Frontiers of Computer Science

Computer scientists work in an interdisciplinary realm between mathematics and engineering. Subsequent to the revolution in personal computing, the tangible benefits of this fertile interaction have tended to concentrate public attention on quantitative progress: cheaper memory, higher resolution graphics, increased storage capacity, and so forth. But computer scientists have long been toiling at the conceptual frontiers. The articles in this issue give a sample of the qualitative as well as quantitative progress being made.

One outpost at the frontier is the effort to build autonomous robots that function in unpredictable environments. The present generation of machines works well, but only if their tasks are well-defined (painting) and their surroundings are more or less unchanging (they are bolted to the factory floor). Traditional robotics relies on central modules that sequentially perceive the world, model it, and then plan and execute some action. Brooks describes a robot architecture in which the overall behavior of the machine is controlled by a network of simple computational elements operating concurrently. The addition of elements layer by layer yields the complex behavior needed for useful work. An advantage is quicker responses to a changing world. As the author himself points out, much work is needed before the performance of these systems can be fully evaluated.

Microprocessor performance is a subject dear to the hearts of computer users everywhere, yet how much more power can be squeezed out of present and future chips? One route is to use many independent processors in parallel. Another avenue, discussed by Fisher and Rau, is instruction-level parallelism. In this case the most basic operations in each processor—addition, multiplication, loading registers—are executed in parallel. The goal is a new generation of microprocessors that will carry out these operations in parallel, even when given a program written with a single sequential processor in mind. More background on this subject is provided in the research news story by Freedman.

Vision is the "remote sensing" that some biological organisms and robots use to guide themselves through the world. What makes the job hard for robots is that the world is a cluttered place, riddled with obstacles and mirages. The goal of computer vision research is twofold: to understand the principles behind vision and to construct robots that can best make use of these principles. Aloimonos and Rosenfeld review past work in computer vision and recent efforts to implement it. Building on early efforts to analyze static two-dimensional images and controlled three-dimensional "blocks-worlds," researchers are extending the theoretical foundations as they take on real-world scenes.

What would it be like if computers could be taught to understand the languages of humans, instead of the other way around? And what might we learn about human language as a result? Joshi discusses the field of natural language processing, an undertaking at the intersection of linguistics and computation. Much theoretical effort is devoted to grammars and parsers that specify and analyze natural languages. Another school of thought emphasizes the use of statistical information about real languages to guide the construction of useful systems. A practical application, multilingual translation by machine, has recently undergone a revival that will draw heavily on natural language processing.

The cutting edge in computer science has also found its way into the domain of molecular biology. Prediction of protein conformation from its amino acid sequence is a largely unsolved problem. A less explored but analogous task is the prediction of the twists and turns of RNA from its nucleotide sequence. Major *et al.* have tackled the RNA puzzle by viewing it as a constraint-satisfaction problem. By a combination of symbolic and numerical computation, the authors have accurately predicted three-dimensional RNA structures.

No brief collection of articles can hope to cover any field, especially one as vigorous as computer science. In this issue, entire branches of research in hardware and software have been unforgivably slighted; this will be rectified in issues to come. Still, these articles show that there is more to computer research than further decreases in the cost per megabyte, important as that may be. It is clear that in the future those megabytes will be put to novel and innovative purposes.—DAVID VOSS