Dataflow Supercomputing

Michael J. Flynn

Over the past decades parallel processor speedup has been an elusive quantity for a broad class of applications. Yet with the end of performance scaling for single processors the need has never been greater. The problem is not technology but programming models.

One answer to this speedup problem is to create an idealized data flow machine that exactly correspond to the application and stream data through the resulting machine. This approach can be emulated with FPGAs, providing more than an order of magnitude speedup even as executed as an emulation of the data flow machine.

About Michael Flynn:
Michael Flynn is well known for early work on parallel processing. He directed the Architecture and Arithmetic group at Stanford University for more than 20 years where he is now Emeritus Professor of Electrical Engineering. He is also Chairman of Maxeler Technologies, a UK and US based company dedicated to maximum performance computing. He is a fellow of the IEEE, the ACM and the Inst. of Engineers of Ireland.

Beyond Moore. Beyond Programmable Logic.

Steve Trimberger

Moore's Law continues, but for how long? Many are predicting the end, or at least the slowing, of semiconductor scaling even as FinFETs are being introduced. New technologies, such as 3D integration, offer new opportunities for silicon vendors and customers. These technology trends are converging on Programmable Logic. FPGA vendors are changing the way they build their product and those changes are being reflected in the architecture and tools that surround those products. This talk describes the technological pressures and opportunities for programmable logic vendors and highlights recent product trends and what they indicate for the future of the programmable logic industry.

About Dr. Steve Trimberger:

Dr. Steve Trimberger has been employed at Xilinx since 1988. He was the technical leader for the XC4000 design automation software, led the architecture definition of the Xilinx XC4000X device families and developed a dynamically-reconfigurable multi-context FPGA. He designed the bitstream security functions employed by Xilinx FPGAs and his research led to the development of the Xilinx 2.5D Stacked Silicon Technology. He is currently a Xilinx Fellow heading the Circuits and Architectures Group in Xilinx Research Labs in San Jose, California. He has served as Design Methods Chair for the Design Automation Conference, Program Chair and General Chair for the ACM/SIGDA FPGA Symposium and on the technical programs of numerous Workshops and Symposia. He has published three books and dozens of papers on design automation and FPGA architectures. He has more than 190 patents in IC design, FPGA and ASIC architecture, CAE and cryptography. His innovations appear today in nearly all commercial FPGA devices. He is a four-time winner of the Freeman Award, Xilinx's annual award for technical innovation. He is a Fellow of the ACM and a Fellow of the IEEE.

Compiling OpenCL to FPGAs: A Standard and Portable Software Abstraction for System Design

Deshanand Singh

Today's FPGAs have logic capacities that are steadily increasing. The FPGA is a large array of fine-grained programmable elements that can be configured in such a way to efficiently solve many complex problems. For many applications, FPGAs are a tremendously efficient computational fabric; however, the primary method of design entry for FPGAs is through Hardware Design Languages (HDLs) such as VHDL or Verilog. These languages model the FPGA at an extremely low level where the programmer is expected to understand cycle-accurate details of how data is moved and transformed through the FPGA. While this programming model is required to achieve the highest possible efficiency from FPGAs, it is akin to "assembly language" programming for processors. In this talk, we explore techniques that allow us to program FPGAs at a level of abstraction that is closer to traditional software-centric approaches using an emerging parallel language : OpenCL. The field of high level synthesis has evolved greatly in the last few decades; however, several fundamental parts were missing from the complete software abstraction of the FPGA. These include standard and portable methods of describing HW/SW codesign, memory hierarchy, data movement and control of parallelism. We believe that OpenCL addresses all of these issues and allows for highly efficient description of FPGA designs at a higher level of abstraction.

About Deshanand Singh:

Desh Singh is a Supervising Principal Engineer at Altera's Toronto Technology Center. Desh leads Altera's OpenCL-to-FPGA project and his charter is to develop high level design tools which allow designers to create applications for FPGAs with a higher level of productivity than traditionally possible. Previously, his group was responsible for a number of optimization algorithms in Altera's Quartus II CAD tool. These include Logic Synthesis, Line-level incremental compilation, Physical Synthesis, Metastability Analysis, and IP core optimizations. Desh holds a PhD from the University of Toronto in the area of timing closure techniques for high speed FPGA designs and has authored over 50 patents and publications on FPGA technology.

Going beyond the FPGA with Spacetime

Steve Teig

The idea of dynamically reconfiguring programmable devices fascinated Turing in the 1930's. In the early 90's, DeHon pioneered dynamic reconfiguration within FPGAs, but neither his nor numerous subsequent efforts, both academic and industrial, resulted in a useful and usable product. Over the last several years, we have significantly advanced the hardware, architecture, and software for rapidly reconfiguring, programmable logic: going beyond the FPGA using a body of technology we call Spacetime. Spacetime represents two spatial dimensions and one time dimension as a unified 3D framework: a powerful simplification that has enabled us to deliver in production a new category of programmable devices (3PLDs) that are far denser, faster, and more capable than FPGAs yet still accompanied by software that automatically maps traditional RTL onto these exotic fabrics. In developing Spacetime, we encountered and resolved many complex, technical challenges that any dense and high-performance reconfigurable device must face, many of which seem never even to have been identified, much less addressed, by any prior effort. In this talk, I will identify some key limitations of FPGAs, introduce Spacetime as a means of addressing them, enumerate some of the many challenges we faced, and present solutions to a couple of them.

About Steve Teig:

Steve Teig is the President and CTO of Tabula and the inventor of Tabula's Spacetime 3-Dimensional Programmable Logic Architecture. Prior to founding Tabula, Steve was CTO of Cadence Design Systems, having joined Cadence through its acquisition of Simplex Solutions, where he was also CTO. At Simplex, Steve invented and led the technology development for the X Architecture, which radically improves chip design by pervasively incorporating diagonal wiring. Before joining Simplex, Steve co-founded two successful biotechnology companies: CombiChem, where he was CTO, and BioCAD, where he was CTO and, later, CEO. In the 1980's, he developed key logic simulation and place-and-route technologies that continue to have far-reaching influence. Steve received a B.S.E. degree in Electrical Engineering and Computer Science from Princeton University and holds over 240 patents. In 2011, he was awarded the World Technology Award for IT hardware innovation.

FPGA development in Norwegian Industry

Espen Tallaksen

A number of advanced and interesting products with FPGAs are being developed in Norway. Some of these FPGAs are being pushed to the limit, with complex designs and extreme performance. Hence methodology, experience and knowhow is really important.

Our experience from more than 15 years of design consultancy shows that the main challenges with respect to project efficiency and product quality may be grouped into a few specific problem areas. These seem to apply world-wide, – but unfortunately the FPGA vendors and academia do not really focus on the most obvious improvement potentials for a huge majority of the industry.

We need more focus on practical approaches, – more systemised common sense, more structured methodology, and more awareness on critical aspects of FPGA development. We need to share our experience and knowledge. The annual Norwegian FPGA-forum gathers FPGA designers, project managers, vendors, academia and consultants for this purpose, and has become the most important event for FPGA-technology in Norway.

This talk will give an introduction to FPGA activity in Norway, briefly discuss some important FPGA development issues and suggest some unique approaches to improving the industry.

About Espen Tallaksen:

Espen Tallaksen is the managing director and founder of Bitvis – an independent design centre for embedded software and FPGA. He graduated from the University of Glasgow (Scotland) in 1987 and has 25 years experience with FPGA and ASIC development from Philips Semiconductors in Switzerland and various companies in Norway, including his earlier founded company Digitas. During twenty years Espen has had a special interest for methodology cultivation and pragmatic efficiency and quality improvement. One result of this interest has been a verification system that is currently being used by several Norwegian companies. He has given many presentations in Scandinavia on various technical aspects of FPGA development – including a keynote at FPGAworld and a two-day course on FPGA development Best Practices. He is also the initiator and chair of the Norwegian FPGA-forum.