Engineering Emergence

Fiona Polack
Department of Computer Science
University of York York, YO10 5DD, UK
fiona@cs.york.ac.uk

April 7, 2005

Artificial emergence: visions and nightmares

The visions of Drexler and others of molecular nanotechnology will be realised sometime. The vision is appealing — trillions of tiny autonomous machines creating useful products or cleaning up after us, using the collective effects of some minimal built-in functionality. However, the nightmare scenarios are ever present: a pollution buster that goes rogue and busts the planet may seem a little extreme, but a host of medical nanites that accidentally block an artery as a side-effect of clearing it is not so farfetched.

Conventionally, we might see *engineering* as the way to control risk in technological advances; we argue that, in order to exploit the emergent properties of (artificial) complex systems safely, we need to be able to engineer these systems.

Engineering and assurance

Engineering is a quality-enhancing activity. The goal is to produce robust systems, providing the necessary assurance of functionality and safety. Side effects include being able to define the safe operating conditions for systems, and building in safe responses if the system moves outside its intended operational envelope.

There is considerable expertise, at York and elsewhere, in constructing and certifying conventional critical systems. A typical approach takes assumptions and evidence (facts and statistics) and uses these to construct an argument that a given system will operate safely within stated environmental parameters. Assurance is conventionally constructed using evidence. For example, confidence in a system's reliability might be enhanced by using materials with a recorded history of performance, or by using techniques (and people) that have produced reliable systems before.

We are seeking to establish engineering principles for emergent systems, such that we could construct assurance arguments. This work is conducted in part under the auspices of the TUNA project[1].

A key problem is decomposition — a conventional incremental-component development does not preserve emergence, or, more accurately, makes no guarantee of emergence in the end system. We need to be able to argue the validity of each part of the system, alone and in combination. We cannot rely on a nanite assembly where the nanites do not survive long enough to do anything, or where the nanites in the environment in which assembly is required do not have the resources to perform the assembly.

Emergence and Architecture

An initial observation is that emergence depends crucially on representation. To take an trivial example, the cells of a cellular automaton (CA) are finite state automata; the update rule determines the next state of each cell from the current states of the cell and some collection of other cells, referred to as its neighbours. Such a system can be programmed very easily, and is utterly uninteresting; the value of the state fluctuates, and at some point may stop fluctuating. The CA displays emergence (pretty patterns) only when the cells and their neighbourliness are translated on to some representation. Change the representation, and the emergent structures change or vanish [3].

We cannot define the automaton and the representation using a single set of concepts. Automata are state-and-operation systems; a representation is a visualisation built up from some discretised model

[1] http://www.cs.york.ac.uk/nature/tuna. **T**heory **U**nderpinning **N**anite **A**seemblers, EPSRC grant EP/C516966/1, is a two-year feasibility study at the Universities of York, Kent and Surrey. The participants are S. Stepney, S. A. Schneider, P. H. Welch, J. C. P. Woodcock, A. L. C. Cavalvanti, H. Treharne and F. A. C. Polack

of space. From this observations, we devised a three-part architecture for emergent systems [4]. We speculate that any complex system with required emergent properties comprises,

- elements (cells, nanites, ants, etc), described by simple state behavioural properties
- detectable emergence (patterns, constructions etc), described in terms of the physical reality or physical properties
- a representation or environment, supporting mappings between, and possibly interacting with the other parts.

We speculate that some architectural parts can be developed by direct application of existing engineering techniques. Part of TUNA, for example, is investigating formal modelling and refinement of aspects of the platelet system.

Concepts and Layers

Having established an engineering-friendly architecture, we are now exploring more thoroughly patterns within each part.

CA models employ only "upward causation": the automaton determines the perceived emergent structures. Layers and causation have long been studied in the context of emergence [1], and, when we move beyond simple CAs, we too find a need to incorporate "downward causation".

In TUNA, we are looking at theoretical designs for artificial blood platelets. One simulation uses a stochastic CA with upward causation to a representation. The model is basic; as platelets cluster, they move very slowly, and rifts open in the clusters as they move. However, by migrating the control of clustering and movement out of the CA, we achieve a higher-layer model that is capable of responding to environmental influences; in keeping with reality, the higher layer exhibits relative location and relative motion. At the lower level, the CA update, still synchronised at the CA level, occurs in response to signals from the higher layer. A nice feature of this model (under development by P. H. Welch and F. R. M. Barnes, University of Kent, using occam-π [2, 6]) is that visual representations can be attached either to the low level (to observe absolute locations and the way platelets pass through them over time), or the higher level (to observe the emergent behaviour of clusters of platelets) [5].

The Future

Engineering of emergent systems is clearly crucial. In our work, we are continuing to explore layers, causation and the three-part architecture. We are starting to consider argumentation, and what is needed to argue that a designed or simulated emergent system will operate correctly in an uncertain world. We plan to explore the capture existing knowledge (or emergence, complexity, argumentation etc) as patterns.

References

[1] P. B. Andersen, C. Emmeche, N. O. Finnemann, and P. V. Christiansen, editors. *Downward Causation. Minds, Bodies and Matter.* Åarhus University Press, 2000.

[2] F.R.M. Barnes and P.H. Welch. Communicating Mobile Processes. In I. East, J. Martin, P. Welch, D. Duce, and M. Green, editors, *Communicating Process Architectures 2004*, volume 62 of *Concurrent Systems Engineering Series*, pages 201–218. IOS Press, September 2004.

[3] M. Hill, S. Stepney, and F. Wan. The effect of irregular grids on the long term behaviour of Life. submitted to ECAL2005, April 2005.

[4] F. Polack and S. Stepney. Emergent properties do not refine. In *REFINE 2005, BCS-FACS Refinement Workshop*, ENTCS. Springer, April 2005.

[5] Fiona Polack, Susan Stepney, Heather Turner, Peter Welch, and Fred Barnes. An architecture for modelling emergence in ca-like systems. submitted to ECAL2005, April 2005.

[6] P.H. Welch and F.R.M. Barnes. Communicating mobile processes: introducing occam-pi. In A.E. Abdallah, C.B. Jones, and J.W. Sanders, editors, *25 Years of CSP*, volume 3525 of *LNCS*, pages 175–210. Springer, 2005. To appear.