Interaction, Simulation and Invention: a Model for Interactive Music

Jonathan Impett
School of Music, University of East Anglia, Norwich NR4 7TJ, UK
j.impett@uea.ac.uk

Abstract

This paper describes the incremental development of a model for interactive music – music instantiated in real-time on the basis of local performance and environmental information. Music is understood as a dynamical complex of interacting situated embodied behaviours. These behaviours may be physical or virtual, composed or emergent, or of a time scale such that they figure as constraints or constructs. All interact in the same space by a process of mutual modelling, redescription, and emergent restructuring. The model is implemented as a complex adaptive system in the *Swarm* simulation environment.

1.1 A View of Music

This paper presents a model for interactive music composition and performance – music instantiated during performance on the basis of stored programs and materials, and performance and environmental information. The model is predicated on a view of musical activity as a distributed, situated embodied phenomenon. Each event on the musical surface is the trace of a unique node in a tapestry of threads – structural, cultural, personal and technological. The balance of this distribution might be said to characterise a particular style, cultural context or kind of experience. These strands aggregate hierarchically into structures, materials, works and practices, adapting individually and forming the context for each others evolution. New constructs at any level are the emergent result of their interaction. Their tendency to interaction and self-organisation is the motor of musical activity and experience, and is further constrained by the more slowly evolving nature of human cognition and action, technology and cultural practice.

The phenomenological complexity of musical experience – not identical with its number of elements or “difficulty” – is seen as a critical parameter in all areas of musical activity. Its generation and dynamic maintaining is essential to the affordance of interaction, emergence and self-organisation, both internal to the musical construct and in the context of its composition, performance and reception. The behaviours which comprise this complex may be cultural, personal or “musical”. They unfold in continuously reforming dynamical hierarchies, the self-organisation of the whole adapting to maintain contextually critical values of energy, interaction and complexity.

The vital role of complexity in music has been observed by Meyer (1956) and Narmour (1977, 1992), both of whom point to prediction and hence modelling as crucial mechanisms in its generation and reception. However, the emergent phenomena they describe are inaccessible to reductive analysis or by tracing the evolution of decontextualised strands. Here it is proposed that the principle of predictive modelling obtains not only in the cognitive domain, but also in the structural, technological and cultural spaces through which these behaviours interact. In a dynamical context of multiple and changing alliances, behaviours adapt to the implicit models of the contexts they mutually generate. To facilitate this adaptation, we must also posit *reflexive* models, which allow for self-simulation and proliferation. It is proposed that this complex of perpetual mutual mediation is generated by a process of representational redescription, as described in the cognitive domain by Karmiloff-Smith (1992). By this means, constructs can cross domains – from, for example, the physically motoric to the formal – and move nonlinearly between levels of abstraction. Emergence is thus understood as an interactive phenomenon (Hendriks-Jansen 1996); a contextually construed, structurally significant act of redescription. The wide range across which the process of modelling, mediation and redescription operates allows for the multiplicity of levels of engagement which is characteristic of satisfactory musical experience of all types.

1.2 The Status of the Interactive Work

In the contemporary cultural context – a *tele-graphic* (Lyotard 1991) or *prosthetic* (Lury 1998) culture – the conventional modes in which these processes operate are changed in two significant ways. Firstly, the ranges of the spatial and temporal distribution of musical activity are vastly extended, such that causality and memory may appear nonlinear, subject-object relations are transformed (Baudrillard 1994), and previously discrete activities can be brought into new relationships. Secondly, in the absence of common cultural currency, "meaning" must be generated locally and provisionally. In the interactive work, the act and environment of performance can themselves become the materials. The system by means of which it is instantiated is therefore not an autonomous input/output processor, but an adaptive entity evolving in the context of an environment which it in turn transforms. By the same token, the process of composition is contiguous with those of performance and reception, as they participate in and shape the same narrative.

The internal and external dynamics of the interactive work must therefore be modelled in the same way; they exist in a single space. The work itself exists as a metastable entity. Various writers have pointed to simulation as the most appropriate contemporary cultural paradigm (Lyotard 1991, Baudrillard 1994, Simon 1996). Here it is proposed that an interactive work might simulate itself, adapting dynamically to each new context, and that its constituent elements – embodied (composed), emergent or real – evolve and interact by the same process.

1.3 Modelling Principles

The *complex adaptive system* as formulated by Gell-Mann (1994) and Holland (1995) corresponds to the understanding of musical activity expressed here, and in particular to the status and functionality internal and external of the interactive work. It is an emergence-generating complex displaying coherence in the face of change, comprising multiple behaviours which are heterogeneously situated in the environment and temporality of the whole. Individual behaviours, both composed and emergent, are embodied as dynamical systems. Port and Van Gelder (1995) put forward the strong hypothesis that cognitive and hence cultural processes *are* dynamical systems; any artefact which would engage with such processes must embody theories as to their functionality and relationship. Each behaviour or aggregate exists in a context the parameters of which may be equally physical or virtual. Such parameters are quantised, rather than represented symbolically. Physical and virtual parts thus coevolve as coupled dynamical systems (Langton 1989, Beer 1995a). Each dynamical system maintains its own temporal behaviour relative to changes in its environment, which may include its neighbours or itself, and its past trajectory and activity is implicit in its current state. There is therefore no need for planning or searching.

The interaction and aggregation of behaviours is afforded by their openness as systems and driven by their nonequilibrium nature (Kelso 1998, Prigogine and Boon 1998). To persist and to avoid infinite attractors they require energy exchange with their environment at every level within the work – a property consonant with intuitive understanding of musical creativity. Difference is thus both the engine and the trace of mediation. Each element in the system is both a musical object and a mediating processor of construal, information and energy, with the goal of generating an autopoietic whole.

2.1 Environment and functionality

The *Swarm* complex systems simulation environment was selected for the present project (Minar et al, 1996). As well as affording the necessary functionality and extensibility, *Swarm* is particularly adapted to implementing the dynamic internal reconfiguration characteristic of complex adaptive systems. By virtue of using the relatively unusual *Objective C* language, *Swarm* has a crucial property in addition to those of other object-oriented environments: that a generic placeholder can be used in the definition of one class to represent a reference to all possible classes. The specific class can then be decided at run-time, in context. Such a situation has clear advantages in a system which is intended to be allowed to develop freely and with which interaction might occur at any level.

The aggregate of objects is the *swarm* itself. Multiple swarms can exist, and objects can be members of more than one. Both swarms and objects can be created and destroyed and their relative memberships changed during the run of the simulation. Each swarm maintains a schedule of actions to be performed on or by its members at each system cycle (or at every *n* steps); the whole simulation is thus updated synchronously. In this sense, such an environment embodies the paradigms of both symbolic and parallel distributed computation. Swarms can themselves be members of other swarms; the conventional system architecture establishes an *observer Swarm* which runs, probes and displays those which perform the abstract simulation, with the coordination and any central management tasks handled by a single *model Swarm*. All parameters and variables are available for enquiry, and variables can be reset from any point in the simulation space. The reconfigurability and recursion of *Swarm's* architecture affords the dynamic self-simulation and hierarchies discussed above.

2.2 Real times

It is necessary to be explicit as to the temporal domain over which a system acts. In purely pragmatic terms, there is no continuum between “event” and “audio” rates of computation. Note or control events may be represented using the MIDI protocol, and in the best case might be accurate to millisecond precision. Audio rate performance requires the generation of at least 44,100 sample values per channel per second. Potential complexity and critical design decisions arise not from the generation of this amount of data, but from the keeping track of the discrete sound events it represents. In this case it was decided to focus entirely on the event level; the object of conventional compositional attention, and of perceived conscious activity in performance.

Swarm was run on a Silicon Graphics O2 computer, under *IRIX 6.5*. Information from the physical environment (human performance, environmental sensing or raw data) was received via MIDI. The chief sources were a STEIM *Sensorlab* “real world-to-MIDI” digitising interface, pitch-to-MIDI converters, and a Yamaha *Disklavier* MIDI-enabled piano. More extended experiments were conducted with the *meta-trumpet*, a conventional instrument fitted with a range of physical sensors (ultrasound, pressure, inclination, acceleration) and the sound of which is represented in the same space as MIDI performance data. The meta-trumpet was designed with the intention of exploiting both physical and sound performance information as compositional material, rather than adding extraneous techniques or arbitrarily substituting components of a rich instrumental practice, such that instrument, computer and composition are folded together (Impett 1994).

The external activity of the *Swarm* system is also expressed as MIDI. IRCAM’s *jMax* and Miller Puckette’s *Pd* were compiled and installed on the same computer, communicating with *Swarm* via internal virtual MIDI ports. Both programmes offer real-time MIDI and sound processing tools, a graphical interface and user extensibility. Basic MIDI functionality - input, output, memory and scheduling - is incorporated in a particular simulation as a sub-swarm *SwarmMIDI* created by the main observer swarm and inserted into its scheduler. Incoming MIDI is filtered for relevant streams of commands, and the values are stored in a blackboard structure, accessible to the whole system, from which agents can read any parameters which form part of their particular world.

As a generic simulation environment, *Swarm* has no claims to real-time performance. Most simulations serve precisely to compress or expand physical time to an experimentally observable scale. However, the principle of *situated* activity suggests that the temporal behaviour of a specific system in specific circumstances would be the functional reference, rather than an external, absolute timing grid. As with any other musical artefact, the system *as it performs* is the source of mediation and interaction. The future of a dynamical system can be understood as a function of its present state and its future conditions; each behaviour is located in the real time of its own world.

The issue of scheduling resurfaces, however, in the capacity of a system for interaction, in actions which imply later actions (such as turning a note off), and in its internal modelling of the temporal dynamics of another system – a human performer, for example. An “absolute” time reference, and some means of knowing when a particular period has elapsed, are necessary to establishing aspects of behaviour which are conditioned by rhythm, ballistics or abstract periodicity “perceived” in the outside world. A scheduler was therefore implemented within *SwarmMIDI*.

2.3 Behaviours as Agents

Each primitive of musical behaviour is implemented as an agent derived from the generic class *SwarmObject*. This allows for the establishment of its *world* – a set of parameters which may represent some aspect of the physical world, be generated by other agents within the system, or be some value calculated at a higher level which reflects an aggregate system property. A field of activity is similarly defined, as well as internal parameters including the energy level of this agent, and the definition of its behaviour – generally as a dynamical system. A set of flags indicate the disposition of the agent to form alliances or to be aggregated into hierarchies with other agents, according to the principles established in (Holland 1995).

3.1 Initial Models

As a test of the generalisability of this architecture, initial experiments implement interactive music models from other sources. Dynamical systems have frequently been identified as an appropriate means of generating music by virtue of the potential richness of their behaviour, the effectiveness with which they can be controlled, and the simplicity of their definition. In the first experiment a single agent – *DynSys* – is run by the modelSwarm, and embodies the dynamical system in question as an iterative process. The environment in which it runs is managed by the *modelSwarm*, and consists in this case of any constant values and any substitute for the previous output value of the system, updated by new pitch input. Constant values and iteration frequency are controlled from the meta-trumpet by movement in two-dimensional space or as a function of note inter-onset time and attack density. The self-similar structures of *fractal* processes could also be implemented in this model, passing parameters from *modelSwarm* in the same way.

The succinct formulation and potential for universal computation of cellular automata (CA) has led several researchers to explore their musical application, including the navigation through their behavioural space by genetic algorithm. The model *CA1d* uses a population of sixty CA existing on a one-dimensional grid. It allows for the random seeding of initial states to a given density. Further pitch input then resets the CA states; any pitches received between update times will set to 1 the CA onto which they are mapped on the next cycle. The model plays all positive-state sites, or only those newly activated, according to a software switch. The rule can be changed during run-time by means of a graphical probe.

3.2 Evaluation

The Swarm-based architecture proposed here is able to reproduce the essential functionality of models which map musical input onto the environment of dynamical systems, whether simple or complex, and which map the system behaviour back onto musical space. There are no formal constraints on the complexity of the dynamical system, the number of agents, or the richness of the embedding of the model in real-world physical events. *DynSys3* maintains an independent dynamical system for *each* parameter of its behaviour, avoiding the parameter-locking of the simpler models. The relationship of its activity with performance input is also less linear: it now occupies a space shaped entirely by the performance, without tracking any aspects of that activity.

However controlled or complexified, the system behaviour remains monolithic. Any emergent structure is the result of the particular dynamical system embodied by the user, and whilst it may be put to interesting use, it has no necessary relationship with its context. An unimaginably fast machine might search the vast space of possibilities for a dynamical system to match its environment, but this would be impersonation, not interaction. The potential relationships are those of *mapping*. However many strands or parameters, any interaction remains unidirectional and one-dimensional. The model is passive; it can be depicted but has no voice. The aspirations of adaptive, dynamically hierarchical system behaviour, and of multi-level engagement with the musician are not yet met: there is no emergent *supplement* from within the model.

4.1 Aggregation, Hierarchy and Emergence: Sharing the Same World

Whilst the architecture of *CA1d* embodies multiple parallel dynamical systems and generates situated, nonlinear behaviour of an arbitrary degree of complexity, it does not afford the multiplicity of levels of *engagement* identified above. The dynamical systems evolve in a homogeneous space, each with an identical situatedness, at the same rate. Nonequilibrium must be embodied at the highest structural level, as well as the lowest, if the system is to be able to self-organise with its environment.

We therefore return to the autonomous dynamical system *DynSys* as the basis for a complex structure in which many such behaviours can interact individually with each other and their environment, at different rates and at different levels. Each parameter of this elemental musical behaviour – pitch, volume and time – is now the trace of an independent internal dynamical system. However, the range of possible components of its environment is extended to include the behaviour itself, so parameter-locking or internal feedback loops are possible. The pitch values are mapped linearly, volume and time exponentially with variable ranges, curves and offsets.

If these worlds of these behaviours are to be non-exclusive and dynamical – they may overlap and change – they can no longer be structurally identical with their environment. The whole set of possible physical and virtual parameters is therefore embodied in another *Swarm* object - *World* - from which all are potentially available to all behaviours including, by the *probe* protocol, the *observer.* The dimensionality of *World* is determined by the number of virtual and physical individuals and the number of parameters each presents to the others. Two additional values represent a qualifying index, such as a salience parameter, and an identifying number so that behaviours may know whether they have already taken account of a value without the need for individual storage. Each *CDynSy*s is situated in its own sub-world at initialisation. This *environment* may incorporate virtual and physical parameters, and feedback loops. As the *world* itself now exists independently, these allegiances may overlap or change arbitrarily. At each simulation step, behaviours update their internal parameters in accordance with the present state of their world. Having calculated their next state and performed any actions, they then post their parameters to the *world* for access by other behaviours.

4.2 Energy

As an indicator of the time-varying state of a behaviour, each *CDynSys* maintains an *energy* variable representing the degree of change over time in its own sub-world. This is in effect a simple internal model of the dynamics in which the behaviour finds itself involved; not, however, of its own behaviour, of its own construal of the world, or of the success of either. In its crudest form, this variable can be used to deactivate an individual when it falls below a threshold, such that the activity of the whole reflects the density of interaction of its component parts. Its calculation is implemented as a time-variant filter, so that previous activity is embodied in the current state.

4.3 Coupling and Bifurcation

In the model *CDS*, eight *CdynSys* agents of the type described above are situated in a *World* comprising themselves and a human performer. Each of these contributes four parameters to the *World*. The virtual *CDynSys* behaviours each embody the same dynamical system with the following parameters: current value (mapped to pitch), a velocity value, constant *mu*, and a time value *metro*. The time between subsequent state calculations of a particular *CDynSys* behaviour is determined by its *metro* value, which may change with the new state like any other parameter. Real time is therefore read at each simulation step to maintain timing accuracy. In this way behaviours which are otherwise situated in the same part of the world evolve at different rates, and because of the real time evolution of their environment plot different trajectories through their phase space. The four parameters of the behaviours own *environment* are read from its unique set of four locations in the world before the next state is calculated, and written to the *world* after. Physical performance is also described in terms of four parameters: with the keyboard - key, velocity, inter-onset time and a control wheel; with the meta-trumpet – pitch, volume, and position in two-dimensional space.

For experimental purposes, *CDynSys1* is locked to the physical parameters of the performance. The environments of the others are distributed randomly at initialisation in the total world of thirty-six parameters. The relative dynamics in the energy of each behaviour in an arbitrary run shows that they exhibit the types of interaction and coordination described by Kelso. Temporary parallel or phase-locked development is indicative of the coupling of critical parameters of two or more behaviours, such that the "free" parameters are brought into stability. Certain behaviours find themselves outside any stable structure for the duration of the run. The coupling of behaviours may cease when one enters a critical or unstable state, or the common critical parameter enters a new area of its phase space. Critical instability can be seen before such bifurcations. *Multistability* is observable where the dynamics of a single behaviour switch allegiance periodically.

A striking and positive property of this model is the nonlinearity of the relationship between agents: specifically, in that between the performer and the system as a whole. At certain moments, a particular performance parameter or sequence of states – i.e. some quasi-periodic behaviour – appears to "control" critical parameters of the entire virtual model, either as a whole or in some crucial respect. Even when this situation appears temporarily stable, it may change abruptly as these critical relationships are transformed in the evolution of the internal dynamics of the system.

Three of our theoretical design aspirations are addressed in *CDS*: the lines of engagement within the "composition" are *emergent* – that is, they cannot be reductively associated with the structures in the abstract, independently of their situated interaction – they appear at *multiple levels* in terms of the kind, degree and permanence of the interaction afforded; and there is *energy exchange* both internally and with the environment. These internal interfaces are nowhere written into the structure of the model, and pass through the *physical* environment as the performer is inevitably drawn in to search for them and test their limits, complexity and elasticity – their *thickness.* As emergent phenomena of varying dimensionality, these interfaces in the dynamics of the system do not lend themselves to simple quantification or analysis. The crude measure of *energy* derived in this model is a collective variable in Kelso's sense (Kelso 1995: 44). It gives an indication of the dynamical relationships, and these parameters could be incorporated in the *World* so as to become potential elements of each others environments; an implicit form of internal modelling. There is emergence in the dynamics and behaviour of the system and in the interaction it affords but not, as yet, in its internal structure.

5 The Autonomy of Aggregation

If the aggregations of behaviour are to be dynamical, arbitrarily multiple and potentially overlapping, and if they themselves are to take on some of the properties of a behaviour to allow their *hierarchical* evolution, they must have some degree of structural autonomy within the wider system. The properties of the *Swarm* environment afford a dynamical solution: that the simulation itself observes or is alerted to negotiation between behaviours and creates a new temporary aggregate structure corresponding to the emergence or potential for emergence that has been identified.

5.1 Instability

As an indicator of the disposition of a behaviour to some kind of cooperation as a component of wider self-organisation, we adopt the parameter of *instability* identified by Kelso as the critical property in such circumstances (Kelso 1995: 45). This value is derived individually for each parameter of each behaviour, using a *linear prediction* algorithm to generate an index of prediction error (Press et al. 1988: 568). It is, in effect, an inverse index of linearity. Experiments were also made using *autocorrelation*, following (Tanguiane 1993), but proved too costly to implement in this number in real time. Only the highest degrees of correlation/lowest prediction errors proved to be intuitively useful, because of the high potential dimensionality of more complex patterning.

Linear prediction is applied to the maximum depth of a circular buffer, now added for each parameter. Values are combined to derive a representative index for the behaviour as a whole. Correlation of energy and instability values indicates at each point which behaviours have moved into a stable attractor, and which into chaotic regions. Of most interest are the intermediate states which can be used to harness the “free” parameters of other behaviours, either to release them from a locked attractor or to draw them towards some degree of linearity. At some points, the total instability is clearly correlated with the energy of a behaviour, at others an increase in energy marks the move into a region of stability or instability. In general, instability diminishes over the course of a run, as parameters find their way into stable regions of phase space.

5.2 Reconfiguration

By relating the indices of energy and instability we can exploit the nonequilibrium behaviour of *DynSys* and provide it with a means of energy exchange with its environment, as in the self-organising systems described by Prigogine and the mechanism of synergetic coupling proposed by Kelso (Prigogine 1980, Prigogine and Boon 1998, Kelso 1995). By this means, the internal shape of the model can evolve autonomously, but conditioned by changes in its environment, as the connections between constituent elements are continuously reformed. Environment and simulation are transforming *each other*, and therefore both the world and the behaviours inhabiting it need to be able to change structurally.

The *DynSys* behaviours are now allowed to reconfigure their individual world, one parameter at a time, if the instability of that parameter falls below a given threshold. We might force additional change to keep the instability within a set of bounds. The overall instability of the system can thus be maintained at some desired critical level, rather than diminishing as the system finds its way into stable regions from which it is difficult to emerge. Certain behaviours enter a region of phase space which results in continuous reconfiguration, despite maintaining a level of instability barely above the threshold. Others reconfigure rarely, but have a higher general instability. The degree of emergent restructuring is highly dependent on individual conditions and the characteristics of particular regions of phase space.

5.3 The MultiBehaviour

The next model – *HCDS* - adds functionality at the *ModelSwarm* level. Instances of a new Swarm object – *Mube* - can now be created by the simulation; it can both "see" the World as a whole and negotiate bilaterally with individual behaviours. Its own behaviour is inserted into the simulation schedule. The use of such a *template* for emergent structures is suggested by Bonabeau et al. (1999) as being an efficient strategy with roots in natural behaviour. As a piece of embodied scaffolding, it corresponds to a mature stage of the emergence of *building blocks* proposed by Holland (1995: 166). The low-level behaviours are now silent. Together with performance information, they provide the context for the emergence of a *Mube*, and shape the world in which it acts and evolves. The simulation layer of the model keeps check of the energy and instability of each behaviour. If a *DynSys* behaviour reaches a variable threshold of instability such that it affords emergence, and has sufficient energy to drive this event, a new *Mube* is activated. This new *audible* construct adopts the environment – the local sub-world - of its instigating behaviour. Like the silent behaviours, the *Mube* monitors its own energy and instability, and continuously reconfigures its environment to maintain critical levels. Now, however, the global *World* is extended by the behaviour of the *Mube* itself, opening up new spaces for all of the agents to explore, *Mubes* and behaviours alike. The trajectory of the *Mube* through its phase space is different to that of the continuing instigating behaviour because of a *transfer of energy* that takes place from the latter to the new construct as it is initiated. A *Mube* remains active as long as it can maintain the necessary level of energy in its environment. Its participation in the *World* allows the same process to happen recursively: as a *Mube* ceases to exist, any construct in whose environment it had figured looks for other sources of energy and instability in the usual manner.

One might imagine a far more complex set of mechanisms of exchange, collaboration and reproduction in the generating of emergent hierarchical structures. These mechanisms could be extended arbitrarily and would rapidly become issues of the embodiment of *style*. It was therefore decided to implement the simplest possible mechanism for the affording of emergence, the transfer of energy and the asymmetrical sharing of worlds.

With *HCDS*, performance with the model is becoming a relationship of interaction rather than control. Human and virtual behaviours search their common space for regions of potential activity to sustain the joint system; as they do so, both space and system are transformed. The spaces are *temporal* as well as *parametric*: the traces of quasi-linear change one inevitably construes as intentional suggest time spans over which they may develop. The extent to which the reaction of the model to new events is predictable is itself an emergent and nonlinear property. It affords interaction, but the level on which engagement appears to take effect is a function of that interaction.

HCDS2 complexifies this situation. As a *Mube* is initiated, it not only takes on the properties of its instigating behaviour as a dynamical system, but also creates a further set of sub-behaviours (*subDS*) which adopt the characteristics of all other behaviours which fall into a wider band of coupling-affordance at that moment. The audible behaviour of these *subDS* agents is transformed by that of the *Mube* to which they are attached, such that they leave the trajectories of the behaviours from which they were cloned. Their speed, pitch, volume and dynamical system variables are modulated by those of the *Mube*. In their own efforts to maintain energy and instability, therefore, their environment departs from that of their initial model. In the last two models, a greater level of management of the overall system behaviour could be achieved by dynamically determining the thresholds of instability and energy which govern structural emergence, in the area of *World* which represents the physical performance.

6 Indices of Emergence: *Invention*

The model described thus far is emergent in its structure, behaviour, and in the interaction it affords. The *redescription* happens implicitly, in the mutual reshaping of *World*. However, the internal relationships or those of replication rather than *simulation*: there is as yet no internal state change by virtue of an act of construal. The very possibility of interaction implies an act of construal or interpretation (Hendriks-Jansen 1996; Bogdan 1997), itself predicated on the *predictive* processes of modelling or simulation. If the model itself is to make dynamic decisions as to the *fact* of emergence, it requires some further index by which this may be measured.

The predictive modelling of nonlinear behaviours is not a trivial question, even when the number of parameters has been reduced for experiment, as here. This is particularly the case when, as in this model, the event trace – the musical surface – is itself the product of the interaction of an unknown number of agents and constraints. This is not an artificial issue: the suggestion here is that this is precisely the case in musical engagement of all kinds. Just as their causality is distributed, the temporal spans in which a particular event participates are multiple and arbitrarily wide. This event trace constitutes a nonlinear time-series, except that the events are unevenly spaced, and there is no unambiguous answer as to the continued contribution of each after the moment of its arrival. For this reason, conventional methods of nonlinear time series analysis expand exponentially to deal with the additional complexities.

After experiments with state-space reconstruction (Kantz and Schreiber 1997) suggested its impracticality in this real-time context, a connectionist approach to issue of behaviour description was adopted. The simple recurrent networks described in (Elman et al, 1996) embody a form of non-representational memory of the evolution of the system, by feeding the current state back with new input; a phase-space filter, in effect. Beer (1995b) has shown that the simplest two-node continuous recurrent neural network (CRNN) is theoretically capable of modelling any nonlinear behaviour. The connection weights of 2-node CRNNs as described in (Haykin 1999: 752) were evolved by genetic algorithm (Mitchell 1996: 76), externally to the *Swarm* model.

Accounts of emergence tend to not lend themselves readily to a computational context; they are judgements based on the contextual construal of a particular situation. Crutchfield (1994) has proposed the incorporation of this property in a computational implementation, in a definition which embodies the principles that have been developed here. Emergence, he suggests, is an interpretative act of redescription. A working description of an object behaviour is qualified and complexified *to the limit of the capability of the describing agent*. At that point, this description is no longer instrumentally useful; the object behaviour must be *redescribed*, integrating as much of the accumulated understanding as seems appropriate, and allowing further room for elaboration of the new model. Emergence then becomes situated, enacted knowledge.

In an attempt to embody this principle, the CRNN modelling agents bred externally were incorporated in a new version of the *Mube* from *HCDS2*: *Invention*. Each of the potential emergent hierarchical structures now contains two CRNNs. One attempts to model the environment, the other the *invention's* own behaviour. At each simulation step, each is clamped to its part of the *World*, and allowed to stabilise its response. The invention now has a complex time behaviour – a function of three systems of greater or lesser periodicity. Both CRNNs calculate indices of *error* - a time-filtered function of the difference between their predictions and the actual successive states of inner and outer environments. If this error reaches a threshold value, redescription is forced and new values are chosen for the CRNN. The maintenance of this error has a *cost* for the invention, and conditions the expression of its own behaviour and those coupled to it. The performance of *Invention* now combines the windows of linearity of *CDS* with the nonlinear emergence of *HCDS2*. The system generates internal horizons for its own construal of internal and external behaviours. Tension is generated in the attempt to integrate linear description with nonlinear behaviour, and modulates both. Figure 1 illustrates the system architecture at this stage of its evolution.

Figure 1: System architecture

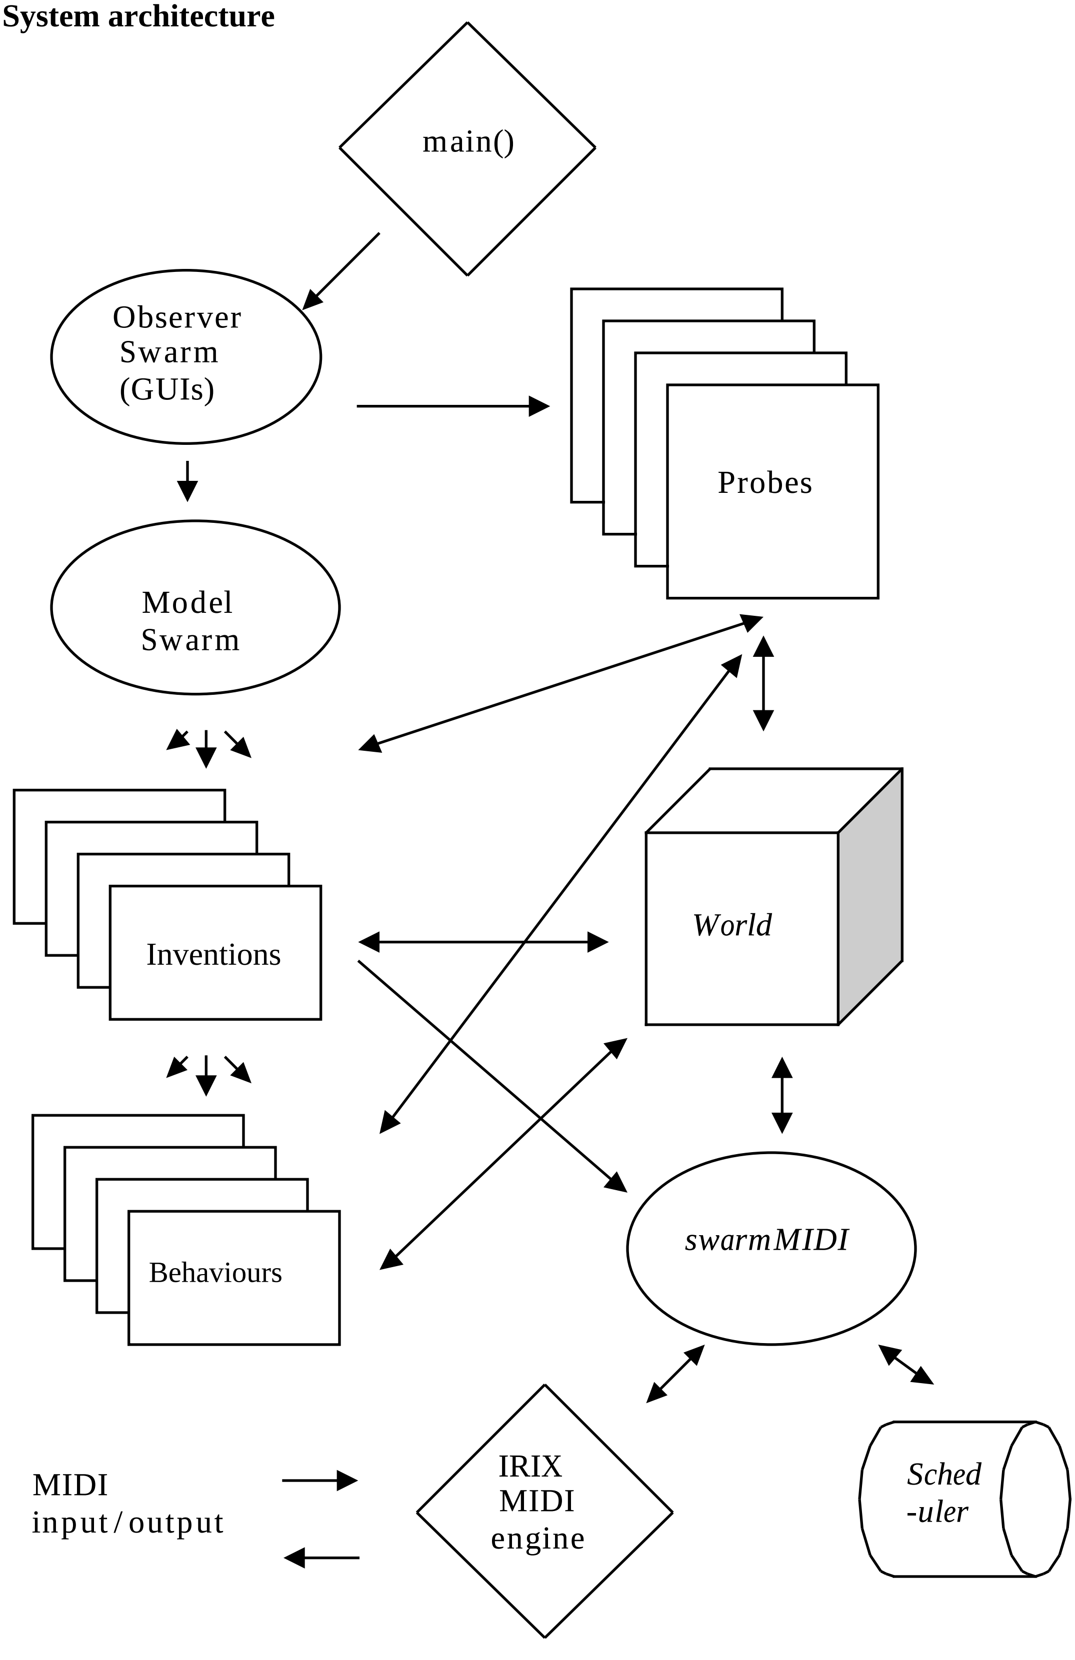

7 Conclusion

The invention is therefore proposed as a unit or level for the compositional engagement with an interactive work. The elements of an *interactive invention* might include melodic behaviours inferred from aspects of a live musician's performance, harmonic or stylistic constraints evolved over many performances, or architectonic structures which build themselves contextually on the basis of a composer's rules. Equally, the driving parameters might be environmental or historical - some representation of musical material formed at an earlier moment. If the unit is too small - at too low a level - the complexity is intractable and its existence imperceptible. And yet this construct has no unique face, no prime form. It exists throughout a work, but only becomes perceptible at points of interaction with other dynamics. It cannot therefore be identical with the work itself, or even sections of a work. The least we can say is that it has behavioural characteristics in *time* - a duration, a periodicity - and has multiple modes of connection to other systems and behaviours, whether cultural, environmental, performance, or note-to-note. Several must be maintained in parallel; they are reconstituted in the course of their interactions; and their relationship to one another is dynamically hierarchical. In the case of interactive music, the blurring of the boundaries between composition and performance, work and environment, is an essential characteristic. It could even be considered, as we have seen, the material itself. The invention is a nexus for the *re*-distribution of musical activity. Inventions might be considered local phase transitions in cultural dynamics, emerging into autonomy from the interaction of those dynamics; an autonomy characterised by the capacity for self-simulation. Situated in the environment of its own creation, the interactive work differentiates itself from the infinitely-extending threads of the fabric of that environment by a process of self-organisation.

The model described here addresses some of the theoretical issues identified as significant in an emerging practice of interactive music. The experimental context requires that such a model generate behaviour which affords tracing and analysis. Nonetheless, this system demonstrates some of the emergence, dynamical restructuring and multiple levels of engagement which were discussed above. In a purely creative context, these properties could be complexified, constrained and mediated arbitrarily. The activities of model development, composition and performance could be folded into each other in a more organic process of development. In a conventional musical context, the work could be seen as an interface between “inner” and “outer” musical worlds which is effectively transparent – a figure drawn on one side may be construed as such on the other. In an interactive work – or arguably any contemporary work - this interface has a dynamical *thickness*; it must be able to adapt itself to its context. This model proposes a means for such adaptation.

References

Baudrillard, J. (1994). *Simulacraa and Simulation* (trans. S.F.Glaser). Michigan: University of Michigan Press.

Beer, R. D. (1995a). Computational and dynamical languages for autonomous agents. In R. Port and T. van Gelder (Eds.) *Mind as Motion*. Cambridge MA: MIT Press.

Beer, R. D. (1995b). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behaviour* 3 (4): 469-509.

Bogdan, R. (1997). *Interpreting Minds: The Evolution of a Practice.* Cambridge MA: MIT Press.

Bonabeau, E., M. Dorigo and G. Theraulaz (1999). *Swarm Intelligence: From Natural to Artificial Systems*. New York: Oxford University Press.

Crutchfield, J. (1994). The calculi of emergence: computation, dynamics and induction. *Physica D* 75: 11-54.

Elman, J. L., E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi and K. Plunkett (1996). *Rethinking Innateness: A Connectionist Perspective on Development*. Cambridge MA: MIT Press.

Gell-Mann, M. (1994). Complex adaptive systems. In G. Cowan, D. Pines and D. Meltzer (Eds*.)* *Complexity.* Reading MA: Addison Wesley/Santa Fe Institute.
Haykin, S. (1999). *Neural Networks: A Comprehensive Foundation* (Second Edition). Upper Saddle River NJ: Prentice Hall.
Hendriks-Jansen, H. (1996). *Catching Ourselves in the Act*. Cambridge MA: MIT Press.
Holland, J. (1995). *Hidden Order: How Adaptation Builds Complexity*. Reading MA: Addison Wesley.
Impett, J. (1994). A meta-trumpet(-er). In *Proceedings of the 1994 ICMC*, San Francisco: ICMA.
Kantz, H. and T. Schreiber (1997). *Nonlinear time series analysis*. Cambridge: Cambridge University Press.
Karmiloff-Smith, A. (1992). *Beyond Modularity: a Developmental Perspective on Cognitive Science*. Cambridge MA: MIT Press.
Kelso, J. A. Scott (1995)*. Dynamic Patterns: The Self-Organisation of Brain and Behaviour*. Cambridge MA: MIT Press.
Langton, C. (1989). Artificial life. In C. Langton (Ed.) *Artificial Life*: 1-48. Reading MA: Addison Wesley/Santa Fe Institute.
Lury, C. (1998). *Prosthetic Culture: Photography, memory and identity*. London: Routledge.
Lyotard, Jean-François (1991). *The Inhuman* (trans. G. Bennington). Cambridge: Polity Press.
Meyer, L. B. (1956). *Emotion and Meaning in Music.* Chicago: University of Chicago Press.
Minar, N., R. Burckhart, C. Langton and V. Askenasi (1996). *The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations.* Available at: http://www.santafe.edu/projects/swarm.
Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. Cambridge MA: MIT Press.
Narmour, E. (1977). *Beyond Schenkerism: The Need for Alternatives in Music Analysis.* Chicago: University of Chicago Press.
Narmour, E. (1992). *The Analysis and Cognition of Melodic Complexity: The Implication-Realization Model*. Chicago: University of Chicago Press.
Press, W. H., S. A. Teukolsky, W. T. Vetterling and B. P. Flannery (1988). *Numerical Recipies in C: The Art of Scientific Computing* (Second Edition). Cambridge: Cambridge University Press.
Prigogine, I. (1980). *From Being to Becoming: Time and Complexity in the Physical Sciences*. New York: Freeman.
Prigogine, I. and J-P. Boon (1998). Le temps dans la forme musicale. In E.Darbellay (Ed.)*Le temps et la forme*. Geneva: Editions de l'Université.
Simon, H. (1996). *The Sciences of the Artificial* (Third Edition). Cambridge MA: MIT Press.
Tanguiane, A. (1993). *Artificial Perception and Music Recognition*. Berlin: Springer.
Van Gelder, T. and Port, R. F. (1995). *Mind as Motion.* Cambridge MA: MIT Press.