

# How did the slime mould cross the maze?

## Adrian Woolfson

EMERGENCE: THE CONNECTED LIVES OF ANTS, BRAINS, CITIES AND SOFTWARE  
BY STEVEN JOHNSON

Allen Lane, 288 pp, £14.99, October 2001, ISBN 0 7139 9400 2

THE MOMENT OF COMPLEXITY: EMERGING NETWORK CULTURE BY MARK  
TAYLOR

Chicago, 340 pp, £20.50, January 2002, ISBN 0 226 79117 3

**I**N London Labour and the London Poor (1861), Henry Mayhew recorded seeing a watercress girl who, eight years old and 'dressed only in a thin cotton gown and a threadbare shawl wrapped round her shoulders', walked the streets crying 'four bunches a penny'; and mudlarks, principally young boys, girls, old men and many old women, who spent the day with 'their trousers tucked up, groping about picking out pieces of coal, iron, wood and copper nails' from the mud on the banks of the Thames. Unlike the street-sellers, whose lives were solitary, the mudlarks formed organised communities. But they were loosely connected. Indeed, Mayhew noted their lack of cohesiveness, observing that they 'peered anxiously about' and held 'but little converse with one another'. Chimney-sweeps, on the other hand, formed distinct and coherent communities, their filthy appearance and offensive smell forcing them into collective isolation. This re-

sulted in the acquisition and nurturing of 'unique habits and peculiarities'. In some communities, the tendrils of continuity reached far back into the past. The fish market at Billingsgate, for example, whose official records date back to the 11th century but whose origins most likely lie in antiquity, remained the geographical and economic centre of fish-selling. While operating within the context and constraints of higher-order organising influences such as the common law, town-planning and the dictates of city guilds, the detail of London's physical and social structure originated according to a dialectic between such 'top-down' principles and less well understood 'bottom-up' self-organisation emerging from the life of the city itself.

That the workings of a system as complex as a city might be computable stretches the limits of credulity. There was nothing inevitable about the emergence of mudlarks by the river's edge, or the watercress

girl loitering near Farringdon market. But much like a living organism, a city is defined by an underlying dynamic that helps generate and perpetuate its core features. If we were able to obtain even a rudimentary insight into these processes we might be able broadly to predict the types of phenomenon likely to emerge within such complex systems, while having to acknowledge that the specifics will remain beyond our grasp.

Jane Jacobs's *The Death and Life of Great American Cities* (1961) fundamentally changed the way in which cities are perceived, arguing that they were 'learning machines', with 'marvellous innate abilities for understanding, communicating, contriving and inventing'. Jacobs made no attempt to express the idea in a formal notation, but is it possible that the broad workings of this complex and distributed 'intelligence' could be captured mathematically? And, if so, could the insights gained enable us to predict a city's likely future? Might we be compelled to take such principles into consideration when planning new cities, if not to dictate their exact form, then to manage their behavioural repertoire and development? Could we in fact work out a tentative formal mathematical science of city structure and dynamics?

In its simplest sense, a city is an example of a 'complex adaptive system'. The fledgling mathematical laws underpinning the behaviour of such systems are indifferent to the material nature of the networks of components that make them up. In this respect, the emergent aspect of a city's existence, which appears to transcend the ac-

tions of any individual, warns against descriptions of the informational content of living organisms that begin and end with DNA. Although the information encoded in DNA can account for organisation at the level of protein molecules, it doesn't explain how those molecules interconnect to form metabolisms, or how such networks can self-organise to form the structures of living things. Whether the system involved is biological or non-biological, what we need are mathematical laws of emergence, laws that would apply as well to a network of neurons in the brain of a porcupine as to the individuals within a specific community or organisation, or to the intricate dynamics of global stock exchanges.

The low-level rules that generate higher-level complex behaviour and define the laws of complexity are the subject of both these books. 'Emergence' refers to the generation of higher-level behaviour or structures within systems made up of relatively simple components. Rather than being programmed into the system from above, such behaviours are generated from the bottom up. Emergence upsets the top-down causality to which we're accustomed; instead, it appears eerily organic, perhaps accounting for the 'missing essence' we intuit in the idea that 'the whole is greater than the sum of its parts.' These books are important, in that they apply the mathematical logic of complex systems to cultural – and, in the case of Taylor, cognitive – phenomena. They do so in a metaphorical rather than explicitly mathematical manner, and focus on the transformative and generative power of self-organising princi-

ples in contexts as diverse as the World Wide Web, global civilisation, artificial life, computer software, architecture, the organisation of cities, artistic creativity and human thought.

‘When I imagine the shape that will hover above the first half of the 21st century,’ Johnson begins, ‘what comes to mind is not the coiled embrace of the genome, or the etched latticework of the silicon chip.’ It is instead the images of computer programs that use the underlying logic of emergence to generate complex, and in some instances unpredictable, behaviour without the help of traditional algorithmic programming. An example are the programs written by Danny Hillis, whose ‘emergent’ software represents a qualitative change from the inflexible programs that preceded it. Hillis’s programs enact the principle of evolution by natural selection: in Johnson’s words, ‘emergent’ software is ‘software that you grow rather than engineer’. These programs begin with nothing more than broad preconceptions of how the numerical tasks that Hillis sets them might be solved; they are in fact programs for generating further, higher-order programs which can then be mutated randomly to produce a new set of related variants. Each successive tier of higher-order programs constitutes a repertoire of possible programs from which the one best adapted to solving the task can be selected. By repeating this process for generation after generation, the program eventually evolves higher-order programs capable of solving the problem. Interestingly, the way in which they achieve a solution is often un-

intelligible to the programmer.

These programs inhabit a silicon universe with a ‘dry’ logical infrastructure, as against the ‘wet’ protein and DNA components of living entities. Johnson argues that the deciphering and harnessing of principles of emergence, which such programs begin to demonstrate, will result in a paradigm shift as fundamental as that brought about by Darwin. And, as with evolutionary theory, the influence of emergence will eventually permeate every area of our intellectual endeavours.

Having considered a range of complex adaptive systems, from the human brain and colonies of ants to the city of Manchester and the guilds of Florence, Johnson addresses his key question – one that was anticipated by H.G. Wells in his *World Brain* (1938). Is it possible that the World Wide Web, a complex artificial system, might one day exhibit emergent behaviour or develop autonomous ‘higher-level consciousness’? Wells speculated that a ‘World Mind’ held out the utopian promise of ‘World Peace’. In *The Impact of Science on Society* (1952), Bertrand Russell similarly conjectured that a scientific society could be stable only if there was a single world government. The World Wide Web currently lacks the formal structure required to achieve higher-level consciousness, but might it be re-engineered to facilitate more modest examples of emergent behaviour? If so, how would such behaviour manifest itself, and should we encourage it? Would we be prepared to ‘let go’ our artificial emergent systems, leaving them to explore the uncharted landscapes of their own potential and to

adopt their preferred modes of behaviour? Furthermore, if the Web were one day to exhibit emergent behaviour or become self-aware, would we even know that it had happened?

Johnson argues that the principal change needed to transform the Web into a 'virtual city of minds' would be the introduction of the types of positive and negative feedback circuits that enable complex phenomena to be generated within self-organising systems. It would also be necessary to order the sea of HTML documents that populate it so as to form communities, linked thematically in a manner analogous to districts in cities. Organisation is a prerequisite; connectivity alone is not enough. The human brain, for example, is not a uniform web or net, but is composed of distinct anatomical regions that are interconnected along discrete pathways. Finally, it is not enough simply to generate complex patterns. For the Web to become truly intelligent it would have to become a complex adaptive system, so that the patterns it generates could form the raw material for a process of cultural evolution.

There are a number of ways in which such changes could be realised. The American company Alexa Internet, for example, archives the Web and its patterns of user traffic. If enough individual surfing histories could be tracked using the information held in such archives, related websites could in principle be hived off into virtual 'neighbourhoods'. (Amazon, the online bookstore, uses this approach to link consumers to related products.) The recently demonstrated phenomenon of 'parasitic

computing', in which a computer uses the Internet's infrastructure to instruct other computers how to perform tasks that are beyond their immediate scope, is another example of how emergent behaviour might one day, in principle, be generated within the Internet.

In *The Soul of the White Ant* (1937), the South African naturalist Eugène Marais painstakingly documented the natural history of the termites that build terminaries or 'ant-hills' on the open veldt. He astutely concluded that the terminary is a 'distributed organism': its masterplan, and the unique behaviour of the community it houses, does not reside within any individual, but is instead an emergent property arising from the ants' interactions. Marais lacked the complexity theorists' formal mathematical language with which to explain such phenomena. Instead, he attributed it, as best he could, to what he called a 'group soul'. Although the structure of ant societies is described using a hierarchical terminology, the command structure of a colony is in fact flat and decentralised, not totalitarian. Indeed, given the architecture of a terminary, it would be physically impossible for the queen to direct the actions of every worker. In Johnson's words, ants 'think locally and act locally, but their collective action produces global behaviour'. The 'macrointelligence and adaptability derive from local knowledge', from individuals that lack a global perspective.

So how might this extraordinary force, which can crystallise order out of chaos and produce structure without the agency of algorithms, be rationalised? How can pat-

terns, whether biological, social or cultural, be generated without programs? Stuart Kauffman, among others, has attempted to simulate the emergent behaviour of abstract networks by using computers, for which purposes it is irrelevant whether the interconnected components are ants, people, communities, molecules or competing multinational companies. Because the results are generic, they offer a starting point for the elucidation of general laws of complexity. The results are remarkably simple: the principal determinants of network dynamics are the number of inputs each component receives, the nature of the rules they use to process these inputs and compute the corresponding outputs, and the bias that each rule introduces into the overall response. A network's behaviour can be 'tuned' by altering any one of these features. Networks in which each component receives only one input are frozen into a small set of inflexible uniform behaviours. Those in which components receive multiple inputs are chaotic and incapable of generating order. The really interesting behaviours emerge when the ratio of inputs is tuned to around two per component. At this level of interconnectivity, order suddenly crystallises and 'emerges' from the simulated networks. If the number of inputs to components is more finely tuned, the system can be taken to the edge of chaos, where it 'quivers on the edge of order' in a largely coherent state, while at the same time producing occasional flurries of chaotic, exploratory activity. The stability provides the robustness needed for continuity, while the chaos generates the bursts of un-

predictable creativity leading to mutations on which processes of natural selection can act.

If we were planning to construct a highly interconnected system from first principles, such as a city, a World Wide Web or an artificially intelligent robot, we would do well to be aware of these fundamental laws. Will Wright's hugely successful computer game SimCity was one of the first to incorporate these bottom-up principles. In SimCity the player constructs virtual towns on a computer screen, but, as Johnson says, rather than actually 'playing' SimCity, 'users grow their virtual cities,' as a player's control over them is always indirect, using the same bottom-up 'swarm logic' that underpins the global intelligence of ant colonies. Anyone who has played SimCity will know that a small linear change in one neighbourhood can have far-reaching non-linear consequences, felt in distant regions of the virtual metropolis. The important point is that although the underlying rule systems determining local behaviour and development are simple and well characterised, they generate global behaviour that could not have been predicted from a consideration of the local rules alone. The emergence of complex phenomena in non-biological, cultural systems was anticipated, as it happens, by Hobbes. In *Leviathan*, he describes the state as a super-organism, 'an Artificiall man; though of greater stature and strength than the Natural, for whose protection and defence it was intended'.

In an ant community the components, or 'nodes', are individual ants, the nature of

whose interconnections depends on their species. Communication between workers in colonies of the fire ant *Solenopsis invicta* relies on a vocabulary of ten signals, nine of which are based on pheromones. The complexity of their semantics is increased by the ants' ability to detect pheromone gradients and to respond to the frequency of their release. These signals enable them to feed information back into the colony 'network'. Consequently, an event detected at one end of the colony reverberates throughout the rest of it. Michael Kreiger and Kenneth Ross have recently shown that a change in a single pheromone-binding gene called Gp-9 can have a major impact on the social behaviour of fire ants.

One can imagine a situation in which the ants in a given colony were genetically engineered in such a way that both the frequency and extent of their interactions would be increased many times over. Each ant would be bombarded with information, resulting in chaotic emergent behaviour that would undermine the colony's ability to transmit and process information efficiently, ultimately threatening its ability to survive. Similarly, a colony with an insufficient number of ants would be unable to generate emergent phenomena, and might become frozen in very restricted and inflexible modes of behaviour. It is likely that evolution has tuned ant networks to the edge of chaos, a mathematical region where they exhibit an optimal mixture of connectivity, order and the type of complexity on which natural selection can continue to act.

The nodes might equally well, however,

be individual workers in a commercial organisation. In this case, a catalogue of the interconnections between nodes might include how many telephone calls were made by each worker, the number of meetings that they participated in each day, and the number of tea breaks they had with their colleagues; or again, the number of pavements, highways, subways, coffee shops and shopping arcades in a city might influence its dynamics and the 'tuning' of its networks.

The slime mould *Dictyostelium discoideum* remains one of the best examples of how intelligent behaviour can be produced by these decentralised means. Slime moulds are composed of swarms of thousands of distinct single-celled units, which for much of the time exist independently of one another. Under certain conditions, however, they coalesce to form a single organism. Toshiyuki Nakagaki has shown that despite lacking a centralised brain, these creatures can learn to navigate simple mazes along the shortest possible routes. Although for many years it was thought that there were specialised slime mould cells that paced or orchestrated the behaviour of the remaining 'slave' cells, experimental data do not support this view. Instead, it appears that the emergent 'intelligence' of slime moulds originates not from a central structure, but rather 'from the ground up', emerging from within the distributed network of interacting cells itself.

Johnson notes that cities are much like slime moulds in that they have an intrinsic 'intelligence' independent of any internal or external 'pacemaker'. Such considera-

tions are not purely academic. When the New York City Planning Commission announced that it planned to raze large sections of the West Village in October 1961, Jane Jacobs's argument that cities needed to be understood from an organic perspective, as having an organised complexity, was crucial in preventing this from happening.

The Moment of Complexity provides a lucid and scholarly summary of a wide range of issues relating to complex adaptive systems and promises fundamental new insights. Unfortunately, these do not materialise. Taylor has an excellent understanding of the debate, but little new to contribute to it. There are some good things scattered throughout the text, however, such as his discussion of Frank Gehry's Guggenheim Museum in Bilbao, whose design resulted partly from the emergent-style computer programs used to transform his drawings into physical structures. Taylor's assertion that the hyper-real paintings of Chuck Close 'provide graphic illustrations of emerging complexity' is similarly intriguing, but also irritating. Exactly how does a Chuck Close painting demonstrate the formal mathematics of complexity? Metaphors alone can't provide a sufficient foundation for such wide-reaching speculations. 'The mix swirling in my mind becomes dense and diverse,' we read, 'like some primordial soup slowly heated to boiling point,' but there is no evidence to support the notion that cognitive processes might incorporate the principles of complex adaptive systems.

What is clear is that the mathematics of

complexity hovers ghostlike over our every action. We may choose to be its masters or its slaves, but irrespective of our choices, it will continue to weave its logic in the core of our biological and cultural existence. Indeed, if H.G. Wells was right, it may become the agency by which we are eventually marginalised, as the world enters a post-biological age dominated by a machine intelligence with an agenda quite different from any concocted within the proteinaceous brains of its creators.