Biological Inspiration in the Design of Computing Systems

BY GIANLUCA TEMPESTI
Department of Electronics, University of York, York, U.K.

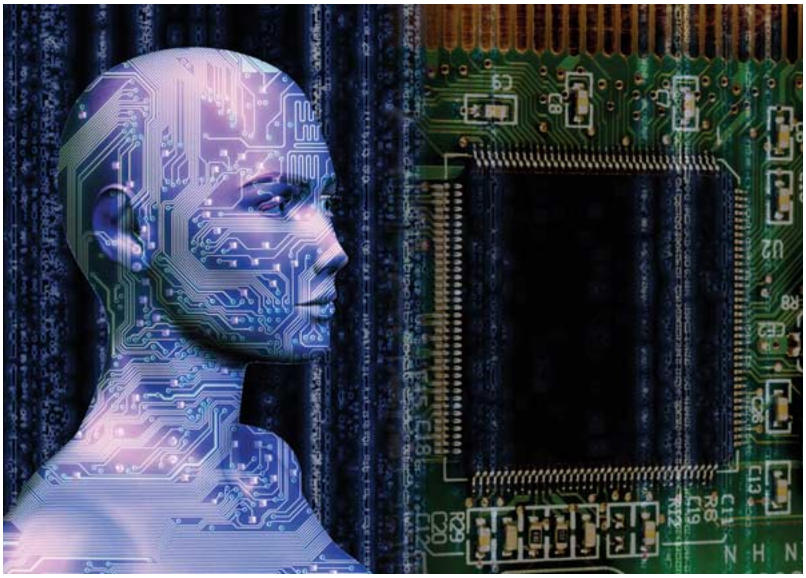

The field of computer engineering has ties to the world of biology that date back to its very dawn. The "artificial brain" paradigm inspired pioneers such as Turing and von Neumann to design what we now call computers. Even if this paradigm has lost much of its value in the decades that followed, as more details on the operation of the brain (and indeed of computers) were discovered, biological inspiration nevertheless lies at the very heart of all computing machines. It is therefore reasonable to wonder whether this kind of inspiration can still be useful to define the next generation of computing systems, a question that becomes all the more relevant as the complexity of hardware substrates slowly begins to approach that of biological organisms.

I. COMPLEXITY IN DESIGN

Few researchers would deny that facing this complexity is one of the main issues in the design of computing systems today and in the foreseeable future. This issue concerns most aspects of computing, and notably it concerns the design of computer hardware. Gone, for example, are the times when the limited amount of resources dictated the need for clever architectural tricks in processor design (microprogramming, to name but one example). The field has gone full circle and now the main challenge in the design of computing systems is represented not by the lack of resources, but rather by their abundance. Design tools and environments increasingly struggle to exploit the number of transistors available through the latest submicrometer technologies and these same technologies introduce layout and fabrication issues that have been solved (so far) only through massive investments in extremely costly processes and factories.

For the moment, this trend shows no sign of a slowdown. Material scientists are performing amazing feats to further shrink the transistor, while the nascent field of molecular-scale electronics promises to introduce unprecedented amounts of computing material. Likewise, the projected costs of the necessary layout and fabrication tools increase at an alarming rate.

II. THE NEXT GENERATION

Whether next-generation circuits will be based on smaller and smaller transistors or on drastically different

Digital Object Identifier: 10.1109/JPROC.2006.890088

molecular components is still a matter of contention, but largely irrelevant: system engineers will have access to more and more resources, but this evolution will introduce a set of major design and fabrication issues. To put it bluntly, we have no idea how to design systems that can operate on the next generation of substrates in an efficient and cost-effective manner. Network-on-a-chip approaches probably represent the future of system architectures, but this observation does not answer the question, since we do not really know how to design and program these systems.

The crux of the problem is that these new substrates, independently of the particular technology, will have properties that are very much unlike what is available today. First of all, the massive amounts of resources and the imperfections of the substrates imply that the top-down design flows that are in use today will be highly inefficient (irrelevant might actually be the correct word). Second, the increased sensitivity to soft errors will require the introduction of much more advanced fault tolerance techniques than what we have today.

III. BIO-INSPIRED SYSTEMS

In this context, enter bio-inspired systems, defined as systems that try to find inspiration from (and not, it is worth pointing out, to imitate) the world of biology to find solutions for the problems facing the design of computing systems. Nature has found ways to cope with complexity and fault sensitivity. A human being (an admittedly complex example) consists of approximately 60 trillion (60×10^{12}) cells. At each instant, in each of these 60 trillion cells, the genome, a ribbon of 2 billion characters, is decoded to produce the proteins needed for the survival of the organism. Faults occur at a very high rate, but are (in the majority of cases) successfully detected and repaired with little or no effect on the operation of the organism.

Even with all the necessary caveats (it is difficult to directly compare a biological organism with a computing system), it is therefore not surprising that an increasing number of researchers are turning to nature to try and find inspiration in the design of highly complex computing machines. This approach can take different forms, depending on which of the many natural mechanisms is chosen as a source of inspiration and on which of the many design and programming issues is chosen as a target. Evolutionary approaches use techniques inspired by the evolution of species to search highly complex solution spaces and go beyond the kind of hardware designs that can be obtained by following a methodological approach. Growth-based approaches try to tackle design and layout issues by observing how a genome codes the instructions for both the construction and the operation of the organism and exploits molecular self-assembly properties to simplify these processes. Learning--based approaches are seen as an alternative way to look at computation in highly parallel cellular systems and design adaptive systems able to tackle complex, unpredictable environments for tasks ranging from robot control to fault tolerance.

IV. BIO-INSPIRED ARCHITECTURES

These are just a few examples of how bio-inspired approaches are being used in computer design: analogies between the world of computer engineering and that of biology can be drawn, explicitly or implicitly, on many levels. And yet the field of bio-inspired hardware design is still far from mature, and plenty of major unsolved problems remain. Evolutionary approaches have hit a scalability wall that needs to be somehow circumvented to obtain useful results and go beyond what are normally called “toy” applications. Growth appears to be a useful paradigm for the layout of systems (see, for example, DNA scaffolding), but is still very much in its infancy and the extent to which it can be applied in practice still has to be proven. Learning is hampered, in the general case, by the very unpredictability that makes it so powerful for some specific applications: the computational behavior of very large learning networks is not well understood and has not been verified in a hardware implementation. Nevertheless, the promise of this kind of approaches remains and sources such as the International Technology Roadmap for Semiconductors place bio-inspired architectures and devices as emerging vectors for next-generation technologies. If some of the critical problems facing the field can be solved, the potential rewards could be vast: nature has been devising molecular-scale systems of astounding complexity using parameters, tools, and mechanisms that are very different from the ones we currently use in the design of computing machines. Is it not reasonable to think that some of these techniques, with all the necessary adaptations, could be useful to determine how to design our own man-made molecular-scale systems?

The need is there: the complexity and fragility of next-generation hardware will require novel design approaches and tools. In this context, biological inspiration in the design of computing systems has never lacked supporters, as a long-term research area. Maybe, just maybe, it is not that long-term any more. ■