Multi-core DSP for Base stations: Large and Small

Doug Pulley

picoChip
Bath, UK. BA1 5JD
e-mail : doug@picochip.com

Abstract - The baseband implementation in cellular base stations for the latest standards specification always presents one of the most exacting challenges in real-time digital signal processing. This paper examines a practical approach to multi-core DSP which addresses the limitations of other technologies that are used to solve the problem and how it not only readily scales up to large high-complexity basestations using the latest techniques in multiple access and multiple antenna systems but down to the lowest cost, high volume systems such as "Femtocell".

I Introduction

Over the last two decades international standards bodies have created a succession of ever more capable but complex air interface specifications. Their implementation, particularly at the base station, has consequently put increasing demands on the baseband digital signal processing. Implementation technologies available to base station architects, Application Specific Integrated Circuits (ASICs), Field Programmable Gate Arrays (FPGAs) and Digital Signal Processors (DSPs), each have their strengths and weaknesses which will be reviewed shortly.

The sheer quantity of processing specified in each generation of standard has tended to be set consciously at the cutting edge of technology available at the time – or what is anticipated within the next year or so as the assumption is always made that continued technological innovation (particularly Moore's Law) will ensure the demands are first met and then reduced in cost over time.

II Conventional implementation technologies

A. ASIC

Of the all the possibilities, ASIC has the longest design cycle, the highest development costs, but the lowest cost per unit shipped. The high development costs may be acceptable if sufficient volumes of units are shipped to amortize these over. In most cellular base station markets the volume is orders of magnitude less than typical consumer electronics volumes and over the past few years development costs have risen such that the *per unit* development cost is very significant.

ASICs do yield the lowest power consumption and offer the ability to take on the most challenging computation taks with best chance of success. However the consequence is a lack of flexibility in device and the time between the release of new features in specifications is reducing (e.g. Baseline UMTS to UMTS HSDPA, UMTS HSDPA to UMTS HSPA) to the point where that time is of the order of, or shorter than, the ASIC design time. This can render the finished article obsolete by the time it is generally available. Add to this the risks of implementing any new standard correctly first time and the need to incorporate learning from the field from early deployments, it is no surprise that the base station ASIC for conventional cellular base stations is almost extinct.

B. FPGA

FPGAs have improved significantly in price/performance terms in recent years and, with the demise of practical ASIC approaches, represents the only "backup plan" for the implementation of functions beyond the capability of conventional DSP devices. They have a particular signal processing niche in performing functions of high computational density but are much less suitable for performing control and management tasks which are far better described in software. The design flow has a lot in common with ASIC and the development effort and time is very significant. Whilst they offer very good design flexibility the debug cycle time can be dominated by place and routing of the design in large or full devices. Power consumption can be very significant in the large devices and costs may still dominate the total in a mixed technology solution.

C. DSP

Conventional DSPs compliment FPGAs very well since they offer a home for the high maintenance, high complexity functions that are unsuitable to FPGA implementation. Performance has continued to improve thanks to Moore's Law, specialist instruction sets and market-specific architectural improvements but these devices alone fall short of being sufficient to implement the full baseband of the latest standards. Completely software defined solutions to wireless baseband have been confined to legacy standards and techniques which are more readily within the capability of DSPs and general-purpose processors.

D. Mixed technology solutions

In the field under examination here – "the latest cellular standards release" - complete solutions for baseband signal processing have required a combination of technologies: FPGA or ASIC hardware assistance for high throughput simpler functions and DSP software implementations of

978-1-4244-1922-7/08/$25.00 ©2008 IEEE

complex control function

The use of mixed technology platforms presents a challenge in architectural design and implementation execution. Each has its strengths and weaknesses that have to be accommodated in functional partitioning of the overall system, each requires a design team with skills in that particular technology and each requires its own particular design tools and development flow. The successful co-ordination of these activities and integration to deliver the final system is a lengthy, risky process.

II. picoArray approach to multi-core processing

A. *picoArray architecture*

Parallel processing (or "multi-core" as it is now referred to) is not a new concept and has been extensively studied for many years. It has always offered the potential for scalable high performance computing but not achieved sufficient momentum to get significant deployment in a mainstream application. Massive parallel systems have typically struggled in general purpose computing applications from the point of development flow and system debug.

Fig 1 shows a parallel architecture intended for a specific class of applications real-time embedded digital signal processing. It shows an array of processing elements *P* connected together by a high speed time-division multiplexed bus. For now it is not important what *P* actually is just what *P* does as observed from the outside. The elements *P* need not be the same and could be complete processors (with internal code and data storage) or pieces signal processing hardware provided they observe this paradigm. The connections between processors are made on a cycle by cycle basis and are statically allocated at design time automatic by the development tools. Since the allocation to of processing to processors (*P*) is also made statically at the same time, all

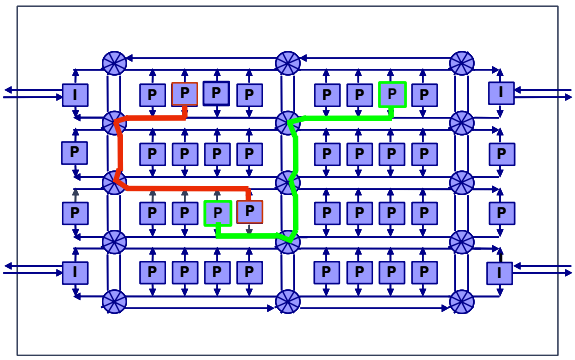

Fig. 1. picoArray approach

resources in the device (processing, storage and communications) are allocated prior to runtime so at runtime there is no arbitration or re-allocation as the result of the application resulting in entirely deterministic behaviour of the system.

The number of elements *P* is limited only the practicalities of device manaufacture as the bus behaviour is highly scalable. Further more, the bus behaviour can be transported from device to device so the cascade of *N* devices each with 300 processors yields what the designer can treat as a single platform of 300*N* processors. In this way an almost arbitrarily large signal processing platform can be build to deal with the most exacting of requirements of the biggest base stations.

B. *Development flow*

Much effort has been expended over the years in development of "parallelizing compilers" – compilers that can take large amounts of sequential code (typically written in C language), partition into parallel functions that can be simultaneously executed and target and array of processing devices with the resulting code. These efforts have not yielded significant results and are not required for to use a picoArray-based device.

Debugging is achieved using conventional single stepping, breakpoints and memory/register content viewing on a processor by processor basis where all processors or single step or stopped on breakpoint together in synchronism. However, this is code level debug, much of the effort in integration of complex base station systems is spent on algorithm system level debug. A consequence of the provision of very high bandwidth interconnect is that even in a complete system or full device, spare bandwidth exists to transport significant quantities of algorithm or system behaviour information out of the device for run-time observation of key signals such as channel estimation values, closed loop feedback system values etc. The flexibility and ease of adding this non-intrusive diagnostics is a significant advantage over achieving similar in FPGA/DSP hybrid systems.

IV. Scaling down

There are emerging applications such as "Femtocell" or "Home Base station" which call for high volumes of small cheap base stations. In this environment cost is of paramount importance and every effort needs to be made to minimize the cost of the equipment. It has already been discussed how ASIC yields the cheapest per unit cost but is only practical in high volume. The baseband processing of 3GPP UMTS Node Bs is sufficiently complex that even with stable specifications and equipment definitions that a significant proportion of the design can only be described in software but the rest to less or greater degree could be considered for ASIC implementation (but as part of the same device). If those items are chosen to be replaced with ASIC equivalent which communicate on the bus in the same manner as their processor counterparts then they become "drop-in replacements" for the groups of processors and software whose function they are replacing. These components can occupy orders of magnitude less silicon area depending on the function they perform considerably reducing the cost of the device though reducing the flexibility by making it more application specific. Given the granularity of the ability to pick and choose functions to "harden", the device architect can pick with fine control the trade-off of flexibility and cost.

V. Summary and Conclusions

An architecture for multi-core digital signal processing has been presented and the benefits of its development flow described in the context of the strengths and weaknesses of conventional approaches for cellular base station baseband implementation. In particular the scalability of the approach is shown to satisfy not only the largest and most complex base station systems but the smallest and most cost sensitive by the architecture's natural ability to allow arbitary mix of software and hardware to satisfy the right cost verses flexibility trade-off.