1.GENERAL BACKGROUND

Advances in integrated circuit technology continue to open up exciting opportunities for the creation on new electronics, communications and computing systems. For example, the new era of System-on-Chip (SoC) technology will in future allow single chip systems to be created that contain billions of digital logic gates. This is prompting innovative developments for example, in low cost-embedded computing systems and leading to the creation of new products that were perhaps unforeseen a decade ago. Parallel research is also increasingly allowing analog circuitry to be integrated on the same silicon substrate - mixed signal devices - with important advances (including research in this university) suggesting that this will increasingly extend to RF front-end systems. This potentially enables the creation new generations of mobile and wireless connected systems with a wide range of multi-media capability e.g. transmission and reception of video (including digital TV), audio, speech and images, as well as other types of broadband information. Critical issues also include network access, security and for fixed networks the ability to process, in real-time, packets of information in the terabits range (e.g. the transmission and reception of multiple HDTV channels, Voice over IP, video-on-demand etc.).The successful design of these and other future systems presents important research challenges, it also emphasises the need to recruit and train a new generation of electronics system and chip designers whose capabilities cover a spectrum of disciplines such as embedded systems design, multi-media digital SoC architectures, mixed signal design and RF design. In addition, expertise is also needed at system level in a variety of areas including real-time network packet processing architectures and network security.

2. TOWARDS A UBIQUITOUS PROCESSOR

The types of potentially mass market electronic devices and systems described typically today require the creation of heterogeneous SoC's that today combine general purpose embedded processors with dedicated datapaths. The former are used to implement programmable functionality, the latter to achieve the levels of real-time performance and/or power dissipation required to create these cost effectively. Examples include SoCs for digital video encoding/ decoding, for multi-media applications, for communications signal processing, for encryption/decryption etc. Indeed it is the case that, even with this or similar approaches, many of these types of applications are pushing the limits of current technology. An example, might be a portable multi-media system with HDTV/ broadband and Voice-over-IP capability connected securely to a wireless networks.

The limitation here that with NRE costs increasing exponentially it increasingly becomes difficult to justify the investment needed to develop these types of customised SoCs in all but applications where the numbers required are very significant. Even then time to market issues can be a major constraint. In an ideal world it would be much more desirable to be able to implement such systems on processors with a high level of programmability and reconfigurability, not least to accommodate the increasing rate of introduction of new communications or video compression standards, for example. The challenge, however, is to achieve the computational performance requirements demanded within stringent power dissipation budgets, acknowledging that various studies have shown that the intrinsic computational efficiency (for example as measured by the ATE product) of special purpose computational architectures can be in the region of 3- 7 orders of magnitude greater than programmable e.g. RISC or DSP processors.

We believe that this is achievable through the creation of systems, which replace the dedicated hardware with a reconfigurable fabric in the first instance with this fabric tailored to a particular class of application (e.g. video compression, graphics processing, voice processing/recognition, encryption etc.) Whilst existing FPGA technologies indicate the potential for doing this, the hardware overhead involved is still much too expensive and power hungry, particularly for high volume and/or portable applications. It is proposer's view that the achievement of such goals is probably best undertaken using a two-step approach. The first is to focus on a range of domain specific applications, such as those described and undertake research within each of the types of domain described to derive suitable electrically configurable (or mask programmable) architectural fabrics within each of these. The assumption here is that what might be strongly suited to one domain of applications (e.g. for multiple standard video compression such as H.264, MPEG4, WM9, MPEG2 etc) is probably not the same solution that might be derived for graphics processing, or cover a wide range of cryptographic standards and protocols. In the case of the latter the underlying mathematics tends to be dominated by non standard arithmetic (e.g. Finite Field arithmetic) and thus electrically reprogrammable architectures that are suited to this are perhaps unlikely to be suited to video compression applications where motion estimation and transform coding (based on sum of products computations) tend to be dominant. Similar arguments, we believe hold in other domains. In the case of network processing, the emphasis is on packing switching and the use of different queuing algorithms, for example. Memory requirements also vary considerably within application type and thus in the first instance the likely solution is the creation of families of scalable platform architectures tailored to various domains. The figure below schematically gives some initial thoughts on the type of architecture that could be used in image and video processing applications.

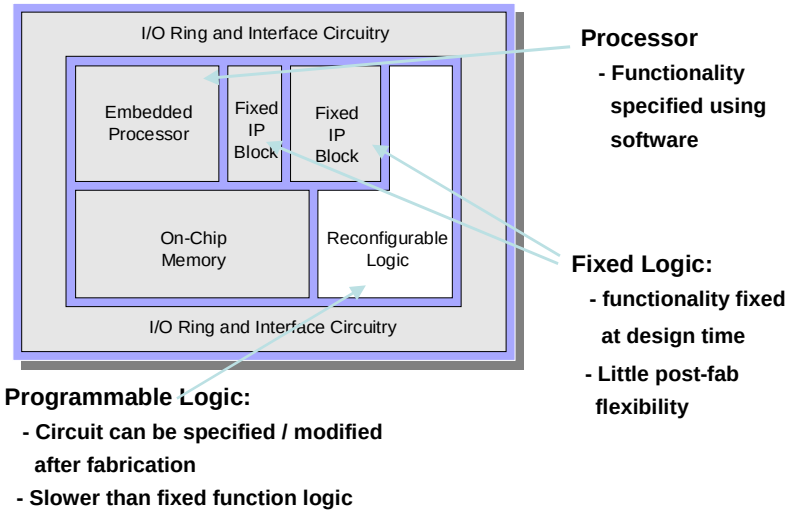

The ultimate objective and thus the second phase of such a Grand Challenge should we believe to attempt to define create and define the ubiquitous General Purpose SoC processor of the future i.e. the "Pentium Processor of the next decade" which could be produced in cost effectively in high volumes and meet the challenging performance and power constraints demanded as more and more real-time processing becomes and internal requirement within future generations of general purpose SoCs. Indeed and interesting up-front exercise might be to try and define what functionality a General Purpose SoC created in 2015 might be expected to contain and in doing so work back to where we are today to ascertain is achievable and what research barriers must be crossed to achieve this.