The Sound Gallery: Recent Developments

Sam Woolf and Adrian Thompson
Interact Lab,
School of Cognitive and Computer Sciences,
Sussex University
Brighton BN1 9QH
England.

samw@cogs.susx.ac.uk

00-44-1273 877221

Abstract

"The Sound Gallery" is an ongoing experiment in art and science that grew out of research in the field of Evolutionary Electronics at the University of Sussex, England.
The aim of this experiment is to utilise ideas and techniques derived from the science of Artificial Life to create an interactive and adaptive installation artwork that displays aesthetically interesting behaviour. This paper gives a brief overview of the 'sound gallery' project, and discusses some of the recent developments that have been made. Recent work on the ‘sound gallery’ project has mainly involved the development of ultrasonic range-finder devices for use as a portable sensory apparatus and novel interface to the sound gallery system.

Overview:

The field of Artificial Life is one that has been explored by a number of recent artists. It has spawned a diversity of new media, and has proven a rich source of inspirations. There now exists a substantial body of work belonging to the canon of A-life art, and a growing number of artist/scientists who are increasingly active in the field. "The Sound Gallery" adds a new artwork to this lineage, drawing inspiration from other works of A-Life art, but also making some notable advances. It achieves interactive and adaptive behaviour in novel ways, making use of reconfigurable hardware technology that (as far as we are aware) has not before been appropriated for artistic ends.

What is music? The composer Arnold Schoenberg claimed that music could be defined as "repetition and variation"(Retallack 1996). The aim of "The Sound Gallery" is to take this statement literally, and to its logical extreme: A sound-source is transformed into 'music' through repetition and distortion as it plays through four separate speakers simultaneously. The signal to each speaker is endowed with variations by the intervention of separate signal processing circuits. The electronic character and make-up of these circuits, and hence the precise distortion effects they entail, are perpetually in a state of flux and change. The circuits are constantly manipulated and reconfigured through a process of artificial evolution driven by aesthetic selection pressures.

The four speakers are to be positioned in an open gallery setting. Audience members, passers-by and curious parties will be encouraged to interact with this gallery space in whatever ways they feel fit. They should allow themselves to drift freely through the gallery-space, exploring the soundscape generated by the speakers within it, seeking out those areas they find the most aesthetically pleasing, or at worst, those that are the least repellent...

Sensors in the gallery track the movements of those present over time. Fitness values for each of the four speakers are derived from this information and passed to a genetic algorithm controlling the evolution of the four reconfigurable circuits. The fitness value of a particular speaker will be calculated from the movements of the audience and the ways in which individuals interact with that speaker. The fitness definition used has been designed to place the four speakers in direct competition, each one struggling against its adversaries for the possession of a scarce resource... the attention and appreciation of human spectators.

A working prototype of "The Sound Gallery" system was developed at The University of Sussex during the summer of 1999. Recently a few refinements have been made, and various other improvements are still planned. I shall discuss the behaviour and substance of this prototype model, and also consider some of the refinements still to be made. But first, we speak of 'Evolvable Hardware' and its practical and aesthetic pertinence to this project.

Evolvable Hardware

Although many of the more mechanistic aspects of electronics design have been automated, at present an inventive skilled human must still be at the helm of the Computer Aided Design (CAD) tools that are now in common use. This person still experiences the pleasure of engaging in a creative activity. However, over the past five years, some researchers have begun exploring the idea of using the apparently ingenious, opportunistic, and innovative aspects of artificial evolution to aid the creativity of human electronics practitioners. (Higuchi, Iwata and Weixin 1997; Sipper, Mange and Perez-Uribe 1998; Miller, Thompson, Fogarty and Thomson 2000)

In natural evolution there is no clear long-term objective. Adaptations are made according to environmental pressures, and these change over time – especially since some of the most important parts of the environment are other life forms doing the same thing. Even when there is an obvious demand for a particular short-term improvement, such as frustrating a new predator, it is often hard to see what changes to the body or behaviour would be best. Artificial evolution for electronics design has usually been much more constrained, with a pre-specified target behaviour as the only objective. Even given such a rigid goal, if the conventional rules of electronics design are not enforced, evolution can produce circuits that seem alien to a trained designer. When such ingenious but bizarre circuits are presented to a group of engineers, in our experience one half is delighted and the other is disgusted. (This can even happen with only one engineer in the group…)

Perhaps the first example of this (Thompson 1996) is shown in Figure1. The circuit has a single input and a single output, and was required to discriminate between a high-pitched audio tone (10kHz) and a lower-pitched one (1kHz). The output should be a steady high voltage whenever one tone is present at the input and a steady low for the other. At the right of the picture we can see some of the peculiar behaviours produced during evolution before finally exactly the target one was achieved. Even though the circuit is small and the external behaviour is simple, it took about two week's hard effort from two skilled designers to analyse roughly how it worked (Thompson and Layzell 1999).

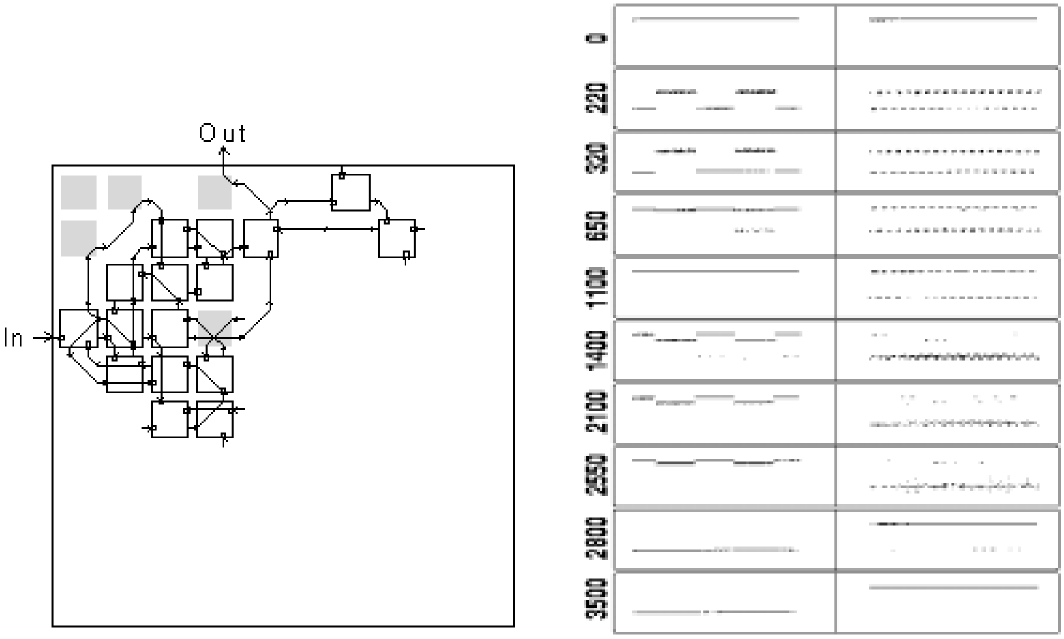

FIGURE 1.

The inspiring evolved tone discriminator circuit. Only the wiring between the components is shown, though their behaviours were also subject to evolution. The components shaded grey are exploited through parasitic coupling, even though they are not directly wired to the rest of the circuit. At the right, the responses of some of the circuit's ancestors are shown, starting with an initial random population (generation 0), and finishing with a near-perfect result (3500 generations).]

Some of the tricks used by the circuit are so unusual that sometimes the analysts' experience of conventional design was positively unhelpful. Some of the investigative techniques needed were more akin to neuroscience than to normal electronics. It turns out that the central part of the circuit, stumbled upon by chance early in evolution, is an oscillator with two different stable modes of oscillation. Which of the two modes it falls into is very sensitively determined by the conditions on the silicon chip at that time. In fact, to discriminate between low and high tones, this sensitive dynamics is used to transduce and amplify a tiny and normally undetectable 'parasitic' side effect of the way the chip is made. This mechanism was fine tuned during the rest of evolution and other components were added to give the high/low output rather than an oscillation. The whole is a finely balanced system, exploiting many of the precise time delays and physical properties of the silicon components. Such exploration beyond the scope of conventional design (Thompson, Layzell and Zebulum 1999) could have exciting engineering implications, especially when steps are taken to ensure 'robustness' over a variety of conditions of temperature, fabrication variation, and so on. (Thompson and Layzell 2000)

These experiments within the field of 'Evolvable Hardware', and the strikingly unconventional qualities of the circuits they produced, set the stage for the Sound Gallery, and inspired one of its key aims: To enable random and untrained members of the public to participate, perhaps unwittingly, in the design of complex and subtle electronic structures, hence demystifying and subverting the design process whilst simultaneously shrouding its final product with an almost indecipherable and 'alien' complexity. The meandering public are to become expert electronics designers and alchemists, the sum of their base drives and motivations transformed to silicon gold.

The direct exploitation and manipulation of the physical properties and behaviours of real interconnected components embodied on a silicon chip provides more freedom to the processes of evolution than some other possible implementations of the Sound Gallery system (for example, the evolutionary 'tweaking' of the parameters of a digital signal processing (DSP) algorithm). It potentially allows for a far greater range of unpredictable behaviours, as evolution may exploit subtle properties of the silicon to unforeseeable and unrepeatable affect. It also renders this behaviour tangible, as each evolved circuit exists as a true physical entity, and not merely as a range of parametric values stored in a computer memory buffer.

On a more practical note, the use of 'reconfigurable' chips avoids the need to evaluate the evolving circuits in a software simulation. Although important in the developing field of 'evolutionary electronics', simulations of circuit behaviour inevitably constrain possibilities, or miss nuances of physical reality, allowing only the smallest and most simplistic of circuits to be simulated fast enough to accommodate real-time interactive behaviour.

Reconfigurable chips consist of an array of components interspersed with wires. There are electronic switches distributed throughout, that control the behaviours of the components and how they connect to the wires. In the types most useful for evolution, the electronic switches are made from transistors turned on or off according to the 1 or 0 contents of adjacent bits of an on-chip RAM memory. Any computer can be interfaced to the device, and the computer's software can write to this configuration memory just like a normal memory chip. By doing so, a particular circuit is physically instantiated in silicon by the reconfigurable switches. This circuit then behaves in real-time according to the laws of semiconductor physics, rather than as a computer program stepping through a sequence of instructions.

Many kinds of reconfigurable chip exist, such as Field-Programmable Gate Arrays (FPGAs) intended for digital use (Oldfield and Dorf 1995) (as abused in the tone-recognition experiment above), or arrays of analogue components. The Sound Gallery uses the Zetex Totally Reconfigurable Analogue Circuit (TRAC) (Zetex group1999), with a separate TRAC020 chip manipulating the sound on its way to each loudspeaker. The four TRAC020 chips are housed on a Zetex development board that connects to the computer via a parallel cable. The genetic algorithm software, controlling the individual circuits configured onto each chip, runs on a single PC as will be described later.

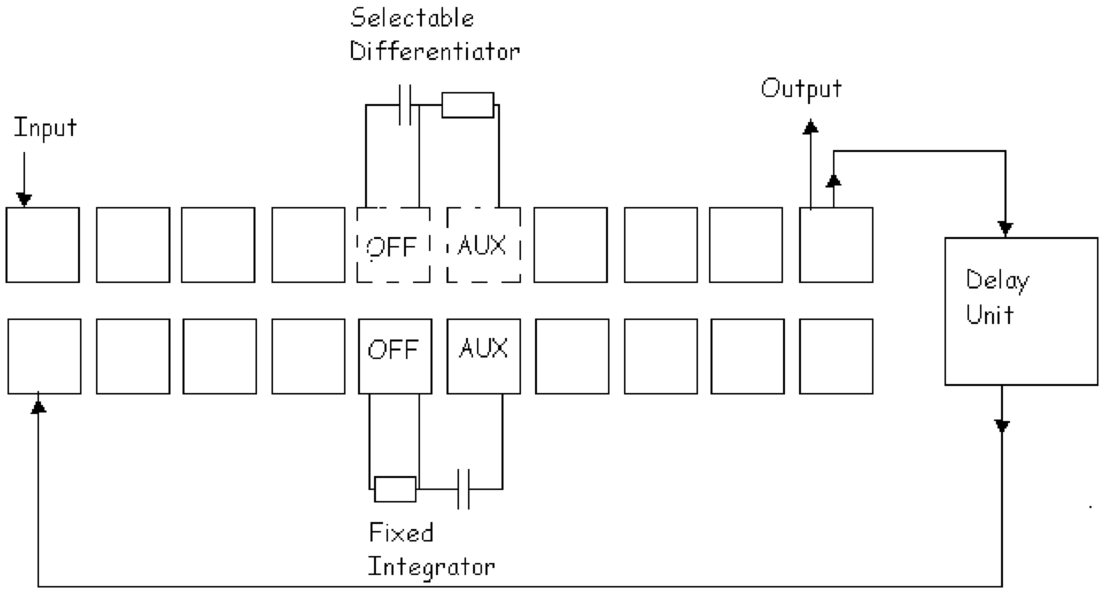

FIGURE 2: The layout of the TRAC020 chip, with external components.

The layout of the TRAC chip, along with its connections to some external components useful in our application, is shown in Figure 2. The component cells, or 'configurable analogue blocks' (CABs), are arranged in two rows of ten. The basic signal flow is from left to right, with each CAB taking an input from the two cells immediately to its left, one from each row. The inputs and output of each CAB are also connected to pins of the chip package. In our case most of the external pins are left unconnected apart from to supply signals to the leftmost cells, to extract an output from the top rightmost cell, and to provide some external resistors and capacitors, as shown.

Each CAB contains an operational amplifier surrounded by a few components with reconfigurable connections. Nominating a cell's input from the same row A and from the other row B, then the repertoire of possible CAB behaviours is: A+B, -B, A, log(A), exp(A), Half-Rectify(exp(A)), Aux(external components), and High-Impedance-OFF. For some of these there are scaling factors and offsets, given in the datasheet. The topology and CAB repertoire would in itself give very limited acoustic distortion effects, so some external resources were provided. For each TRAC chip, a separate digital delay module was connected as shown in the figure. The additional external capacitors and resistors shown provide the possibility of differentiation or integration of the signal if the pair of CABs they are connected across are configured in the combination (OFF, Aux).

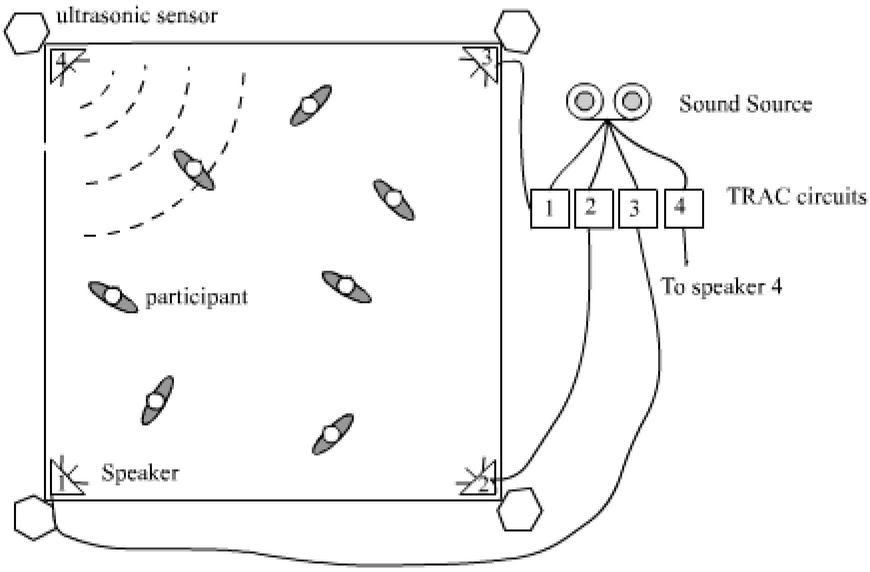

FIGURE 3: The gallery setup.

Figure 3 is a sketch of the gallery arrangement. A central sound source is transmogrified by a different evolved circuit, embodied in the reconfigurable TRAC chips, on its way to each speaker. Participants' movements in the surrounding gallery space are influenced by their aural experience. A computer (not shown) observes the ostensible appreciation of each speaker via an ultrasonic sensing system, and assigns evolutionary fitnesses to the circuits accordingly. The computer maintains a sub-population of circuits for each speaker, and schedules individuals for evaluation, whereupon they are sent as configurations to the chips.

Sensing System.

Up until now, experiments with the Sound Gallery system have been conducted using an ad-hoc sensing system that relies on a human operator directly observing the movements of participants and entering values representing their relative positions into the computer by hand for each new generation of evolution (see the description of the early experiments below). This was at best only a temporary solution, and recent work on the development of ultrasonic range-finder devices to fully automate the system is now nearing completion.

The choice of what kind of sensor system to implement has been a difficult one.
The sensor system needs to be tailored to the environment of the installation, and as it is hoped that the Sound Gallery may eventually be exhibited in various locations, this factor is to a large extent an unknown quantity. There are several choices for sensors that can detect the distance or position of humans. Which is best depends on cost, time, the size and furnishing of the space, indoors or out, and any need to withstand attack. Cost-effective candidates are floor pads, passive infrared (sensing movements of body-heat sources within an area, as done for burglar alarms), sonar (ultrasonic time of flight), electric field proximity sensing , and image analysis from an overhead camera. At greater cost, sensors based on microwave Doppler-shift or laser time of flight might be explored.

Finally, in a well-controlled setting, if the participants are in on the game they can each be given some sort of reflector, transponder, or beacon, allowing their position to be triangulated from base-stations by radio or ultrasound.

It is also possible to envisage that in some environments the positions of participants may not be the most useful data. In a crowded night-club setting, for example, freedom of movement may be limited and the distance of participants from the speakers may prove to be a poor indicator of their preferences. In such circumstances crowd enthusiasm levels could perhaps instead be betrayed by the vigor of the participants' bouncing in the vicinity of each speaker, which could be measured with pressure pads under the floor covering Since people may be unaware of their participation, and should not necessarily have to be conscious of the preferences they express, asking them to score the speakers more directly is not an option.

The choice of an ultrasonic rangefinder or sonar sensor system was eventually arrived at after more obvious candidates like pressure pads and image analysis of overhead camera footage were ruled out on the basis that these would not be viable solutions in many potential environments (such as outdoor locations or rooms with low ceilings). It is hoped that the ultrasonic sensors should work reasonably well in most situations. In addition, they can be built cheaply, are light and portable, and they do not require intrusive overhead cameras or participant tagging.

Description of Ultrasonic Sensor System

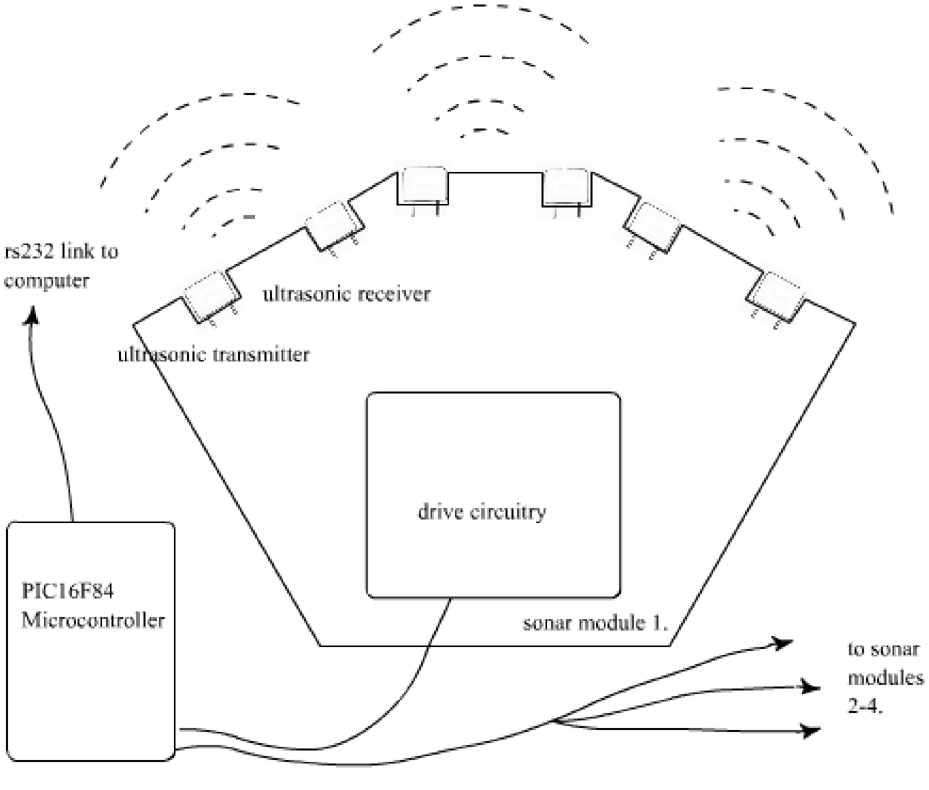

FIGURE 4: The Ultrasonic Sensor System.

The ultrasonic sensor system consists of four separate sonar modules, one of which is to be positioned in each corner of the gallery space. Each sonar module houses three pairs of ultrasonic transducers, each pair consisting of a transmitter and a receiver. The transmitters emit pulses of ultrasonic (40 kHz) energy, and the receivers detect the echoes of these pulses as they bounce back off objects within range. By counting how much time has elapsed between the sending of the ultrasonic pulse and the return of its echo, it is possible to calculate how far away objects (people) in the room are. After each set of ultrasonic pulses has been transmitted, the receive section waits for the first three returning echoes. This means it is possible to detect the distances of the three nearest objects to the active transducer pair with every set of pulses transmitted .

The ultrasonic transducer pairs have been chosen so that their receptive fields approximate a narrow beam of roughly 30 degrees and range of around three metres. By angling each transducer pair so their receptive fields cover different 30 degree slices of the gallery space, it becomes possible to receive not only distance, but also approximate directional information about the positions of nearby obstacles.

A programmable microcontroller chip called a PIC16F84 controls and coordinates the sonar modules and communicates with the main computer via a rs232 link.

Description of "The Sound Gallery" Programs and Algorithms:

The algorithm actually works in two distinct phases, as shall be described in more detail below. First a hillclimb phase is implemented, and then a steady-state island model genetic algorithm. The hillclimb algorithm is used to initialise the four islands (sub-populations) associated with the chips. Hillclimbing may start from random initial circuits, from hand designed circuits, or from randomly generated circuits that have been selected out by the user. Hillclimbing creates initial populations for the genetic algorithm comprising of reasonably fit diverse solutions that are already related to each other to some degree.

The second phase of the algorithm begins with the implementation of an island model genetic algorithm. An island model GA, as defined by Whitley et al (Whitley 1997) is a GA that uses several processors to explore the same problem space simultaneously and in parallel, whilst periodically exchanging genotypes between machines to enable the sharing of information. Our GA only uses one processor, but does maintain four independent sub-populations between which individuals migrate at times. The island model GA was chosen so that each speaker would be driven through circuits evolved from different sub-populations, each having their own individual genetic characteristics, hence reducing the likelihood of all the speakers producing similar sounds. It was also hoped that the migration operator used with island model GAs would cause interesting effects to occur; allowing, for example, aesthetically interesting sounds to 'jump' from one speaker to the next.

The Hillclimb Phase:

Four initialisation genotypes are generated, one for each of the four sub-populations. Hillclimbing then commences, with each sub-population working in parallel to the other three. The initialisation genotypes undergo repeated mutations, generating new genotypes that represent new TRAC configurations. Each new genotype is evaluated and assigned a fitness value. When a mutation is evaluated to be fitter than, or equally fit to, the parent genotype from which it was derived, then this mutant genotype is stored as the next member of its sub-population, and will be used as the source for subsequent mutations.

The hillclimb algorithm is run until the four island populations are filled to capacity. This provides the genetic algorithm with its initial populations, each individual of which is already associated with a fitness value. These fitness values could not have been calculated in the conventional manner, as they represent a conglomeration of the subjective opinions and random wanderings of those present in the gallery, and cannot be mathematically derived from any aspect of individual genotypes or phenotypes. From this initial condition, the steps of the island model GA are followed and repeated (forever).

The Island Model Genetic Algorithm Phase:

Linear rank-based selection is used to select two 'parent' genotypes from each island sub-population. Child genotypes are derived from each pair of parents through the application of crossover, mutation and replication genetic operators. The TRAC development board is then reconfigured so that the circuit specifications represented by each of the new child genotypes are physically manifested in silicon. Each of the four circuits on the TRAC development board are then allocated fitness values, and the new genotypes replace the least fit members of their respective island sub-populations. This sequence of events represents one iteration of the algorithm.

The Sound Gallery GA also implements a migration scheme similar to that developed by Whitley et al. Every X iterations, each island copies an elite group of the fittest individuals from within its own population to one of the other islands. The receiving island avoids overpopulation by deleting the same number of the least fit individuals from within its own ranks. Islands are paired differently for each iteration of this migration cycle, so that they send and receive individuals from different sources throughout the duration of the evolutionary run.

The need for ageing:

Early experiments with small island populations showed that the design of the genetic algorithm allowed very fit genotypes to remain in the population indefinitely. Less fit more recent additions to the population could be eradicated by the periodic migrations of fitter strings from neighbouring islands. This situation could lead to population convergences, an unwelcome prospect, as apart from the obvious desire to maintain genetic diversity in the populations, it also seemed likely that some of these fitter genotypes could be the products of freak evaluation conditions. There could be times, for example, when a particular genotype receives a very good fitness value simply because of some freak fluctuation in audience movements, and not because it represents a circuit of particular merit. This could easily occur, for instance, at times when there are only one or two persons present in the gallery setting. Since genotypes are only evaluated once, and as this evaluation may lose validity with increasing temporal distance from its context, it was felt necessary to allow the genetic algorithm some means of recovery from freak convergence conditions. A periodic 'ageing' of fitness values was therefore introduced. The fitness values of each member of all island populations are periodically incremented by a small amount, making them more costly. This has the effect of making those genotypes that have been in the population for a large number of generations increasingly less fit, so that freakishly 'good' genotypes cannot linger indefinitely in the highest echelons of the population.

They are instead quietly retired, to make room for fresh young strings.

Description of encoding scheme:

A very simple encoding scheme has been used. Each genotype is a string of integer values, each sequentially corresponding to a single CAB of a TRAC chip (see figure 2). Each node of the genotype is allocated a value corresponding to the function type to which its associated CAB is to be configured. This encoding scheme enables all possible TRAC configurations to be represented as a string of 20 integers whose values may vary between 0 and 7.

The Fitness Function:

Conventional genetic algorithms are often used to find an optimal solution for some well defined problem domain. A traditional favourite GA domain is "The Travelling Salesman Problem" (TSP), where the goal is to find the most efficient path of travel between a number of unevenly distributed nodes, where each node is visited once and once only. In such optimisation problems, finding the fitness of a solution normally involves evaluating the optimality of some property of its phenotype. For the TSP, for example, the fitness of a solution is normally calculated as the total length of the path encoded by the evolved genotype.

The GA utilised for this project perhaps has more in common with natural selection than it does with algorithms used for optimisation problems. Ultimately whether or not a particular genotype is selected for or not depends not on some calculated value, but on how suited that genotype is to its current environment. This will depend on the process of interactions of all in the space concerned, and individual motivations may vary. Fitness of a genotype will depend on the positions of human participants in the gallery. However these participants may not be making conscious decisions and choices, but will register their votes in a more subliminal level. People in the gallery space need not even be aware of the fact that their each and every movement has a direct bearing on the evolution of new circuits, as omniscient sensors track their changes in position over time...

The fact that at each new step of the algorithm, four new child genotypes will be simultaneously evaluated means that the GA effectively places these new genotypes in direct competition. Each new genotype is evaluated in the context provided by the others. If the same genotype re-emerges in a different context, it is possible that it may be interpreted differently by the human participants, and hence receive a different fitness score. The fact that the geographical location of participants controls fitness evaluation also makes possible a kind of implicit 'fitness sharing', as a participant positioned equidistantly between two rival speakers will reward the genotype of the circuits driving each of these speakers equally.

Fitness evaluations may also be affected by noise in the measurements provided by devices tracking positions of human participants, by the time lengths allowed for each evaluation, and by idiosyncrasies in the sound source.

The Experiment:

An experimental run of the prototype Sound Gallery system developed for this project was executed in a small room in the Centre for Computational Neuroscience and Robotics (CCNR). The four speakers were positioned in each corner of the room, each one connected to one of the TRAC020 chips on the development board. Each chip received an identical input signal, a repetitive audio loop comprised of low frequency drum beats, high frequency cymbals, and synthesised notes of various frequencies, produced with drum-machine simulation software running on a laptop computer. The Sound Gallery program was running on a second laptop computer, and this computer was connected to the development board by means of a parallel cable. The room was divided into quadrants, with the dividing lines marked clearly on the floor of the room. Each speaker was allocated its own quadrant. Volunteers were instructed to explore the sounds generated by each speaker, and to settle in the quadrant allocated to whichever speaker they felt produced the most aesthetically pleasing sounds; this process to be repeated each time the circuits were reconfigured.

Before the start of the experiment, each chip was configured to a randomly generated circuit selected for its aesthetic value. The hillclimb algorithm then used these circuits as its starting point, and hillclimbing continued for ten steps. After ten steps of hillclimbing, the island-model genetic algorithm was implemented, and continued running for a further 40 steps.

Fitness values for each evolved circuit were entered into the computer running the Sound Gallery program by hand. To make this task executable, the fitness of a circuit was defined as (T-Q), where T was the total number of volunteers participating in the experiment, and Q was the number of volunteers who settled in the quadrant allocated to the speaker driven by the circuit under consideration.

The experiment ran for about an hour and a half. During this time, the number of volunteers present fluctuated between five and eight.

The volunteers responded well to the concept of exploring the different sounds emerging from different parts of the room, and seemed to enjoy and be genuinely enthusiastic about the experience. The Sound Gallery proved itself capable of producing a fairly large repertoire of interesting distortion effects, and when each chip distorted the sound to each speaker differently, a pleasing cacophony of related but highly individualised reverberations filled the room. To describe the resultant aural chaos as 'music' may be to somewhat labour the point- indeed, one of those present thought it more redolent of the sounds that brainwash Harry Palmer in "The Ipcress Files" than of anything more euphonious. However, with the four TRAC chips perpetually improvising new adaptations to their most popular configurations, and with the four speakers continually emitting their distorted variations to the sound-source's leading theme, it seems fair to claim that the "repetition and variation" required to fulfil Schoenberg's definition was achieved.

Some interesting artefacts of the way the system was implemented became apparent during our experiment. About half way through, speaker D stopped producing sounds, and speaker C became very quiet. This revealed an interesting idiosyncrasy: Participants were actually drawn to the quieter speaker, and crowded around it in an attempt to hear the sounds it produced. This effectively gave the quiet speaker an abnormally good fitness value, and so in the next stage of evolution, that speaker became quieter still, until in fact it became barely audible.

This meant that only Speakers A and B were producing interesting sounds, so for the rest of the run, people generally boycotted the quadrants allocated to speakers C and D, only moving between the quadrants allocated to A and B. This meant that every individual generated by the genetic algorithm from the populations of island C and island D effectively received the same fitness score. This meant that proper selection could not take place, and the sounds produced by speakers C and D did not improve. The only time C and D could have benefited from evolution would have been when receiving high fitness migrants from other islands. In practice, such migrations were too few and far between to have much beneficial effect. Our desire to retain the genetic diversity of each island had driven us to set the migration rate to just one individual every twenty steps. This, it would
The experiment described above stimulated both enthusiasm and scientific curiosity, and kept a roomful of people entertained for an hour and a half. During this time the dingy back-room of the CCNR was transformed into a perpetually mutating soundscape that at every moment offered up new aspects of itself for exploration. The Zetex TRAC020 chip, working in concert with the digital delay units, proved itself capable of producing a reasonably large and interesting repertoire of distortion effects. It was particularly good at picking out very high and very low frequencies from the source signal, and of distorting these extreme frequency bands in interesting ways. A major weakness was the fact that if a speaker repeatedly failed to produce interesting sounds, became very quiet, or produced no sound at all, and hence was continually boycotted by the participants, aspects of the design of the GA made it very difficult for such speakers to recover. Steps are to be taken in future implementations of "The Sound Gallery" to exploit or eradicate the effects of these idiosynchracies.

Acknowledgements

Many thanks to my supervisors Adrian Thompson and Yvonne Rogers. This work has been funded by EQUATOR Interdisciplinary Research Council.

References and Bibliography:

- Borenstein, J et al [1994] “Where Am I?” Sensors and Methods for Autonomous Mobile Robot Positioning, University of Michigan. Technical Report UM-MEAM-94-21.
- De Garis, H [1995] EVOLVABLE HARDWARE Workshop Report:, ATR, Kyoto http://www.cwi.nl/~ehrenbur/EvolvableHardware/WorkshopReport.html
- Flockton S, Sheehan K [1998] ”Intrinsic Circuit Evolution Using Programmable Analogue Arrays” in Evolvable Systems: From Biology to Hardware. Proceedings Second International Conference, ICES 98 Lausanne, Switzerland.
- Flockton, S and Sheehan, K [1999] “A System for Intrinsic Evolution of Linear and Non-Linear Filters” in Proceedings of The First NASA/DoD Workshop on Evolvable Hardware, Pasadena, California,1999.
- Goldberg , D.E [1989] “Genetic Algorithms in Search, Optimisation and Machine Learning”. Addison-Wesley Publishing Company.
- Higuchi, Iwata and Weixin (Eds.)[1997] "Proc. 1st Int. Conf. on Evolvable Systems: From Biology to Hardware (ICES'96)" LNCS 1259. Springer-Verlag 1997

- Holland, J [1975] “Adaption in Natural and Artificial Systems”.The University of Michigan Press, Ann Arbor.
- Husbands, P[1991] “Genetic Algorithms in Optimisation and Adaption” The University of Sussex CSRP -??
- Miller, J [1999] “On the Filtering Properties of Evolved Gate Arrays”, in Proceedings of The First NASA/DoD Workshop on Evolvable Hardware, Pasadena, California, 1999.
- Miller, Thompson, Fogarty and Thomson (Eds.)[2000] "Proc. 3rdInt. Conf. on Evolvable Systems: From Biology to Hardware (ICES2000)" LNCS tba. Springer-Verlag 2000. In press.
- Motorola [1997] “Motorola Semiconductor Technical Summary. Introducing Motorola’s Field Programmable Analogue Array”, Motorola Inc. 1997. http://www.mot.com/SPS/
- Oldfield and Dorf [1995]"Field Programmable Gate Arrays: Reconfigurable logic for rapid prototyping and implementation of digital systems." Wiley 1995.
- Ozsvald, I [1998], “Short-Circuiting the Design Process: Evolutionary Algorithms for Circuit Design using Reconfigurable Analogue Hardware”, KBS MSc Thesis, September 1998, School of Cognitive and Computing Sciences, University of Sussex.
- Retallack, J ed. [1996] “‘Musicage’- John Cage in Conversation with Joan Retallack”,Wesleyan University Press.
- Schwefel, H and Rudolph G [1995], “Contemporary evolution strategies” in Advances in Artificial Life (Lecture Notes in Artificial Intelligence, vol. 929), F. Moran et al Eds. Berlin, Germany: Springer-Verlag.
- Sims, K. [1991]. “Artificial Evolution for Computer Graphics," in Computer Graphics (Siggraph '91 proceedings), Vol.25, No.4, July 1991.
- Sims, K [1994] “Evolving 3d Morphology and Behaviour by Competition” Artificial Life IV Proceedings, MIT Press.
- Sipper, Mange and Perez-Uribe (Eds.)[1998] "Proc. 2nd Int. Conf. on Evolvable Systems: From Biology to Hardware (ICES'98)" LNCS 1478. Springer-Verlag 1998.
- Spector, L., and Alpern, A [1994]. “Criticism, Culture, and the Automatic Generation of Artworks”. In Proceedings of the Twelfth National Conference on Artificial Intelligence, AAAI-94, Menlo Park, CA and Cambridge, MA: AAAI Press/The MIT Press.
- Spector, L., and Alpern, A. [1995]. “Induction and Recapitulation of Deep Musical Structure”. In Working Notes of the IJCAI-95 Workshop on Artificial Intelligence and Music.
- Thompson [1996]"Silicon Evolution" in "Genetic Programming 1996: Proc. 1st Annual Conf. (GP96)" J. R. Koza, D. E. Goldberg, D. B. Fogel and R. L.Riolo (Eds.), pp.444-452. Cambridge, MA: MIT Press, 1996.
- Thompson, A [1997] “Artificial Evolution in the Physical World”, in Evolutionary Robotics: From Intelligent Robots to Artificial Life (ER'97), T.Gomi (Ed.), pages 101-125. AAI Books, 1997.
- Thompson, Layzell and Zebulum[1999]"Explorations in Design Space:Unconventional electronics design through artificial evolution" IEEE Trans.Evol. Comp. 3(3), pp.167-196. 1999.

- Thompson and Layzell[1999]"Analysis of Unconventional Evolved Electronics" Communications of the ACM 42(4), pp.71-79. 1999.
- Thompson and Layzell [2000] "Evolution of Robustness in an Electronics Design" in J. Miller, A. Thompson, T. Fogarty and P. Thomson (Eds.) "Proc.3rd Int. Conf. on Evolvable Systems: From Biology to Hardware (ICES2000)" LNCS tba. Springer-Verlag 2000.
- Todd & Latham, W [1992] "Evolutionary Art and Computers" Academic Press
- Tomassini, M and Sipper, M [199?] "An introduction to Evolvable Hardware" http://www.evonet.polytechnique.fr/Coordinator/evonews/news3/ehard.htm
- Whitley, D et al [1997] "Island Model Genetic Algorithms and Linearly Separable Problems" in Proceedings of the AISB Workshop on Evolutionary Computation 1997.
- Woolf, S [1999] "Artificial Life and Art", in Vol.8 Proceedings of the 3rd World Multiconference on Systemics, Cybernetics and Informatics and the 5th International Conference on Information Systems Analysis and Synthesis, Orlando, Florida 1999.
- Xilinx. Inc [1996]. XC6200 Advanced product specification V1.0. June 1996. In The Programmable Logic Data Book. 1996. http://www.xilinix.com
- Zebulum R, Pacheco P, Vellasco M [1996] "Evolvable Hardware Systems: Taxonomy, Survey and Applications", in Proceedings of the First International Conference on Evolvable Systems, Tsukuba, Japan, 1996,pp. 344-358.
- ZETEX TRAC020 [1999] "Totally Re-configurable Analogue Circuit" -TRAC Issue 2- March 1999, Fast Analogue Solutions, Zetex group. http://www.fas.co.uk