

EMBEDDED INTELLIGENCE and processes of self-organization: the case of intelligent vehicle/highway systems.

MANUEL DE LANDA

One useful way of viewing the evolution of computer technology is as a slow migration of problem-solving skills from the human body to formal systems, and from there to electro-mechanical devices. That is, when Aristotle created his famous syllogistic logic, he in effect transferred some elementary skills from humans to a mechanical recipe (or algorithm). Later on, nineteenth century logicians (Boole, Frege) enlarged the capabilities of these algorithms to encompass other deductive logic skills. When Alan Turing created his imaginary machine to execute these mechanical recipes, and then, under the pressure of WW II, he, and others like John Von Neumann, finally embodied this abstract device in a concrete machine, one more link in this transfer of mechanical intelligence took place. Thus, it was a slow but real migration from the human nervous system to computers in three steps. Problem-solving skills, which began as informal heuristics (or rules of thumb) embodied in flesh, end up as algorithms in silicon, via the intermediate step of combinatorial rules working with physical inscriptions on pieces of paper. Nowadays, this second step has been eliminated by knowledge engineers, who transfer human heuristics directly to the machine, through a process of intense interviewing of human experts. This is a process in which informal, half-conscious skills are brought to the surface, articulated and formalized, and then compiled together to create a so-called "Expert System", the most successful product of Symbolic Artificial Intelligence to date.

A second migration is taking place today. Mechanical intelligence, so far confined to specific devices like the personal computer, is beginning to move outwards into the human environment, into its buildings and appliances, into its roads and vehicles. And it is this second migration that promises to alter social patterns of interaction in a more radical way. As Mark Weiser of the Xerox Palo Alto Research Center has said, "the most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it." He gives the example of the electrical motor, which begins its evolution as a very visible entity in the workshop, the central source of power for all tools, and then goes on to vanish from view, as miniaturization allows each tool to have its own motor. The idea is that, once motors or computers become part of the environment, they cease to command our attention and allow us to concentrate on the tasks we want to accomplish. In this sense, invisible computing is the exact opposite of immersive virtual reality technology which tends to focus our full attention on the medium itself, instead of letting us concentrate on what we want to achieve. As Weiser remarks, "only when things disappear in this way are we free to use them without thinking and so to focus beyond them on new goals."¹ However, embedded intelligence raises a host of issues regarding the potential misuse of these technologies. In a world where most display surfaces around us have become smart tabs, pads and boards, and where these devices need to know the location and function of every human in a building, the spectre of a Panopticon surveillance system arises. The problem of protecting individual privacy in these new environments is already an intensely debated issue in the context of the Internet as well as in the world of credit cards. The solution proposed in all three cases involves the use of cryptography to make sure that only the information relevant to the system (transaction amount and not purchasing habits, in the case of credit cards) flows through the circuits. (Yet, as the Clipper controversy in the U.S. demonstrates, privacy through cryptography may be a thing we will need to struggle for, and not something that we can take for granted). Beyond privacy-related issues, the social implications of embedded intelligence will revolve around the question of what kind of structures we will build with these technologies. If we imagine that, thanks to miniaturization,

the very materials that we use to build social structures will become smart, the question is whether we should build command-and-control hierarchies with this "intelligent stuff" or should we rather build self-organized market-type structures out of this "programmable matter". In one case, we proceed by sorting out people into internally homogenous ranks, and allocate control at the top rank. In the other, we allow control to remain decentralized and we aim at articulating people with heterogenous goals by a mutual meshing of their skills and needs. Human history shows that it is much easier for us to build hierarchies than it is to create dynamic mesh-works. We do not even seem to have an adequate theory of markets as self-organizing structures. Part of the blame for this must be attached to the discipline of economics, which has failed to come to terms with the real dynamics of markets. This comes from partly subscribing to the obsolete notion of an "invisible hand", partly by its failure to see monopolies and oligopolies for what they always have been, control hierarchies operating outside self-organized markets. Fortunately, the situation is beginning to change. Theories of self-organized processes are now thirty years old, and the early insight that non-linear matter-energy flows are capable of spontaneously generating order out of chaos, has now developed into an international research program which is beginning to have applications in the social sciences. Specifically, the idea that, in the right conditions, wholes that are more than the sum of their parts (i.e. synergistic wholes) can emerge out of the local dynamics of their components, is starting to find expression in the field of management science. Management should aim, according to these theories, not at imposing a preconceived plan upon subordinates but at catalysing the creation of a self-renewing network of decision-making processes. Instead of rigidly defining behaviours from above, the aim is to gently constrain them so that they mesh with one another forming an auto-catalytic loop, a self-renewing set of human actions, whose coherence emerges from its local dynamics, due to reciprocal, anticipatory adaptation.² It is my belief that the most innovative uses of embedded intelligence will come from an interaction with non-linear group dynamics. That is, invisible computing needs to be designed to support and make easier the emergence of self-organized mesh-works of activity in human institutions. The particular institutions I would like to explore here are those related to the organization of transportation functions in society. As historians have recently discovered, cities throughout history have emerged around geographical points of intense motion of people, goods and information, such as the intersection of two rivers or of ancient trade routes.³ It is almost as if early cities emerged as a kind of "mineralization" of a trade post, the growth of a mineral infrastructure of roads and buildings around a flow of goods, further spawning the emergence of other urban settlements by setting into motion other flows. Hence, transportation technology, however primitive, is one of the basic ingredients in the creation of urban life. At the same time, as cities grow in size and population, the shortcomings of this technology may contribute to their death, by causing crowding, congestion and pollution beyond the abilities of institutions to handle. Today we are rapidly approaching this point, as many cities have become unable to cope with the transportation problems created by the technology of motorized vehicles and a centralized traffic signal system. In the 1970's, the U.S. government began to consider the possibility of using Artificial Intelligence (AI) to solve these problems, such as the Automated Highway System pioneered at Ohio State University.⁴ Though the 1980's saw a virtual freeze of funding in this area, research has been intensified in the last few years, with budgets increased from two to two hundred million a year.

Here too we find a dual approach to intelligent vehicle/highway systems, one stressing centralized decision-making, the other self-organized, emergent behaviour.

Interestingly, each of these approaches is reflected in the kind of AI used to frame and define the problem. When one approaches a given system, whether natural or human-made, as if it

was made of homogenous parts (i.e. a hierarchy), the easier thing to do is to take a top-down approach. One analyses or dissects the whole into its homogenous components, and then puts the system back together again. However, when the system under consideration is a meshwork of heterogenous parts, analysis fails because the synergistic properties of the whole arise out of the interactions between its components, and these interactions are lost when one dissects them. Therefore a synthetic or bottoms-up approach is what is needed here, as exemplified by the discipline of Artificial Life, where a population of abstract animals are allowed to live and reproduce in the context of other animal populations, with the hope that a heterogenous meshwork resembling a real ecosystem will emerge spontaneously. The key here is that only local rules of interaction must be defined explicitly, all global behavior must self-organize. In the case of AI, the symbolic approach takes the top-down, analytical route, while the neural net and animal approaches take the bottoms-up, synthetic road. Symbolic AI has been successful at modelling evolutionary late skills, such as chess playing and theorem proving, while Behavioral and Connectionist AI have excelled in the simulation of more basic skills, such as face recognition and pattern-based reasoning. Philosophers, pondering the implication of this dichotomy, have concluded that the human mind may indeed be a hybrid of hierarchical and meshwork structures. Much as Expert Systems were forced to use a serial computer (with its centralized processing unit) to simulate a parallel computer (e.g. a decentralized system of production rules), the human brain may be a parallel computer that at some point in its evolution became capable of simulating a serial one. It is this latter one that we experience as our stream of consciousness.⁵ And much as the human mind may be a hybrid structure, most of our institutions may also be pictured as varying mixtures of centralized and decentralized control, with pure command hierarchies and pure mesh-works being the exception.

Current attempts at embedding mechanical intelligence into the non-linear dynamical system formed by motorized vehicles, in interaction with the road and traffic signal infrastructure, have mostly taken the top-down approach. Again this is not surprising since hierarchical thinking seems to be much more entrenched in our habits and organizations. However, the actual history of our public transportation systems has involved both central planning by government authorities and self-organized processes arising from complex interactions between users, fares, running costs and number of lines. Public transportation within a city and that between cities and suburbs exhibit mutual enhancement, and this leads to non-linear threshold effects, as when adding one more bus line leads to a sudden and dramatic increase in bus travellers.

Attempts by central authorities to increase use of a given mode of transport typically fail when the amount of investment does not reach the critical threshold. However, when the whole dynamical system is poised near a bifurcation, central commands may have effects that go well beyond the investment made. Hence, city transportation is another example of a hybrid system, with the centralized component being much more visible than the self-organized one, which has only recently been revealed by computer simulations.⁶

The scientific study of vehicular traffic dates from the 1950's, when researchers began to approach it as a continuous fluid, and discovered that near bottlenecks traffic shock waves may form which propagate throughout the fluid. A more bottoms-up approach, in which this fluid is decomposed into its car-driver units, has revealed that traffic flows owe their coherence to the drivers' psychological reaction to proximity to other vehicles. These dynamic flows have at least two stable states. At low vehicular concentrations, we have a dynamical regime governed mostly by the drivers' desired speed. In this stable regime, the patterns that form are of the platoon type: short trains of cars at short distance, with larger distances

between platoons. On the other hand, as the concentration of vehicles reaches a critical threshold, the fluid switches to another regime, a collective flow in which each car moves at a speed totally constrained by its neighbours. Centralized decision-making must operate within these dynamical limitations, that is, it can only effect switches between regimes. (A particularly important parameter which may be manipulated is the percentage of time cars spend stopped at traffic lights.)⁷

What these two examples illustrate is that centralized control always operates within the constraints of a larger, self-organized system. The efficiency of its decisions will depend on how close to a threshold the system is. Away from that threshold, an inefficient amount of effort will have to be spent to effect a switch in the mode of transportation people use (e.g. car vs. bus), or in the state of the fluid of vehicles. Near the threshold, relatively small amounts of invested effort will be amplified by the system itself, resulting in greater overall efficiency. However, benefiting from these recent discoveries will entail a change in our thinking, since complete planning and control must be given up, and the dynamics of the traffic system itself need to be made part of a decision-making process. One approach to this question, pioneered by researchers studying the self-organization of insect societies, is to decentralize decision-making and to allow the environment to contribute to the solution of the traffic problem. J.L. Deneubourg and his colleagues call this approach one of "collectively self-solving problems". They point out that in the case of ant colonies, a kind of "swarm intelligence" emerges out of a few behavioural rules and the possibility of strong interactions among the ants. For example, if one ant finds a food source it recruits another ant via chemical communication. Because this recruiting operation is repeated by the second ant, and then repeated again and again, the effects of the interactions are amplified, and trails of ants emerge. If two sources of food are discovered simultaneously, the one at a shorter distance from the nest will form a trail more rapidly, and win over the rival source. No centralized decision needs to be made to pick one source over the other. In a sense, the closer food source selects itself. Hence the distribution of food in the environment contributes to the solution. These researchers go on to point out that unlike ants our motorized vehicles interact only weakly, and hence no exploratory trails can form. Embedding computers on the cars, so that they could attract one another with varying strength, may lead to the emergence of such "swarm intelligence" in the traffic flow. For example, if the attraction between cars diminished with traffic congestion, bottlenecks would be automatically avoided. Similarly, a vehicle that discovers a new uncongested route would automatically share the benefits with others, since it would attract others to form a trail behind it.⁸

What will perhaps be the first instance of AI applied to traffic problems has decidedly not taken the decentralized approach. I am referring here to the so-called "Smart Corridor", a thirteen mile stretch of the Santa Monica freeway, which is supposed to begin operating sometime in 1994. The old centralized decision-making structure has been kept intact, only augmented via Expert System technology, both for traffic and surveillance management. The heuristic know-how of traffic experts has been painstakingly extracted from their bodies, and converted into rules, which are then added together into a knowledge base, and brought into action via an inference engine. The system itself, however, is not capable of learning. Heuristic know-how develops by interacting with reality, and human experts constantly update their knowledge from the results of these interactions. The Smart Corridor, on the other hand, cannot itself learn new things as it aids in the solving of real problems, that is, it cannot automatically update its knowledge base with new rules, and it must have human analysts do that by hand. It cannot deal with unexpected events occurring simultaneously, or with time-dependent consequences of a chain of events, or even perform truth-maintenance in real-time, a necessary skill since the validity of deduced facts changes overtime.⁹

It may be objected that these are not real limitations, since an Expert System is not supposed to take over traffic management, but simply assist human managers as they make their own decisions. Hence, the human operators could compensate the shortcomings of the digital assistant. Yet, by not decentralizing control to the vehicles themselves, even the human operators may become overwhelmed by the non-linearity which exists in any real traffic situation: bad weather conditions, uneven driver performance, non-cooperating pedestrians, stalled vehicles and so on. What is needed is a traffic system that learns, perhaps involving a fast genetic algorithm capable of breeding solutions to unforeseen situations, and a new approach to management that does not command the form the traffic should take but attempts to catalyse the formation of efficient patterns, such as long lasting platoons. Ideally, the whole vehicle/ roadway/traffic signal system should become an emergent optimizer.

I would like to conclude this presentation with some general remarks. Although most human institutions have been dominated by the hierarchical command component of their mixtures, they ultimately operate within vaster systems, where the meshwork component predominates. Cities as a whole are mesh-works of heterogenous elements in which dynamics severely constrain the decisions and actions of their centralized governments. The same seems to be true of smaller scale structures, such as commercial firms. Embedded intelligence could allow small entrepreneurs to create firms in which catalysing the formation of mesh-works of decision-making processes would be relatively easier. Invisible computing would function as part of the non-linear dynamic system, which a firm uses in its day to day operation, facilitating the emergence of self-organization. Mesh-works of such small firms could then have access to different economies of scale than their larger, centralized counterparts, and therefore would be able to compete more successfully. In a world that is becoming increasingly homogenized, and in which much of the homogenization is being caused by command hierarchies (central governments, large corporations), it could not be but healthy to add more meshwork to the mix. However, the antidote to homogeneous articulation is not disarticulated heterogeneity. The current situation in the Balkans is proof of the dangers of that. Rather, we must learn to create structures without homogenizing, learn to articulate the heterogenous as such. And, of course, the methods needed to do this cannot be centrally developed or commanded into existence. In this regard, the main contribution of embedded intelligence would be to help create the environments where these methods could evolve themselves, where humanity could learn to self-organize.

FOOTNOTES:

(1) M. Weiser. The Computer for the 21st Century. (Scientific American, Vol 265, Number 3.) Page 94

(2) F. Malik and G. J. B. Probst. Evolutionary Management in Self-Organization and Management of Social Systems. H. Ulrich and G. J. B. Probst. (Springer Verlag, Berlin 1984). Page 109

(3) Fernand Braudel. Capitalism and Material Life. (Harper Colophon, NY 1973), Page 389

(4) D. Rack, D. Hoskins and D. Malkoff. Intelligent Road Transit. (AI Expert, Vol 9, Number 4) Page 16

(5) Andy Clark. Microcognition: Philosophy, Cognitive Science and Parallel Distributed Processing. (Bradford, Cambridge Mass, 1990). Page 135

(6) D. Kahn, J. L. Deneubourg and A. De Palma. Public Transportation: A Dynamic Model of Mode Choice and Self-Organization. In Robert Crosby ed. Cities and Regions as Nonlinear Decision Systems. (AAAS, Washington DC, 1983) Page 63

(7) R. Herman. Remarks on Traffic Flow Theories and the Characterization of Traffic in Cities. In P.M. Allen and W. C. Schieve eds. *Self-Organization and Dissipative Structures: Applications in the Physical and Social Sciences*. (Univ. of Texas Press, Austin 1982). Page 266

(8) J. L. Deneuborg, S. Goss, R. Beckers and G. Sandini. Collectively Self-Solving Problems. In A. Blaboyantz ed. *Self-Organization, Emerging Properties & Learning*. (Plenum Press, NY 1991). Page 271

(9) D. Rock, D. Hoskins and D. Malkoff. *Op. cit.* Page 21