time: the preferred timeslot at which the consumption starts
- Duration: the preferred number of slots the consumption lasts
- Probability: the probability that the consumption event will occur. This is used for non-daily used appliances, to overcome the daily structure of the preference list.
- Appliance details: the characteristics of the appliance; wattage, usage pattern (continuous vs. intermittent) and usage rate.

Meanwhile, whereas the preference lists outline specific start times and durations for each given household, the constraint list, $C = \{C_1, C_2, ..., C_{20}\}$, outlines a set of constraints $C_i$ for each household $i$, indicating how much that household is prepared to operate outside its current habits. A specific household's constraint list indicates the limits within which a consumption event can vary. Each constraint is attached to a consumption event, but not all consumption events necessarily have constraints. There are two types of constraints, Start time constraints and Duration constraints. A Start time constraint defines the maximum accepted value for the attached consumption event's Start time field, and the Duration constraint defines the minimum accepted value for the Duration field.

### 5.1.2 Estimating the 'Performance' of the simulation

As discussed, optimisation requires repeatedly running the simulation with different successive configurations, where the choice of next configuration is guided by the quality of previous configurations (essentially, most stochastic optimisation methods work by trying out new configurations that are close to the better-performing previously tested configurations). The quality function is rather crucial to this process. As should be clear, a 'configuration' in this set of experiments amounts to a particular daily schedule of consumption events (within the constraints) for each household. The quality function, as noted, is simply a measure of the amount of external energy consumption that is recorded in the simulation, when run with a given configuration.

In some more detail, the quality function of a configuration is calculated in the following way. Given a particular configuration to evaluate, the ABM model runs in time steps of one (simulated) hour. In each such hour, energy consumption is calculated for each household (according to the schedule of consumption of events in that household, for that given configuration), and the available energy supply from the microgrid is calculated for that household (its portion of the supply emerging from the wind turbine, plus that emerging in that hour from the household's solar panel). Consequently, for each household, and for each hour of the simulation, we have a difference value – the difference between microgrid supply available to a household, and the electricity consumption in that household. Each of these difference values is squared, and then they are all summed to arrive at an overall quality value for the configuration. The aim is then to *minimise* this measure.

It is worth considering this issue with care: in both sets of experiments, the quality measure is defined so that we obtain ideal quality by matching the curve of consumption levels over time as closely as possible to the curve of microgrid supply levels over time. Differences – either positive (excess demand, so the national grid must be used to plug the gap), or negative (excess supply, wasted energy generation in the village) are penalised. This reflects a suitable target for ecovillages that do not have viable means to store, or otherwise capitalise on, excess energy, and was a suitable experimental design issue. However, we remark that it is trivially simple to explore alternative quality functions for alternative scenarios in which, for example, REALL was able to usefully store a limited amount of its supply, perhaps at a certain cost, which could itself be incorporated into the quality function.

### 5.1.3 Notes on Statistical Confidence

Recalling our discussion in section 3.2, since the ABM simulation is stochastic it is clear that it would not be sensible to evaluate the quality of a configuration on the basis of a single run of the simulation with that configuration. Ideally we would take as our quality evaluation the mean result from several runs with the same configuration. However, simulations have an appreciable time cost, and we need to minimise the number of repeat runs, but at the same time attain a useful level of statistical confidence in the evaluation.

We address this matter in the following way. Whenever a configuration is to be evaluated, we run the simulation five times to obtain a mean quality value, and we also note the variance of this value over the five runs. If the variance is such that we cannot form 95% confidence in the result (on the basis of a T test), a further simulation is run, and the mean and variance re-calculated now over the six runs. This is repeated until we have obtained a mean quality value that has suitably high confidence, although we stop at ten iterations irrespective of confidence level.

## 5.2 Exploiting the Social Network

The second set of experiments was the same as the first set in all respects, but had the additional characteristic of a social interaction network and its exploitation. Ecovillages are communities in which we can naturally expect a high level of social interaction will feature, and we model this in a simple way by defining a network in which each household is a node, and a link between nodes represents a social link between the two households. The aim of this is to investigate to what degree the sociability of the community can lead to reduced energy demand. It is assumed that neighbours who are linked with a strong social tie are more likely to share or combine resources. For example, if household A is linked socially to household B, then it is more probable that it would, for instance, invite residents of household B for dinner. Household B would then have less demand for cooking, lighting, TV and heating for that night. If similar types of interaction are actively and regularly carried out in the community then considerable savings might be achieved. The second set of experiments investigates this hypothesis, and also explores the effect of modifying the density of the social network.

The density of the social network is characterized by $m$, which is related to the maximum number of links that an individual household can have. Each household is assigned randomly, at the start of the simulation, a random number of links (between 0 and $m$) to other households. Notice that the special case of $m$=0 corresponds to the first set of experiments. During a simulation, sharing worked as follows. Each simulation day, a household can choose to share, or not to share (with probability 0.5), its resources. If the decision is to share, then one linked neighbour on the social network is chosen randomly, and sharing takes place between the original household (the inviter) and the linked household. The sharing is then manifest as a (potentially only small) increase in the inviter's consumption for the shared appliances only, and zero consumption for the invitee household for the same appliances. The increase in consumption of the inviter is a randomly chosen multiplier between 1 and 1.5; this range reflects the cases of devices that do not cause extra-consumption when shared (e.g. TV) and others which may be increased (lighting, cooking, etc.). It is worth noting that a boundless variety of sharing 'models' could be implemented; more realistic approaches, for example, would bias the sharing events according to strength values on the links in the social network, and would be informed by available survey data. However the approach described was deemed sufficient for this experiment.

.

Finally, while in the first set of experiments we compared the three exemplar black-box optimisation methods (HC, SA, and EA), we used only the best of these (according to results on he first set of experiments) for the social network experiments – this turned out to be the EA.

## 6. Simulation and Optimization Results

### 6.1 Optimizing Energy Balance

Each of HC, SA and EA were tested on for a variety of algorithm parameter configurations. We do not reproduce a full set of results here, but this is available by request. In this section we briefly summarise the relative performance of the three optimization methods, and focus more on the physical interpretation of the ABM simulation results.

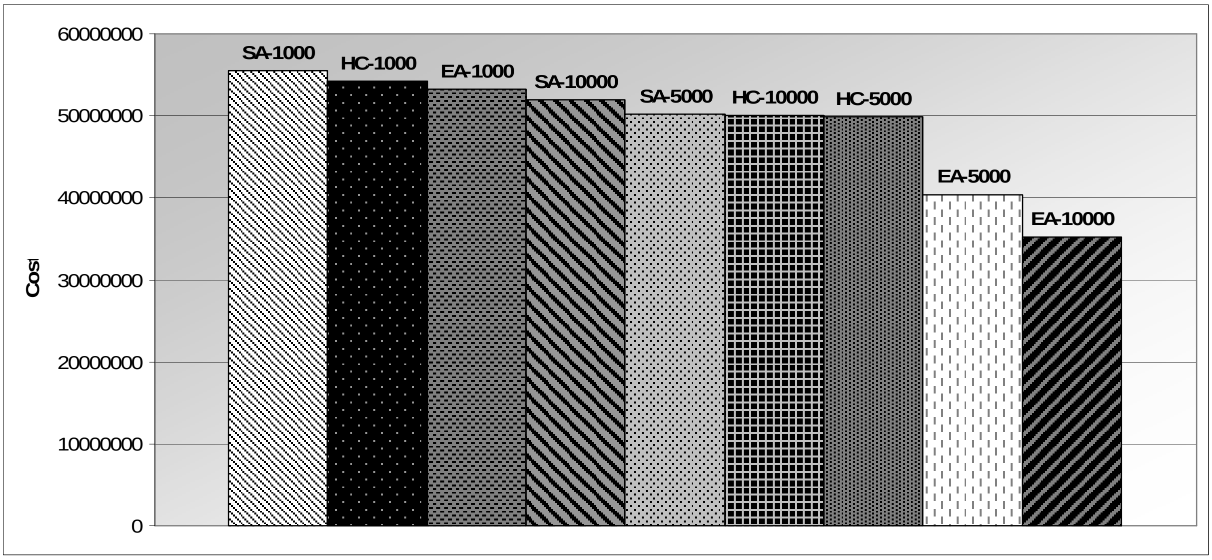


Figure 2: Algorithm performance comparison for different iteration counts

Figure 2 summarises the performance of HC, SA, and EA for this optimisation task, after many experiments with varying parameterisations of the algorithms. In this figure, lower is better, but the numbers on the vertical axis can be ignored, to be considered as arbitrary units for current purposes, but is roughly proportional to squaring the amount of energy required from the national supply. A relatively low value indicates a configuration of household schedules that, while each operating within their own comfort zones, amounted overall to an energy consumption curve that was relatively well aligned with the availability of energy from the village's microgrid. The label at the top of each bar indicates the algorithm, and the number of iterations for which it was run (which can be taken as directly proportional to the time taken). In this way the figure also includes a broad summary of the speed/quality tradeoffs for the different optimisation methods on this problem. Interestingly (but not unusually) the EA is the best quality algorithm, whether at 1,000, 5,000 or 10,000 iterations. An 'iteration' corresponds to a single run of the ABM simulation on a single configuration. The EA also clearly benefits more than the others (in terms of its ability to find better configurations) when given extra time. The superior solution quality found by the EA is particularly significant when experiments continue for 10,000 iterations, so the EA was chosen as the sole optimiser to use in the next set of experiments, and is also the method that led to the solutions we interpret below.

Figures 3 and 4 show visualisations of the preference lists at steps 0 and 5000 respectively, from a 5,000 iteration run of the EA. That is, figure 3 shows a preference list that results from an unoptimised situation – in which households simply follow their natural preferences, without reference to any community-based planning, and figure 4 represents an ideal situation that might emerge by households attempting to align their electricity consumption in co-operation with and to the benefit of the community as a whole, but still within their own constraints of acceptability.

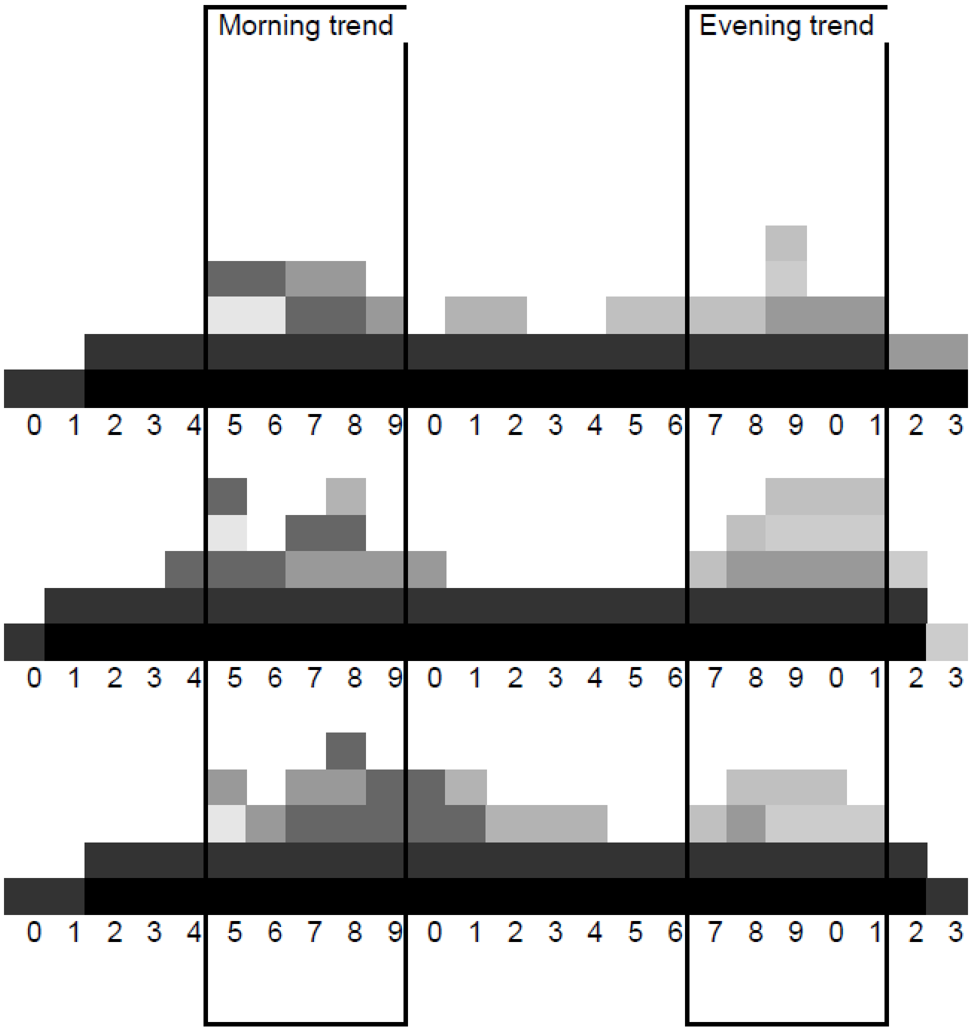


Figure 3: visualisation of the preference lists for three households – step 0

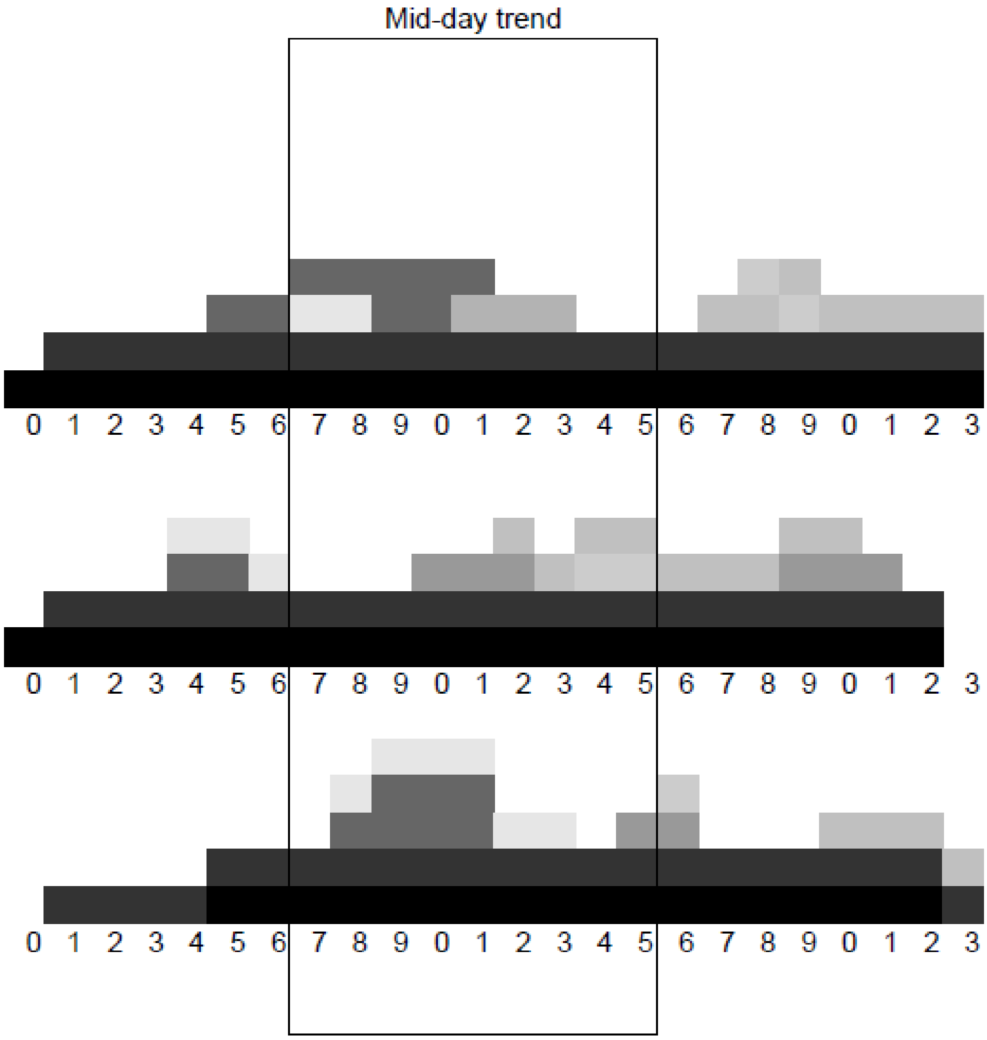


Figure 4: Visualisation of the preference lists for three households – step 5000

In figures 3 and 4, each colour denotes an appliance type. The horizontal axis holds the time of the day, while the vertical axis indicates consumption in three different households. The figure 3 case shows clear trends of consumption in hours of morning (breakfast time) and evening (dinner time). However, presumably influenced by the fact that the dominant source of energy is solar, and the availability of solar energy is concentrated around midday, the output of step 5000 shows more concentration of consumption towards the middle area, with less in the morning and the evening. Also, the optimizer enforces a kind of co-operative balance among different households consumption, taking into consideration their constraints, to achieve the least possible degrees of dissatisfaction.

A consideration of the energy savings available by optimising the preference lists in this way is given in the next section. Recall that this set of experiments is the ‘no friends’ specialisation of the set of experiments that investigate the exploitation of the social network.

## 6.2 Optimizing Energy Balance via Exploiting the Social Network

In this set of experiments, the same objective function was used, and separate optimization runs were done for different levels of social network density. Three values were tested for *m* (the maximum number of links per household), reflecting different levels of connectivity, in addition to the null case where no social network is incorporated (corresponding to the experiments summarised in section 6.1). Summarised results, following the use of the EA as the optimiser, running for 5,000 iterations per experiment, are shown in Figure 5. Again, in this figure the vertical ‘cost’ axis is in arbitrary units, but can be considered directly related to squared excess reliance on the national supply, rather than the village’s microgrid.

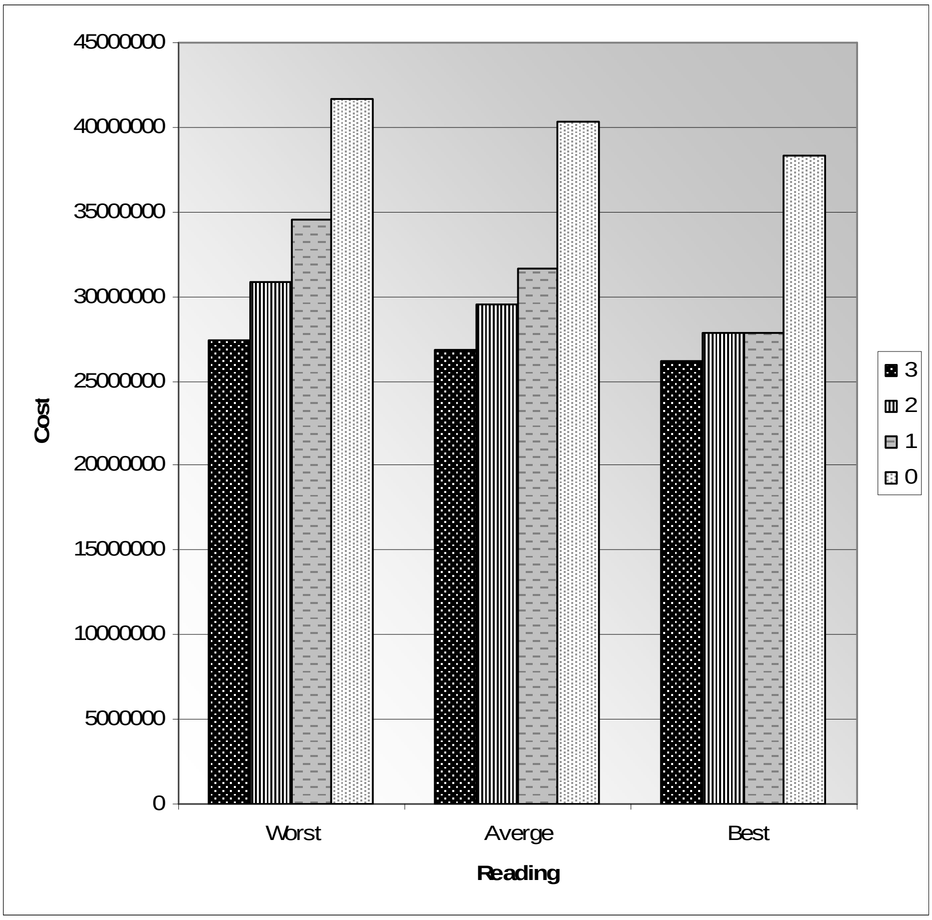


Figure 5: Comparison of cost for different levels of social network density

Figure 5 summarises the results of several optimization runs for each social network density, and record the best, worst and mean values that emerged from those runs. It is quite clear that the introduction of the social network leads to significant improvement in the cost values.

Recalling Figure 2, we note that the results achieved by the 5,000-iteration EA were in the region of 40M cost units. This already represents a considerable improvement over configurations where no optimisation has been done (i.e. each household simply operates unchanged according to their initial preferences), in which the cost is around 75M units (Hawasly, 2009). In figure 5, if we simply consider the average results, the exploitation of a low density social network already leads to a sharp improvement (33M cost units), while networks with density level 3 (certainly not unreasonable for an eco-village community) achieve around 28M cost units.

When we correctly interpret these results in terms of direct amounts of electricity, rather than the units of the fitness function – or, equivalently, if we interpret these results correctly in terms of potential financial savings), we find that the configurations found by optimisation using the density 3 social network represent 40% savings over an unoptimised, `un-managed' community.

If we consider 5,000-iteration optimisation only, we note that simply optimising household schedules without exploiting the social network (the first set of experiments) can lead to ~25% savings, and this increases to 40% with the most dense social network studied (in which each household is friends with roughly 15% of its small community). However it is worth pointing out that longer iteration runs (the 10,000 iteration EA) were able to find schedules with ~32% savings over the unoptimised case, without exploiting the social network. For pragmatic reasons, such longer runs were not done for the 'exploiting social networks' experiments in this paper, but the implication is that better than 40% savings could well be achievable.

# 7. Conclusions

## 7.1 A Summary of the model, experiments and results

We have described a relatively simple agent-based model (ABM) of the planned Riccarton ecovillage, in which households' preferences for the their daily use of electrical appliances were simulated, in tandem with realistic time-varying availabilities of solar and wind energy. Via optimisation, we explored the potential benefits that could be gained in such a community if it adopted a simple and acceptable community management strategy, in which, informed by evidence from simulations, households would be individually requested to lean towards particular times for their regular energy consumption activities. Further, we also explored the potential accrued benefits available if the community exploited its social network by regular resource sharing. Our findings suggest that quite significant savings are achievable via such measures. Simply attempting to establish a mutually acceptable, and mutually beneficial schedule of consumption activities, with no assumed resource sharing, can lead to 25% savings in cost (reliance on the national supply), while exploitation of the social network seems capable of raising that to beyond 40%. Meanwhile it is entirely conceivable that more optimisation effort could find solutions with appreciably more savings; such requires computational expense, but by no means undue such expense.

We find the implications of this, for the design and management of such communities, to be both intriguing and exciting. As discussed, simulation models such as the REALL model allow 'what if' explorations, but the impact of these explorations can be rather limited without some way to discover how to organise or manage the community to achieve a specific goal. Via optimisation, however, we can automatically explore the space of strategies, and find ones that achieve our targets. As a tool for managing communities, this general approach has been very little explored to date, but the applicability (in the face of no sensible alternatives) and the potential of such an approach seem clear. Naturally, such models can incorporate considerably more complexity than we have shown here; for example, in one thread of continuing research we plan to enhance the REALL model to reflect more real-world complexity by switching from modelling households into modelling people with intentions and beliefs, incorporating more appliances, investigating real households' consumption habits and adaptability, better representing time, and the development of a complex environment where the agents would interact. Such enhancements, among others, allow for a greater range of experiments, and a consequently greater range of implications for design and management strategies for the REALL community. It must be noted of course that such enhancements also carry a computational burden that tends to constrain the amount of optimisation that could be done, however, the ever-increasing availability of affordable high-performance computing resources comes to our rescue in this respect. In particular, if we view ABM-based models as essential management tools for such communities, the wise use of which can promise substantial benefits beyond shorter sighted approachers, then the requirement of reasonably costly computational resources becomes acceptable.

### 7.2 Discussion

The growing need to build resilience in our societies at a time of rapid social, economic and environmental change with better choices for, and management of the built environment and the human capacity to deal with change is clear

This paper has shown that new Agent Based complex models are now capable of informing choices of technological solutions with the impacts of attitudes, values and behaviours to optimize the energy and environmental benefits of decisions and minimize their costs. It offers the best method to date of ensuring that the best choice is made within a wide range of available adaptive opportunities derived from an extremely complex range of iterations produced from relationships between interacting forces.

Systems such as the REALL model offer a tool to introduce credible 'fresh thinking' and one that has significant potential to engage and educate the public, raise environmental awareness and personal 'buy-in' to the implementation of 'socially viable' solutions, in the transformed markets. It also may answer a wide range of questions such as:

- A question arises of what level are such solutions best fostered at?
- As a result of the outputs of such simulations, would ensuing local community based changes be best led by governments, dedicated to stimulating economic growth?
- Will local communities find it easier to agree to lower standards of living for their own futures if they can clearly see the clear, equitable, benefits using such models?
- Can such models be used to avoid 'collapse' events being the only effective driving force for markets?
- Is it inevitable that the only sustainable future is a simpler one that will work adequately in the approaching eras of scarcity, of growing climate impacts and increasingly unaffordable capitalism?

- Are there new paradigms of ‘smart development’ where man and machines can work side by side to automate, motivate and decarbonise economies while maintaining standards of living and qualities of life?
- People have been shown to be motivated by environmental concerns while also holding strong material concerns which motivate them to purchase new products and increase their environmental impact. Can the use of tools such as this to “trade-off” gains and losses clearly demonstrate the effectiveness of solutions to promote goals of long-term sustainability that would appeal to, and be adopted by, local residents?
- Are imposed step changes imposed by society and their politicians, or periodic system collapses, the only way to affect the stringent cuts required by emerging greenhouse gas reduction targets?
- To what extent are people’s expectations of quality of life and comfort amenable to modification with information and “ethical persuasion”?

These tools can be used interactively with real populations to test the water on a case by case basis, to explore what are their core values and what they are willing to sacrifice in the bid to build social and economic resilience in the face of rapid change. Pathways to deep emissions reductions in commercial buildings require not only increased consideration of passive and low energy architectural principles and renewable energy technologies but also the buy-in of populations in doing so.

Substantial building sector energy and emissions reductions are available through advanced technologies and controls but even in the very high tech buildings systems are only capable of reaching their full potential when coupled to behavioural techniques to manage the expectations and perceptions of building users.

There appears to be a pressing need to develop new ways of engaging society to participate in the necessary changes ahead. It is clearly necessary to have a clear and scientifically valid focus on the targets, in time, ahead and the tools to hand to practically achieve them, but radical transformations of the type necessary to meet global reduction targets are unlikely to progress without the engagement of the key stakeholders involved in the built environment in making the decisions to change. Models such as this may well provide a new language to the dialogue that engagement may require as well as ensuring that the best possible solutions are adopted.

**References**

Baker, T.J.A., Botting, M., Berryman, M.J., Ryan, A., Grisogon, A.M. and Abbott, D. (2005), 'Adaptive Battle Agents: Emergence in Artificial Life Combat Models', Proceedings of the SPIE, Smart Structures, Devices, and Systems II, Bellingham, WA, pp. 574-585.

Boccara, N. (2004), Modeling complex systems, Springer.

Bossomaier, T., Amri, S. and Thompson, J. (2007), 'Agent-Based Modelling of House Price Evolution', Proceedings of the 2007 IEEE Symposium on Artificial Life (CI-ALife 2007), Honolulu, HI, pp. 463 - 467.

Corne, D, Frisco, P. (2008) Dynamics of HIV infection studied with cellular automata and conformon-P systems, Biosystems 91(3): 531-544.

dos Santos, R., Coutinho, S. (2001) Dynamics of HIV Infection: A Cellular Automata Approach, *Phys. Rev. Lett.*, **87**(16): 168102-1, 4pp.

Edmonds, B. (2005) 'Simulation and Complexity - how they can re- late'. In Feldmann, V. & Mhlfeld, K. (Eds.) Virtual Worlds of Pre- cision - computer-based simulations in the sciences and social sciences. Lit Verlag, 5-32.

Fu, M.C. (2002), 'Optimization for Simulation: Theory vs. Practice'.
INFORMS, Vol. 14, No. 3, Summer 2002 pp. 192215.

Fu, M.C., Glover, F.W. & April, J. (2005), 'Simulation Optimiza- tion: A Review, New Developments, and Applications'. Proceedings of the 2005 Winter Simulation Conference.

Garcia, J. S. D. , Avila, S. L. & Carpe, W. P. (2006), 'Intro- duction to Optimization Methods: A Brief Survey of Methods'. IEEE Multidisciplinary Engineering Education Magazine, volume 1, issue 2.

Gilman, R. (1991), 'The Eco-village Challenge', IN CONTEXT, 29, Page 10.
(http://www.context.org/ICLIB/IC29/Gilman1.htm)

Gimblett, H. R. (2002), 'Integrating Geographic Information Systems and Agent-Based Technologies for Modeling and Simulating Social and Ecological Phenomena', in Gimblett, H. R. (Ed.) Integrating Geo- graphic Information Systems and Agent-based Techniques for Simulating Social and Ecological Processes, Oxford University Press, Oxford.

Gray, P., Hart, W., Painton, L., Phillips, C., Trahan M. & Wagner, J. (1997), 'A Survey of Global Optimization Methods', [online], Sandia National Laboratories. viewed 15 July 2009. (http://www.cs.sandia.gov/opt/survey/)

Griva, I. ,Nash, S.G. & Sofer, A. (2009), Linear and Nonlinear
Optimization, 2nd Edition. SIAM, PA.

Hartmann, S. (2005), 'The World as a Process: Simulations in the Natural and Social Sciences', PhilSci Archive, archive.pitt.edu/archive/00002412/).

Hawasly, M. (2009) Optimization of Complex Systems via Simulation**,** MSc Dissertation, Heriot-Watt University, August 2009.

Energy Academy of Heriot-Watt University, Website, viewed 28 July 2009, (http://www.energy.hw.ac.uk/projects/ecovillage.cfm).

Hood, L. (1998), 'Agent based modelling', Proceedings of Greenhouse Beyond Kyoto Conference, Bureau of Resource Sciences.

John Gilbert Architects, 'Riccarton Ecovillage - Feasibility Study'.

Kasper, D.V.S. (2008), 'Redefining Community in the Ecovillage', Research in Human Ecology, Vol. 15, No. 1, 2008. (http://www.humanecologyreview.org/pastissues/her151/kasper.pdf )

Laguna, M. (1997), 'Optimization of Complex Systems with OptQuest', OptQuest for Crystal Ball User Manual, faculty.colorado.edu/laguna/articles/optquest.pdf )

López-Sánchez, M., Noria, X., Rodríguez, J.A. and Gilbert, N. (2005), 'Multi-Agent Based Simulation of News Digital Markets', International Journal of Computer Science & Applications 2(1): 7 – 14.

Olafsson, S. & Kim, J. (2002), 'Simulation Optimization'. Proceedings of the 2002 Winter Simulation Conference.

Railsback, S.F., Lytinen, S.L., & Jackson, S.K. (2006), 'Agent- based Simulation Platforms: Review and Development Recommendations'. Simulation, 82, 9 (Sep. 2006), 609-623.

REPAST, Website. viewed FEB 2009.(http://repast.sourceforge.net/).

Roaf, S., D. Crichton and F. .Nicol (2009). Adapting buildings and cities for climate change, Architectural Press, Oxford.

Roaf, S. (2009),'Riccarton Eco-Village: Towards Zero Carbon Housing'.

Shalizi, C.R.(2006). Complex Systems Science in Biomedicine, Springer.

Stephan, C. and Sullivan, J. (2004), 'Growth of a Hydrogen Transportation Infrastructure', in Macal, C.M., Sallach, D. and North, M.J. (eds.), Proceedings of the Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence, Chicago, IL, pp. 731-742, Available at http://www.agent2005.anl.gov/Agent2004.pdf

Stern, H. (2008) The accuracy of weather forecasts for Melbourne, Australia, Meteorolgical

Applications, 15(1), 65-71.

Takama, T. (2005), Stochastic Agent-Based Modelling for Reality: Dynamic Discrete Choice Analysis with Interaction, University of Oxford, Oxford, UK, Available at http://www.tri.napier.ac.uk/Events/TDM/prestonpaper.pdf.

Weise, T. (2009), Global Optimization Algorithms - Theory and Application, 2nd Edition. E-book, (http://www.it-weise.de/).

Wikipedia (2009), 'Complex System', wiki article, viewed 11 July 2009, (http://en.wikipedia.org/wiki/Complex system).

Wikipedia (2009), 'Ecovillage',wiki article, viewed July 2009, (http://en.wikipedia.org/wiki/Ecovillage).

Wikipedia (2009), 'Optimization (mathemat- ics)', wiki article, viewed July 09, (http://en.wikipedia.org/wiki/Optimization (mathematics)).

Wikipedia (2009), 'System',wiki article, viewed July 09, (http://en.wikipedia.org/wiki/System).

Zambonelli, F., Jennings, N. R., & Wooldridge, M. (2003), 'Developing multiagent systems: The Gaia methodology'. ACM Trans. Softw. Eng. Methodol. 12, 3.