Preprint for: McCormack, J. (2003) Evolving Sonic Ecosystems, *Kybernetes* (to appear)

Evolving Sonic Ecosystems

Jon McCormack

Centre for Electronic Media Art
School of Computer Science and Software Engineering
Monash University, Wellington Road
Clayton, Victoria, 3800 Australia

`jonmc@csse.monash.edu.au`

Keywords *Artificial Life, Genetic Algorithms, Computer Music Composition, Human-computer interaction, Generative Art.*

Abstract

This paper describes a novel type of artistic Artificial Life environment. Evolving agents, who have the ability to make and listen to sound, populate a synthetic world. An evolvable, rule-based system drives agent behaviour. Agents compete for limited resources in a virtual environment that is influenced by the presence and movement of the artwork's audience. Through a link between the real and virtual spaces, virtual agents evolve implicitly to try to maintain the interest of the human audience.

1. Introduction

> *»One thing that foreigners, computers and poets have in common is that they make unexpected linguistic associations.«*
> Jasia Reichardt (Reichardt, 1971b).

Music and art are undoubtedly fundamental qualities that help define the human condition. While many contrasting and parallel discourses contribute to our understanding of art theory — and hence our interpretation of art itself — there are two implicit themes that connect all artworks. The first is the act of creation. Even the most abstract or conceptual artworks cannot escape the fact that, as ideas, objects or configurations, they must be *made.* Secondly, the importance of novelty, either perceived or real, is a fundamental driving force behind any creative impetus or gesture.

Artificial Life (AL) methodologies can play an important role in developing new modes of artistic enquiry and musical composition. For artists, AL offers significantly different contributions than those traditionally provided for the creative arts. For the first time in the history of art, AL offers, in theory at least, that it may be possible to create artificial organisms that develop their own autonomous creative practices — to paraphrase the terminology of Langton (Langton, 1989), *life-as-it-could-be* creating *art-as-it-could-be.*

In addition, AL has important contributions to make in our understanding of genuine novelty[1], often referred to under the generalized term *emergence* (Cariani, 1991; McCormack and Dorin, 2001; Emmeche et al., 1997).

---

[1] The concept of novelty is a vexed one with many different interpretations in the literature, and could easily occupy an entire paper in itself. Some authors argue that novelty and emergence have no relation (Nagel, 1961), whereas others see them as fundamentally the same (Cariani, 1991). In the sense the term is used in this paper, novelty suggests that which has never existed before, hence the issues sur-

1.1. Artificial Life Art

Techniques from Cybernetics and Artificial Life have found numerous applications in the creative arts. General contemporary overviews can be found in, for example (Bentley, 1999; Bentley and Corne, 2002; Sommerer and Mignonneau, 1998; Wilson, 2002).

Cybernetics has a rich and often overlooked history in terms of computing and the arts. The seminal ICA exhibition *Cybernetic Serendipity*, held in London in the summer of 1968 was one of the first major exhibitions to look at connections between creativity and technology (Reichardt, 1971a). Even the title suggests notions of novelty and discovery, a key theme for many works and critics in the decades that have followed the exhibition. Interestingly, the exhibition's curators shunned distinctions between artists and scientists and instead focused on ideas and methodologies that connected the two.

One particularly relevant concept from Cybernetics is that of *open-ended behaviour*, what Ashby referred to as *Descartes Dictum:* how can a designer build a device which outperforms the designers specifications (Ashby, 1956). Cyberneticist, Gordon Pask, built an 'ear' that developed, not through direct design, but by establishing certain electro-chemical processes whereby the ear formed and developed in response to external stimuli (Cariani, 1993).

The goal of the work described in this paper is an AL artistic system, titled *Eden*, which tries to capture (in spirit at least) the concepts of an open-ended system that is *reactive* to its environment. In order to address this issue, two important questions were asked during the design and development of the work. Firstly, how a designer can create a virtual AL world that evolves towards some subjective criteria of the audience experiencing it, without them needing to explicitly perform fitness selection. Secondly, how the relationship between real and virtual spaces is acknowledged in a way that integrates those spaces phenomenologically.

In terms of software, *Eden* is an AL environment that operates over a cellular lattice, inhabited by agents who have, amongst other capabilities, the ability to make and 'listen' to sound. Agents use an internal, evolvable, rule-based system to control their behaviour in the world. The virtual environment that the agents inhabit develops in response to the presence and movement of people experiencing the system as an artwork. This software system will be described more fully in sections 2 and 3 of this paper. Interaction with the work is detailed in section 4, with a summary of results and brief conclusion in sections 5 and 6.

1.2. Related Work

In terms of the specific software system described in this paper, much technical inspiration and methodology is drawn from John Holland's *Echo* (Holland, 1995), particularly in the use of rule-based methods for the internal decision-making system of agents. Many others have used evolutionary systems as a basis for musical composition, but in the main for compositional *simulation* (Wiggins et al., 1999; Todd and Werner, 1998), rather than as a new form of creative tool for the artist and audience.

The *Living Melodies* system (Dahlstedt and Nordahl, 1999) uses a genetic programming framework to evolve an ecosystem of musical creatures that communicate using sound. *Living Melodies* assumes that all agents have an innate 'listening pleasure' that encourages them to make noise to increase their survival chances. The system described in this paper, *Eden,* contains no such inducement, beyond the fact that some sonic communication strategies that creatures discover should offer a survival or mating advantage. This results in the observation that only some instances of evolution in *Eden* result in the use of sonic communication,

---

rounding novelty are connected with determinism (Emmeche et al., 1997). For art, almost every new artwork is in some sense novel, however we may at least be able to apply criteria that suggest a degree of novelty, such as descriptive causality and explainable causality. Moreover, in an AL sense, we require not only the artwork to be novel, but the behaviour of the virtual agents to be novel as well.

whereas in *Living Melodies*, *every* instance evolves sonic communication. Living Melodies restricts its focus to music composition, whereas *Eden* is both a sonic and visual experience.

2. *Eden*: an Artificial Life Artwork

Eden is a 'reactive' artificial life artwork developed by the author. The artwork is typically experienced in an art gallery setting, but in contrast to more traditional artworks, is designed as an *experiential* environment, whereby viewers participation and activity within the physical space have important consequences over the development of the virtual environment.

The artwork is exhibited as an installation, and experienced by multiple users simultaneously. It consists of multiple screens, video projectors, audio speakers, infrared distance sensors, computers, and custom designed electronic systems. Figure 1 shows a floor plan and simulated visualization of the work. As shown in this figure, physically the work consists of two semi-transparent screens suspended from the ceiling of the exhibition space. The screens are positioned at 90º to each other, forming an 'X' shape when viewed in plan. The ambient lighting is minimal; making the screens and the light they reflect and transmit the predominant source of visual interest in the space. The screens' transparency enables them to be viewed from either side. Multi-channel audio is provided by a number of speakers placed on the periphery of the main screen area.

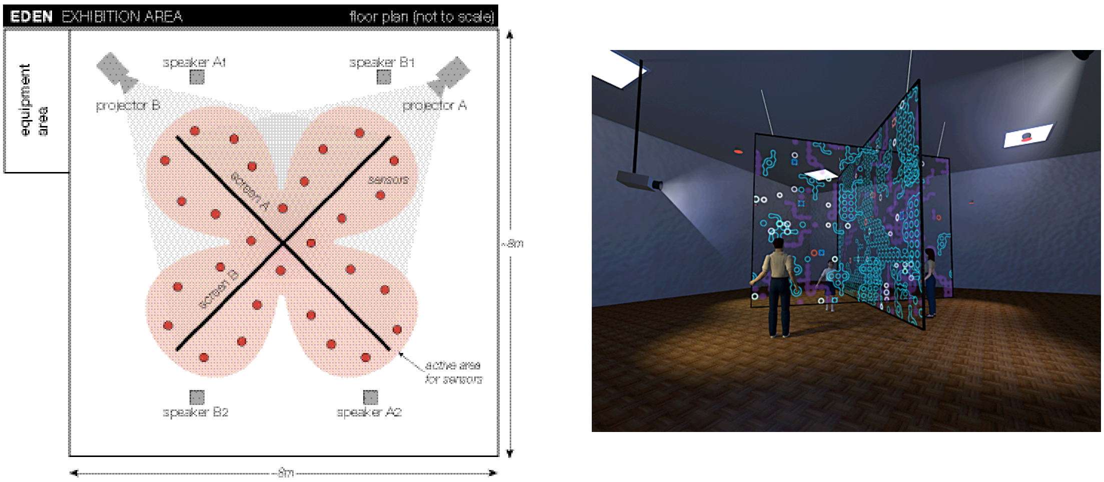

Figure 1: Floor plan of *Eden* (left) showing the layout screens, speakers, projectors and sensors. The active sensor area is shown in red. The image on the right is a simulation of the work running in a typical gallery environment, illustrating the effect of using transparent screens to visualize the work.

In addition to this audio-visual infrastructure, a series of infrared range sensors are placed around the screen area. The purpose of these sensors is to measure the position and movement of people experiencing the work. The sensors themselves are not visible when viewing the work. They function as an environmental stimulus for the virtual agents' world and ultimately contribute to selective pressures that aim to encourage a symbiotic relationship between people experiencing the work and the agents populating the virtual world. The role of the sensors and their effect on the development of the virtual environment portrayed in the work is detailed in section 4.

3. Agents and Environments

This section gives technical details on the major software components of the system, with particular emphasis on the mechanisms that facilitate development of sonic agents within the system. Further details, particularly the *payoff* and *bidding* processes for rule selection may be found in (McCormack, 2001).

3.1. The Eden World

The environment projected onto the screens is known as the Eden *world*. In implementation terms, the world consists of a two-dimensional cellular lattice that develops using a global, discrete, time-step model — a popular AL model based on the theory of cellular automata (Ulam, 1952; Conrad and Pattee, 1970). Each cell in the lattice may be populated by one of the following entities:

- *Rock:* inert matter that is impervious to other entities and opaque to sound and light. Rock is placed in cells at initialisation time using a variation of the *diffusion limited aggregation* (DLA) model (Witten and Sander, 1981). Rocks provide refuge and contribute to more interesting spatial environmental behaviour of the agents.
- *Bio-mass:* a food source for evolving entities in the world. Bio-mass grows in yearly[2] cycles based on a simple feedback model, similar to that of *Daisyworld* (Watson and Lovelock, 1983). Radiant energy (in 'infinite' supply) drives the growth of bio-mass. The amount of radiant energy falling on a particular cell is dependent on a number of factors, including the local absorption rate of the bio-mass and global seasonal variation. Probabilistic parameters can be specified at initialisation time to control these rates and variations. The efficiency at which bio-mass converts radiant energy into more bio-mass is also dependent on the presence of people in the real space of the artwork. This dependency is detailed in section 4.
- *Sonic Agents:* mobile agents with an internal, evolvable *performance system.* Agents get energy by eating bio-mass or by killing and eating other agents. More than one agent may occupy a single cell. Since these agents are the most complex and interesting entity in the world, they are described in greater detail in section 3.2.

A real-time visualization of the world produces images that are projected onto the screens, as illustrated in Figure 1 (in this case there are two worlds, each running on a separate computer, but connected as a single logical world running over two computers). The visualization process is described more fully in section 3.3.

3.2. Agent Implementation

Sonic agents are the principle evolving entity in the world. Essentially, the agent system is similar to that of Holland's *Echo* system (Holland, 1995). An agent consists of a set of *sensors*, a rule-based *performance system* and a set of *actuators.* This configuration is illustrated in Figure 2. Sensors provide measurement of the environment and internal introspection of an individual agent's status (detailed in section 3.2.1). The performance system relates input messages from the sensors to desired actions (detailed in section 3.2.3). The actuators are used to show intent to carry out actions in the world. The success or failure of an intended action will be dependent on the physical constraints in operation at the time and place the intent is instigated. Actuators and actions are detailed in section 3.2.2.

---

[2] An *Eden* year lasts 600 *Eden* days, but passes by in about 10 minutes of real-time.

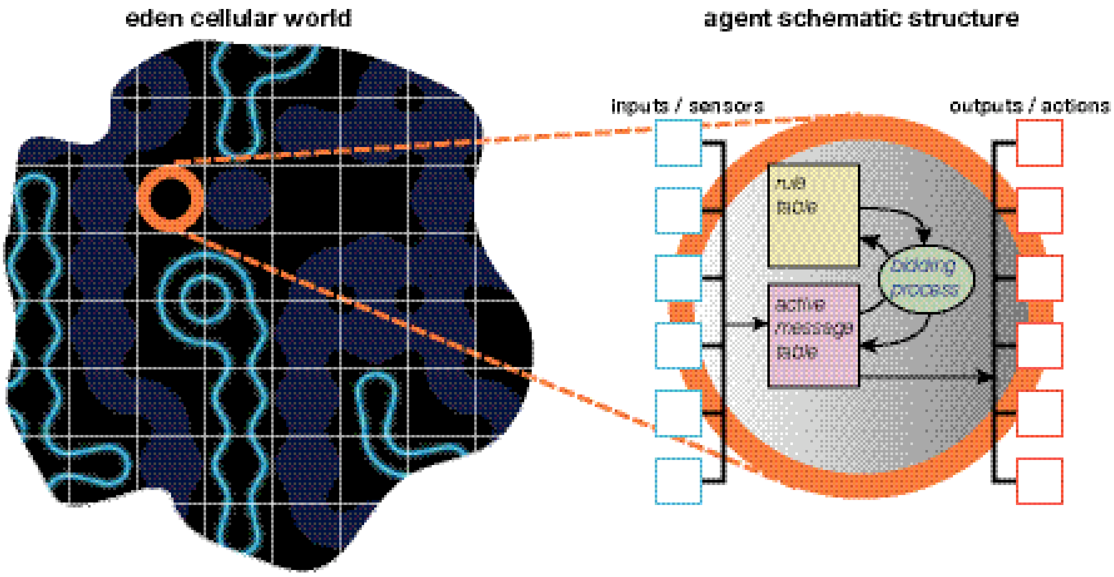

Figure 2: Shows a section of the *Eden* cellular lattice in visual form (left). To emphasise the lattice structure, grid lines have been layered over the image. The image shows rocks (purple), biomass (light blue) and an agent (orange). The diagram (right) shows the agent's internal schematic structure, consisting of a number of sensors, a performance system that evolves, and a set of actuators.

At initialisation of the world, a number of agents are seeded into the population. Each agent maintains a collection of internal data. This data includes:

- *Current age*, an integer measured in time-steps since birth. Agents live up to 100 years and cannot mate in their first year of life.
- *Health index:* an integer value indicating the overall health of the agent. A value of 100 indicates perfect health, if the health index falls to 0, the agent dies. An agent can loose health via a sustained negative *energy level* differential (explained below); by bumping into solid objects, such as rocks; or by being hit by other agents. In addition, the loss in health from being hit by another agent depends on both its mass and health index.
- *Energy level:* a measure of the amount of energy the agent currently has. Agent's gain energy by eating bio-mass or other agents. Energy is expended attempting to perform actions (regardless of their success); a small quantity of energy is expended even if no action is performed at a given time-step. If an agent's energy level falls to zero, the agent dies and its body is converted to bio-mass in a short number of time-steps.
- *Mass*: an agent's mass is linearly proportional to its energy level, plus an initial 'birth mass' that varies uniformly over the population.

3.2.1. Sensors

Sensors provide a way for an agent to *measure* itself and its environment (Pattee, 1988). Sensor data is presented as bit strings constructed from local environmental conditions and from the internal data structures held by the agent. Sensor data is updated once each time-step. An agent can use a range of sensor types, but sensors themselves do not undergo any evolutionary process and are therefore fixed in function, sensitivity, and morphology. It is up to an individual agent's performance system to make use of any given sensor, hence sensor data will only be used in the long term if it provides useful information that assists the agent's survival or mating prospects. Sensor use does not incur any cost to the agent.

Sensor information available to an agent includes:

- A simple local vision system that detects the ‘colour’ of objects on facing and neighbouring cells (the range is fixed at a single cell). Rocks, bio-mass, and agents all have different ‘colours’, thus enabling an agent to distinguish between them.
- A sensor to detect the local cell nutritional value. Cells that contain bio-mass or dead agents have a high nutritional value.
- A sound sensor that detects sound pressure levels over a range of frequency bands. Sound can be detected over a much larger spatial range than, for example, vision and also with greater fidelity.
- An introspection of *pain.* Pain corresponds to a negative health index differential and would usually indicate attack by another agent or that the agent is bumping into rocks.
- An introspection of the current energy level.

3.2.2. Actuators

Actuators are used to signal an agent’s intent to carry out an action in the world. The physical laws of the world will determine whether the intended action can be carried out or not. For example the agent may intend to ‘walk forward one cell’, but if that cell contains a rock, the action will not be possible. Furthermore, all actions cost energy, the amount dependent on the type of action and its context (e.g. attempting to walk into a rock will cost more energy than walking into an empty cell).

As with the sensors, the set of actuators is fixed and does not change as the performance system evolves. However, actions will ultimately only be used if they benefit the agent in the long term and analysis of actions used by agents shows not all agents make use of the full set of actuators.

Possible actions an agent may perform include:

- *Move* forward in the current direction;
- *Turn* left or right;
- *Hit* whatever else is in the cell occupied by the agent. Hitting another agent reduces that agent’s health level using a non-linear combination of the mass, health, and energy level of the agent performing the hit. Hitting other objects or being hit will cause pain and a loss of health.
- *Mate* with whatever is currently occupying the agent’s cell. Obviously only useful if another agent is in the same cell. In addition, mating is only possible if the age of both agents is greater than one year.
- *Eat* whatever is currently occupying the agent’s cell. Agent’s can only eat bio-mass, or dead agents (which turn into bio-mass shortly after death).
- *Sing*: make a sound that can be heard by other agents. Sound is detailed more fully in section 3.4.

Performing an action costs energy, so agents quickly learn not to perform certain actions without benefit. For example, attempting to eat when your nutritional sensor is not activated has a cost but no benefit. Attempting to move into a rock has a cost greater than moving into an empty cell.

Agents may also choose not to perform any action at a given time step (a ‘do nothing’ action), but even this costs energy (although less than any other action).

3.2.3. Performance System

The performance system connects an agent’s sensors to its actuators (Figure 2). It is based on the classification system of (Holland, 1995). Sensory data arrives from the sensors in the form of a *message,* a binary string of fixed length[3]. Messages are placed in an *active message table*, a first-in, first-out (FIFO) list of messages currently undergoing processing. Each agent maintains a collection of *rules*, stored in a database or *rule table.* Rules consist of three components: a *condition string*, an *output message* and a *credit*. Condition strings are composed from an alphabet of three possible symbols: {1,0,#}. At each time-step, the message at the head of the active message table is processed by checking for a match with the condition string of each rule in the rule table. A ‘1’ or ‘0’ in the condition string matches the corresponding value in the message at the same index. A ‘#’ matches either symbol (‘0’ or ‘1’). So for example, the message ‘10010111’ is matched by any of the condition strings ‘10010111’, ‘10010##1’, ‘########’. The condition string ‘#######0’, however, would not match.

Rules whose condition strings match the current message bid for their output message (also a bit string of the same length as sensor messages) to be placed in the active message table. This bid is achieved by calculating the rule’s *strength.* Strength is the product of the rule’s credit (detailed in a moment) and it’s *specificity.* Specificity is a unit normalized value, inversely proportional to the total number of ‘#’ symbols in the condition string. So for example, a condition string consisting entirely of ‘#’ symbols has a specificity of 0; a string with 75% ‘#’ symbols has a specificity of 0.25; and so on.

For each rule that matches the current message under consideration, it’s strength is calculated. The rule with the highest strength is selected and then places it’s output message into the active message table. If more than one rule has the highest strength, then a uniform random selection is made from these rules. The selected rule places its output message into the active message table. Most output messages are *action messages*[4], i.e. they trigger an actuator. Action messages are removed from the table once they have been translated into actuator instructions.

The process outlined in this section is illustrated below in Figure 3.

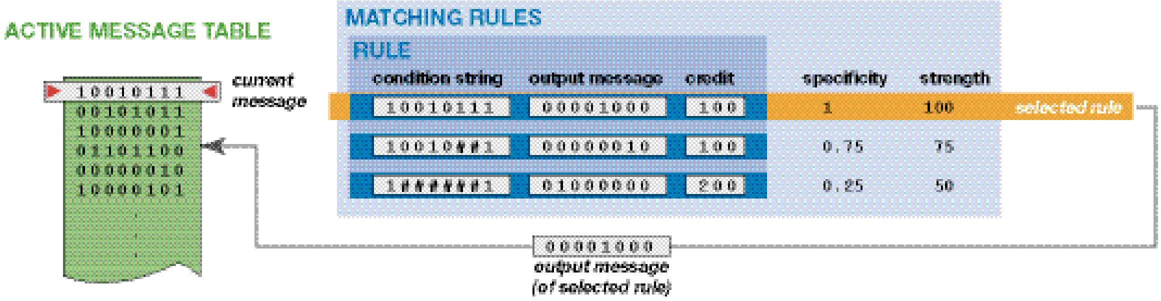

Figure 3: The rule matching and bidding process. The top message from the active message table is selected and becomes the *current message.* Rules whose condition string matches the current message have a strength calculated as the product of their credit and specificity. The rule with the highest strength then becomes the selected rule and it’s output message is added to the active message table. The current message is then discarded and the process repeats. Some messages are action messages and trigger actions.

3.2.3.1. Credits and Payoffs

Each rule maintains a credit, essentially a measure of how useful this particular rule has been in the past. Rules begin with a default credit value and earn or loose credit based on how

---

[3] Currently a message length of 32 bits is used, but the actual length does not concern the processes described. Larger message lengths allow more bandwidth in sensor messages, but require more storage.

[4] Action messages are distinguished from other messages by a marker bit in the string being set— all other message types are guaranteed not to set this bit.

useful the rule is in assisting the agent to live and mate in the world. As described in section 3.2, agents maintain an energy level and health index. The differentials of these quantities are monitored, and when they reach a certain threshold, a *credit payoff* is performed. The credit payoff rewards or punishes rules that have been used since the last payoff (held in a separate list), by altering their credit according to frequency of use and the magnitude of the change in energy since the last payoff. Further details regarding this process may be found in (McCormack, 2001).

The credit payoff system enables rules that, over time, assist in increasing health and energy to be rewarded; those that decrease health and energy will decrease in credit. The rational being, that next time rules have a chance to bid, if they have been useful in the past, they'll probably be useful in the current situation.

The number of time steps between successive payoffs will be dependent on how quickly or slowly the agent's health is changing. For example, if a creature is being attacked and losing health quickly, payoffs will be more frequent. The rules involved in letting the agent get hit will also decrease in credit quickly (hopefully soon being outbid by other rules that may prove more successful, if the agent is to survive).

Maintaining a list of rules that have been used since the previous payoff allows rules that indirectly increase health to receive appropriate credit. For example, while the rule to 'eat when you find food' is a good one, you may need to walk around and look for food first to find it. The rules for walking and turning, although they decrease health in the short term, may result in finding food. This increases health in the longer term. Overall, if such rules are helpful in increasing health, their credit will increase. A rule whose strength falls to zero, will be automatically removed from the agent's rule table, since it is on longer able to bid to be used.

As specified in section 3.2.3, a rule's strength is the product of its credit and specificity. This is necessary, since rules that are more specific will be used less often as they match fewer messages. More specific rules will have less chance to receive credit payoffs, but still may be useful. Hence when two or more rules with the same credit match a message, the more specific rule will have greater strength and thus will be selected over the more general one.

3.2.4. Agent Evolution

The credit payoff system allows rules that have contributed to the agent's survival to be used more often. However, this will only select the best rules from the currently available set. The problem remains as to how the agent can discover better rules than those it currently uses.

Genetic algorithms follow a Darwinian metaphor in that they operate as a search method over the phase space of possible phenotypes in a given system — searching over the *fitness landscape* for individuals of higher *fitness* (Goldberg, 1989; Mitchell, 1996). In the *Eden* system, a rule functions as the genetic unit of selection and new rules are brought into an agent's genome via the standard operations of *crossover* and *mutation* (See the references for explanations of these terms).

Recall from section 3.2.2 that mating is a possible action an agent can perform. If two agents successfully mate, they produce a new agent whose rule table is a combination of the parent's tables. A proportion of rules from each parent are selected, based on the strength of the rules — the rules of highest strength from each parent being selected. These selected rules undergo crossover and mutation operations, as per the schema system of Holland (Holland, 1992) resulting in the creation of new rules. Mutation rates vary according to the behaviour of people experiencing the artwork in the exhibition space. This is further detailed in section 4.

Since rules of highest strength are selected from each parent, the evolutionary process is Lamarckian (Bowler, 1992), in that the rules that have been most useful to an agent's survival in it's lifetime are selected. This design decision was used to allow more rapid adaptation to changing environmental conditions — a necessary feature if the agent's in the artificial ecosystem are to adapt to the behaviour of people experiencing the work.

3.3. Image

Representation of the entities of *Eden* is achieved using tiling patterns, loosely based on Islamic ornamental patterns (Grünbaum and Shephard, 1993). Only the representation of bio-mass will be considered here. The visual representation of the bio-mass is based on a 16 tile set. A tile is chosen for a particular cell based on the neighbour relationships of adjacent cells. For the purposes of tile selection, tiles are selected based on the binary occupancy condition of the cell's neighbours, considering only cells of the same type. For the 16 tile set, only immediate orthogonal cells are considered — thus there are 16 possible configurations of neighbouring cells. Figure 4 shows the individual tiles and the neighbour relation necessary for the tile to be used. The resultant images formed by a grid of cells (illustrated in Figure 5), form a continuous mass of substance, as opposed to squares containing individual entities. These minimalist geometric textures suggest abstract landscapes rather than the iconic or literal individual representations that are common in many Artificial Life simulations. This design decision forms an integral aesthetic component of the work.

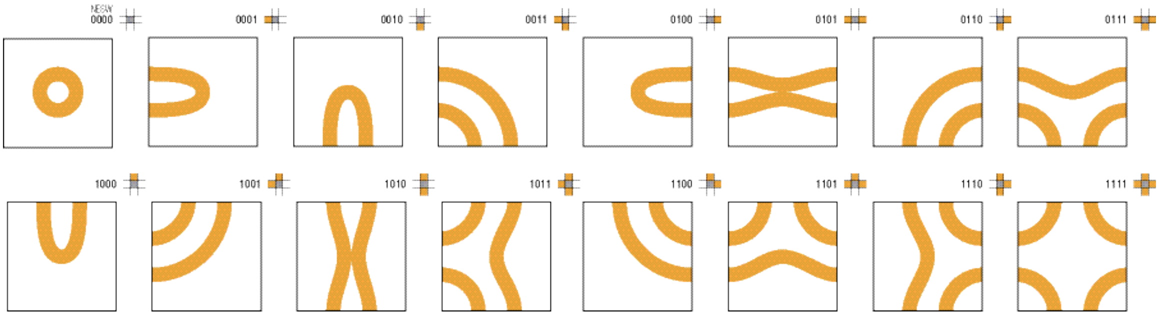

Figure 4: Cellular tiling set for *Eden*'s bio-mass. Each cell considers it's four immediate neighbouring cells (north, south, east and west). The neighbouring relations determine the image used for the cell. A function returns the bit pattern representing the neighbourhood state for the cell and the tile is selected based on the supplied index. Four bits are required, each representing the four directions. The bits are encoded NESW (from MSB to LSB). The symbols above each cell pattern shown here illustrate the bit pattern and corresponding neighbourhood relationships.

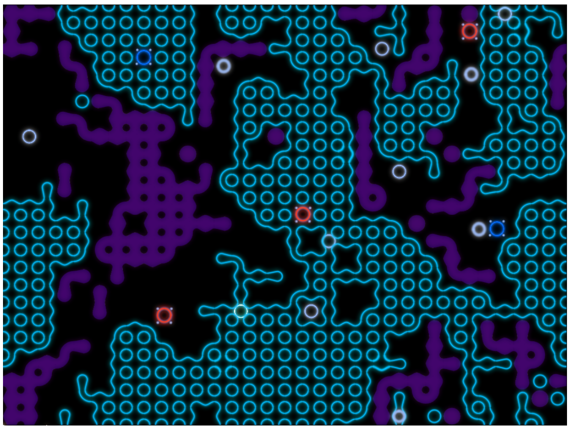

Figure 5: Illustrates the visualization of the *Eden* world, showing rocks (purple), bio-mass (light blue) and agents (circular elements, red and white).

3.4. Sound

One of the key elements of *Eden* is the ability of agents to create and listen to sound. A large proportion of sensor bandwidth is devoted to sound, allowing orthogonal sensing of both frequency and sound pressure (volume). Some basic physical modelling is performed on sound pressure levels, however many physical sound propagation aspects are simplified in the interests of efficiency.

3.4.1. Sound Generation

Actuator messages requesting sound generation need to be converted into a generated sound. As described in section 3.2.2, actuator messages are bit strings. A portion of the string encodes the sound generation command ('sing'), the remainder the sound generation data (energy levels over a range of frequency bands). In the current implementation, there are three distinct frequency bands, each occupying one third of the total of the sound generation data for the 'sing' actuator message (see Figure 6).

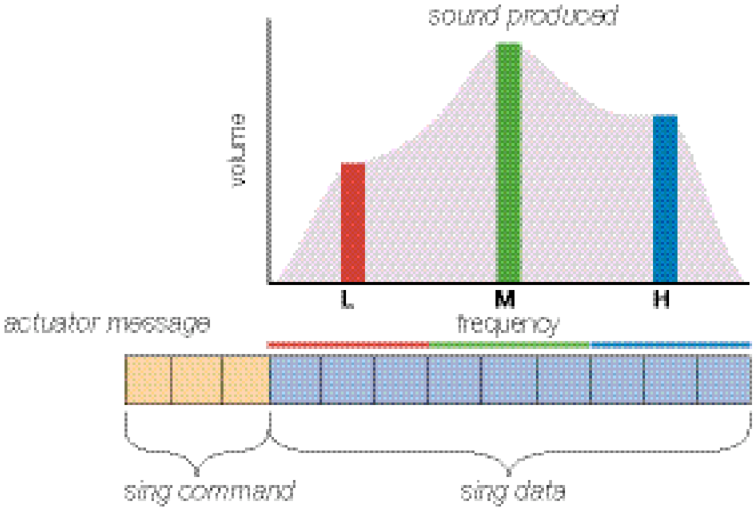

Figure 6: The 'sing' actuator message contains two parts. The first is the command requesting the agent to perform a sing operation; the remainder contains the sing data: volume levels for three distinct frequency bands. Using three bits per frequency band results in 2^9 or 512 distinct sounds.

When an agent 'sings' the spectral signature determined by the sing data in the actuator message is registered for the current time step. In addition, the same signature is used to drive a sonification process, so that people in the exhibition space can hear sounds that correspond to the 'singing' activities of the agents. To drive this sonification process, the three frequency bands are assigned labels 'L', 'M' and 'H' corresponding to low, medium and high pitched sounds (for example the majority of spectral energy in the 100, 1000 and 10,000Hz regions respectively). When an agent makes a sound, the corresponding selection from a pre-computed library of sounds is triggered and sent to the audio subsystem. The audio subsystem does basic sound spatialization using the four channel audio system that is part of the artwork. Sounds are spatialized according the position of the agent that is making the sound on the screen. Thus, as an agent making sound moves across the screen that sound will appear to move with the agent to human observers. The audio sub-system allows many agents to be making sound simultaneously.

3.4.2. Sound Reception

Agents have a significant amount of sensor bandwidth devoted to sound reception. An agent's sound reception area is a forward facing conical pattern that, like the sound generation, is sensitive across three separate frequency bands (see Figure 7). Each band has the same propagation and reception pattern, i.e. there are no frequency dependent differences in the modelling.

At each time-step, the conical reception area for each agent is checked for any other agent that is singing within that area. A simple physical model (Roederer, 1975) controls the propagation of sound through the environment[5]. Sounds arriving at the agents cell are summed on a per-frequency basis and the resultant sensor message instantiated.

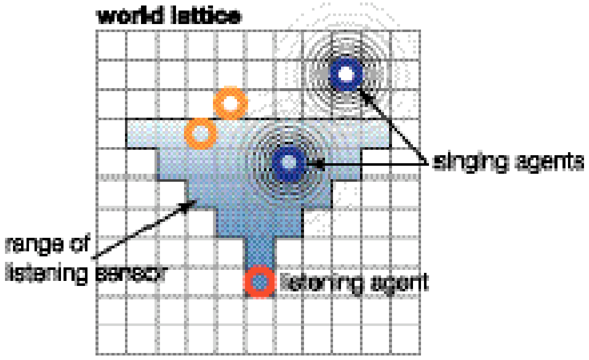

Figure 7: The reception area of an agent. The listening agent will hear agents who are making sound within the blue area only. A simple physical model controls the perceptual volume levels for the agent.

4. Interaction

The *Eden* system has a unique relationship between the physical and virtual components of the system. As shown in Figure 1, a series of infrared sensors[6] are placed around the screens in the exhibition space. These sensors measure distance. Each sensor has a range of approximately 100cm. Data collected from individual sensors is digitized in a data-collection subsystem that is then used to infer the presence and movement of people in the space. This data is used to drive environmental parameters in the virtual simulation. Before discussing the details of the mappings between sensor data and the simulated environment, a background discussion on the rational for such mappings will be presented.

4.1. The Problem of Aesthetic Evolution

Typically, genetic algorithms evolve towards finding maxima in *fitness*, where fitness is some criteria that can be evaluated for each phenotype of the population. Many systems define an explicit *fitness function* that can be machine evaluated for every phenotype at each generation (Mitchell, 1996).

Aesthetic evolution or *aesthetic selection* is a popular technique that replaces the machine evaluated fitness function with the subjective criteria of the human operator. Aesthetic evolution was first used by Dawkins (Dawkins, 1986) in his ‘Blind Watchmaker’ software to evolve two-dimensional, insect-like shapes. Aesthetic selection has been used to successfully evolve images (Rooke, 2002; Sims, 1991a), dynamic systems (Sims, 1991b), morphogenic forms (Todd and Latham, 1992; McCormack, 1993), even musical patterns and structures (Bulhak, 1999). Regardless of the system or form being evolved, aesthetic selection relies on the user to explicitly select phenotypes at each generation. Users typically evolve to some subjective criteria — often described as ‘beautiful’, ‘strange’ or ‘interesting’ — criteria that prove difficult to quantify or express in a machine representable form (hence the use of the technique in the first place).

However, aesthetic evolution has two significant problems:

---

[5] When sound propagates in a medium such as air at standard temperature and pressure the relationship between distance and perceived levels is The perceptual mechanism for loudness behaves in an exponential way, as it does for humans, $L = 20 \cdot \log_{10}(P/P_0)$. Where L is the sound pressure level in decibels (dB), P_0 a reference pressure corresponding roughly to the threshold of hearing in humans (Roads, 1996).

[6] The use of the term ‘sensors’ here referrers to physical devices and should not be confused with the sensors described in section 3.2.1, which are virtual (i.e. simulated).

- The number of phenotypes that can be evaluated at each generation is limited by both screen area (in the case of visual representation) and the ability of people to perform subjective comparisons on large numbers of objects (simultaneously comparing 16 different phenotypes is relatively easy, comparing 10,000 would be significantly more difficult).
- The subjective comparison process, even for a small number of phenotypes, is slow and forms a bottleneck in the evolutionary process. Human users may take hours to evaluate many successive generations that in an automated system could be performed in a matter of seconds.

What we would like is a system that combines the ability to subjectively evolve towards phenotypes that people find ‘interesting’ without the bottleneck and selection problems inherent in aesthetic evolution.

4.2. *Eden* as a Reactive System

The solution used to the problem described in the previous section is to map the presence and motion data of people experiencing the artwork to the environmental parameters of the virtual environment. Thus the virtual world in which sonic agents live and evolve is dependent not only on the simulated qualities discussed so far, but the presence (or absence) of people experiencing the work and their behaviour within the exhibition space.

Eden has no explicit fitness function. Agents continue to be part of the system based on how well they can survive and mate in the current environment. If certain selection pressures are applied, such as food becoming scarce, only those agents who can adapt and find food will prosper. By driving environmental conditions from the presence and movement of people in the exhibition space, agents must implicitly adapt to an environment that includes aspects of the world outside of the simulation.

In the current system, the following mappings are used:

- *Presence in the real environment maps to bio-mass growth rates*. The presence of people around the screen area affects the rate of bio-mass growth in that local area of the *Eden* world. Areas with no people correspond to a barren environment — little bio-mass will grow without the presence of people in the real environment.
- *Movement in the real environment maps to genotype mutation rates*. The greater the movement of people in the space the higher the mutation rate for rule evolution (see section 3.2.4).

These mappings are based on the following assumptions. Firstly, people will generally only spend time experiencing something if it interests them. In the context of experiencing an artwork, generally people may spend a short time evaluating their interest in an artwork, but after some time, if it no longer interests them, they will leave. Of course there may be other reasons for leaving, but in general the duration of stay will have some relation to how ‘interesting’ the experience is.

Agents require food to survive. If people are in the real environment, then food will grow at a more rapid rate. An agent who is making ‘interesting’ noises for instance, would have a better chance of keeping a person’s attention than one who is not. Moreover, an agent making a progression of sounds, rather than a just a single, repeating sound, is likely to hold a person’s attention even longer. Agent’s who encourage and hold a person’s attention in the space implicitly give the environment a more plentiful food supply.

The movement of people in the space mapping to mutation rates is based on the assumption that people will move over an area looking for something that interests them, and when they find it, will stay relatively still and observe it. Hence, the movement of people within the real space serves to inject ‘noise’ into the genome of agents who are close to the source of movement. Higher mutation rates result in more variation of rules[7]. If an agent or group of agents are holding the viewer's attention, then less rule discovery is needed in the current environment. Whereas, if people are continually moving, looking for something 'interesting', this will aid in the generation of new rules.

Further details on the dynamics of this component of the system can be found in (McCormack, 2002).

5. Results

At the time of writing, a number of exhibitions of the work have been completed. Images from an exhibition of the work are shown below in Figure 8. A typical exhibition may last several weeks, giving plenty of opportunity for the agent evolutionary system to take into account the behaviour of people experiencing the work. Certain factors have a marked effect on this behaviour and need to be compensated for. For example, when the gallery is closed there will be no people in the space anyway[8].

Analysis of the rules used by agents shows that sound is used to assist in mating as would be expected (Werner and Dyer, 1991) and with the influence of people, sound is used in other ways as well. Once the environmental pressures from audience behaviour are incorporated into the system, the generation of sound shows a marked increase and analysis of the rules discovered shows that making sound is not only used for mating purposes.

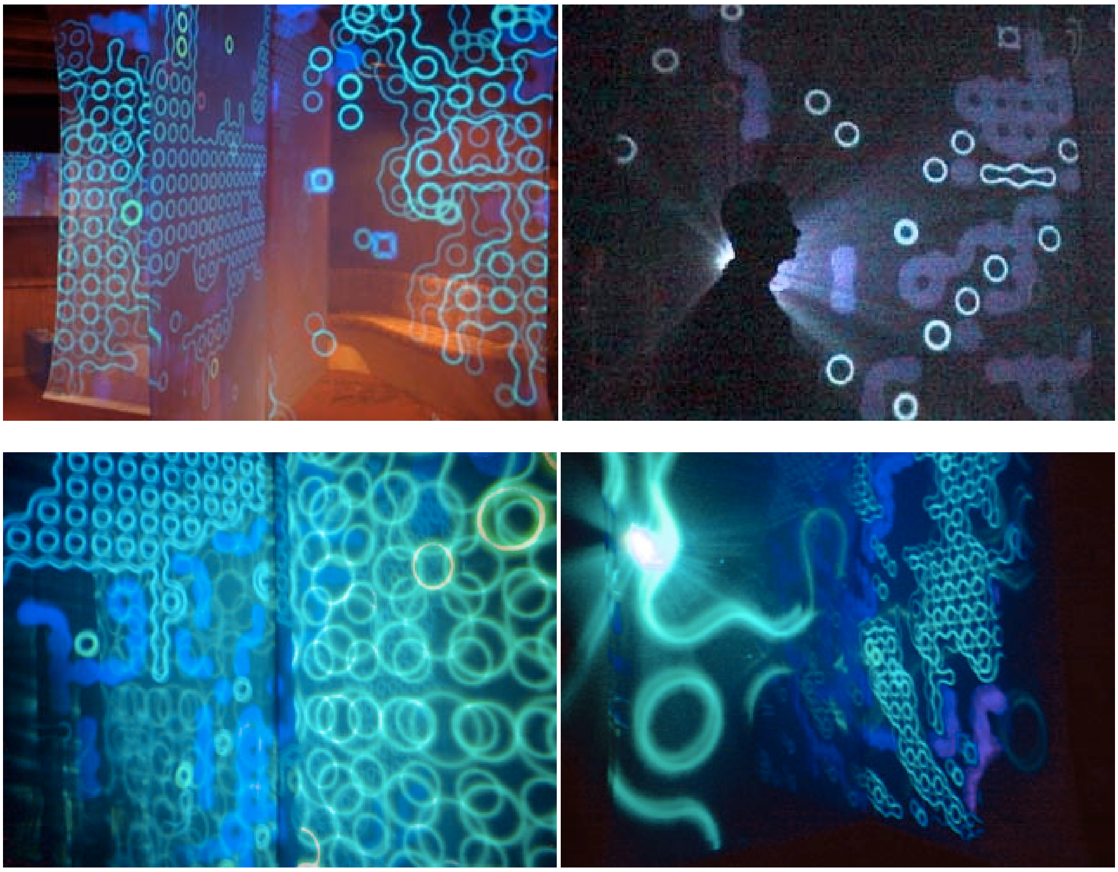

---

[7] Most child rules that mutate will not be 'better' than the parent rule, but in general, the use of mutation does provide the possibility for the system to discover rules that would not be possible by crossover alone.

[8] Without compensation for gallery opening hours, the entire population dies out each night!

Figure 8: Images of *Eden* in operation.

6. Conclusion

This paper has described a novel evolutionary system, where agents make use of sound to assist in survival. While the main impetus and methodologies are based around the development of an artistic system, it is hoped that some of the ideas presented here may be of interest to those interested in Artificial Life from other perspectives or with different agendas and applications.

In summary, a system has been produced that attempts to integrate the open-ended nature of synthetic evolutionary systems into an interactive virtual space. The approach used in this paper has been to measure components of the real environment, incorporating them into that of the virtual one, thus enabling an evolutionary relationship between virtual agents and the artwork's audience, without need for explicit selection of phenotypes that engage in 'interesting' behaviour.

Further information, including sample sound recordings and video documentation of the work are available on-line at: http://www.csse.monash.edu.au/~jonmc/projects/eden.html.

7. Acknowledgements

The author would like to thank the reviewers for their helpful comments.

8. References

Ashby, W.R. (1956), *An Introduction to Cybernetics,* Chapman & Hall, London.

Bentley, P.J. (1999), *Evolutionary Design by Computers,* Morgan Kaufmann Publishers, San Francisco, Calif.

Bentley, P.J. and D.W. Corne (eds.) (2002), *Creative Evolutionary Systems,* Academic Press, London.

Bowler, P.J. (1992), Lamarckism, *in* Keller, E.F. and E.A. Lloyd (eds), *Keywords in Evolutionary Biology*, Harvard University Press, Cambridge, MA, pp. 188-193.

Bulhak, A. (1999), Evolving Automata for Dynamic Rearrangement of Sampled Rhythm Loops *in* Dorin, A. and J. McCormack (eds), *First Iteration: a conference on generative systems in the electronic arts*, CEMA, Melbourne, Australia, pp. 46-54.

Cariani, P. (1991), Emergence and Artificial Life *in* Langton, C.G., et al. (eds), *Artificial Life II, SFI Studies in the Sciences of Complexity*, Vol. 10 Addison-Wesley, pp. 775-797.

Cariani, P. (1993), To Evolve an Ear: Epistemological Implications of Gordon Pask's Electrochemical Devices, *Systems Research* **10**(3)**,** pp. 19-33.

Conrad, M. and H.H. Pattee (1970), Evolution Experiments with an Artificial Ecosystem, *Journal of Theoretical Biology* **28,** pp. 393.

Dahlstedt, P. and M.G. Nordahl (1999), Living Melodies: Coevolution of Sonic Communication *in* Dorin, A. and J. McCormack (eds), *First Iteration: a conference on generative systems in the electronic arts*, Centre for Electronic Media Art, Melbourne, Australia, pp. 56-66.

Dawkins, R. (1986), *The Blind Watchmaker,* Longman Scientific & Technicial, Essex, UK.

Emmeche, C., S. Køppe and F. Stjernfelt (1997), Explaining Emergence: Towards an Ontology of Levels, *Journal for General Philosophy of Science* **28,** pp. 83-119.

Goldberg, D.E. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning,* Addison-Wesley Pub. Co., Reading, Mass.

Grünbaum, B. and G.C. Shephard (1993), Interlace Patterns in Islamic and Moorish Art, *in* Emmer, M. (ed) *The Visual Mind: Art and Mathematics*, MIT Press, Cambridge, MA.

Holland, J.H. (1992), *Adaptation in Natural and Artificial Systems : An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence,* MIT Press, Cambridge, Mass.

Holland, J.H. (1995), *Hidden Order : How Adaptation Builds Complexity,* Addison-Wesley, Reading, Mass.

Langton, C.G. (1989), Artificial Life, *in* Langton, C.G. (ed) *Artificial Life, Sfi Studies in the Sciences of Complexity*, Vol. 6, Addison-Wesley, pp. 1-47.

McCormack, J. (1993), Interactive Evolution of L-System Grammars for Computer Graphics Modelling, *in* Green, D. and T. Bossomaier (eds), *Complex Systems: From Biology to Computation*, ISO Press, Amsterdam, pp. 118-130.

McCormack, J. (2001), Eden: An Evolutionary Sonic Ecosystem *in* Kelemen, J. and P. Sosík (eds), *Advances in Artificial Life, 6th European Conference, ECAL 2001*, Springer, Prague, Czech Republic, pp. 133-142.

McCormack, J. (2002), Evolving for the Audience, *International Journal of Design Computing* **4**(Special Issue on Designing Virtual Worlds).

McCormack, J. and A. Dorin (2001), Art, Emergence and the Computational Sublime *in* Dorin, A. (ed) *Second Iteration: a conference on generative systems in the electronic arts*, CEMA, Melbourne, Australia, pp. 67-81.

Mitchell, M. (1996), *Introduction to Genetic Algorithms,* MIT Press, Cambridge, MA.

Nagel, E. (1961), *The Structure of Science : Problems in the Logic of Scientific Explanation,* Routledge, Lond.

Pattee, H.H. (1988), Simulations, Realizations, and Theories of Life, *in* Langton, C.G. (ed) *Artificial Life*, Vol. VI, Addison-Wesley, pp. 63-77.

Reichardt, J. (1971a), *Cybernetics, Art and Ideas,* New York Graphic Society, Greenwich, Conn.

Reichardt, J. (1971b), Cybernetics, Art and Ideas, *in* Reichardt, J. (ed) *Cybernetics, Art and Ideas*, Studio Vista, London, pp. 11-17.

Roads, C. (1996), *The Computer Music Tutorial,* MIT Press, Cambridge, Mass.

Roederer, J. (1975), *Introduction to the Physics and Psychophysics of Music,* Springer-Verlag, New York.

Rooke, S. (2002), Eons of Genetically Evolved Algorithmic Images, *in* Bentley, P.J. and D.W. Corne (eds), *Creative Evolutionary Systems*, Academic Press, London, pp. 339-365.

Sims, K. (1991a), Artificial Evolution for Computer Graphics, *Computer Graphics* **25**(4)**,** pp. 319-328.

Sims, K. (1991b), Interactive Evolution of Dynamical Systems *First European Conference on Artificial Life*, MIT Press, Paris, pp. 171-178.

Sommerer, C. and L. Mignonneau (eds.) (1998), *Art@Science,* Springer-Verlag, Wien, Austria.

Todd, P.M. and G.M. Werner (1998), Frankensteinian Methods for Evolutionary Music Composition, *in* Griffith, N. and P.M. Todd (eds), *Musical Networks: Parallel Distributed Perception and Performance*, MIT Press/Bradford Books, Cambridge, MA.

Todd, S. and W. Latham (1992), *Evolutionary Art and Computers,* Academic Press, London.

Ulam, S. (1952), Random Processes and Transformations *Proceedings of the International Congress on Mathematics*, Vol. 2, pp. 264-275.

Watson, A.J. and J.E. Lovelock (1983), Biological Homeostasis of the Global Environment: The Parable of Daisyworld, *Tellus* **35B,** pp. 284-289.

Werner, G.M. and M.G. Dyer (1991), Evolution of Communication in Artificial Systems, *in* Langton, C.G. (ed) *Artificial Life II*, Addison-Wesley, Redwood City, CA, pp. 659-682.

Wiggins, G., et al. (1999), Evolutionary Methods for Musical Composition *Proceedings of the CASYS98 Workshop on Anticipation, Music & Cognition*.

Wilson, S. (2002), *Information Arts : A Survey of Art and Research at the Intersection of Art, Science, and Technology,* MIT Press, Cambridge, Mass.

Witten, T.A. and L.M. Sander (1981), Diffusion-Limited Aggregation, a Kinetic Critical Phenomenon, *Phys. Rev. Letters* **47,** pp. 1400-1403.