



What's Emergent in Emergent Computing?

Klaus A. Brunner

Institute of Design and Technology Assessment

Vienna University of Technology

Favoritenstr. 9–11, A 1040 Vienna

Austria

email: k.brunner@acm.org

Abstract

Emergent Computing (EC) has become a widespread concept in computer science, particularly in AI und A-Life. In this paper, the relation of EC to the philosophy of emergentism, as developed by Alexander, Broad, Morgan and used in modern theories of self-organisation, is discussed using examples from the theory of cellular automata. Additionally, fundamental issues of emergentism and (in-)determinism in relation to EC are brought up for discussion.

1 Motivation and Goals

The goal of this paper is to describe ongoing work to examine the notion of “emergent computing” in relation to the concept of emergence as known from emergentist philosophy, and to show some issues that came up in the course of that work. The relevance is seen in the fact that “mind” and “consciousness” are often considered emergent properties of the physical world (thereby touching on central issues of AI and AL), and that emergence can also be seen as one of the cornerstones in the understanding of “information” [Hofkirchner 2001].

Questions to be examined include: what does the concept “emergent computing” mean considering philosophical ideas of emergence; and is it a well-defined, commonly agreed on concept? More fundamentally, in which ways can “emergence” occur in computing, if at all?

2 The Concept of Emergence in Philosophy

2.1 Roots and Proponents

Emergentistic thinking in the modern sense was developed primarily by Samuel Alexander, Charles Dunbar Broad and Conwy Lloyd Morgan, who centred a great deal of their scientific work around concepts of emergence, and in the 1920s built up what came to be known as American and British Emergentism [Stephan 1999].

Emergentism established a middle ground between the extremes of mechanism (as exemplified in the idea

of the Laplacean daemon) and vitalism, introducing a degree of indeterminism while remaining a fully naturalistic philosophy by rejecting supernatural influence. Today, emergentistic thinking can be found in the areas of cybernetics and systems theory. It is also commonly associated with the modern concept of self-organisation [Krohn and Küppers 1992].

2.2 Principles

In an exhaustive survey of emergentism, Stephan [1999] identifies a number of definitive properties of emergentistic theories. Common to all these theories is the notion of systemic properties: these are properties that the system has, but none of its parts have. In Stephan's taxonomy, this is one of the minimum requirements of “weak emergentism”. Additionally, “strongly emergentistic” theories postulate the existence of systemic properties that are irreducible, unpredictable, or both:

- irreducibility: a systemic property is considered irreducible if it cannot be deduced from the properties of the system's constituents. This is emergence in the analytical or synchronous sense: the system's emergent properties cannot be (fully) explained even with complete knowledge of the structural arrangement and the individual properties of its parts. This is a more formal definition of the idea that “the whole can be more than its parts”.
- unpredictability: a systemic property is considered unpredictable if it is in principle (that is, even with “perfect knowledge”) impossible to predict its appearance in the future. It follows that an irreducible systemic property is unpredictable if it has not been instantiated before, or if the structure exhibiting said systemic property is in principle unpredictable. This is emergence in the diachronous or evolutionary sense, as it emphasises the idea that the properties of new structures are not fully foreseeable (the epistemic interpretation) or not fully determined (the ontological interpretation) before their realisation.

Emergentist theories often employ a system of levels, such as physical, chemical and biological levels. However, unlike reductionists, emergentists con-

sider “higher” levels to be more than just macro-descriptions of the level(s) beneath them, as can be deduced from the first item in the above list: irreducible properties exist on higher levels of reality, not only of description. The emergentist concept of downward causation expresses the idea that higher levels can have a causal influence on lower levels. This idea comes in different varieties, as pointed out by Emmeche et al. [2000] in their distinction between weak, medium and strong downward causation.

3 The Concept of Emergent Computing

3.1 A Few Definitions

EC is generally characterised by the interaction of relatively simple entities, forming a system that as a whole is said to exhibit emergent properties. EC systems are decentralised (or, “bottom-up”) in the sense that there is usually no central point of control governing the entire system’s behaviour. Examples typically cited as EC applications include cellular automata, neural networks, genetic algorithms, and agent-based systems.

Forrest [1991] draws a distinction between emergent properties in general, and emergent properties that are considered computations (“[a] natural interpretation of the ephenomena as computations”). Only the latter are considered “emergent computations”. Notably, Forrest’s definition of EC rules out irreducibility and downward causation, as she clearly designates emergent computations as epiphenomena. This, however, appears to be in contradiction with the statement that “[g]lobal patterns may influence the behavior of the lower-level local instructions, that is, there may be feedback between the levels.” [Forrest 1991, p. 2]

In a different interpretation of Forrest’s text, Hordijk [1999] understands emergent properties as functionality of, and indeed for, the system:

Emergent pattern formation in decentralized spatially extended systems often entails an important functionality for the system as a whole. In other words, the emergent patterns give rise to some form of globally coordinated behavior, or global information processing, which is used by the system to sustain itself or make certain decisions. . . . This global information processing in decentralized spatially extended systems, mediated by emergent pattern formation, is known as emergent computation . . . [Hordijk 1999]

Remarkably, Hordijk claims that the system’s own emergent behaviour may be used to sustain that system itself, a common theme in theories of self-organisation¹.

Holland [1998] appears to focus on the unpredictability aspect of what he calls emergent properties,

¹cf. the concept of autopoiesis [Maturana and Varela 1987], which emphasises the constant self-creation of autonomous systems

citing the vast number of distinct states that complex systems can exhibit through myriad possibilities for interaction. However, at least in the context of computer based systems, this is a notion of practical unpredictability, and it is argued that prediction is always possible to some extent with appropriate models of the system:

[B]y attending to selected details, we can usually extract recurring patterns. When these recurring patterns are regularly associated with events of interest, we call them emergent properties. [Holland 1998]

Holland provides a valuable framework for the modeling of the building blocks of typical EC systems with his concept of constrained generating procedures (cgps). A direct, concise description of the term “emergence” can not be found in Holland’s “Emergence” as the author himself acknowledges explicitly. However, the relation (or rather, non-relation) of his ideas to emergentism becomes clearer when the cgp model is applied to various EC methods, as I shall point out below.

Cariani [1991] puts “emergence” to the eye, and subsequently, to the cognition and actions of the beholder by creating a concept of relative emergence: “The emergence-relative-to-a-model view sees emergence as the deviation of the behavior of a physical system from an observer’s model of it”, concluding that “the interesting emergent events that involve artificial life simulations reside not in the simulations themselves, but in the ways that they change the way we think and interact with the world”. This view of emergence is completely contrary to even weak emergentism, as it leaves out and actually denies emergence as an ontological concept in computer systems.

3.2 Cellular Automata as an Example of EC

Cellular automata can be considered a prime example of typical EC setups: a potentially large number of extremely simple elements (“cells”) are interconnected to form an array. Each cell’s state is fully determined by the states of its adjacent cells, and updated repeatedly. The strengths of CAs lie in their relative simplicity, ease of implementation on computers, and the ease of representing them graphically in a simple, yet impressive manner. John Conway’s “Game of Life” CA is frequently quoted as an example:

These computations [in CAs] may or may not be emergent. For example, the soliton-like computation in CAs is explicitly constructed by using different CA cell-states to encode the absence or presence of, and the interactions between, different kinds of soliton-like particles with which the computation is performed. Thus, this computation would not be considered emergent. In the ‘Game of Life’ emergent structures called gliders and glider guns are used to create explicitly a particular initial configuration such that the CA dynamics mimics computation with logical

gates. This seems to be somewhere in between emergent and explicitly programmed. The intrinsic computation embedded in a CA's dynamics, however, appears to be truly emergent. [Hordijk 1999]

It is notable that Hordijk draws a distinction between "emergent" and "explicitly programmed". That distinction appears to be based on the designer's intention: while certain effects are intended (such as the computation of a single cell's state, explicitly formulated in some programming language), others are not ("the intrinsic computation in a CA's dynamics"), although they may prove useful in some way. Yet, those effects are completely determined by the explicit programming at the bottom level of the system. They are not irreducible, and therefore completely epiphenomenal.

However, it can be argued that the so-called truly emergent computations are unpredictable in the sense that accurate prediction would require exactly the same calculations that are executed when the system is run: thereby, predicting the system's emergent runtime behaviour is functionally identical to its first execution. Prediction using a simplified model of the system could result in drastically different behaviour, a situation similar to models of deterministic chaos. In that epistemic sense, the system is actually showing emergence in the evolutionary or diachronous meaning of the word.

Holland's [1998] application of his cgp model to cellular automata goes to show that his notion of emergence is a purely descriptive one, contrary to emergentist thinking:

If we turn reductionism on its head, we add levels to a basic description. More carefully, we add new laws that satisfy the constraints imposed by laws already in place. ... these new laws apply to complex phenomena that are consequences of the original laws; they are at a new level. [Holland 1998, p. 190]

New laws satisfying the constraints of laws already in place are reducible to laws (axioms) already in place: a new law is not "new" in the sense of emergentism, as it is neither irreducible nor unpredictable.

In the cgp framework, a new level is derived from an aggregation of existing cgps, forming a new meta-cgp. Applied to a two-dimensional cgp such as Conway's Life, the automaton may be viewed as an array of tiles made up of 3x3 arrays of cells, as opposed to the original view of an array of single cells. In Holland's words: "There is a one-to-one correspondence between the behaviors of the two arrays, though the behavior of the tiled automaton *appears* to be much more complex." (emphasis added)

Neither irreducibility nor unpredictability (in principle) are issues in Holland's cgp example, and it follows from the above citations that the notion of levels clearly refers to levels of description, not of ontology.

3.3 Other Examples

In this paper, Cellular Automata serve as the prime example of EC systems. As mentioned before, EC spans a number of different approaches from the fields of Artificial Intelligence and A-Life. Neural networks can be considered as the next step from CAs: they can be modeled with a richer, more complex set of interaction parameters that are themselves variables: such as time-varying thresholds, changes in synapse weights, fatigue, and so on [Holland 1998].

The complexity of the system's components, and the variability of the coupling between them become even greater with agent-based systems. The general observations found in the examination of CAs should, however, apply to these more complex systems as well.

4 Some Issues in EC versus Emergentism

4.1 The Role of Indeterminism in Emergence

Is emergence (in the strong, emergentistic sense) possible in computer systems such as CAs, where indeterminism is generally absent? Fleissner and Hofkirchner describe the situation in systems undergoing self-organisation as follows:

Inputs and outputs are not related in a way which can be plotted as bijective mapping. Different inputs may lead to the same output, and the same input may lead to different outputs. So causes and effects are not coupled unambiguously. Due to mathematical short cuts not being applicable, emergent phenomena cannot be predicted in detail. There is no mechanistic transformation which turns the cause into the effect. There is an activity of the system itself which selects one of the several possible ways of reacting. [Fleissner and Hofkirchner 1997]

This situation (also described as "less than strict determinism" in [Hofkirchner 2001]) is hypothesised to be one the hallmarks of self-organisation and emergence.

If a degree of indeterminism, or less than strict determinism, is a required ingredient of emergence, can it be injected from the outside—indeterminism being defined as an effect that is, at least to some degree, not caused by endogenous variables? It is tempting to simply introduce a measure of indeterminism by mixing in pseudo-random (exogenous) input parameters to the decisions of an EC system's components: for instance, in a CA, the cell's state transition function (taking the cell's neighbours' states at t as input, and yielding the state in $t+1$ as output) could be changed so that it is influenced by an additional, pseudo-random parameter.

Whether this leads to any change in the quality of the resulting phenomena is to be doubted: first of all, the idea of incomplete determinism as a necessary condition of emergence is associated with a certain global state of the system where self-organisation is

about to occur (a “critical point”). In that sense, it is related to the macrolevel of the system. An indiscriminate mixing-in of random influence at the microlevel, which could be considered “background noise”, may result in different states of the system, but there is no logical reason to expect the sudden emergence of a new quality (or, the sudden emergence of structures with new qualities) from that alone. Additionally, this “indeterminism” at the lowest level would not entail or enable any form of downward causation. It is indeed unrelated to events at the system level.

Is there a way of introducing indeterminism at the microlevel (individual constituents) dependent on the macrolevel (the system)? For instance, the mixing-in of pseudorandom parameters at the microlevel could be triggered by certain measures of the system’s activity as a whole (such as “far from thermodynamic equilibrium”, in an EC application modeling physical systems). A similar idea can, in fact, be found in the philosophy of emergence: Popper [Eccles and Popper 1977] refers to situation-dependent “propensity” (likelihood of realisation) of events in the world as a crucial element in emergence.

4.2 Inside Out: Computing Systems in the Real World

In many examples of emergent computing, a common trait is that they are essentially closed software systems (such as CAs). What if the “real world” becomes part of a system’s feedback loop, essentially forming a hybrid system? Can a fundamentally different quality of emergence be expected when the behaviour of an ant colony is not simulated in computer memory, but in the form of robots ([Krieger et al., 2000])?

From a purely mechanistic world-view, this question is meaningless, or at least not very interesting: the realms of computing and the “real world” are viewed as equally deterministic. However, if real-world physical systems are expected to exhibit something like incomplete determinism (if only under special circumstances such as “distant from thermodynamic equilibrium”), then the possibility appears not as easy to rule out.

5 Conclusion and Outlook

Concepts of emergence and self-organisation are often viewed as the building blocks for truly interdisciplinary research on systems. The recurrence of “emergence” in computer science publications may appear as a promising sign for progress in this interdisciplinary undertaking. However, a sampling of relevant literature shows that the concept of emergence is open to a wide range of different interpretations even within just one discipline. In many cases, emergence is used in the “weak” sense as defined by Stephan, in other cases it seems to be merely a figure of speech. The more fundamental question—whether emergence in the “strict” sense, marked by either irreducibility, or unpredictability, or both, can be attained in computing systems—opens up a variety of sub-questions, such as those on the role of indeterminism. Approach-

ing emergent computing from this angle could provide added insight on a theoretical level (such as in the understanding of “information”) and new ideas for empirical research on EC systems.

References

- [Cariani, 1991] Peter Cariani. Emergence and Artificial Life. In: Langton *et al.*(eds.), *Artificial Life II*. SFI Studies in the Sciences of Complexity X 775-797
- [Eccles and Popper, 1977] John C. Eccles and Karl R. Popper. *The Self and Its Brain: An Argument for Interactionism*. Springer, 1977.
- [Emmeche et al., 2000] Claus Emmeche *et al.* Levels, Emergence, and Three Versions of Downward Causation. In: Peter Bøgh Andersen *et al.*(eds.), *Downward Causation: Minds, Bodies and Matter*. Aarhus University Press, 2000.
- [Fleissner and Hofkirchner, 1997] Peter Fleissner and Wolfgang Hofkirchner. Actio non est Reactio: An Extension of the Concept of Causality towards Phenomena of Information. In: World Futures 3-4(49) & 1-4(50)/1997, 409-427
- [Forrest, 1991] Stephanie Forrest. Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks. *Physica D*, 42 (1990) 1-11
- [Hofkirchner, 2001] Wolfgang Hofkirchner. The Hidden Ontology: Real-World Evolutionary Systems Concept as Key to Information Science. In: Emergence 3.3 (forthcoming)
- [Holland, 1998] John H Holland. *Emergence: From Chaos To Order*. Springer, 1977
- [Hordijk, 1999] Wim Hordijk. *Dynamics, Emergent Computation, and Evolution in Cellular Automata*. Ph.D. thesis, University of New Mexico, 1999.
- [Krieger et al., 2000] Michael J.B. Krieger, Jean-Bernard Billeter and Laurent T. Keller. Ant-like task allocation and recruitment in cooperative robots. *Nature* 406 (Aug 31):992
- [Küppers and Krohn, 1992] Günter Küppers and Wolfgang Krohn (eds.). *Emergenz: Die Entstehung von Ordnung, Organisation und Bedeutung*. Suhrkamp, 1992.
- [Maturana and Varela, 1987] Humberto R. Maturana and Francisco J. Varela. *The Tree of Knowledge: The Biological Roots of Understanding*. Shambala, 1987.
- [Stephan, 1999] Achim Stephan. *Emergenz: Von der Unvorhersagbarkeit zur Selbstorganisation*. Dresden Univ. Press, 1999.