

Invited Address

Living with Failure: Lessons from Nature?

Steve Furber
University of Manchester, UK



Steve Furber is the ICL Professor of Computer Engineering in the School of Computer Science at the University of Manchester. He received his B.A. degree in Mathematics in 1974 and his Ph.D. in Aerodynamics in 1980 from the University of Cambridge, England. From 1980 to 1990 he worked in the hardware development group within the R&D department at Acorn Computers Ltd, and was a principal designer of the BBC Microcomputer and the ARM 32-bit RISC microprocessor, both of which earned Acorn Computers a Queen's Award for Technology. He took up the ICL Chair at Manchester in 1990 where he now leads the EPSRC-funded Advanced Processor Technologies Portfolio Partnership. <http://www.cs.manchester.ac.uk/apt/>

Steve is a Fellow of the Royal Society, the Royal Academy of Engineering, the British Computer Society, the Institution of Electrical Engineers and the IEEE. In 2003 he was awarded a Royal Academy of Engineering Silver Medal, and in 2004 became the holder of a Royal Society-Wolfson Research Merit Award.

Abstract - The resources available on a chip continue to grow, following Moore's Law. However, the major process by which the benefits of Moore's Law accrue, which is the continuing reduction in feature size, is predicted to bring with it disadvantages in terms of device reliability and parameter variability. The problems that this will bring are underlined by the predictions from an Intel commentator: within a decade we will see 100 billion transistor chips. That is the good news. The bad news is that 20 billion of those transistors will fail in manufacture and a further 10 billion will fail in the first year of operation.

What does a 20-30% device failure rate mean for designers and what does it mean for production test?

As a designer, I have some idea how to design for very low device failure rates. Redundancy, fault-tolerance and ECC are all approaches that can cope with very low failure rates. The basic assumption is that faults are infrequent so we only have to cope with one at a time. But a 20-30% failure rate will clearly violate this assumption, and the bottom line is that I have no idea how even to begin to design useful circuits that can cope with this level of failure. For an example of a functional device that can cope with this level of failure, we have to look to nature. Brains can cope with very high levels of neuron failure. But we have no idea how they work, let alone how they keep working

after these failures. What might we be able to learn from biology about building systems that continue to function as components change and fail? Will manufacturing test change from being primarily about checking that every device on the chip works to checking that enough devices are working to ensure that the chip functions correctly and is likely to continue to do so even after many more devices have changed or failed over the early operational life of the chip?

In this talk, I will describe a proposed chip multiprocessor system that is being developed primarily to help understand how the brain works, but which will also present the sorts of challenges that will increasingly dominate the future of production test. The chip does not need to be fully functional to be useful, so how can the production test establish that enough works for the chip to be useful, even after further early-life failure?