

# FLEET, INFINITY & MARINA

by

Adam Megacz and Ivan Sutherland

July 2009

## ABSTRACT

This talk describes a radically different architecture for computing called Fleet. Fleet accepts the limitations to computing imposed by physics: moving data around inside a computer costs more energy, more delay, and more chip area than the arithmetic and logical operations ordinarily called "computing." Fleet puts the programmer firmly in charge of the most costly resource, communication, instead of in charge of the arithmetic and logical resources that are now almost free. Fleet treats arithmetic and logical operations as side effects of where the programmer sends data.

Fleet achieves high performance through fine grain concurrency. Everything Fleet does is concurrent at the lowest level; programmers who wish sequentiality must program it explicitly. Fleet presents a stark contrast to today's multi-core machines in which programmers seek concurrency in an inherently sequential environment.

The Fleet architecture uses a uniform switch fabric to simplify chip design. A few thousand identical copies of a programmable interface connect a thousand or so repetitions of basic arithmetic, logical, input-output, and storage units to the switch fabric. The uniform switch fabric and its identical programmable interfaces replace many of the hard parts of designing the computing elements themselves.

Both software and FPGA simulators of a Fleet system are available at UC Berkeley.

Berkeley students have written a variety of Fleet programs; their work helped to define what the programmable interface between computing and communication must do. A simple compiler now produces the programs required at source and destination to provide flow-controlled communication. We expect work on a higher-level language to appear soon as a PhD dissertation.

A recent 90 nanometer TSMC test chip, called Infinity, demonstrated switch fabric performance at about 4 GHz. A new test chip, called Marina, has just gone out for fabrication. Marina will test the programmable interface, and if successful, will give us confidence to build a complete Fleet. We seek participation from sponsors, programmers, and designers of basic computation elements.

## BIO

Ivan Sutherland is a Visiting Scientist at Portland State University where he and Marly Roncken have recently established the "Asynchronous Research Center" (ARC). The ARC occupies both physical and intellectual space half way between the Computer Science (CS) and Electrical and Computer Engineering (ECE) departments at the university. The ARC seeks to free designers from the tyranny of the clock by developing better tools and teaching methods for design of self-timed systems. Prior to moving to Portland, Ivan spent 25 years as a Fellow and Vice President at Sun Microsystems. A graduate of Carnegie Tech, Ivan got his PhD at MIT in 1963 and has taught at Harvard, University of Utah, and Caltech.