Development of On-Chip Signal Processing for Audio/Video/User interfaces

Edited by
Karlheinz Brandenburg, FhG IIS
Patrick Dewilde, DIMES, TU Delft

Contributors

Core group

Karlheinz Brandenburg, FhG IIS

Patrick Dewilde, DIMES, TU DELFT

Ahmed Jerraya, TIMA/INPG

Rudi Lauwereins , IMEC

Fernando Pereira, IST

Additional experts

Andrea Cuomo, STM, who presented the enclosed figures at the MEDEA+ Worksop, Leuven, 4 June, 2003

Hugo De Man, IMEC

Table of Contents

Executive Summary

This report identifies several major future trends in the applications of microelectronics for digital signal processing. The task given to the group was to report on the "Development of On-Chip Signal Processing for Audio/Video/User Interfaces". In the process of writing this report, the scope was somewhat broadened to include other major developments, in particular application scenarios. The task was carried out not as a complete overview of our vision of future technology, but to provide examples in particular in areas where European research, both basic research at university level and applied research, is already strong and where further strengthening would help to maintain or gain a world leading position in these application areas.

We advice MEDEA+ to take the following recommendations into account:

1. Make aggressive moves towards the new G4 paradigm. Take the lead by putting strong efforts into the open standardisation of its components (including operating systems).

2. Support the implementation of new standards. A significant amount of good basic work has been carried out and standardised with the involvement of European companies, institutes and universities who now have to compete with proprietary solutions. In a lot of cases reference designs are needed for complex new standards (e.g. MPEG-4, MPEG-7, MPEG-21).

3. Develop an energy-efficient platform for 'intelligent system design', consisting of embedded processors, modules for decision making, modules for pattern recognition, local networking devices and a variety of networking protocols that can be logically integrated. Move aggressively in extending the capabilities of System-C in this direction. Ensure the completeness of the design trajectory from system design, verification all the way down to mapping on hardware (chip, package) and creation of low voltage, low leakage IP blocks with the emphasis on low leakage embedded memory. This requires efforts in the front-end (architecture design, mapping) but also at circuit level (e.g. control of leakage current and its impact on the front-end). This can best occur in a network of industry, universities and research centres (see item 8).

4. Start research efforts at higher system levels. The emerging 'user's control' paradigm requires new ways of handling interactions with a community of user's that get more and more demanding, desire more services, wish to exercise more control and have expanded levels of computing and memory facilities as well as requiring increased security. These scalability, architecture and security problems are solvable but not without a major effort. Europe should play a strong role in this domain.

5. Explore aggressively the new possibilities in the multimedia era: new types of transducers, interaction with the Internet, user's control, new services. A lot of opportunities for new products exist here. 'Ambient intelligence' is an important step, but more is needed in terms of applications, services, Quality of Service and roaming user support.

6. Commence new programs to bridge the large gap between basic research and readily available reference implementations to jump start the development of new products, since higher semantic levels of meta-data and user interface are one big opportunity for future systems-on-a-chip.

7. Start programs in 'body area networking' and place extra emphasis on the development and system integration of devices surrounding the individual.

8. Develop stronger co-operation with team-driven network of universities (and research institutes). The strength of the US microelectronic environment has been greatly enhanced by the co-operative attitude of American companies w.r.to a large collection of universities based all over the country. Many ideas and many experiments are required, as success is based on small probabilities. Develop broad regional involvement, try to get many universities and research groups interested and contributing new ideas.

1. Introduction

This report identifies several major future trends in the applications of microelectronics for digital signal processing. The task given to the group was to report on the "Development of On-Chip Signal Processing for Audio/Video/User Interfaces". In the process of writing this report, the scope was somewhat broadened to include other major developments, in particular application scenarios. In order to identify major trends and to establish where future progress in IC technology for on-chip signal processing is required, a certain procedure was followed leading us to:

- Identify the main trends in applications of very complex signal processing. These trends are on purpose "apples and oranges" in the sense that they include both important requirements for future systems and trends in applications.

- Identify a few "killer applications". We defined killer applications as applications which will not just make use of new enabling technologies (including silicon technologies and new algorithms), but will turn up as appliances which will then be sold in the order of many million highly complex devices.

- Derive research issues from the main trends and killer applications. Again this list of research issues is not meant to be complete, but includes very important issues for future systems-on-a-chip.

- Derive consequences for future systems on a chip. This chapter is the closest to microelectronics technology issues.

- Derive recommendations for MEDEA+

The task was carried out not as a complete overview of our vision of future technology, but to provide examples especially in areas where European research, both basic research at university level and applied research, is already strong and further strengthening would help to maintain or gain a world leading position in these application areas.

2. Overview

The following chapter briefly outlines a listing of principal key words for the main paper.

1. A short overview of main trends

- paradigm change towards peripheral control and ambient intelligence,
- large collection of 'intelligent devices' surrounding the user and interacting with him,
- exponential increase of demand for signal processing on the chip both between user and background network (video on demand) and between users (roaming),
- emphasis on a scaled 'Quality of Service',
- new importance of security and privacy protection while allowing easy access,
- upcoming of 'personal networks',
- need for large wireless bandwidth,
- enormous increase of computational complexity of signal processing algorithms,
- upcoming 'immersive experiences',
- seamless roaming access,
- need for support of 'context awareness',
- renewed importance of energy considerations,
- augmented reality.

2. Overview of selected potential killer applications

We have identified the following three:

1. The Personal Information Centre or Wearable Digital Assistant,
2. Surrounding multimedia and video-phony,
3. Ubiquitous communications.

3. Overview of important research Issues

Here are the most important ones that are further developed later in the note:

1. Techniques for (broadband) wireless communications,
2. User's access and interfacing,
3. Next generation multimedia as truly immersive,
4. Networking issues,
5. Increase of 'true' intelligence,
6. New types of distributed architectures,
7. New design technology needs.

4. Consequences for future systems on a chip

- Networks on a chip techniques will be needed,
- Real time operating systems to be intensively developed,
- Combination of heterogeneous signal processing resources,
- New dynamic and adaptive techniques for task management,
- Power-speed trade-offs.

5. Recommendations for MEDEA

- See the Executive Summary

3. Main trends

The massive increase in chip density and functionality will have a tremendous impact on how electronic systems will be organised in the future, ranging from multimedia systems to systems for telecommunications. This may be considered a true **'paradigm change'**: due to the massive availability of light weight (and cheap) large scale memory and processing power the 'centre of control' can now move from servers at fixed locations to the 'customer's pocket'. As new measures for the effectiveness of 'wearable control' become available, we can start using terms as 'Gigabytes/gram' or 'Gigamips per Watt' to indicate space and power performance. The consequences of the paradigm shift are that fixed resources that are now dedicated to a specific user (PC's, displays, web access) become 'general purpose' while the control of them is transferred to the roaming user. A large collection of service providers and provisions will be offering their goods mediated via this mobile user, and the web will be a public place where all sorts of activities take place, very much like an open city but with a much greater range of accessibility. The phenomenon is already visible in the fast emerging WiFi (Wireless Fidelity) environment.

Although this main trend has been going on for a while, it has not yet been realised to its most pervasive consequences, nor have the necessary steps been taken to facilitate it in a way that will satisfy both producers and consumers. In our view, however, it is inevitable much as the web has been unpredicted and inevitable and will have major consequences for the design and realisation of equipment The effects will be felt across a wide range of categories, which are detailed to a certain extent below. To start with signal processing and multimedia:

- Between the roaming (wireless) user a few orders of magnitude of multimedia type information will be transferred, ranging from 'video on demand' and 'net-meeting', to information on any items the user demands or his system needs, in particular information concerning the properties of the wireless environment the user is roaming in;
- QoS (Quality of Service) becomes an overriding issue; the user wants quality in correspondence with the availability of resources, the properties of his environment and his system. The QoS must be adaptive so as to provide a performance that is as optimal as possible given the concrete situation;
- Accessibility will be the first requirement, but also security and the correct handling of paying customers, hence an added importance of watermarking, encryption, shielding from eavesdroppers, authentication etc. And all this to be handled in close interaction with the roaming user already mentioned;
- A large increase of signal processing demands is to be expected both for the user's 'personal area network' (PAN) and for the user's 'personal network' (PN) which extends throughout the web or the Internet.
- Games are evolving to 'immersive experiences' - but these could be useful in other circumstances as well. There is no reason to restrict the user's direct environment to just the classical communication experiences we have now: TV type broadcasting;
- The grey area between mass media (information on demand) and personal communications gets blurred, the user is also gradually becoming a supplier of multimedia information, e.g. generating video while on the run - a fundamental

change for service provision as well as for signal processing in the user's environment.

The main 'ubiquitous wireless communications trend' paradigm has more consequences for present day system design and the conception of new types of products. We mention:

- Seamless roaming access across a variety of networks, WAN, LAN, Fibre, Cellular requires redesigns of wearable devices, base stations, routers and networks in general;
- The electronic environment has to change. Not only will there be an increased amount of sensors and actuators, transducers etc., leading to sensitive 'context awareness', but the electronics will also become more invisible, active in the background, self-explanatory and incorporated in the immersion.
- Energy supply becomes an extra point of attention, ranging from the source of energy to making consumption as minimal as possible;
- The natural environment gets supplemented with 'artificial' information, added virtual reality, possibly customised. It may also be attached to the environment of the user. This may happen in many circumstances, e.g. in a shopping mall, but also in a car that is driving along a 'smart road' - that is a road that gives him information on conditions further ahead.

Taken together, all these elements boil down to a fundamental reconsideration of system design technology in general and signal processing in particular, leading to new types of chips that better exploit the emerging requirements. Before diving into a more detailed account of the consequences, the new paradigm has for technological research, we consider three future potential so called 'killer applications'; defined as new products that in our view have the potential to attract a mass market.

4. Potential killer applications

As outlined in Andrea Cuomo's presentation at the MEDEA+ Workshop in Leuven on 4 June 2003, a variety of applications are expected in the future:

From this wide range of applications we have chosen the following killer applications:

4.1. The 'Personal Information Center' / 'Wearable Digital Assistant'

The foremost application is the 'Personal Information Center' (PIC) - the successor of the PDA/Mobile phone which everybody wears in his or her pocket or handbag. Its importance is driven by the enormous increase of memory and computing power available in a very compact space. We will witness a true paradigm shift with the centre of control moving from static devices (PC's and servers) to mobile. You will carry an information centre around with you from which you will be able to direct all your personal needs and wishes, as far as information and communication is concerned, when you wish to do so. When you enter a room, the static devices available there (computers, displays, television set, telephone), and any other 'intelligent' surrounding devices will be at your disposal, controlled wirelessly from your PIC. Everything will be concentrated into a single device. It is clear that we are already witnessing a strong evolution towards this state - the present PDAs have an increasing amount of functions, but the paradigm shift, which in our view is inevitable, must still occur.

Compared to currently available devices in the field, future PICs will have many Gbytes of local storage and processing power several orders of magnitude beyond current capabilities of low power systems.

4.2. Surrounding multimedia and videophony (home server / home gateway)

In future homes there will be many electronic devices most of them with large amounts of memory and with universal standardised communication interfaces thus enabling them to communicate with each other. For instance it will be possible to listen to a song on the digital radio, whose data stream contains metadata about the song. In this way it will be easy to find background information from other sources, such as which group produced the song, where they will have their next concert, or to purchase the song over the Internet.

The communication will be wired and wireless, and will include devices like the home entertainment centre, speakers, displays, the PICs, and PC's.

Displays and speakers will be separate intelligent devices. Displays will be flat and on the wall, with very high resolution, requiring a high communication data rate controlled by users. Speakers will also be flat, and their number will be much higher than today. The number of speakers will be tens or hundreds to create a three-dimensional immersive sound impression, much like a holographic image. They will be integrated into the wall as a "sounding wall paper" and controlled wirelessly by the home entertainment centre. Access to content will be over digital audio and video broadcast, and through broadband Internet connections. Accessible content will include large virtual libraries, which can be employed to search for a wide variety of properties. For example, it will be possible to hum a tune and to have a search engine suggesting songs with that exact tune. The personal entertainment centre will learn the taste of its users and suggest songs they may like. The entertainment centre will also be used to make telephone calls, but more in the sense of teleconferencing. To talk to somebody remotely will require no handset; it will be a hands free operation with a very high audio quality and with a virtual position in the room for each remote participant. This improves the intelligibility if a conversation involves more than 2 people.

The centre piece of this scenario is a "home server" or "home gateway" connecting different applications and serving as the gateway to remote services. These home servers will not only require vast amounts of storage for holding audio and video data, but also processing capabilities (e.g. for deriving high level metadata from new incoming multimedia essence) well beyond current capabilities of even high end computing servers. The wave field synthesis subsystem (see below) will need to run acoustic room simulations and render the signal separately (including high complexity digital filters) for each loudspeaker system connected. Each loudspeaker (there might be hundreds of them) has its separate rendering, D/A conversion and digital amplifier circuit. Only progress in systems-on-a-chip can make this technology affordable.

Counterpart of the wearable: "immersive home"

The fast evolution in plasma and LC displays will gradually make them ubiquitous. There will be one in every room - aside from the fact that you have already one on your PIC to fall back on if nothing else is available. A variety of other devices must complement the display, sound reproduction and capture in particular, but also text capture. All these devices are to react wirelessly to directions received from your PIC. This is the essence of the 'ad hoc' networking environment. The local environment will support all the global communication needs of the person, in particular videophone connections (net-meeting etc.) but also the need to have direct access to an assortment of (chosen or surreptitiously imposed) centres of information, much like we now access the internet to order tickets, find information on topics of interest to us, carry out banking operations etc. Surround media also implies the use of multimedia for a wide variety of public interactions: e-ticket, access control, interaction with traffic management etc. All these devices are part of the 'killer application'.

Andrea Cuomo demonstrated Multimedia over Wireless, as an example:

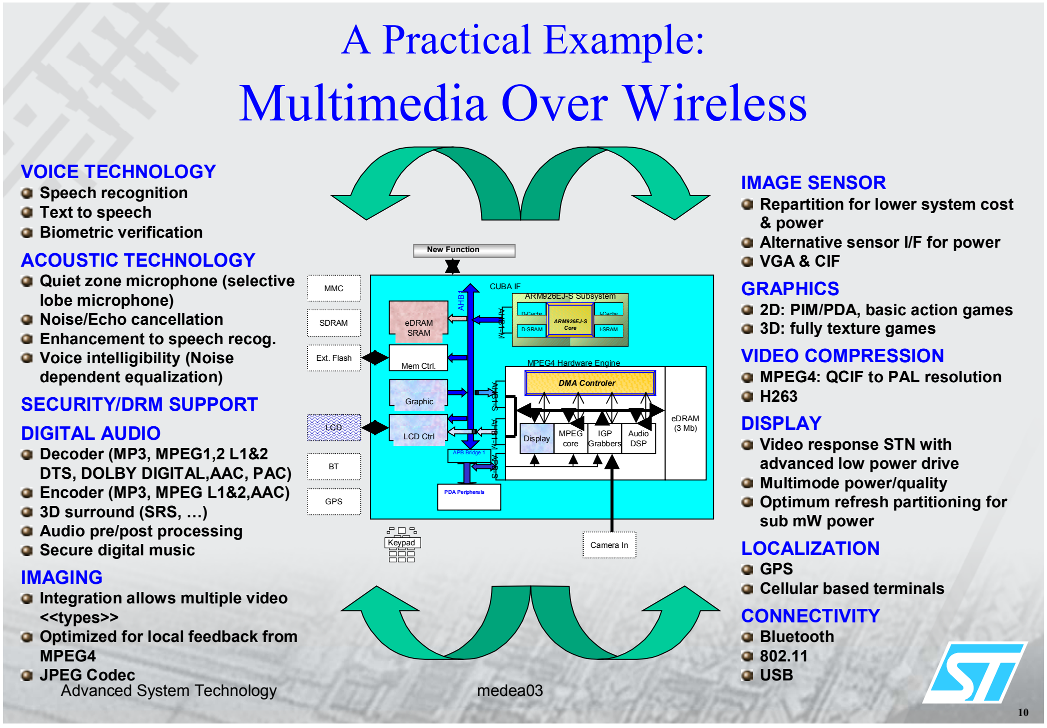

© Andrea Cuomo, STMicroelectronics

4.3. Ubiquitous communications

Each piece of equipment in the user's environment needs to communicate with any other piece and the user's controller. This is the 'bluetooth ideal' put to its logical consequences. The user is surrounded by a 'personal area network' (PAN) that collects and distributes all the information pertinent to his local environment. The PAN controllers, both at the location of the user and locations of sensors and actuators will be a primary piece of equipment to be reproduced in large quantities. The precise architecture of such a PAN is not known at present, present day systems such as WiFi or Bluetooth are either too rigid, offer too little control and hierarchy or are not powerful enough. There are a number of options, all of which show promises and shortcomings. The ideal would be a system that is digital almost all the way to the transmission/receiving circuit, that does not depend on over precise electronics, that can handle a wireless broadband environment, use space division multiplexing among other schemes, is scalable, self initialising and has a hierarchical definition (i.e. can handle simple devices all the way up to the PAN server). Such a device would have the potential to be found in almost any future piece of equipment. There is clearly a need for a new standard here governing the future generation of ad-hoc networks and PAN's.

5. Research Issues

5.1 Techniques for wireless communications

Some of the major trends in digital telecommunication systems are growing data rates, the evolution towards optimal receivers and increased mobility. The high data rates are needed for applications such as fast Internet access, video-on-demand, distance learning, scalability to large collections of users accidentally present etc. and require advanced modulation schemes that take care of the broadband and diverse nature of the communication scene. Recent evolution in VLSI brings the integration of advanced error coders, multi-user receivers and smart antennas on silicon within reach. These techniques will allow for building receivers whose performance is very close to the theoretical (Shannon) limits. Finally, the increasing mobility places some difficult requirements on the synchronisation of the broadband receivers.

The above evolutions are found in many application fields, ranging from high-speed digital subscriber lines and cable modems over the universal mobile telecommunication system to satellite communication systems. One of the most challenging is the high-speed (above 100 Mbps) wireless LAN or access that will be the basis for future mobile broadband systems. The goal of the wireless systems research should be the definition of innovative future wireless systems and the demonstration of their key functionalities. The focus will be on the physical layer of broadband multimedia communication systems. The main research challenges that can be identified in this area are: fast synchronisation for burst mode communication, space-time processing with multi-antenna front-ends, multi-user receiver architectures, multimedia communication over a wireless link with QoS provisions, high-performance error-coding based on turbo-codes and low density parity check codes, integrated broadband front-end architectures, modulation techniques for broadband communication over different channels, and software defined radio.

Although a lot of novel algorithmic techniques in all these domains are emerging from the telecommunication community and although everyone speaks about the “software defined radio”, very little attention has been paid by the silicon community to the consequences of all of this for energy-efficient implementation, usually very compute intensive algorithms. One thing is for sure, technology and voltage scaling alone will not be sufficient to reach a reasonable power autonomy for the future wearable assistant.

On the other hand it becomes clear that, over the coming decade, the battle for new standards will keep raging because the “ultimate” form of multisphere communication is not yet in sight. Hence we will be forced to create wireless terminals capable of adapting to novel standards by software downloads and to provide flexibility for the terminals and base-stations to provide end-to-end optimisation of the energy budget depending on de QoS requested by the end user.

Hence the challenge will be to come up with a global hardware-software architecture with the necessary and sufficient set of parameters that are automatically and dynamically adjusted to operate at the highest level of energy efficiency possible. This must be done in addition to the creation of VDSM novel programmable, configurable and parameterised communication IP blocks operating at 0.5 Volt level with new circuit techniques to reduce leakage power in standby and active mode. This is only possible if a multidisciplinary, concerted research is organised with clear and concise goals driven by an appropriate roadmap for the future market.

This requires a quick and decisive European approach as the competition from the military driven US or the game and multimedia driven Japanese industry is fierce. E.g. Intel's WiFi drive is again likely to catch us by surprise. Europe already lost the Hyperlan-802.11 battle in spite of better standards…

5.2 User's access and interfacing

5.2.1 Introduction

The current standardisation landscape shows that many visual coding problems (as we see them today) are solved or are being solved, e.g. by MPEG (at least until the next technological jump arises). While minor refinements and improvements based on well-known technology are still possible, they are not expected to bring significant changes. It is however clear that standard audiovisual representation technology – coding and description - is not enough to deploy powerful multimedia applications. The technology in the standards is just the enabling factor to which non-normative technology has to be added to deploy advanced applications. Good MPEG-4 content and applications will require powerful tools for object-based rate control, error concealment, natural and synthetic object composition, object-based authoring, audiovisual segmentation, 3D modelling, etc. In the same way, powerful MPEG-7 applications will require adequate technology for feature extraction, action classification, low-level to high-level information mapping, advanced querying, matching criteria, etc. All these technologies are non-normative: this means that they do not need to be mature at the time of the development of the standard but their performance will largely determine the success of the standards since they will determine the quality and functionality of the applications.

5.2.2 Powerful analysis, understanding and abstraction tools

Much of the non-normative technology referred to above is related to the analysis and understanding of the visual data to better code and describe it as well as the analysis and understanding of the compressed data and of the descriptions, in order to better and more easily use the content available. In other words, it is only with powerful analysis, understanding and abstraction tools that it will be possible to make good use of the tools made available by the standards. The better the visual content can be analysed and understood, the better it can be coded and described but also manipulated and used. Moreover analysis and understanding will be more and more multi-modal and cross-modal, exploiting multi-sensory cues to improve performance and allow additional functionality. For example, it is possible to improve video story unit detection by using both video and audio information and it should be possible to look for a Pavarotti's picture using a Pavarotti's song.

While there are a lot of unsolved problems finding the best algorithms for feature extraction and image/sound understanding, even for algorithms which are known today, the computational complexity is a limiting factor for large scale introduction of such systems. In other words, progress in microelectronics will be needed in order to make these technologies economically feasible.

5.2.3 Intelligent sensors, interfaces and displays

In addition to advanced processing and representation technologies, the deployment of advanced multimedia applications will be largely impacted by the developments in sensors, interfaces and displays. Smart sensors will ease the analysis and understanding tasks, collecting more data from the real world, e.g. depth information for a visual scene (see the report of the scientific committee of MEDEA+ on "Heterogeneity on Si or in a package for future system innovation". Pushing the limits, it is possible to imagine a future where analysis and understanding, e.g. for segmentation or scene description, will deal with environments where at least some of the objects will themselves spontaneously or on request report some of their characteristics: smart sensors will be complemented with smart objects. Also interfaces shall allow a more natural and powerful relation of the user with the content, e.g. navigation and interaction within virtual multimedia worlds. Finally, advanced multi-sensory high resolution 'displays' will allow creating fully immersive environments in which users will enjoy rich interactive experiences.

5.2.4 Hybrid synthetic / natural worlds

Although big steps towards the smooth integration of real and synthetic audiovisual data have been made, e.g. with MPEG-4, there is still a long way to go before hybrid real-virtual worlds behave as a continuum. The trend to extend, complete and substitute the reality by the virtuality has strong motivations, still growing in the future, such as the shortage of natural resources and ecological awareness. The real-virtual continuum will be present in augmented worlds, fully virtual or fusing real and virtual, with high degrees of realism and interaction providing a powerful sensation of tele-presence. For example, a big neutral room may become in the same day a sports pavilion, a concert hall and a meeting place with all the necessary equipment, just by means of virtual decoration, saving huge amounts of materials and human resources. The achievement of the real-virtual continuum will rely on many interdisciplinary areas such as audiovisual processing and representation, smart sensors, advanced (very likely wearable) displays and powerful interfaces.

5.2.5 Digital item adaptation

In order for the user to obtain the best possible multimedia experience, it is essential that the content is adapted to take into account as many relevant usage environment parameters as possible. Typical relevant usage environment dimensions are the network, the terminal, the natural environment and the user; each of these dimensions may be characterised by multiple parameters. Standard usage environment description tools, just like content description tools, are needed to ensure a high degree of interoperability between terminals and servers.

MPEG-21 Digital Item Adaptation (DIA) is a relevant development in this area; in part, it targets the specification of standard description tools for the usage environment. Since multimedia experiences are centered on the user, the incorporation of the user themselves into the adaptation and delivery processes needs to be at a deeper level than today's typical approach. As an example, the MPEG-21 DIA specification only considers a narrow set of user preferences for multimedia selection and adaptation. With the future emergence of continuously wearable devices, e.g., with the capability of measuring several human parameters such as blood pressure, cardiac rhythm, temperature, etc., it is possible to take a much wider variety of data into account during the multimedia adaptation process. The result will be the provision of a multimedia experience adapted precisely to the user's current characteristics, e.g., in terms of mental and physical mood or circumstance. For example, the Internet-radio's wake-up music could be selected based on these parameters; this would allow an improved multimedia experience since it would not only be based on static (or slowly varying) user preferences but also on instantaneous measures of the user's condition. This type of capability would require sensors to not only perform the physiological measurements (wireless devices with this type of capability are already available) but also a standard way to adequately describe the features. This standard solution could simply be an extension of the MPEG-21 DIA framework (currently under development). In addition to the selection and standard description of the relevant features, the provision of the capabilities in question would require (non-standard) solutions to map the physiological sets into psychological dispositions or moods with which certain types of multimedia content are associated. This would likely involve complex psycho-physiological studies so as to enrich the user/human dimension that multimedia experiences require.

5.2.6 Transcoding / cross modal adaptation

A central role in adapting multimedia content to different usage environments is played by content adaptation. This encompasses a wide range of processing mechanisms, namely single object/data type transcoding of video or audio, structured content adaptation, e.g., filtering of MPEG-21 Digital Items, and cross-modal adaptation, e.g., conversion of video to images or text to speech. Content adaptation may be performed at various locations in the delivery chain or even distributed across several nodes. Moreover, it may be performed in real-time or off-line providing a set of variations from which to select at adaptation time. While adaptation through trans-coding is well known, cross-modal adaptation is a more challenging process which may range from a simple key-frame extraction process (to transform video into images) to more complex conversions such as text to speech, speech to text or even video to text or speech. A video to speech cross-modal conversion may be as simple as using the information in the textual annotation field of the associated metadata stream, e.g., MPEG-7 stream describing MPEG-2 or MPEG-4 content, or as complex as analysing the content, recognising some objects, e.g., for which specific models are available, and synthesising some speech based on the recognised objects. The adaptation of structured content such as portals or MPEG-21 Digital Items may be divided into two stages. Firstly, filtering the content components to give a set suitable for the consumption environment, and, secondly, after a certain content component is selected, adapting that component in terms of transcoding or cross-modal conversion.

5.2.7 IP management and protection

A key factor in the delivery of multimedia content today is an increasing desire by intellectual property owners to establish, control and protect their rights. A number of schemes for protecting music and video content have been used; most notably the protection of DVD content and more recently attempts by the music industry to protect CDs from copying. However, such techniques have generally been proven to be limited in their protection. Thus, new technologies and mechanisms for content protection are continually being introduced. In fact it may be that it is the platforms that must be altered to ensure protection of content from copying. Currently, various bodies are creating standards (an example is the work in MPEG on intellectual property management and protection, a rights expression language and a data dictionary) which will provide a framework for digital rights management. However, the enforcement technologies are likely to be a significant area of research and development for the foreseeable future. For the best multimedia experiences, the expression of rights to access and alter content metadata, perform (or prevent) certain adaptations and the enforcement of those rights will be critical if content providers are to become willing to distribute material into new user environments. Further, intellectual property management and protection will be an issue not only for the content but also for usage environment information. Usage environment information could reveal personal information such as location and a user's state of health; it is unlikely that users would be happy to have such information freely available on the network.

5.2.8 Peer-to-peer networking architectures

The traditional content delivery model is one where users access content from a set of servers. The 'master' of the content is thus the administrator of the server and, in general, they control who has access to the content. An alternative vision of content delivery is one where every user is interconnected and any user can act as a server of content for any other user. Such peer-2-peer networks are currently being heavily used for music and file-sharing on the Internet. While such use is currently shrouded in controversy due to legal issues, there is no doubt regarding the power of the model. Peer-2-peer networking has been a long-term form of communication; face-to-face conversation, telegraphy, telephony and the postal system are all examples. It has thus been natural for peer-2-peer to quickly grow in electronic forms on the Internet and in cellular systems. At first email was the prime example but now we see Instant Messaging, the file-sharing networks and, among cellular users, the rapid growth of Short Message Services (SMS). The latter is already being enhanced by Multimedia Messaging Services as users demand an improved peer-2-peer multimedia experience. It thus seems likely that legitimate content delivery (and our current infrastructures) will evolve to one where peer-2-peer transfer is commonplace. This will require changes in the way we consider rights and security of access to content for users. However, while peer-2-peer will clearly be very important, and perhaps dominant, server delivered content is likely to remain significant. One reason for this is the reliability of information and experience gained from legitimate and commercial content providers. Thus, we will see a mixture of the two architectures mirroring today's society, with the original news content (rapidly exchanged via email and SMS) likely coming from a broadcaster. The impact of a mixed peer-2-peer and server based architecture is that adaptation will be required both in 'professional' delivery from servers as well as between individual users. Since the latter may not have the processing capability or software to deliver complex adaptations, a new breed of network services may emerge to provide content adaptation to peer-2-peer users. This further emphasises the necessity of transparency – users wish to exchange content as easily as they can converse in the street (already this is a reality in the text world of SMS), but most will not have or even wish to acquire technical skills.

5.3 Next generation multimedia: true immersive environments

Next generation multimedia services will even more than is the case today blur the distinction between personal communications and mass media or between synthetic worlds and recordings of the real world.

The reproduction of real world scenes will allow true immersive environments. Led by the world of computer games, more realistic scenes can be rendered both in audio and video.

Three-dimensional sound rendering systems can be obtained with the so-called Wave-Field Synthesis. Here, individual channels are not reproduced (as in a stereo system), but the original wave filed is reproduced, much like in optical holography. This requires many speakers to be built into the room. Hence it will be necessary for the entertainment centre to find and measure the position of the speakers, which can be obtained by sensors (acoustic and other).

In order for a wide variety of devices to communicate with each other, it is necessary to have a standardised interface. Hence it is necessary to develop a suitable standard. Many devices will be very complex, and a network of these devices will be a complex entity in itself. This requires suitable system- and software-design technologies, and again standards to give each component or device a well-defined common interface.

To purchase content online, both a working and accepted rights management system and security technologies have to be available. To identify songs, for example on the radio, or to find songs similar to a hummed tune, suitable recognition technologies are needed. The basis for the teleconferencing scenario are audio coding schemes with high quality and very low end-to-end delay, and effective echo-cancellation technologies which can handle three-dimensional sound.

Content can be downloaded to a wide variety of devices, ranging from the home entertainment centre to small portable devices, like players, and small computers or the PIC. Since they have different connection speeds and computing power, it is necessary to offer content scalable, suitable to the range of devices. This needs coding which enables the scalability, and also a production or postproduction to support the scalability. This can be, for instance, the inclusion of several different coding "objects" in the coded streams.

In conclusion, standards for systems and protocols and signal processing algorithms are required.

We must also act quickly in the domain of wearable multimedia. Future terminals will have to cope with much higher quality pictures especially when the user also becomes a content provider of real live video pictures that, at the receiver side may be received on high quality displays. Future MPEG-4 implementations scaled linearly from today's best (Japanese) implementations lead to power budgets of several Watt's even after applying known transformations to the software. Hence new architectural research is needed to come up with appropriate architectures that take the dynamic nature of the algorithms into account such that energy consumption is linked to average behaviour instead of worst case. All of this must be done together with a projection towards future low voltage. Low leakage VDSM libraries where low leakage embedded memory plays a central role. Attention must be paid to the co-design of multiprocessor networks on a chip and the compilation and OS environment that goes with it. If we combine that with the wearable communication system, we must also provide a global PHY/MAC/QoS architecture which must be extremely energy efficient to reduce the overhead it otherwise represents to the system of which it tries to optimise the overall energy consumption.

Quick action here is also required in view of the enormous push of the Japanese multimedia and game hardware and software implementations (see e.g. ISSCC 2003 and http://zdnet.com.com/2100-1103-948493.html).

5.4 Networking issues: how will multimedia reach customers

Networking issues take place at several levels of the networking hierarchy: at the level of the user's PAN (Personal Area Network), the wireless and/or wired and/or cellular environment that connects him with the outside network, the home network, the 'first mile' and the background network. Each merit further consideration:

The PAN
The structure of the personal area network capable of handling all the desires a user may have in his direct vicinity, what connects his WDA or PIC to its immediate environment is at present not at all settled. We need the successors of Bluetooth and WiFi offering improved performance, greater flexibility, better bandwidth coverage. Research in this area should also lead to a new standard that should become ubiquitous for the user's environment. Specific research problems are how to handle the division between broadcasting and streaming, multimedia and multicasting in the ad-hoc environment, taking care of scalability (catering to a substantial variety of local users in the same environment) and of security issues pertinent to the wireless environment. The new user-centred paradigm is also taking place at this level (see also the section on 'intelligence' and on 'parallel architectures'). We envision a large and flexible number of 'intelligent agents' interacting with each other and embodying both the functionality of the environment, their ability to communicate and the interaction with the controlling user.

The connection of the user with the Internet
A number of issues arise where the user's environment connects to the Internet (assuming that the user will be surrounded by a personal environment). Authentication and security is an important one. The interaction strategy between personal environments and the chaining of information is another, as well as the efficient usage of overlapping wireless space. The adaptive catering of Quality of Service. The integration of ad hoc protocols with the Internet protocols.

The Internet itself
New generations of IP and their standardisation will remain an active research area if only because of increasing demands on the protocols. Lower and higher level protocols have also to be revised. Voice and video over IP will stay prominent in the interest of system designers. It is not hard to demonstrate that it can work, but it is also not hard to predict that an extremely large amount of traffic will have to be handled putting high pressure on every link in the chain, including signal processing aspects. Hierarchies of networks are becoming inevitable which will also require new standards.

All of this requires a concerted European action. We are confronted with the huge power of the cash rich companies such as Intel, IBM and Microsoft that are de facto determining in a very pragmatic way what the future standards will be. How will European system and semiconductor companies act in a concerted way to respond with the same vigour?

5.5 Gradual increase of 'true' intelligence

We may talk of an 'intelligent system' as soon as there are several independent but interacting operating systems active in one overall system environment. The fast increasing chip complexity is already bringing in the concept of 'network on a chip'. However, besides the network on the chip, we saw above that we also have the personal area network that in itself brings into play a potential multitude of autonomous OS's or at least a multitude of parallel threads. An operating system is a system that is aware of its resources, capable of monitoring its environment and make decisions on the handling of information that comes in and is being issued out. These are already the basic ingredients of an 'intelligent environment'. A step further is whereby the various composing systems are capable of developing 'context awareness', e.g. they have a good idea of the communication environment they are in, are able to customise is, set up co-operation with neighbouring systems, react adequately to user demands, handle power in a 'clever' way, perform graceful recovery etc. On one level higher, these systems might be able to set up co-operations to achieve common goals they are sharing and that have been negotiated with users. In short: because of the availability of large scale and distributed processing power, a whole hierarchy of 'intelligent' layers will appear. Research has to be devoted to the development of design templates or design patterns covering these layers in which designers can express what they expect from the system they are conceiving. In view of the intrinsic difficulty of handling operating system type behaviour, not to mention a large collection of co-operating systems in one environment, and in view of the known difficulties with programming multi-agent systems, this may be an arduous design issuein addition to the 'change in paradigm', designers have to adapt to (it may be that traditional hardware designers are better equipped to handle interactions between autonomous systems than their software counterparts who are used to strictly structured hierarchical designs!). In addition to these architectural questions, classical signal processing tasks connected to intelligence such as pattern recognition, voice identification and biometry may become much more powerful using the new environments we are considering. And last but not least, we may expect that 'semantic information' will move down from the web into the local environments and up again. This means that the local environment will have to be able to handle these internet layers of intelligence, and in particular the signal processing aspects of this such as extracting semantic information from images, speech and other media of information. The 'intelligent system design level' will become of paramount importance in the future.

5.6 Architectures for this all: what will they look like?

The architectures should offer flexibility to accommodate new applications (maybe downloaded over the network connection), adaptability to changes in the environment, and should offer maximal QoS to the user. Flexibility can be provided by the use of software running on embedded processors. Emerging fine-grain and coarse-grain dynamical and partial reconfigurable architectures open the path to dynamic “task” creation in hardware. Future SoC platforms will consist of tiles containing combinations of application-specific components, various types of instruction-set processors, reconfigurable hardware, and memories, all communicating through an on-chip network. On these platforms various tasks can be created dynamically either in software or hardware and relocated at run-time from software to hardware and back, according to the requested QoS of the various simultaneously running applications. An additional challenge is the incorporation of 'intelligence'. The first emanation of this is the ubiquitous presence of autonomous systems, a large collection of operating systems handling their resources and interacting with each other in the local environment and the Internet. Because of their compact nature and the special requirements placed on them, these operating systems may need special realisation techniques. They may be subdivided into subsystems that are acting as agents in their own right (e.g. a subsystem that is continuously gauging the properties of the surrounding communication environment) or as subsystems that execute dedicated and complex tasks such as recognition and detection, rule base consultation, signal decoding, resource polling etc. One stage higher, situation abstraction can be performed and shared between agents, division of tasks, merging of information and sharing of data. In extreme cases the very presence of specific agents and information may be doubtful requiring redundant and fault tolerant designs distributed over collections of agents. The whole field of design of co-operating intelligent agents is still wide open for research, and will impact the development of new integrated systems.

An example for such architecture was demonstrated by Andrea Cuomo:

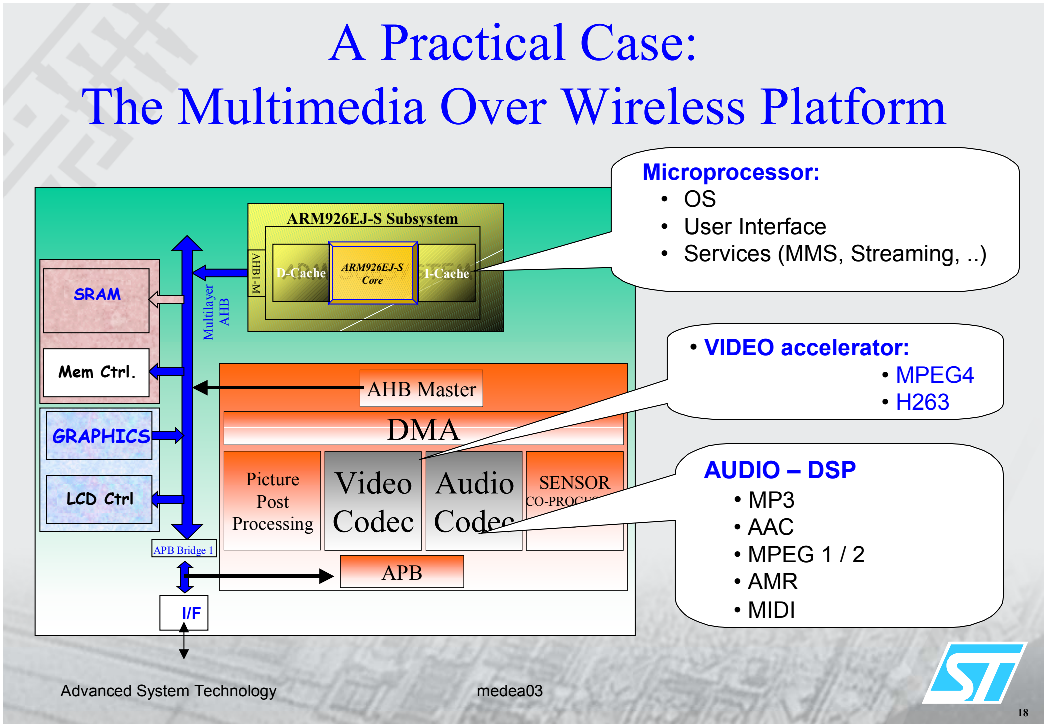

© Andrea Cuomo, STMicroelectronics

5.7 Design technologies

We must revise what we understand by design technology. If it is the creation of system level design tools to be marketed in the EDA back-end way, there is little hope. Design technology makes sense if it is developed as a methodology that supports the design of the type of killer applications discussed above. In fact what we see is that, what used to be called design technology, now becomes part of the implementation of the system. An example of this is the end-to-end optimisation of a communication link for a given QoS.

Nanoscale integration implies that larger portions of an electronic system of great complexity can be integrated on a single chip or on a single package. What used to be regarded as a system in itself will in the future become a component of a much larger system. The complexity of the design problem will scale up to limits that are unthinkable today. Current tools and methodologies will not scale or be capable of handling the new level of complexity. This phenomenon has been coming up in recent years and has too often been discarded by the microelectronics community. Only a few years ago it was stated, most notably in Europe, that the 'MP3 algorithm is too complex for integration'. Now that we can buy MP3 chips and coding/decoding systems of even higher complexity for less than 10 Euro (designed in the Far East…) we know better! Europe cannot be ambitious enough in these matters. Clearly, we need a radical change in our design paradigms. Multidisciplinary approaches will be the basis for this change. Even today there are only a few examples of successful systems on chip that have been taken to market. The challenges clearly lay in design methodologies that rely upon re-use, on design for manufacturing methods, and on the careful analysis of price/performance barriers that will guide the decisions of full integration versus multi-chip solutions.

Nano-scale Integrated Systems design including MEMS will not only be electronics design but also involve several diverse physical aspects, such as electricity, magnetism, optics, radio, quantum, thermal properties, mechanics, chemistry, and biology; i.e. it will be a *multi-physics* design. Smart and wireless devices (e.g. PDA) require RF and digital mixed signal design. Advanced fabrication technology requires Integrated Systems designers to cope with all the physical imperfections (e.g. crosstalk, transmission line effects, power supply noise, and leakage) by considering the coupled interaction. Currently, very few multi-physics effects are considered in the SoC design phase. No standards are available for the exchange of information describing multi-physics at different modelling levels. Very few methods are available for measuring nano-scale multi-physics material parameters.

The nanoscale CMOS and non-CMOS technologies will require new concurrent optimisation of application platforms, circuit architectures, and design methods in order to reduce the cost and resource demands of System-of-Chip designs and thus enabling a larger scale industrial deployment. We anticipate that the main obstacle is not the technology but the design and usage community. The technical challenges are: (i) computational complexity within SoCs, (ii) communication complexity, and (iii) physical complexity. New architectural and design paradigms need to be established for nanoscale SoC. In addition to the current computational complexity issues, on-chip communication complexity and physical complexity are influencing all levels of the design abstraction.

We need to raise the SoC abstraction level from conventional RTL (register transfer level) and ISA (instruction set architecture) level to platform API (application programming interface) level. ‘Network-on-chip’ is a new area aiming to enable SoC design to start from a SoC platform that is configurable for both computation and communication (using the network concept). Compared with conventional SoC design which starts from scratch or from a relatively fixed platform, the configurable SoC platform simplifies the SoC design thereby giving a faster design.

Interaction between heterogeneous subsystems is a major bottleneck when dealing with Integrated Systems. Analogue and mixed-signal capabilities are required more and more. Furthermore, since analogue and digital must coexist and co-operate, many new problems arise especially when high-performance is needed and/or when low-voltage, low-power is demanded. Achieving an adequate dynamic range, when biasing is low, demands new circuit techniques which have to be compatible with more conventional approaches. Compatibility with high-level digital design techniques is another important goal as well as a common framework for design and test. Finally, exploring design alternatives that may be considered unconventional is worth a mention. In the context of such a heterogeneous system such as a future SoC, taking advantage of the sensor/actuator properties combined with analogue processing cells and digital memory may be a promising alternative such as in massive visual processing.

As billions of transistors per single SoC are predicted in a near future, power consumption is a dramatic issue, not only for portable devices, but also for high-end chips due to heat dissipation and reliability. New applications, such as wireless ad hoc networks, require extremely low power consumption, as the lifetime has to be several years on a single small battery. Very advanced deep submicron technologies (90 and 65 nanometers, SOI) will present a dramatic increase of leakage power, interconnect problems and crosstalk, requiring new design and circuit techniques, such as very low and multiple Vdd (i.e. 0.4 to 0.6 Volt), complex DVS (dynamic voltage scaling) schemes and new power sources. On the system level, more and more programmable and reconfigurable SoCs including wireless links, which impacts their power consumption, will be required by customers. New high-level design methods are required to trade-off flexibility and low power consumption.

The most important SoC challenges are design space exploration and the validation at the architectural level, implementation productivity, manufacturing costs and the integration of heterogeneous components from multiple implementation fabrics. Exponentially increasing system complexity (algorithmic and platform complexity, functionality etc.) and silicon complexity (parasitic effects, crosstalk, global interconnect, process variability etc.) place long-standing design methodologies at risk.

With SoCs, silicon is becoming the unifying medium into which system design paradigms from RF, A/D interfacing, MEMS, optoelectronics, dedicated signal processing units, programmable and reconfigurable architectures, embedded real time software etc. have to be merged. Thus, SoCs represent multi-paradigm systems and many paradigm shifts may become necessary.

Design methodology issues at higher abstraction levels are frequent and implementation platforms are always very domain specific and cannot be considered independent of the application domain ("pyramid tops are domain specific" according to Prof. Hugo De Man). Moreover SoCs have to be globally optimised within the environment in which they operate – both from an implementation point of view *and* cost point of view. Thus, there is a strong demand for application domain-oriented research. We explicitly refer to the Chapter ' System and Design Issues' of the Report of the MEDEA+ Scientific Committee on "Heterogeneity on SI or in a package for future system innovation", 2003.

5.8 A network of industry and academia

A strategy for research and development, as well as for product definition and marketing for electronic systems in view of the incoming nano era, requires expertise and even excellence in so many diverse areas that even the largest semiconductor industry in the world, INTEL, cannot conceive it in isolation. The worldwide trend is to develop a joint program between industry and academia with public and industrial funding sustained over a rather long time horizon to develop methods and technologies to master the difficult art of Integrated System Design. The coalitions being formed cut across disciplinary areas; system design and manufacturing, semiconductor manufacturing and design, manufacturing equipment and IP creation are all diverse topics that need to be addressed for Integrated System Design.

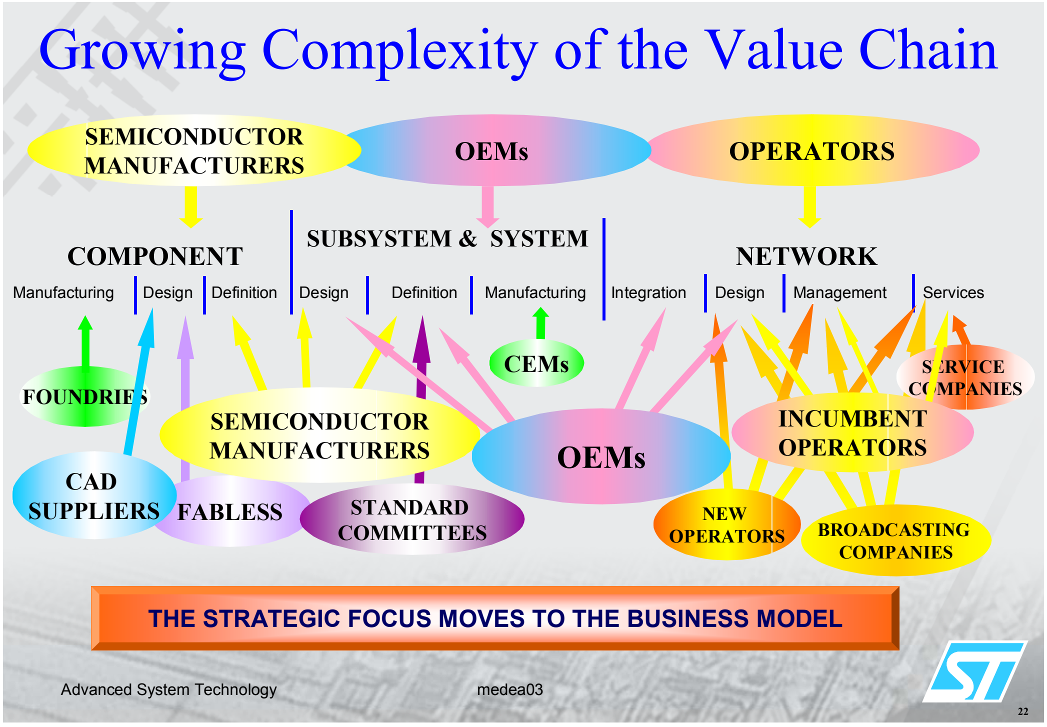

© Andrea Cuomo, STMicroelectronics

For example, the GSRC program is funded by electronic system and semiconductor companies united under the MARCO program and by DARPA. This research consortium has been operating, together with three companion programs, for four years and it is expected that they will be renewed for another three years. The approach seeks to identify the best groups in the US to create a joint research program. The funding is directed towards carrying out specific research and, only as a by-product, towards creating an infrastructure across different Universities to support research for many years to come. In Europe, the research community is certainly more fragmented being more widely dispersed among centres in different countries and, historically, it has been more distant from industrial applications.

We need to change the way research is carried out in the field of Integrated Systems by bringing the different communities together. To do so, it is necessary to secure enough funding and incentives to create a virtuous circle whereby research groups will team up to identify promising research areas and to share knowledge across industry-academia boundaries**.** Researchers from different knowledge domains have to be summoned in a common program where ideas can be freely exchanged and expensive infrastructures can be leveraged. Industry and academia have to come together as never before. Information must be shared. Education must be reorganised and harmonised to accommodate the disciplines that form the basis of Integrated Systems technology.

6. Consequences for future systems on a chip

It is envisaged that these future SoC platforms will consist of tiles containing combinations of application-specific components, various types of instruction-set processors, reconfigurable hardware, and memories, all communicating through an on-chip network. On these platforms various tasks can be created dynamically either in software or hardware and relocated at run-time from software to hardware and back, according to the requested QoS of the various simultaneously running applications.

A real-time operating system will manage all the resources of the future SoC platform and will be responsible for the dynamic creation of tasks in hardware and software, re-location of running tasks from hardware to software and vice-versa, and unification of the communication mechanisms between tasks running in hardware and software. The RTOS should abstract as many details as possible from the underlying architecture and offer the application programmer a software-like abstraction of the platform.

Rapid evolution in sub-micron process technology allows ever more complex systems to be integrated into one single package. Technology advances are however not followed by an increase in design productivity, causing technology to leapfrog the design of integrated circuits (ICs) and consumer markets.

We believe that several key research issues need to be addressed to enable the above vision, with the following believed to be the most crucial:

- Pre-processing application source code transformations,
- Management of dynamically allocated complex data types,
- Instruction and configuration storage and access management,
- Re-configurable processor mapping and handling of concurrent tasks,
- Power-speed trade-offs for embedded memories in interconnect-dominated deep sub-micron technologies,
- On-chip interconnect networks in interconnect-dominated deep sub-micron technologies.

For the design of the front-ends, on-chip active components will be combined with passive components (inductors, capacitors, and even complete RF band pass filters) and antennas that are processed in new technologies. The combination of this technology with silicon technology for the active components will enable the development of new architectural concepts. Further, micro machining will be used to make MEMS components such as high-performance antenna switches, variable capacitors, very sharp filters, etc.

The implementation of a highly integrated mixed-signal design, even in the case of multi-chip solutions, will also typically contain a mixed-signal technology that combines digital functionality with the analogue blocks.

7. Recommendations for MEDEA+

From the detailed recommendations, esp. in chapter 5 ‘Research issues’ we derive the following:

1. Make aggressive moves towards the new G4 paradigm. Take the lead by putting strong efforts into the open standardisation of its components (including operating systems).

2. Support the implementation of new standards. A significant amount of good basic work has been carried out and standardised with the involvement of European companies, institutes and universities which now have to compete with proprietary solutions. In a lot of cases reference designs are needed for complex new standards (e.g. MPEG-4, MPEG-7, MPEG-21).

3. Develop an energy-efficient platform for 'intelligent system design', consisting of embedded processors, modules for decision making, modules for pattern recognition, local networking devices and a variety of networking protocols that can be logically integrated. Move aggressively in extending the capabilities of SystemC in this direction. Ensure the completeness of the design trajectory from system design, verification all the way down to mapping on hardware (chip, package) AND creation of low voltage, low leakage IP blocks with the emphasis on low leakage embedded memory. This requires efforts in the front-end (architecture design, mapping) but also at circuit level (e.g. control of leakage current and its impact on the front-end). This can best occur in a network of industry, universities and research centres (see item 8).

4. Start research efforts at higher system levels. The emerging 'user's control' paradigm requires new ways of handling interactions with a community of user's that get more and more demanding, desire more services, want to exercise more control, have expanded computing and memory facilities and require more security. These scalability, architecture and security problems are solvable but not without a major effort. Europe should play a strong role here.

5. Explore aggressively the new possibilities in the multimedia era: new types of transducers, interaction with the Internet, user's control, new services. A lot of opportunities for new products here! 'Ambient intelligence' is an important step, but more is needed in terms of applications, services, Quality of Service and roaming user support.

6. Start new programs to bridge the big gap between basic research and readily available reference implementations to jump start the development of new products, since higher semantic levels of metadata and user interface are one important opportunity for future systems-on-a-chip.

7. Start programs in 'body area networking' and put extra emphasis on the development and system integration of devices surrounding the individual.

8. Develop stronger co-operation with team-driven network of universities (and research institutes). The strength of the US microelectronic environment has been greatly enhanced

by the co-operative attitude of the American companies w.r.to a large collection of universities based all over the country. Many ideas and many experiments are needed since success is based on small probabilities. Develop broad regional involvement, try to get many universities and research groups interested and contributing new ideas.

In addition, we strongly recommend that MEDEA+ and its industrial and institutional partners follow 'Marketing in Time to Achieve Time to Market approach discussed by Andrea Cuomo at the Leuven-Workshop on 4 June 2003:

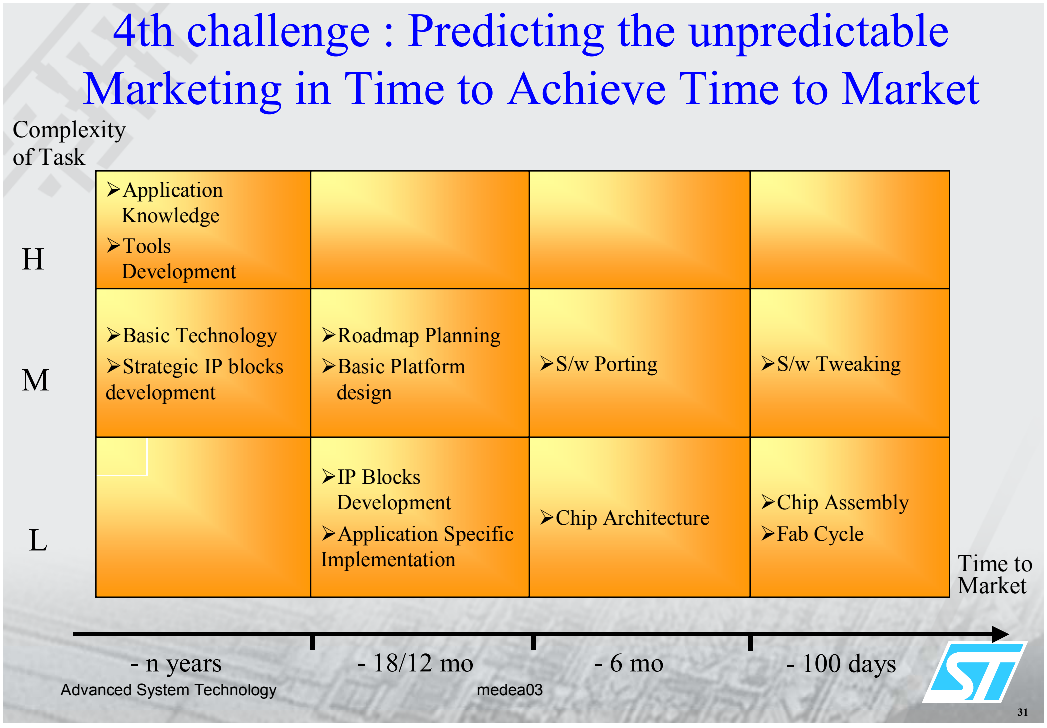

© Andrea Cuomo, STMicroelectronics

References

- ITEA roadmap 'Softec' project http://www.itea-office.org/projects/itearoadmap.htm
- ISTAG (ambient intelligence, ftp://ftp.cordis.lu/pub/ist/docs/istagscenarios2010.pdf)
- http://www.wireless-world-research.org/BoV1.0/BoV2001v1.1B.pdf
- ITRS design roadmap http://public.itrs.net/Files/2001ITRS/Design.pdf
- Embedded systems roadmap http://www.stw.nl/progress/Esroadmap/ESRversion1.pdf
- Report of the MEDEA+ Scientific Committee on "Heterogeneity on SI or in a package for future system innovation", 2003
- MEDEA+ applications roadmap
- ISSCC 2003
- http://zdnet.com.com/2100-1103-948493.html

Glossary

WiFi	Wireless Fidelity: The Wi-Fi Alliance is a nonprofit international association formed in 1999 to certify interoperability of wireless Local Area Network products based on IEEE 802.11 specification.	www.wi-fi.org
PAN	Personal Area Network	
PN	Personal Network	
PIC	Personal Information Center	
PDA	Personal Digital Assistant	
VDSM	Very Deep Sub-Micron effects in IC design	http://www.eedesign.com/story/OEG20000609S0017
MPEG	Moving Picture Experts Group is the name of family of standards used for coding audio-visual information (e.g., movies, video, music) in a digital compressed format.	http://www.mpeg.org/MPEG/index.html
SMS	Short Message Services	
WDA	Wearable Digital Assistant	
QoS	Quality of Service	
EDA	Electronic Design Automation	
RTL	Register Transfer Level	
ISA	Instruction Set Architecture	
RTOS	Real-Time Operating System	
GSRC	Gigascale Silicon Research Center	http://www.gigascale.org/
DARPA	Defense Advanced Research Projects Agency	http://www.darpa.mil/
MARCO	Microelectronics Advanced Research Corp.	http://fcrp.src.org