It's FPL, Jim - but not as we know it! Opportunities for the new Commercial Architectures

Tom Kean

Algotronix Consulting, PO Box 23116, Edinburgh EH8 8YB, United Kingdom
tom@algotronix.com

Abstract. Following the simple Programmable Logic Device (SPLD) and Field Programmable Gate Array (FPGA) generations a third generation of programmable logic technologies is now reaching the marketplace. These new architectures are driven by the move to system level integration and fast expanding markets such as networking and wireless communications which are not addressed adequately by mainstream FPGA's. This paper considers the technologies, business models and chances of success of the third generation companies using the Triscend CSoC and Systolix Pulse DSP architectures as examples.

1 Introduction

The 25 year history of commercial programmable logic shows two important technical generations: simple PLD (SPLD) sum-of-products (PAL) technology and Field Programmable Gate Array (FPGA) technology. Today, we are seeing the beginning of a third-wave of programmable architectures with distinctive features focussed on the challenges of high volume embedded applications (figure 1).

Although PAL technology was introduced more than 20 years ago and there have been few significant architectural improvements in the last ten years it is still an important niche component in the marketplace. No single player managed to dominate the simple PLD (SPLD) market and SPLD's quickly became a commodity component until improvements in process technology made them irrelevant.

Today, FPGA technology is, after 16 years of development, also reaching maturity. Product generations now serve to tune the architecture for improvements in process technology rather than to introduce fundamentally new structures. The last significant improvement in mainstream FPGA architecture was, arguably, the introduction of block RAM by Altera in 1995 which allowed FPGA's to address applications which required medium sized memories. In contrast to SPLD's, the FPGA business has come to be dominated by two strong players Xilinx and Altera and has remained highly profitable.

Given this background of incremental development of a highly profitable but apparently mature technology it is, perhaps, not surprising that a wave of radical new

architectures have emerged from start-up companies. These companies are vying to define a third-generation of programmable logic. This surge of architectural development is pulled by demand from emerging markets which the established second generation companies have chosen not to address and fuelled by the recent flood of venture capital into technology companies.

An important technical theme underlying the new FPL architectures is the merging of FPGA technology with ideas from computer architecture - including DSP processors, parallel computers and VLIW machines. Since many of the new architectures are directed at computational applications it is not surprising to see direct support in the silicon for arithmetic operations and multi-bit words.

Influences from the academic FPL community and research funded by DARPA are clear (notably the context switching proposals of Andre de Hon [1], the processor like operators of the Kress Array [2] and the combination of FPL with microprocessor in OneChip [3] and Napa [4]) as are influences from research on non-mainstream FPGA architectures by the second generation companies [5][6].

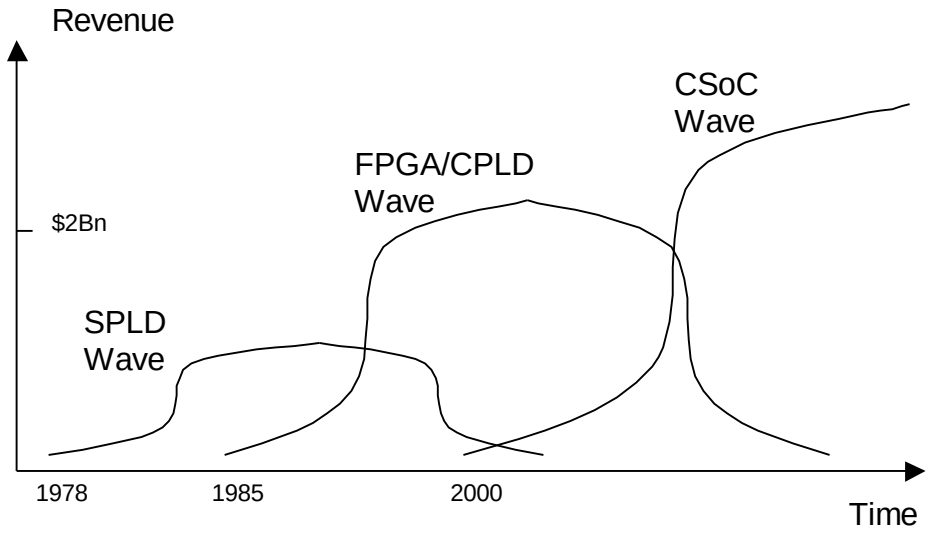

Fig. 1. The history of programmable logic can be viewed as a sequence of waves of development driven by major changes in architectural style

2 Market Drivers

While academic research has largely concerned itself with applying FPL to general purpose computing the explosive growth in the networking, mobile communications and multimedia markets is driving the commercial architectures. These embedded applications demand raw computational performance without sacrificing flexibility, power consumption or time to market. Dataquest estimated the market for function specific signal processing as $6.3 billion in 1999. Table 1 shows some of the third generation FPL companies and their target markets.

Company	Architecture	Business Model	Markets
Adaptive Silicon	Not disclosed	Sell Cores	Embedded DSP
Chameleon Systems	Array of 32 bit processors	Sell Chips	Voice over IP, Software Radio, networking
Malleable	Not disclosed	Sell Chips	Voice over IP
Morphics	Not disclosed	Sell Solutions (chips and software)	Cellular Communications / Software Radio
Systolix	Systolic Array	Sell Cores	Signal Conditioning, Embedded DSP
Triscend	System on Chip	Sell Chips	Communications, Embedded Systems

Table 1. The new FPL companies

2.1 Software Download

A key driver for third generation architectures is the move away from programming the PLD once, during manufacture to downloading improved designs in the field. This ability to upgrade product features after the initial sale is a fundamental benefit of reconfigurable technologies over ASIC's and changes the economics of embedded equipment - moving it to a model more like the computer industry. The majority of software currently running on personal computers was not available on the day they were purchased.

Smith [7] illustrates the economic impact of time to market in ASIC design by showing how a 1 month lead in time to market can result in a 70% difference in product lifetime revenue. Figure 2 updates this analysis to show the effect of software download. Firstly, if a manufacturer has an option to download software in the field he can ship a product at an earlier stage of development. Thus, the product with software download reaches market first (illustrated by the upper curve). Some time later, the competing product reaches market (the lower curve). By this time the first product has established itself and the competing product never gains the same market acceptance. So far this mirrors Smith's analysis of time to market benefits: however, the important difference is that using software download the manufacturer of the first product can extend the functionality of the product and thus its life span in the marketplace. These updates are a further source of revenue from existing customers. Thus, the revenue curve in the case of software download can have multiple peaks.

Today, software download is a compelling technology for embedded products with a communications capability. The largest scale deployment of software download to date was in the 56K bit modem market where two groups of competing companies shipped large numbers of modems before the V90 international standard was approved. Both groups knew that their modem's would eventually have to be standard compliant but neither could afford the time-to-market disadvantage of waiting until the standard was finalized. When the standard appeared both groups supplied downloadable patches via their web sites that upgraded their modems to be standard compliant.

Although time to market and upgrade revenues are more than sufficient reason in themselves to use software download there are three other important applications:

1. Bug fixes. Software download allows serious problems to be fixed in the field without an expensive product recall. This substantially reduces the risk of market introduction and may allow a less expensive and time consuming beta-test program.
2. Customization. Download of small amounts of data can be used to enable product features or securely install customer specific information. When the equipment has a limited user interface it may be convenient for customers to deal with a call center or website which then downloads personal or configuration information rather than input it directly.
3. Remote diagnostics and monitoring. A special version of the FPL configuration which stresses the system and checks for common symptoms can be downloaded allowing more effective product support.

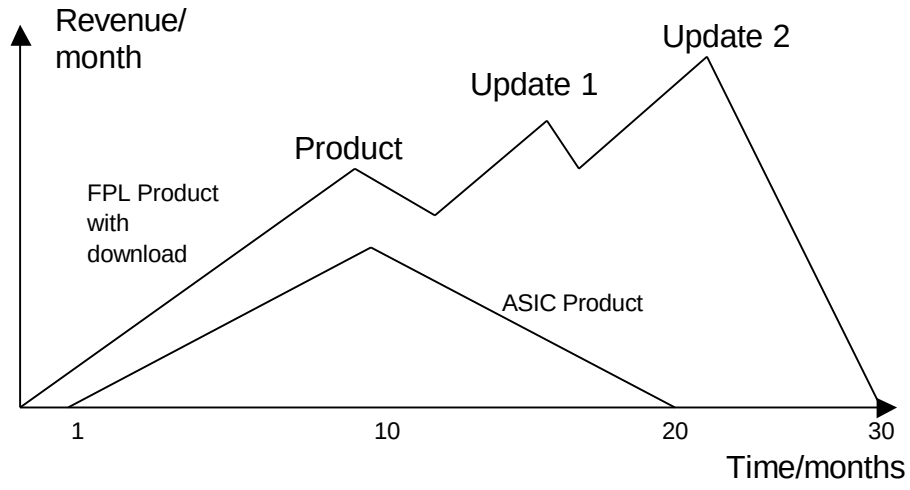

Fig. 2. Economic benefit of software download.

3 Business Models

The second generation FPL companies distinguished themselves over first generation companies not only through new architectures but also by a new business model: the fabless semiconductor company. It is interesting to speculate whether these FPGA companies could have built sustainable businesses based only on the architectural breakthrough or whether their commercial success was determined as much by the new business model which allowed easier access to the marketplace.

Third generation FPL companies are entering the market with innovative architectures which combine aspects of DSP and VLIW processors with FPL technology. These architectures seem 'prima-facie' better suited to the DSP challenges of the target markets than 'conventional' programmable logic. However, there is also a step change in business model as we enter the era of system-on-chip. System on chip represents an opportunity for the new companies and a challenge to the business practices of their competitors.

Some of the new companies are reacting to System on Chip by becoming chip-less as well as fabless: that is they offer Intellectual Property (IP) cores to be integrated onto their customers System on Chip designs rather than physical components. The move to IP cores dramatically lowers the barriers to entry into the marketplace: a typical IP vendor may require only a few tens of employees, where a fabless vendor

may require a few hundred and a vendor with its own manufacturing facility more than a thousand. Thus, using this model a venture funded startup company can compete and win sales from a billion dollar a year established company. What is less clear is whether this model will eventually result in a small number of high profit companies as did the fabless-model or whether the low barriers to entry will create a cluttered playing field of small companies in cut-throat competition with each other. The second outcome might benefit the large systems companies who will create the System on Chip IC's at the expense of the FPL companies.

A second group of companies are keeping the model of selling chips rather than IP but are producing Configurable System on Chip (CSoC) IC's as standard products. These devices contain micro-controllers, memory and configurable logic so they can efficiently implement all aspects of a customer design. Thus, these companies are bringing the benefits of deep sub-micron SoC technology to companies and projects which could otherwise not justify it financially. This is closely analogous to the way the FPGA companies brought the benefits of ASIC gate array technology to companies which could not justify the associated tooling charges and design risk. Historically, this has been a very successful business model. Clearly, this model requires more investment than the IP cores model and it is possibly more vulnerable to competition from established second generation companies.

Of course, there is no reason why a company should not adopt both business models: selling CSoC chips and IP cores. In fact, the technology developed for the CSoC marketplace is directly applicable to the IP cores business. Selling IP cores can be seen as a way of leveraging extra profit from technology whose costs are already covered. From a customer point of view cores provided by a company which also supplies CSoC chips are likely to be better supported and have more developed CAD and library support.

3.1 Horizontal and Vertical Business Models

The second generation FPL companies operated as 'horizontal' companies, supplying the same component to a variety of different industry segments. Recently, a cluster of small IP companies have formed around the FPL vendors to provide 'cores' targeted at important segments (figure 3 (a)). Vendors have also begun some in-house development of segment focussed IP cores and small concessions to the requirements of important segments have been made on the base silicon (such as the introduction of carry-chains and shift register mode for LUT RAM to help DSP applications). However, the second generation companies are not producing distinct product families for industry segments. Instead there are generic 'high performance' and 'low cost' product ranges which are roughly synonymous with 'this years' and 'last years' architectures.

New entrants to the market must show significantly improved price/performance over the mainstream architectures and may have difficulty persuading customers to invest engineering time in creating complex designs for untried architectures. They do not yet have the sales to foster an infrastructure of third party IP companies in the same way as the second generation vendors. One approach to address these competitive difficulties is to focus the architecture on a particular industry segment (figure 3(b)). This allows tailoring of the architecture more closely to that segments requirements - possibly improving price/performance. More importantly, segment focus allows the product vendor to develop a comprehensive supporting range of IP for the programmable architecture. Thus customers can be presented with packaged solutions for their applications which can be compared directly on price and performance with existing solutions: customers are not forced to 'buy-in' to the underlying architecture or to commit engineering resource to designing for it.

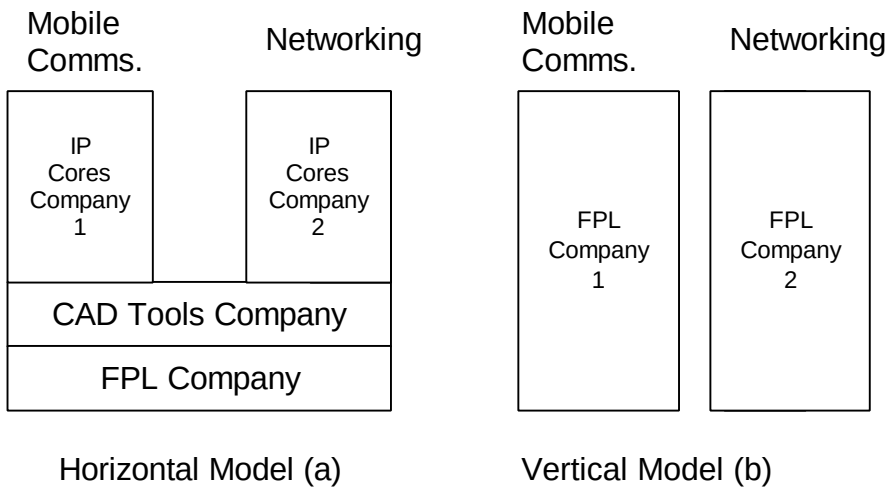

Fig. 3. Horizontal (a) and Vertical (b) Business Models

4 Systolix Pulse-DSP

The Systolix Pulse DSP is a third generation FPL architecture aimed at digital filtering and signal conditioning. The architecture is described in [8] and consists of an array of bit-serial arithmetic processors coupled by a special interconnect framework which transfers multi-bit wide words and error correction information (figure 4). Targeting the architecture at a narrow - but very important - class of signal processing algorithms has allowed the high level textual Pulse Programming Language and the graphical Filter Express design tools to be developed. These allow the customer to capture their design in terms of DSP algorithm parameters without compromising implementation efficiency.

The first product to use the Systolix Pulse-DSP technology is the AD7725 sigma-delta A/D converter from Analog Devices [9]. The Systolix array provides flexible user specified signal conditioning of the converted data. Systolix offers cores of various sizes and performance ratings for integration into third party System on Chip ASIC's. Thus in our categorization Systolix is an example of a fabless, chipless and vertically targeted FPL vendor.

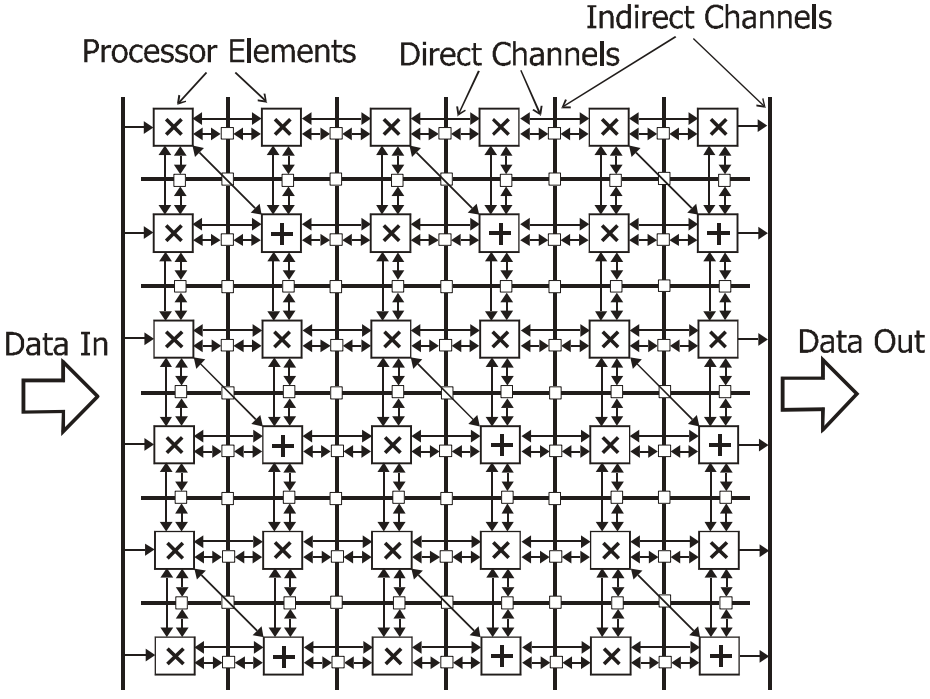

Fig. 4. Systolix PULSE-DSP Architecture (figure courtesy of Systolix Ltd.)

5 Triscend E5-Configurable System on Chip

Unlike Systolix, Triscend sells chips directly to end-users. Triscend's focus is to bring the benefits of system on chip technology to companies or projects for which an ASIC based SoC is not viable - for example, because of low volume requirements. The FPL component of Triscend's System on Chip technology - termed Configurable System Logic (CSL) is itself a parameterisable FPL-core.

Triscend's CSL technology [10] is more general purpose than many third generation FPL architectures. For example, in figure 4, the Triscend chip is shown in a wireless communications application [11] implementing 'glue logic' functions such as display interfaces as well as computational functions like channel filtering and Viterbi.

Triscend has targeted its first family of devices (the 8032 based E5) at microcontroller users. The Triscend FastChip software provides a drag-and-drop user interface through which customers can add peripherals to the central core. Thus the E5 could be viewed as providing the world's largest catalogue of 8032 variants! Naturally, customers can also make use of standard FPGA tools such as Synopsys and OrCad to create their own custom peripherals as required. Triscend's second generation technology will feature the ARM7 TDMI processor and is directed at networking and communications applications [12].

A particularly important feature of the E5 given its target market is the close coupling between the microcontroller core and the programmable logic: the microcontroller can manage the reconfiguration of the programmable logic and directly address registers within the user logic. A header file for C or assembly language programmers containing symbolic references to the resources provided by the design mapped onto the CSL array is generated automatically by the CAD tools.

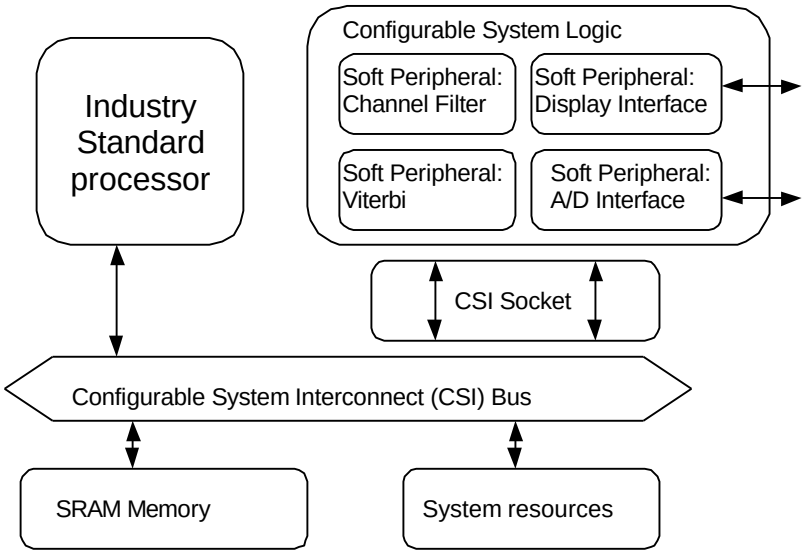

Fig. 5. Triscend CSoC in Wireless Communications (figure courtesy of Triscend Inc.)

6 Summary

The central driving force in programmable logic architecture is, as it has always been, the ability of silicon to implement more and more functionality each year. As this happens a programmable chip can map larger and larger sections of the system and it must address different challenges (table 2).

First wave, sum of products architectures were efficient for small amounts of glue logic but did not scale as the process technology improved and more complex sections of the system had to be mapped to the programmable logic device. Complex PLD (CPLD) architectures borrowed some of the features of FPGA's and managed to delay the inevitable for some time but as process technology improved the underlying deficiencies became more and more apparent.

Similarly, second generation FPGA devices can efficiently map large blocks of logic within a system but do not scale to map entire systems as demanded in the era of-system on chip. Processors, their associated program memories and busses and purely computational functions (such as filter acceleration) do not map efficiently to general purpose FPGA architectures.

Third generation FPL devices, now coming on to the market, address the requirements of system on chip integration. As with earlier generations of programmable logic there is a step change in business model from chip vendor to IP core

vendor. This change addresses the market requirement for lower cost of programmable components and allows easier market entry for small companies. The lower revenues achievable with this model also make competition from established second generation companies less likely.

It will be interesting to see whether the combination of the new architectures, the chipless business model and the rapidly growing application areas constitute a strategic inflection point which reshapes the FPL industry to the same extent as the FPGA architecture and fabless model did 16 years ago.

	First Wave	Second Wave	Third Wave
Defining Company	Monolithic Memories	Xilinx	TBD
Defining Architecture	22V10 PAL	XC4000 FPGA	TBD
Programming Model	Write Once	Write Many	Download in Field
Functionality provided	Glue Logic	Subsystems	Systems
Business Structure	Own Fab. Sell chips	Fabless, Sell chips	Fabless, Sell IP cores and chips
Business Model	Horizontal	Horizontal	TBD

Table 2. Characteristics of FPL generations

References

1. de Hon, Andre : Reconfigurable Architectures for General Purpose Computing, PhD Thesis, MIT Artificial Intelligence Laboratory 1996, AITR1586
2. Kress, R. : A fast Reconfigurable ALU for Xputers, PhD Thesis, Universitaet Kaiserslautern, 1996, D-386
3. Wittig R. et al : OneChip: An FPGA Processor with Reconfigurable Logic, , Proc. IEEE Symposium on Custom Computing Machines, FCCM 97, Napa CA 1996
4. Rupp C.R. et al : The NAPA Adaptive Processing Architecture, Proc. IEEE Symposium on Custom Computing Machines, FCCM 98, Napa CA 1998
5. Churcher S. et al : The XC6200 FastMap™ Processor Interface. Proceedings FPL95, Springer LNCS 975
6. Trimberger, S. et al : A Time-Multiplexed FPGA, Proc. IEEE Symposium on Custom Computing Machines, FCCM 97, Napa CA 1997
7. Smith M. J. : Application-specific Integrated Circuits, Addison Wesley Longman 1995
8. Jones, G. : PulseDSP - A Signal Processing Oriented Programmable Architecture, Field Programmable Logic and Application, Proceedings FPL99, Glasgow UK, Springer LNCS 1673
9. Analog Devices Inc. : AD7725 16-bit Sigma-Delta ADC with Programmable Post Processor, Preliminary Technical Data, Norwood MA, Feb 2000
10. Triscend Inc. : Triscend E5 Configurable System on Chip Family, Mountain View CA, January 2000
11. Kean, T. : Soft RF: new frontier for hackers, Focus on Software Radio, EE Times, August 16 1999
12. Triscend Inc. : Triscend Announces Industry's First 32-bit Configurable Processor will be ARM based, Press Release, November 1998