

Beyond design: cybernetics, biological computers and hylozoism

Andrew Pickering

Received: 17 April 2007 / Accepted: 18 November 2008 / Published online: 19 December 2008
© Springer Science+Business Media B.V. 2008

Abstract The history of British cybernetics offers us a different form of science and engineering, one that does not seek to dominate nature through knowledge. I want to say that one can distinguish *two different paradigms* in the history of science and technology: the one that Heidegger despised, which we could call the Modern paradigm, and another, cybernetic, nonModern, paradigm that he might have approved of. This essay focusses on work in the 1950s and early 1960s by two of Britain's leading cyberneticians, Stafford Beer and Gordon Pask, in the field of what one can call biological computing. My object is to get as clear as I can on what Beer and Pask were up to. At the end, I will discuss Beer's hylozoist ontology of matter, mind and spirit. This material is not easy to get the hang of—but that is what one should expect from an unfamiliar paradigm.

Keywords Cybernetics · Computing · Biological computing · Design · Stafford Beer · Gordon Pask · Martin Heidegger · Hylozoism

We tend to think of technology as a unitary category, as if we could have a single, unitary philosophy of technology. Martin Heidegger (1976 [1954]), for example, thought that the essence of technology is a project of *enframing*, of putting the world at our disposal, and that science is the means of achieving that. We dominate nature through knowledge: science gives us access to the levers of power. And, at the same time, we enframe ourselves, becoming parts of a posthuman assemblage of production and consumption. I think Heidegger was right. The last few centuries can be characterised

Presented at an international conference on the philosophy of technology, Copenhagen, 13–15 Oct 2005.
Revised for submission to Synthese.

A. Pickering (✉)
University of Exeter, Exeter, UK
e-mail: a.r.pickering@exeter.ac.uk

by an ever-growing assemblage of knowledge and power, science and technology, which does enframe the world and ourselves within it. I don't see why we should not have a unitary philosophy of that assemblage: Heidegger is a good place to start.

That is not what I want to talk about. The history of British cybernetics offers us a different form of science and engineering, one that does not seek to dominate nature through knowledge. I want to say that one can distinguish *two different paradigms* in the history of science and technology: the one that Heidegger despised, which we could call the Modern paradigm, and another, cybernetic, nonModern, paradigm that he might have approved of. This essay focusses on work in the 1950s and early 1960s by two of Britain's leading cyberneticians, Stafford Beer and Gordon Pask, in the field of what one can call biological computing. My object is to get as clear as I can on what Beer and Pask were up to. At the end, I will discuss Beer's hylozoist ontology of matter, mind and spirit. This material is not easy to get the hang of—but that is what one should expect from an unfamiliar paradigm.¹

As a first pass through my topic, let me start with some work that is in the same spirit as Beer and Pask's, though less ambitious, and which is going on now. If you look on the web, you can find a project in the Arts, Computation and Engineering Masters Program at UC Irvine undertaken by Garnet Hertz, entitled 'Control and Communication in the Animal and the Machine.'² This title is a direct quotation of the much-cited subtitle of Norbert Wiener's 1948 book, *Cybernetics*, the book which introduced the word 'cybernetics' to the public. We can take this project, then, as a piece of self-consciously cybernetic engineering. What is it? It is a cockroach-controlled robot.

In Fig. 1, a giant Madagascan cockroach stands on the white trackball at the top of the assembly, attached by Velcro on its back to the arm which loops above the other components. Motions of the cockroach's legs rotate the trackball which in turn controls the motions of the cart (much as a trackball can be used to control the motion of the cursor on a computer screen). Infrared sensors detect when the cart is approaching an obstacle and trigger the appropriate light from an array that surrounds the roach. Since roaches tend to avoid light, this causes the roach to head off in another direction. The entire assemblage thus explores its environment without hitting anything or getting stuck—ideally, at least. So what?

What is singular about this assemblage is the cockroach. Over the past twenty years or so, many autonomous robots similar to Hertz's have been built, but their intelligence resides in electrical circuitry, not an insect. The first of these were small robots built by another of my cyberneticians, Grey Walter, in the late 1940s (Fig. 2)—he called them tortoises, or *Machina Speculatrix* (Walter 1953; Pickering 2004b). They had, as he put it, two neurons, each composed of an electronic valve, a capacitor and a relay switch. The present-day descendants of the tortoises have silicon chips instead to supply their intelligence. So here we come to a first contrast: in Hertz' robots, a roach takes the

¹ For more on British cybernetics, including the work of Beer and Pask, see Pickering (2002, 2004a, b, 2005a, b, 2007, forthcoming).

² www.conceptlab.com/control/—and follow the links for more information. Since I first looked at this site (21 July 2005) the title of the project has been changed to 'Cockroach Controlled Mobile Robot,' though the previous title is still listed too. I am extremely grateful to Ellen Fireman for first telling me about this project, and to Garnet Herz for telling me more about it when I visited Irvine in October 2005.

Fig. 1 Cockroach-Controlled Robot. *Source and permission:* Garnet Hertz




Fig. 2 The Robot Tortoise. *Source:* de Latil (1956, facing p. 34)

place held by a computer in conventional robotics—and hence the phrase *biological computing*—the roach-as-computer.

We need to think further about this contrast. We can note that centuries of engineering and science go into the manufacture of valves, chips and electronic circuits in general. The human race has had to accumulate an enormous amount of material infrastructure, knowledge and expertise to put us in a position to build and program computers. Hertz instead needed to know *almost nothing* about roaches to build his machine. The roach remains, in a cybernetic phrase, a *black box*, known only in terms of its performative inputs and outputs—it runs away from light, and that’s it. If the construction of a conventional robot entails ripping matter apart and reassembling it in accordance with our designs, Hertz’s involves a different strategy: *the entrainment of the agency of nature*, intact into a human project, without penetrating it by knowledge and reforming it.

Of course, we have been doing that sort of thing forever—think of the use of natural power sources: the horse pulling a barge, a stream driving a mill. But that is hardly our paradigm for thinking about technoscientific design today, and it is certainly not how we usually think about robots, computing and intelligence. That is why Hertz’s work strikes me, at least, as profoundly *strange*. And I could rephrase this strangeness as a *denaturalisation* of the Modern paradigm itself. In relation to Hertz’ biological computing, the long and arduous conventional history of science and technology, industry and warfare, leading up to mainstream contemporary robotics looks like a massive *detour*—penetrating, knowing and rearranging matter at vast expense of money and effort. Hertz’ robots show that there is another and much simpler way to achieve comparable ends without the detour. We can see two different *stances towards matter* in play here: the conventional one that involves penetrating black boxes through knowledge, and the cybernetic one that seeks to entrain boxes that remain black into our world. And we could understand this contrast ontologically and epistemologically. Cybernetics centres itself on a world of performative black boxes and their interrelations, whereas the Modern paradigm emphasises an intrinsically cognitive relation to matter.

Hertz’ work is very interesting and it leads in nicely to what I want to discuss. Now I want to go back to the more or less forgotten origins of biological computing.

My God, I’m a cybernetician!

Stafford Beer, on first reading Wiener’s *Cybernetics*. *The Falcondale Collection* (1994).

I need to begin with some general history of cybernetics. I am especially interested in British cybernetics, and although the standard story locates the origins of cybernetics in Norbert Wiener’s military research in WWII, British cybernetics (and, indeed, much of American cybernetics) began as a *science of the brain*, in which the brain was approached from a distinctive angle. As Ross Ashby, one of the British pioneers, wrote in 1948: ‘To some, the critical test of whether a machine is or is not a ‘brain’ would be whether it can or cannot ‘think.’ But to the biologist the brain is not a thinking machine, it is an *acting* machine; it gets information and then it does something about it’ (1948, p. 379). Something of the strangeness of cybernetics becomes clear here. We usually think of the brain as a *representational*, cognitive device which we use for thinking. This is the view at the heart of 20th-century

philosophy of science and of postwar research in artificial intelligence. And the point to grasp is that cybernetics did not approach the brain from that angle; it did not conceive the brain as primarily a cognitive organ. The cyberneticians thought of it instead as *performative*, an organ that does something, that acts—as an active and engaged switchboard between our sensory inputs from the environment and our motor organs. And beyond that, cybernetics was interested in the brain as the organ par excellence of *adaptation*, especially of adaptation to environments that had never been encountered before.

Where did cybernetics take this idea? A further defining attribute of early cybernetics was the attempt to understand the performative and adaptive brain by building electromechanical models that could mimic its performance. This was Grey Walter's rationale for constructing his robot tortoises. Like Hertz' roach-robots today, these were adaptive inasmuch as they could find out about and explore environments they had never encountered before without ever constructing any inner representation of their worlds. But here we need to focus on Ashby's notion of adaptation rather than Walter's, and the place to start is with the biological notion of *homeostasis*.

Homeostasis refers to the ability of mammals to keep certain essential variables within narrow limits, independent of fluctuations in their environment—the constancy of blood temperature is an obvious example. Here one thinks of the feedback systems and servomechanisms that were themselves key referents for early cybernetics: the thermostat, for instance, that senses the temperature in a room and turns the heating up or down accordingly. One can say that the thermostat maintains the temperature in a state of *dynamic equilibrium*, tilting it one way if it tends to move in the other. Ashby thought of the body as having a whole array of such homeostatic mechanisms hard-wired into it by the genes. But he argued that we could not be hard-wired to deal with all possible fluctuations; we also have to be able to *learn* about environments we have not encountered before, and such learning was Ashby's paradigm of adaptation.

The brain was, in this sense, the supreme organ of adaptation, and to explore the mechanism of adaptation Ashby also built a machine in the late 1940s—his famous homeostat, which formed the centrepiece of his 1952 book, *Design for a Brain*, on which I can comment briefly. First, the name homeostat is misleading. The homeostat did not simply mimic a thermostat by maintaining some essential variable constant. Instead, it *became* a thermostat—a servomechanism—by randomly reconfiguring its wiring whenever it encountered environments with which it could not maintain dynamic equilibrium. Relays would trip when the current passing through it exceeded preset limits, changing the resistance and reversing the polarity of its circuits until it regained a thermostat-like relation to the environment. This *becoming-thermostat*, as [Deleuze and Guattari \(1987\)](#) might say, was the homeostat's key property, the property that qualified it, in Ashby's opinion, as a model of the adaptive brain. And, again, we should note that the homeostat was a performative rather than a representational device. Like Walter's tortoises, the homeostat contained no inner representations of its world; it *acted*, in a double sense: it transformed inputs into outputs, and it transformed itself in pursuit of an equilibrium between inputs and its outputs (Fig. 3).




Fig. 3 The Homeostat. Source: Ashby (1952, p. 94, fig. 8/8/1)

From the homeostat the history of cybernetics proceeded in several directions. One was Ashby's visionary attempt to develop giant artificial homeostat-like brains.³ But we can go in another direction today. We could see Ashby's construction of the homeostat as a classically scientific detour away from the living brain in pursuit of a representational understanding. And the key to understanding Beer and Pask's work in biological computing is to see that it *undid this detour*. If a primitive adaptive system like the homeostat could be thought of as a model brain, well, the world is already full of much more sophisticated adaptive systems than the homeostat, which could function as much more sophisticated brains—perhaps even more sophisticated, in certain ways, than our own.


This line of thought was most clearly laid out in a paper entitled 'Towards the Automatic Factory' that Beer presented to a symposium on self-organisation, held at the University of Illinois in June 1960 (Beer 1994 [1962a]).⁴ Beer opened with a discussion of the concept of an 'automatic factory,' then attracting great interest, especially in the US (see Noble 1986). This was a vision of industrial automation taken, one might think, to the limit. In the automatic factory, not only would individual machines and productive operations be performed by machines without human interference, but materials would be automatically routed from one operation to the next. In the 'lights out' factory, as it was also called, the entire production process would be conducted by machines, and human labour made redundant.

Beer was not in 1960 a critic of the automatic factory, except that he felt that current conceptions were not automatic enough. He compared them to a 'spinal dog'—a dog

³ At various times throughout the 1950s Ashby attempted to develop a more sophisticated synthetic brain called DAMS, for Distributed And Multistable System (Pickering forthcoming). This project did not, in the end, get far, though it is perhaps being taken up again just now. See the link at www.ibspan.waw.pl/ICANN-2005/workshops.html to the 'Building a Brain' workshop: www.ibspan.waw.pl/ICANN-2005/CreatingABrain.pdf (accessed the 4th of July). I am grateful to Peter Asaro for bringing this meeting to my attention.

⁴ Page number citations here and below are to the reprints of Beer's writings in Beer (1994).

Fig. 4 Beer in the 1960s.
 Source: Harnden and Leonard (1994). Reproduced by permission of John Wiley & Sons and Allenna Leonard



Stafford Beer as Managing Director
 SIGMA (Science in General Management) Ltd 1961–66

whose nervous system had been surgically disconnected from the higher levels of its brain. The automatic factory (1962a, p. 164):

has a certain internal cohesion, and reflex faculties at the least. [But] When automation has finished its work, the analogy may be pursued in the pathology of the organism. For machines with over-sensitive feedback begin to ‘hunt’—or develop ataxia; and the whole organism may be so specialized towards a particular environment that it ceases to be adaptive: a radical change in the market will lead to its extinction.

Beer’s argument was that to make it adaptive and to avoid extinction in market fluctuations, the automatic factory would need a brain.

At present, such an automatic factory must rely on the few men left at the top to supply the functions of a cerebrum. And . . . the whole organism is a strange one—for its brain is connected to the rest of its central nervous system at discrete intervals of time by the most tenuous of connections. The survival-value of such a creature does not appear to be high . . . This will not do. The spinal dog is short of a built-in cerebrum; and the automatic factory is short of a built-in




Fig. 5 The Cybernetic Factory. *Source:* Beer (1962a, pp. 200–201, Fig. 4). Reproduced with the permission of John Wiley & Sons and Allenna Leonard

brain. The research discussed in this paper is directed towards the creation of a brain artefact capable of running the company under the evolutionary criterion of survival. If this could be achieved, management would be freed for tasks of eugenics; for hastening or retarding the natural processes of growth and change, and for determining the deliberate creation or extinction of whole species (1962a, p. 165).

The cybernetic factory, as Beer imagined it, would be *viable*—a key term for him: it would react to changing circumstances, it would grow and evolve like an organism or species, all without any human intervention at all.

Figure 5 is a schematic of Beer’s cybernetic factory, and for present purposes the key element to think about is the box marked ‘homeostatic loop’ at the bottom centre, which Beer elsewhere referred to as the ‘U-machine’. This sits like a spider in a mesh of information flows—inputs reflecting the state of the factory and its environment, and outputs that determined its parameters. The U-machine was the brain of Beer’s imagined automatic factory. The question was then, just what should the U-machine be? Beer remarked that (1962a, pp. 209, 215):

As a constructor of machines man has become accustomed to regard his materials as inert lumps of matter which have to be fashioned and assembled to make a useful system. He does not normally think first of materials as having an intrinsic

sically high variety which has to be constrained. . . [But] We do not want a lot of bits and pieces which we have got to put together. Because once we settle for [that], we have got to have a blueprint. We have got to design the damn thing; and that is just what we do not want to do.

This critique of design is a striking index of the difference between cybernetics and Modern technoscience. Intrinsic to the latter is a principled rearrangement of otherwise passive matter to bend it to our will; but this was what Beer proposed not to do. Along the same lines, he also talked about ‘new developments in solid state physics,’ for which he had nothing but admiration, but he argued that ‘I am unable to see how “design” can be eliminated from molecular electronics.’ One has to specify the configuration of a chip, say, before one can build it, and this is an impediment to the construction of a U-Machine that could continually reconfigure itself. To be adaptable, the machine would have to know its own design, work variations on that, and then reconstruct itself. And, of course, the ‘techniques so far available . . . involve massive equipment that could hardly be visualized as operated by the U-Machine to change its own internal mechanism’ (1962a, p. 210).⁵

Nevertheless, said Beer in 1960, ‘Before long a decision will be taken as to which fabric to use in the first attempt to build a U-Machine in actual hardware (or colloid, or protein)’ (p. 212). Colloid or protein? This gets us back to biological computing and the adaptive brain. Beer’s idea was simply—if one can say that—to enrol some naturally occurring adaptive system as the U-machine, as the brain of a totally automatic adaptive factory. During the second half of the 1950s, he had, in fact, embarked on what he described as ‘an almost unbounded survey of naturally occurring systems in search of materials for the construction of cybernetic machines’ (1959, p. 162), and in 1962 he published a brief and, alas, terminal report on the state of the art, which makes fairly mind-boggling reading (Beer 1962b). We can glance at some of the systems he discussed there to get a flavour of this work.

Beer’s list of candidates for the U-machine begins with quasi-organic electrochemical systems, but I want to postpone discussion of these for a while. Then comes Beer’s successful attempt to use positive and negative feedback to train young children (presumably his own) to solve simultaneous equations *without* teaching them the relevant mathematics—to turn the children into a performative (rather than cognitive) mathematical machine. Beer then moves on to discuss various thought-experiments involving animals (1962b, pp. 28–29):

Some effort was made to devise a ‘mouse’ language which would enable mice to play this game—with cheese as a reward function . . . In this way I was led to consider various kinds of animal, and various kinds of language (by which I mean intercommunicating boxes, ladders, see-saws, cages connected by pulleys and so forth). Rats and pigeons have both been studied for their learning abilities . . . The Machina Speculatrix of Grey Walter might also be considered (with apologies to the organic molecule) . . . However no actual machines were built. . . By the same token, bees, ants, termites, have all been systematically considered as

⁵ ‘BEER: And the other big point I would like to make, about the big electronic machines, which I think are just dinosaurs . . ./BOWMAN: Subject to the same fate?/BEER: I think so’ (Beer 1962a, p. 220).

components of self-organizing systems, and various ‘brain-storming’ machines have been designed by both Pask and myself. But again none has been made.

Here, I would guess, we are at the origin of the mouse-computer that turns up in both Douglas Adams’ *Hitch-Hikers Guide to the Universe* and Terry Pratchett’s *Discworld* series of fantasy novels.⁶ The most convincing fictional representations of biological computing that I have come across are the visibly organic control systems of alien space ships that featured in various episodes of *Dr Who* and, more recently, in Greg Bear’s novel, *Slant* (1997), which includes, as I recall, colonies of beans and wasps.

Beer had, however, devoted most of his own efforts to systems composed from simpler organisms: colonies of *Daphnia*, a freshwater crustacean (Pask had considered mosquito larvae), of *Euglena* protozoa, and an entire pond ecosystem (Fig. 6):

[P]ure cultures . . . are not, perhaps, ecologically stable systems. Dr. Gilbert, who had been trying to improve the *Euglena* cultures, suggested a potent thought. Why not use an entire ecological system, such as a pond? . . . Accordingly, over the past year, I have been conducting experiments with a large tank or pond. The contents of the tank were randomly sampled from ponds in Derbyshire and Surrey. Currently there are a few of the usual creatures visible to the naked eye (*Hydra*, *Cyclops*, *Daphnia*, and a leech); microscopically there is the expected multitude of micro-organisms. The state of this research at the moment is that I tinker with this tank from time to time in the middle of the night.



Fig. 6 Maverick Machines: Pond as Biological Computer. Source: Pask and Curran (1982, p. 143)

⁶ In the *Hitch-hiker’s Guide*, the earth is a giant analogue computer built by mice-like beings to answer (or was it discover?) the Ultimate Question. On the earth as an analogue computer, see also Blohm et al. (1986).

Some clarification might be needed here. The key point is that all the systems Beer talked about are adaptive systems, capable of reconfiguring themselves in the face of environmental transformations. In a steady state, an ecosystem like a pond, for example, exists in a state of dynamic equilibrium with its environment, homeostatically responding to fluctuations that threaten its viability. And if the environment changes, the ecosystem will reconfigure itself to achieve a dynamic equilibrium with that, just like Ashby's electromechanical homeostats. Beer's idea was that if one could only couple such an adaptive system to a factory, say, making the factory part of the pond's environment, and vice versa, the health of each could be made to hinge on that of the other, in a process that Beer called *reciprocal vetoing*. Disturbances from the factory might trip the ecosystem into a new configuration, which would in turn perturb the operation of the factory, and if the factory in its new state was still unstable, new disturbances would travel back to the ecosystem—and so on until the pond and the factory achieved a collective state of dynamic equilibrium with each other and their outside environments. This is the way in which a pond with some small organisms and a leech could serve as an adaptive brain for the automatic factory. An amazing idea, though I can see no reason in principle why it should not work.

Having said that, of course, there are not, at the moment, any such biological computers. As I hinted already, this project came to naught. The immediate technological problem lay in achieving a coupling between naturally adaptive systems and the systems they were intended to control. From the schematic of the automatic factory it is clear that Beer had analysed what the key input and output variables were. The problem was to make biological systems *care about them*. How could they be translated into variables that would impinge significantly on a biological controller? In his 1962 review Beer mentioned a couple of attempts to do this, and indicated where difficulties had arisen (1962b, p. 29):

Many experiments were made with [*Daphnia*]. Iron filings were included with dead leaves in the tank of *Daphnia*, which ingested sufficient of the former to respond to a magnetic field. Attempts were made to feed inputs to the colony of *Daphnia* by transducing environmental variables into electromagnets, while the outputs were the consequential changes in the electrical characteristics of the phase space produced by the adaptive behaviour of the colony. . . . However, there were many experimental problems. The most serious of these was the collapse of any incipient organization—apparently due to the steadily increasing suspension of tiny permanent magnets in the water.

To put it another way, having consumed the iron filings, the *Daphnia* excreted them and the water filled with rust. Another attempt hinged on the fact that *Euglena* are sensitive to light, and Beer sought to achieve optical couplings to a tank full of them. 'However, the culturing difficulties proved enormous. *Euglena* showed a distressing tendency to lie doggo, and attempts to isolate a more motile strain failed' (Fig. 7).

The collapse of Beer and Pask's biological computing project in the early 1960s did not, then, imply any inherently fatal flaw in its conception. We should not refuse to entertain this strange nonModern paradigm on that ground. The primary obstacle lay in getting adaptive biological systems to engage with us, in finding some practical means of signalling to them in ways that they might respond to. And here, the opening









Figure 2 Effects of *Euglena* upon each other. , *Euglena* culture, with tropism displayed as shown; , stimulus; , sensory receptor; , inhibiting and , stimulating influence of a's sensation on b's stimulus

Fig. 7 The *Euglena* Homeostat. Source: Beer (1962b, p. 30, Fig.2)

sentences of Beer's 1962 review bears attention: 'Everything that follows is very much a spare time activity for me, although I am doing my best to keep the work alive—for I have a conviction that it will ultimately pay off. Ideally, an endowed project is required to finance my company's Cybernetic Research Unit in this fundamental work' (1962b, p. 25).

I quoted Beer above on tinkering with tanks in the middle of the night, evidently at home, and his daughter, Vanilla, has fond childhood memories of weekend walks with her father to collect water from local ponds. And here we run into an absolutely typical sociological feature of British cybernetics: key developments were achieved on an amateur basis, outside all of the usual institutions of support. Beer did his work on biological computing on his own time, for fun, while running what was probably the world's biggest industrial OR group for a living. Obviously one can connect the strangeness of cybernetics to the lack of an institutional base—these two features hung together and reinforced one another—but it would take another essay to develop that thought further (Pickering 2005b). I just want to note that (1) given this hobby-ist social basis, it is not surprising that many of the most interesting cybernetic innovations languished after the 1960s, and that (2) some of them *are* now resurgent, as, for example, Garnet Hertz' roach-controlled robots, though (3) projects like Beer's in the 1960s were much more ambitious than latter-day versions, especially in that (4) sophisticated notions of adaptation running from Ashby to Beer and Pask have got lost along the way to the present. And to elaborate that last thought, I promised to return to the first item on Beer's 1962 list, quasi-organic electrochemical computers. Beer and Pask worked on these together and separately, but they are most extensively discussed in several long publications in the late 1950s and early 1960s from Gordon Pask, so the next section is about Pask's work (Fig. 8).⁷

Self-organizing systems lie all around us. There are quagmires, the fish in the sea, or intractable systems like clouds. Surely we can make these work things

⁷ Pask's publication list begins with 18 papers written between 1957 and the publication of *An Approach to Cybernetics* in 1961. Ten are about teaching machines; the other eight are about cybernetic controllers such as chemical computers.

Fig. 8 Gordon Pask.
Reproduced by permission of
Amanda Heitler



out for us, act as our control mechanisms, or perhaps most important of all, we can couple these seemingly uncontrollable entities together so that they can control each other. Why not, for example, couple the traffic chaos in Chicago to the traffic chaos of New York in order to obtain an acceptably self-organizing whole? Why not associate individual brains to achieve a group intelligence?

Pask, 'The Natural History of Networks' (1960, p. 258)

Figure 9 is a schematic of one of Pask's devices. A set of electrodes dips down vertically into a dish of ferrous sulphate solution. As current is passed through the electrodes, filaments of iron—'threads' as Pask called them—grow outwards from their tips into the liquid: Fig. 10 is a photograph of a stage in this process. Very simple, but so what? Three points need to be understood. First, the threads are *unstable*: they grow in regions of high current density but dissolve back into solution otherwise. Second, the threads grow *unpredictably*, sprouting new dendritic branches (which might extend further or dissolve). We can thus see how such a system might be seen as conducting a search through an open-ended space of possibilities, and we can also see that it has the *high variety*—another cybernetic term of art—required of a controller. Third, as extensions of the electrodes, the threads themselves influence current densities in the dish. Thus the present thread-structure helps determine how that structure




Fig. 9 Schematic of a Chemical Computer. *Source: Pask (1960, p. 247, Fig. 4)*

will grow as a function of the currents flowing through the electrodes, and hence the growth of the thread-structure exhibits a path-dependence in time: it depends in detail on both the history of inputs through the electrodes and on the emerging responses of the system to those. The system thus has a *memory*, so it can *learn*. This was Pask's idea: the chemical computer could function as an adaptive controller, in the lineage of the homeostat.

The question now becomes one of how such a system might be interested in us: how can a chemical computer be induced to substitute for the human manager of a factory, for example? As with Beer's biological computers, the answer is simple enough, at least in principle. Imagine there are two different sets of electrodes dipping into the dish of ferrous sulphate with its thread structure. One set are inputs: the currents flowing through them reflect the parameters of the factory (orders, stocks, cash-flow, etc). The other set are outputs: the voltages represent instructions to the factory (buy more raw materials, redirect production flows). There will be some determinate relationship between these inputs and outputs, fixed by the current thread-structure, but this structure will itself evolve in action and, as Ashby would have said, the combined system of factory plus controller will inevitably 'run to equilibrium.' Like a set of interacting homeostats, the chemical computer and the factory will eventually find some operating condition in which both remain stable: the factory settles down as a viable system, in Beer's terms, and the chemical computer, too, settles down into a state of dynamic equilibrium (until some uncontrollable perturbation arrives, when the search process starts again).

The magic is done—almost. Let me make one comment before I go on with the story. I need to refine my sense of biological computing as avoiding the Modern detour away from the world. Even when we talked about Beer and ponds, some detour away from nature-as-found was required—literally: Beer had to wander down to the nearest pond with a bucket to collect water and organisms and take them home again. In




Fig. 10 Threads Growing in a Chemical Computer. *Source:* Pask (1959, p. 919, Fig. 12)

the case of Pask's organic computers, the detour is more obvious. One does not find electrolytic cells lying around in nature; they have to be made. So this sort of cybernetic engineering does entail a detour. But I still want to note three points of contrast with the Modern paradigm. First, the cybernetic detour is noticeably shorter than those typical of Modern technoscience. Filling a bucket in a pond or setting up an electrolytic cell are immeasurably simpler operations than setting up a semiconductor industry. Second, and very importantly, biological computing *did not involve the same sort of detour* through knowledge as electronic computing. As I said earlier, all that Pask and Beer needed to know about their systems was that they were adaptive black boxes; there was no equivalent here of the knowledge of inner workings that characterises the modern sciences. And third, the other side of that, these black boxes were immediately performative devices: they did not themselves produce knowledge—unlike electronic computers with their representational inner states, for example, they acted, rather than thought. So I want to stress that even when these nuances are added, the contrast

between biological computing and the Modern paradigm remains. Now back to the story.

Pask thought through at least two further aspects of biological computing. First, there was the question of how to get the process of coupling the computer into human affairs going. How should the threads begin to learn about running a factory? One answer was to envisage a ‘catalyst,’ a system that would send current through the ‘least visited’ electrodes, thus fostering a variety of interactions with the factory and enabling the computer to interrogate the factory’s performance on a broad front. Of course, second, the procedure of simply letting the computer and the factory search open-endedly for a mutual equilibrium would almost certainly be disastrous. Who knows what idiotic instructions the computer would issue to the factory before stability was approached, or how quickly the factory would wreck the computer?

Pask therefore imagined that a human manager would be allowed to *train* the controller, monitoring the state of the factory and the machine’s responses to that, and approving or disapproving those responses by injecting pulses of current as appropriate to reinforce positive tendencies in the machine’s evolution, as indicated in Fig. 11. Pask noted that this kind of training would *not* take the form of the manager dominating the controller and dictating its performance; there was no way that could be done. In fact, the interaction would have to take the form of a ‘partly competitive and partly collaborative game . . . after an interval, the structured regions [in the controller] will produce a pattern of behaviour which the manager accepts, not necessarily one he would have approved of initially, but one he accepts as a compromise.’ Thus the manager and the controller come into homeostatic equilibrium with one another at the same time, in the same way, and in the same process as the controller comes into equilibrium with the factory. ‘At this point the structured region will replicate indefinitely so that its replica produces the same pattern of behaviour. The manager may thus be removed and the assemblage will act as an organic control mechanism in the industry’ (p. 169).

My comment now is that here we have a contrast between the nonModern and Modern paradigms that I have not mentioned before. In the Modern paradigm, the

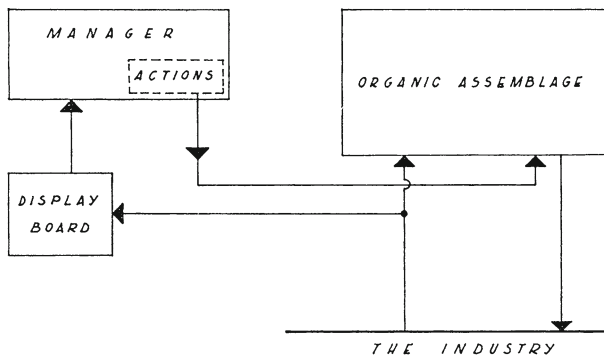


Diagram 2.

Fig. 11 Training a Chemical Computer. Source: Pask (1958, p. 169, diagram 2)

impulse is to bend nature to our will. If a digital computer does not execute the operations we envisage for it, something has gone wrong. Biological computing entailed a much more *symmetric* relation between the human and the nonhuman—a ‘conversation,’ as Pask put it, a ‘compromise,’ in which human performances and goals, the specifics of management, were themselves liable to open-ended transformation—mangling (Pickering 1995)—in negotiation with ponds or electrolytic cells, performative black boxes. Now I should complete the story of Pask’s devices.⁸

I have emphasised that biological computers remained Black Boxes in their inner workings, but a field of knowledge did enter into their construction in relation to their couplings. In Beer’s 1960 vision of the automatic factory, the input and output channels to the U-machine were specified by the designer, Beer himself. Which variables in the factory the U-machine should pay attention to, and which variables it should act on, were hard-wired into the system—as if the brain’s sensory and effector channels were fixed forever in advance. There is nothing surprising about this to the Modern imagination—it seems one of the more unremarkable features of Beer’s vision. But Beer and Pask wanted to go further. Pask imagined a situation in which a transformation in production methods upset Bill Smith, the foreman, whose happiness then turned out to be crucial to the smooth running of the entire factory. In Pask’s scenario, the human manager notices this, and switches some of the production to a project that makes no profit in itself but keeps Bill happy. Could one imagine a biological computer adaptive enough to pick up on new variables that had not been specific in advance—a device which could develop new senses, so to speak?

In a 2001 festschrift for Pask, Beer recalled that in 1956 or 1957 he was visiting London from Sheffield, and spent most of the night with Pask at the latter’s flat in Baker Street, as he often did. They first had the idea of exploring the robustness of Pask’s chemical computers by chiselling out sections of established threads and seeing what happened. It turned out that the systems were very robust—another contrast with Modern computing. ‘And yet these demonstrations, though exciting at the time, were somehow recognized to be trivial’ (Beer 2001, pp. 554–555):

‘Adaptation to the unexpected’ should mean more than this, and yet there must be limits. I was already developing my theory of viable systems, and often used myself as an example. But what if someone pulled out a gun and shot me. Would that be proof that I am not after all a viable system? Surely not . . . Suddenly Gordon said something like, ‘Suppose that it were a survival requirement that this thing should learn to respond to sound? If there were no way in which this [sound] ‘meant’ anything [to the device], it would be equivalent to your being shot. . . We need to see whether the cell can learn to reinforce successfully by responding to the volume of the sound.’ It sounded like an ideal critical experiment. I cannot remember what exactly the reinforcement arrangements were, but the cell already had them in place . . .⁹ And so it was that two very tired young

⁸ What follows is indebted to an essay by Peter Cariani (1993) which first rescued key aspects of Pask’s work from obscurity, and I am very grateful to Cariani for fascinating discussions around the topics at issue.

⁹ As usual, Pask was himself not very forthcoming on the practical details. Pask (1958, pp. 166–167) sketches out a means for monitoring the development of threads in some region using four ancillary electrodes. Two electrodes emit current periodically (presumably to inhibit growth of threads from them) and

men trailed a microphone down into Baker Street from the upstairs window, and picked up the random noise of dawn traffic in the street. I was leaning out of the window, while Gordon studied the cell. ‘It’s growing an ear,’ he said solemnly . . . A few years later Gordon was to write (Pask 1960, p. 261):

We have made an ear and we have made a magnetic receptor. The ear can discriminate two frequencies, one of the order of fifty cycles per second and the other of the order of one hundred cycles per second. The ‘training’ procedure takes approximately half a day and once having got the ability to recognize sound at all, the ability to recognize and discriminate two sounds comes more rapidly . . . The ear, incidentally, looks rather like an ear. It is a gap in the thread structure in which you have fibrils which resonate at the excitation frequency.

There is something truly remarkable about this episode. I can think of no equivalent in the history of Modern science and technology. As Beer put it in 2001, ‘It could well have been the first device ever to do this [develop a new sense], and no-one has ever mentioned another in my hearing’ (Beer 2001, p. 555).¹⁰ It is common enough in the history of electronic computing to add senses to a machine—Grey Walter made it possible for his tortoises to hear by wiring a microphone into their circuits—but this has always been by design, something imposed from the outside by the designer. Beer and Pask, in contrast, simply exploited the inner agency of their adaptive Black Box—a possibility which I think could not even be imagined in the Modern paradigm. One way to see what is at stake here would be to say that the Modern detour through knowledge and away from the world can also be *a block*, a trip that forecloses options that Beer and Pask’s work demonstrates lie actually already at hand. Perhaps we are more impressed by this technoscientific trip than we should be.¹¹

We have come to the end of the technical history of biological computing, apart from the very recent and rudimentary re-awakening that we began with. In the 1960s, both Beer and Pask developed their cybernetics in new directions. In management, Beer articulated his Viable System Model, in which organisational structures and information flows aimed to emulate the adaptive qualities of the human brain and nervous system. Pask focussed on the development of interactive and conversational teaching machines, as well as making very imaginative contributions in the theatre, architecture and robotic art (Pickering 2002, 2007). But to stay with my topic a little

footnote 9 continued

the others register some response. The trick would be to look for changes in response correlated with sounds in the required range (as detected by a filter attached to the microphone). Positive correlations could then be encouraged by injecting more current into the assemblage as they occur.

¹⁰ The question of how we would recognise a new sense if we came across one arises here: see Cariani (1993).

¹¹ I have been seeking to clarify the ontological contrast between Beer and Pask’s biological computers and their Modern counterparts, but Pask also discussed in very interesting ways the epistemological difference between Modern science and cybernetics (for example, Pask 1958, 172–173). He understood the former as having a distanced and intrinsically representational relation to matter going through the experimental testing of pre-given hypotheses. Cybernetics, in contrast, was about establishing productive and performative relations via open-ended and intimate engagements, from which behavioural knowledge was an output.


longer, I will end with a discussion of what one might call the spirituality of biological computing.

In Modernity, science and spirituality have nothing to do with one another. It is part of the Modern Constitution, as Bruno Latour (1993) calls it, that they inhabit disjoint ontological realms. It is just another aspect of cybernetics' nonModernity that it tended to erase that disjuncture. Beer, in particular, was profoundly religious from an early age, and ended his life as a self-described Tantric yogi, so I want here to explore some aspects of his spirituality and how they connect to his work in biological computing (Fig. 12).

We can proceed in two stages. First, in a 1966 essay entitled 'Cybernetics and the Knowledge of God,' Beer began by pointing to the finitude of the human mind: 'Each of us has about ten thousand million neurons to work with. It is a lot, but it is *the* lot. . . [T]his means that there is a strict mathematical limit to our capacity to compute cerebrally—and therefore to our understanding. For make no mistake: understanding is mediated by the machinery in the skull' (Beer 1966, p. 294). What is to be admired and wondered at, therefore, is not our knowledge of the world—or the world-as-known—but the world itself, which Beer referred to as 'an irreducible mystery: that there is anything' (p. 298). And, for Beer, this is where God comes in: '*God is what explains the mystery*' (p. 299).

What Beer points to in this essay is thus a sort of gestalt switch lying between Modern science and his cybernetics. If Modernity proceeds as if nature were exhaustible by knowledge, and as if God were somewhere else, Beer here expresses his

Fig. 12 Stafford Beer, 1975.
Source: Beer (1994, p. 315).
Reproduced by permission of
John Wiley & Sons and Allenna
Leonard



suspicion of representation, as somehow veiling the spiritually-charged thing in itself from us. And in this essay he emphasised this contrast himself (pp. 294–295):

To people reared in the good liberal tradition, man is in principle infinitely wise; he pursues knowledge to its ultimate . . . To the cybernetician, man is part of a control system. His input is grossly inadequate to the task of perceiving the universe . . . there is no question of ‘ultimate’ understanding. . . [I]t is part of the cultural tradition that man’s language expresses his thoughts. To the cybernetician, language is a limiting code in which everything *has* to be expressed—more’s the pity, for the code is not nearly rich enough to cope (pp. 294–295). . . Will you tell me that science is going to deal with this mystery [of existence] in due course? I reply that it cannot. The scientific reference frame is incompetent to provide an existence theorem for existence. The layman may believe that science will one day ‘explain everything away’; the scientist himself ought to know better (p. 298).

Next, we can turn to a book Beer published in 1986 called *Pebbles to Computers* (Blohm et al. 1986). This is a nonlinear history of computing, running from simple counting to digital computers but also embracing, for example, Stonehenge as an astronomical computer and Peruvian *quipus*, beautiful knotted threads, as calculational devices. Here Beer goes beyond his awe at the sheer excess of matter over representation to emphasise an again spiritually-charged wonder at matter’s *performativity* and, especially its computational performance. There are several fascinating passages on this, but to keep things short let me mention just one. Under the heading ‘Nature Calculates,’ Beer comments on a photograph of the Gatineau River (Fig. 13) that catches the endless complexity of the water’s surface (p. 51):

This exquisite photograph of water in movement . . . has a very subtle message for us. It is that nature’s computers *are* that which they compute. If one were to take intricate details of wind and tide and so on, and use them . . . as ‘input’ to some computer simulating water—what computer would one use, and how express the ‘output’? Water itself: that answers both those questions.

Nature does not need to make any detours; it does not just exceed our computational abilities, in effect it surpasses them in unimaginable ways. In a poem on the Irish Sea in the same book, Beer also talks about nature as doing a ‘bigger sum’ than ours, exceeding our capacities in way that we can only wonder at, ‘shocked’ and ‘dumbfounded.’¹² And to emphasise the point, Beer again makes a contrast with the ontological stance of Modern Science (p. 53):

The uneasy feeling that [this poem] may have caused derives, perhaps, from insecurity as to who is supposed to be in charge. Science (surely?) ‘knows the score.’ Science does the measuring after all. . . . But if art is said to imitate

¹² To inject a personal note, these words read very true to me. They remind me of my days as a postdoctoral researcher in theoretical particle physics, spending weeks and months trying and failing to understand mathematically how quarks interact, while being painfully aware that the quarks themselves were doing their own thing all the time, in real-time, throughout the cosmos (or so I believed at the time). That kind of experience leaves one with a feeling for scientific knowledge as a pale simulacrum of the world (or not even that in my case), a simulacrum one nevertheless finds it hard not to mistake for the thing in itself.



Fig. 13 The Gatineau River, Quebec. *Source:* Blohm et al. (1986, p. 51)

nature, so does science. . . Who will realize when the bathroom cistern has been filled—someone with a ruler and a button to press, or the ballcock that floats up to switch the water off? **Nature** is (let it be clear that) **nature** is in charge.

‘Nature is in charge’ is a succinct expression of what I would call Beer’s hylozoism: his spiritually-charged awe at the activity and powers of nature in relation to our inability to grasp them representationally. And I can offer two comments on this. First, while it is not necessary to subscribe to Beer’s hylozoism to be interested in biological computing, one can certainly see how such an ontology hangs together with Beer and Pask’s strange projects. If we are already and inevitably plunged into the infinite computing power of nature, the long trip through chip manufacture and digital computation certainly appears as a massive detour. Second, we can take this line of thought further. Beer’s thought was that mind is everywhere, immanent in matter itself. I focussed earlier on naturally occurring adaptive systems as model brains, but now we can add the Gatineau River, the Irish Sea and toilet-flushing mechanisms (all watery phenomena, interestingly). If we add the cybernetic notion that all of these mind-phenomena are reciprocally coupled to one another in a homeostatic fashion, then we move towards a conception of minds, including the human mind, as intrinsically decentred and merging into one another. And this takes us from the starting point of this section, Beer’s awe at the excess of matter, towards a Buddhist and very nonModern image of the human mind as extending beyond the Modern self in performative engagement with the nonhuman mind-stuff of the cosmos. *Pebbles to Computers* comes to an end with a quotation Beer attributes to the *hsin hsin ming* by Sengstan, the third Zen patriarch (d. 606) (p. 105):

Things are objects because of the mind;
 The mind is such because of things.
 Understand the relativity of these two
 and the basic reality: the unity of emptiness.

In this emptiness the two are indistinguishable
and each contains in itself the whole world.

I should try to sum up. What I have been trying to do is contrast a nonModern cybernetic paradigm with the Modern technoscientific paradigm that Heidegger railed against. I have tried to get this into focus by talking about Beer and Pask's strange and wonderful initiatives in biological computing. This last section of this essay helps me, at least, to get the ontological aspect of the contrast into focus better. Cybernetics placed the inexhaustible powers of nature at its centre, where Modern technoscience has representation instead. Modernity is thus characterisable by its enormous detours through knowledge and industry, while cybernetics stays close to the performative world as found, abstaining from the attempt to unwrap Black Boxes. And if Modernity is defined by projects of domination, then cybernetics is marked by a symmetric accommodation to the ultimately uncontrollable. If I developed this last point, especially as it includes our dealings with other people, it would get us into a discussion of a distinctly cybernetic politics which I find distinctly preferable to the politics of Modernity today. Finally, I outlined a spiritual stance that also differentiates cybernetics from Modernity in its hylozoist collapse of spirit and matter. Cybernetics, one might say, retains the space for wonder that Modern technoscience obscures—and that is not the least of this strange paradigm's virtues.

References

- Ashby, W. R. (1948). Design for a brain. *Electrical Engineering*, 20, 379–383.
- Ashby, W. R. (1952). *Design for a brain* (2nd ed. 1960). London: Chapman & Hall.
- Bear, G. (1997). *Slant*. New York: Tor.
- Beer, S. (1959). *Cybernetics and management*. London: English Universities Press.
- Beer, S. (1960). Retrospect—American diary, 1960, In S. Beer (Ed.), *How many grapes went into the wine? Stafford Beer on the art and science of holistic management* (pp. 229–309). New York: Wiley, 1994.
- Beer, S. (1962a). Towards the automatic factory, In H. von Foerster & G. Zopf (Eds.), *Principles of self-organization: Transactions of the University of Illinois symposium on self-organization, Robert Allerton Park, 8 and 9 June, 1961* [sic: actually 1960] (pp. 25–89). New York: Pergamon. (Reprinted in Beer, *How many grapes went into the wine? Stafford Beer on the art and science of holistic management* (pp. 163–225). New York: Wiley, 1994).
- Beer, S. (1962b). A progress note on research into a cybernetic analogue of fabric, *Artorga*, Communication 40, April 1962. (Reprinted in Beer, *How many grapes went into the wine? Stafford Beer on the art and science of holistic management* (pp. 24–32). New York: Wiley, 1994).
- Beer, S. (1966). Cybernetics and the knowledge of God. *The Month*, 34, 291–303.
- Beer, S. (1994). *The Falcondale collection: Stafford Beer initiates an audience into the world of systems and managerial cybernetics*. Videotapes and transcript. Liverpool: JMU Services Ltd.
- Beer, S. (2001). A filigree friendship. *Kybernetes*, 30, 551–559. doi:[10.1108/03684920110391779](https://doi.org/10.1108/03684920110391779).
- Blohm, H., Beer, S., & Suzuki, D. (1986). *Pebbles to computers: The thread*. Toronto: Oxford University Press.
- Cariani, P. (1993). To evolve an ear: Epistemological implications of Gordon Pask's electrochemical devices. *Systems Research*, 10, 19–33. doi:[10.1002/sres.3850100305](https://doi.org/10.1002/sres.3850100305).
- de Latil, P. (1956). *Thinking by machine: A study of cybernetics*. London: Sidgwick and Jackson.
- Deleuze, G., & Guattari, F. (1987). *A thousand plateaus: Capitalism and schizophrenia*. Minneapolis: University of Minnesota Press.
- Harnden, R., & Leonard, A. (Eds.) (1994). *How many grapes went into the wine? Stafford Beer on the art and science of holistic management*. New York: Wiley.
- Heidegger, M. (1976 [1954]). The question concerning technology, In D. Krell (Ed.), *Martin Heidegger: Basic writings* (pp. 287–317). New York: Harper & Row.

- Latour, B. (1993). *We have never been modern*. Cambridge, MA: Harvard University Press.
- Noble, D. F. (1986). *Forces of production: A social history of industrial automation*. Oxford: Oxford University Press.
- Pask, G. (1958). Organic control and the cybernetic method. *Cybernetica*, 1, 155–173.
- Pask, G. (1959). Physical analogues to the concept of growth. In *Mechanisation of thought processes: Proceedings of a Symposium held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958*, 2 vols (pp. 877–928). London: Her Majesty's Stationery Office.
- Pask, G. (1960). The natural history of networks. In M. Yovits & S. Cameron (Eds.), *Self-organizing Systems: Proceedings of an Interdisciplinary Conference, 5 and 6 May* (pp. 232–263). New York: Pergamon.
- Pask, G., & Curran, S. (1982). *Micro man: Computers and the evolution of consciousness*. New York: Macmillan.
- Pickering, A. (1995). *The mangle of practice: Time, agency, and science*. Chicago: University of Chicago Press.
- Pickering, A. (2002). Cybernetics and the mangle: Ashby, Beer and Pask. *Social Studies of Science*, 32, 413–437. doi:10.1177/0306312702032003003.
- Pickering, A. (2004a). The science of the unknowable: Stafford Beer's cybernetic informatics. In R. Espejo (Ed.), *Tribute to Stafford Beer*, special issue of *Kybernetes*, 33 (2004), 499–521.
- Pickering, A. (2004b). Mit der Schildkröte gegen die Moderne: Gehirn, Technologie und Unterhaltung bei Grey Walter, (transl. by Gustav Rossler), In H. Schmidgen, P. Geimer & S. Dierig (Eds.), *Kultur im Experiment*. (pp. 102–119). Berlin: Kulturverlag Kadmos. (English version: The tortoise against modernity: Grey Walter, the brain, engineering and entertainment. In *Experimental cultures: Configurations between science, art, and technology, 1830–1950*. pp. 109–122. Berlin: Max Planck Institute for the History of Science, preprint 213, 2002).
- Pickering, A. (2005a). A gallery of monsters: Cybernetics and self-organisation, 1940–1970. In S. Franchi & G. Güzeldere (Eds.), *Mechanical bodies, computational minds: Artificial intelligence from automata to cyborgs* (pp. 229–45). Cambridge, MA: MIT Press.
- Pickering, A. (2005b). Cybernetics as nomad science. Annual Meeting of the Society for Literature, Science, and the Arts, Chicago, 10–13 November 2005.
- Pickering, A. (2007). Science as theatre: Gordon Pask, cybernetics and the arts. *Cybernetics & Human Knowing*, 14(4), 43–57.
- Pickering, A. (forthcoming). *Sketches of another future: The cybernetic brain, 1940–2000*. Chicago: University of Chicago Press.
- Walter, W. G. (1953). *The living brain*. London: Duckworth. (2nd ed. Penguin, 1961).
- Wiener, N. (1948). *Cybernetics, or control and communication in the animal and the machine*. Cambridge, MA: MIT Press.