

The Density Advantage of Configurable Computing



An examination of processors and FPGAs to characterize and compare their computational capacities reveals how FPGA-based machines achieve greater performance per unit of silicon area. If we can exploit this advantage across applications, configurable architectures can become an important part of general-purpose computer design.

André DeHon
California
Institute of
Technology

A large and growing community of researchers has successfully used field-programmable gate arrays (FPGAs) to accelerate computing applications. The absolute performance achieved by these configurable machines has been impressive—often one to two orders of magnitude greater than processor-based alternatives. Configurable computers have proved themselves the fastest or most economical way to solve problems such as the following:

- *RSA (Rivest-Shamir-Adelman) decryption.* The programmable-active-memory (PAM) machine built at INRIA (Informatics and Automation Research Institute, Paris) and Digital Equipment Corporation's Paris Research Lab achieved the fastest RSA decryption rate of any machine (600 Kbps with 512-bit keys, and 185 Kbps with 970-bit keys).
- *DNA sequence matching.* The Supercomputer Research Center's Splash and Splash-2 configurable accelerators ran DNA-sequence-matching routines more than two orders of magnitude faster than contemporary MPPs (massively parallel processors) and supercomputers (CM-2, Cray-2) and three orders of magnitude faster than the attached workstation (Sparcstation I).
- *Signal processing.* Filters implemented on Xilinx and Altera components outperform digital signal processors (DSPs) and other processors by an order of magnitude.¹

- *Emulation.* Chip designers use FPGA-based emulation systems to simulate modern microprocessors.²
- *Cryptographic attacks.* Collections of FPGAs offer the highest-performance, most cost-effective programmable approach to breaking difficult encryption algorithms. For example, Berkeley students showed that an Altera FPGA can search 800,000 keys per second, whereas a contemporary Pentium searches only 41,000 keys per second.³

From an operational standpoint, what we see in these examples is a reconfigurable device (typically an FPGA) completing, in one cycle, computations that take processors tens to hundreds of cycles. Although these achievements are impressive, by themselves they do not tell us why FPGAs were so much more successful than their microprocessor and DSP counterparts. Do FPGA architectures have inherent advantages? Or are these examples just flukes of technology and market pricing? Can we expect the advantages to increase, decrease, or remain the same as technology advances? Can we generalize the factors that account for the advantages in these cases?

To attack these questions, we must quantify the density advantage of configurable architectures over temporal architectures—both empirically and with a simple area model. We must also understand the tradeoffs that configurable architectures make to achieve this density advantage. Once we understand the tradeoffs involved in using general-purpose computing

Field-Programmable Gate Arrays

An FPGA is an array of bit-processing units whose function and interconnection can be programmed after fabrication. Most traditional FPGAs use small lookup tables to serve as programmable computational elements. The lookup tables are wired together with a programmable interconnect, which accounts for most of the area in each FPGA cell (Figure A). Many commercial devices use four-input lookup tables (4-LUTs) for the programmable processing elements because they are area efficient.¹ As their name implies, FPGAs were originally designed as user-programmable alternatives to mask-configured gate arrays—the bit-processing elements implementing the logic gates, and the programmable interconnect replacing selective gate wiring.² Increasingly, FPGAs have served as spatial computing devices.

Most of the examples mentioned in the introduction of this article use Xilinx XC4000 or Altera A8000 components as their main computational workhorses. These commercial architectures have several special-purpose features beyond the general model—for example, carry-chains for adders, memory modes, shared bus lines—but they are basically 4-LUT devices.

References

1. J. Rose et al., "Architecture of Field-Programmable Gate Arrays: The Effect of Logic Block Functionality on Area Efficiency," *IEEE J. Solid-State Circuits*, Oct. 1990, pp. 1217-1225.
2. S. Trimberger, *Field Programmable Gate Arrays*, Kluwer Academic, Norwell, Mass., 1992.




Figure A. A three-input lookup table (3-LUT) FPGA. A programmable interconnect wires the lookup tables together to serve as programmable computational elements.

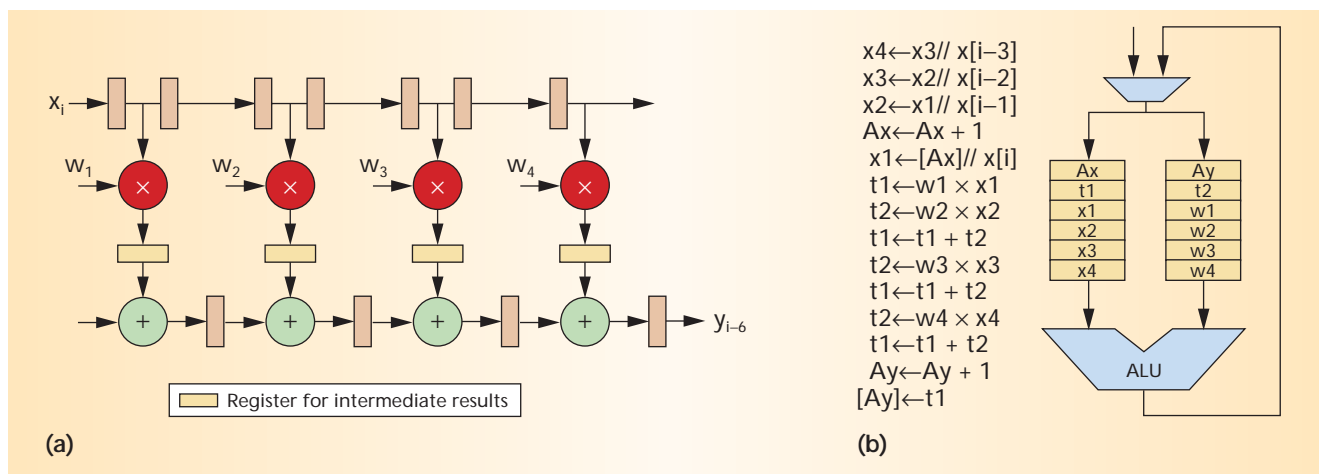
blocks, we can expand the comparison to include custom hardware and functional units. Taking these effects together, we can see how configurable computing fits into the arsenal of structures we use to build general, programmable computing platforms.

CONFIGURABLE COMPUTING

Computing with FPGAs is called *configurable computing* because the computation is defined by configuration bits in the device that tell each gate and interconnect how to behave. Like processors, FPGAs are programmed after fabrication to solve virtually any computational task—that is, any task that fits in the device's finite state and operational resources. This impermanent, postfabrication customizability distin-

guishes processors and FPGAs from custom functional blocks, which are operationally set during fabrication and can implement only one function or a very small range of functions. (See the "Field-Programmable Gate Arrays" sidebar.)

Unlike processors, the primitive computing and interconnect elements in an FPGA hold only a single device-wide instruction. (Here, the term *instruction* broadly refers to the set of bits that control one cycle of operation of the postfabrication programmable device.) Without undergoing a lengthy reconfiguration, FPGA resources can be reused only to perform the same operation from cycle to cycle. In these configurable devices, we implement tasks by spatially composing primitive operators—that is, by linking



them together with wires. In contrast, in traditional processors, we temporally compose operations by sequencing them in time, using registers or memory to store intermediate results (see Figure 1). The single-instruction-per-active-computing-unit limitation in FPGAs provides an area advantage, at the cost of restricting the size of the computation described on the die at any point in time.

EMPIRICAL COMPUTATIONAL DENSITY

As noted earlier, a single reconfigurable device often can compute, in a single cycle, a computation that takes a processor or DSP hundreds of cycles. We can place a simple filter, for example, spatially on a single FPGA, as in Figure 1a, so that it takes in a new sample and computes a new result in a single cycle. In contrast, a processor or DSP takes a few cycles to evaluate even one filter tap, easily running tens of cycles for even the simplest filter structures. The FPGA might require tens of cycles of latency to compute a result, but because it performs the computation as a spatial pipeline composed of many active computing elements, rather than sequentially reusing a small number of computing elements, it achieves higher throughput.

FPGAs can complete more work per unit of time for two key reasons, both enabled by the computation's spatial organization:

- With less instruction overhead, the FPGA packs more active computations onto the same silicon die area as the processor; thus, the FPGA has the opportunity to exploit more parallelism per cycle.
- FPGAs can control operations at the bit level, but processors can control their operators only at the word level. As a result, processors often waste a portion of their computational capacity when operating on narrow-width data.

As examples, consider the Alpha 21164 processor⁴ and the Xilinx XC4085XL-09 FPGA. Both devices were built in 0.35-micron CMOS processes. The 21164 contains two 64-bit ALUs and runs at 433 MHz. As a result, it performs, at most, 2×64 single-bit ALU operations every 2.3 nanoseconds. This gives us a maximum theoretical computational throughput of 128 bit operations

per 2.3 ns, which equals 55.7 bit operations per ns.

In contrast, the XC4085XL-09 consists of 3,136 configurable logic blocks and runs at a peak clock rate of 4.6 ns per cycle. For a rough comparison, we can equate one CLB to one ALU bit operation. (One CLB consists of two 4-input lookup tables. In many cases, we can put more than one ALU bit operation in a CLB, but the conservative estimate suffices for illustration.) The FPGA achieves a computational density of 3,136 bit operations per 4.6 ns, or 682 bit operations per ns, easily an order of magnitude greater than the computational density of the processor in the same process technology.

This crude comparison does not tell the whole story of the useful computation these devices can perform or the factors that prevent them from achieving their maximum theoretical peak performance. Nonetheless, it does illustrate how it is possible for an FPGA to extract more computational capacity from a silicon die than a RISC processor can.

There are challenges to making the FPGA run consistently at its peak rate, just as there are challenges to making the processor issue productive cycles at its peak rate. A big problem with the FPGA is the difficulty of adequately pipelining the design to consistently achieve such a high clock rate. Conventional FPGA architectures and tool methodologies make it difficult to contain interconnect delays and reliably target clock rates near the device's peak capacity. Yet, as a recent reconfigurable design developed at UC Berkeley demonstrates, engineering FPGA designs and spatial computing arrays that reliably achieve high-clock-rate execution is possible.⁵

Figure 2 compares the computational densities of a wide range of processor and FPGA implementations. It shows that the anecdotal density observation just discussed holds broadly across device implementations. That is, FPGAs have an order of magnitude more raw computational power per unit of area than conventional processors. This trend has remained true for many process generations if we consider total device area. As the amount of silicon on the processing die has increased, both FPGAs and processors have turned the larger dies into commensurately greater raw computational power, but the gap between densities has remained.

Figure 1. (a) Spatial and (b) temporal computations for the expression $y[i] = w_1 \cdot x[i] + w_2 \cdot x[i-1] + w_3 \cdot x[i-2] + w_4 \cdot x[i-3]$. These are implementations of a 4-tap FIR filter.




Figure 2. Comparison of processor and FPGA computational densities. These data are based on published clock rates, device organization, and published and measured die sizes.⁶ ALU bit operations/ λ^2 s (bit operations per λ^2 second) is the density of operations per unit of area-time (area \times time). Area is normalized by the technology feature size (λ is half the minimum feature size). Time is given in seconds, an unnormalized unit, since several small feature effects prevent delay scaling from being a simple function of feature size.

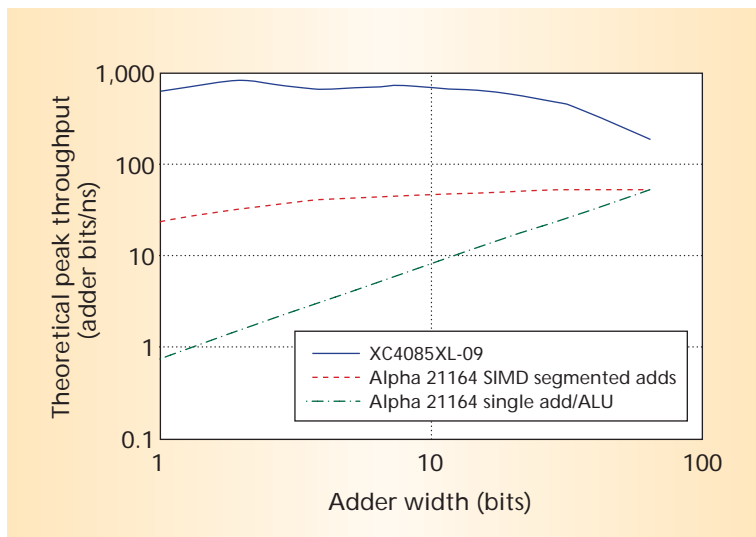


Figure 3. Maximum adder throughput as a function of unpipelined adder word width. Here, we assume the FPGA must complete an entire add of the specified width within a cycle. The FPGA throughput varies because of combinational add latency, granularity issues associated with packing adds into a row, and overhead costs for starting and completing each add. For the processor's single add, we assume only one add of the specified width is performed in the ALU. For the segmented adds, we assume that a single guard bit is left between words, and that data are otherwise perfectly aligned for the operation.

These peak densities only tell us what the architecture can provide when task requirements match architectural assumptions. If the task requires manipulation of small data words, but we are using a large-word CPU, our yield will be only a fraction of the CPU's peak capacity. For example, a 64-bit architecture processing 8-bit data items would realize only an eighth of its peak

processing power. Since FPGAs are controlled at the bit level, they do not suffer this problem. Consequently, when operating on narrow data items, FPGAs have the potential for a second order-of-magnitude advantage in computational density over processors.

The peak densities also only tell us how much throughput these devices can achieve. They do not tell us how much latency a single data set incurs when traversing a complete computational sequence in any of these devices. In most cases, given comparable implementation technologies, hardwired structures in the ALU enable processors to complete a single add operation in much less time than a contemporary FPGA requires for an equally wide add.

For example, if the add itself is not internally pipelined on the aforementioned XC4085XL-09, a single 64-bit add would take a little over 17 ns. Because we can get a maximum of 56 of these adders on the FPGA (using 60 percent of its raw resources), this gives a maximum throughput of 56/18 ns, or 3.1 64-bit adds/ns, compared to the processor's 2/2.3 ns, or 0.9 64-bit adds/ns. To illustrate the combination of these effects, Figure 3 shows the maximum theoretical bit-level adder throughput available on the Alpha and the XC4085 when a single add latency sets the pipeline operating frequency.

SIMPLE MODEL

The preceding section suggested that FPGAs achieve their density advantage and fine-grained controllability by forgoing the deep instruction memories found in processors and DSPs. Simple area accounting is consistent with this view.

Each FPGA bit operator, complete with lookup table, configuration bits, state, and programmable interconnect, requires an area of 500,000 to 1 million λ^2 (see my thesis⁶). A RISC processor instruction is 32 bits long and is usually stored on the processor die in static RAM cells whose bulk area is about 1,200 λ^2 per bit. Thus, one RISC instruction occupies roughly 40,000 λ^2 . Assuming for the moment that the instruction memory is all that takes up space on the processor die, we can put 25 RISC instructions in the space of a single 1-million- λ^2 FPGA bit operator. The RISC instruction typically controls a 32-bit, single-instruction, multiple-data (SIMD) data path, so we can place $32 \times 25 = 800$ RISC instructions in the space of 32 FPGA bit-processing units.

The processor also needs data memory to hold 32-bit intermediate results. Each intermediate result will occupy at least 40,000 λ^2 in SRAM area. Assuming we keep one word of state for each instruction, the area per active computation bit reaches parity when the RISC processor holds instructions and state for 400 operations. So, if we design the RISC processor to support 4,000 instruction words and 4,000 words of

Table 1. Comparison of 16 × 16 multipliers.

Device style	Design	Feature size 2λ (μm)	Area (λ^2)	Time (ns)	Area-time ($\lambda^2\text{s}$)	Ratio
Custom	*Fadavi-Ardekani ⁷	0.63	2.6M	40	0.104	1
FPGA	Isshiki and Dai ⁸ (88 CLBs × 1.25 M λ^2 /CLB, 7.5 ns/cycle × 16 cycles)	0.60	110M	120	13.2	130
DSP	Kaneko et al. ⁹	0.65	350M	50	17.5	170
Processor	Yetter et al. ¹⁰ Magenheimer et al. ¹¹ (66 ns/cycle × 44 cycles)	0.75	125M	2,904	363	3,500

* From a survey of a large number of multiplier implementations,⁶ this example is the densest 16 × 16 multiplier and is implemented in a feature size most comparable to the other devices listed.

state data to keep the 32-bit ALU busy, we require as much area as 320 FPGA bit-processing units. These FPGA processing units, if heavily pipelined, can provide 10 times the per-cycle active computational capacity of the 32-bit RISC data path.

Modern processor designs do allocate space for thousands of instructions and on-chip state data memory per active data path, making the last comparison most relevant. In practice, the RISC processor would require area for its actual data path, its high-speed interconnect paths, and its control. But this makes the FPGA look even better by lowering the actual parity point below 400 operations. This simple accounting clearly demonstrates the trade-off that differentiates processors and FPGAs: *Processor architectures make a large sacrifice in actual computational density to tightly pack the description of a larger computation onto the die.*

The last comparison also underscores the trade-off FPGAs make to achieve their high computational density. By packing a single instruction and state element with each active bit operator, the FPGA stores the state and description of a computation much less densely than a processor. That is, an FPGA bit operator's 1 million λ^2 of area is less dense than a RISC instruction's 40,000 λ^2 by a factor of 25. Consequently, when performance or throughput is not important, the processor often can implement a large computation in less area than an FPGA.

The comparison couples the two main organizational differences between the processor and the FPGA—deep instruction memory and wide, SIMD-controlled data paths. To better understand their contributions, it is worthwhile to separate these factors. Let's look at two intermediate designs: a 1-bit processor data path and a multibit FPGA.

If the RISC instruction controls only a single-bit ALU (and we retain our earlier assumption that instruction and data memory are the only things consuming space in the processor), we see that 25 instructions take the same space as one FPGA cell. Both devices offer one active computational bit operator per cycle in this space. Now, when we have only 250 instructions, the FPGA has more than 10 times the processor's computational density. This example underscores the fact that the processor is using its SIMD control to help mitigate the expense of deep instruction memory.

Commercial FPGAs use approximately 200 bits to specify function, interconnect, and state storage for each 4-input lookup table (4-LUT). In practice, these configuration bits are highly decoded, so their information content is much smaller, perhaps closer to 64 bits,⁶ but we can use the larger number here for illustration. Assuming the same 1,200 λ^2 per SRAM bit as in the earlier example, we could save at most 240,000 λ^2 by sharing instructions in SIMD fashion among FPGA bit operators. This is an upper bound since sharing implies additional wiring between cells, and the bound very generously assumes that interconnect is completely identical between cells. A 32-bit SIMD FPGA data path would occupy $31 \times 760,000 \lambda^2 + 1,000,000 \lambda^2$, which approximately equals 25 million λ^2 , or about the area of 25 bit-controlled FPGAs. Thus, in contrast with the processor, the FPGA, with its shallow instruction memory, does not pay a large density penalty for its bit-level control.

SPECIALIZED FUNCTIONAL UNITS

Previous sections focused on the use of generic processing elements such as ALUs and lookup tables. In practice, modern microprocessors regularly include specialized, hardwired functional units such as multipliers, floating-point units, and graphic coprocessors. These units provide a greater effective computational density when called upon to perform their respective tasks but provide little or no computational density when different operations are needed. The area per bit operation in these specialized units is often 100 times smaller than the amortized area of a bit operation in a generic data path. Therefore, including such functions is worthwhile if they will be used often enough.

Example: hardware multiplier

A hardwired multiplier is often one of the first specialized units added to a processor architecture and is a primary architectural feature of a DSP. Given their regularity and importance, multipliers are among the most heavily optimized computational building blocks. Therefore, they serve as an extreme example of how a hardwired unit's computational density compares with its configurable and programmable counterparts.

Table 1 compares several 16-bit × 16-bit multiplier

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.