UNIVERSITÀ DI PISA

Scuola di Dottorato in Ingegneria "Leonardo da Vinci"

Corso di Dottorato di Ricerca in
Ingegneria dell'Informazione

Tesi di Dottorato di Ricerca

Embedded computing systems design: architectural and application perspectives

Tony Salvatore Bacchillone

Anno 2013

UNIVERSITÀ DI PISA

Scuola di Dottorato in Ingegneria "Leonardo da Vinci"

Corso di Dottorato di Ricerca in
Ingegneria dell'Informazione

Tesi di Dottorato di Ricerca

Embedded computing systems design: architectural and application perspectives

Autore:

Tony Salvatore Bacchillone _____

Relatori:

Prof. Luca Fanucci _____

Prof. Sergio Saponara _____

Anno 2013

SOMMARIO

Questo elaborato affronta varie problematiche legate alla progettazione e all'implementazione dei moderni sistemi embedded di computing, ponendo in rilevo, e talvolta in contrapposizione, le sfide che emergono all'avanzare della tecnologia ed i requisiti che invece emergono a livello applicativo, derivanti dalle necessità degli utenti finali e dai trend di mercato.
La discussione sarà articolata tenendo conto di due punti di vista: la progettazione hardware e la loro applicazione a livello di sistema.
A livello hardware saranno affrontati nel dettaglio i problemi di interconnettività on-chip. Aspetto che riguarda la parallelizzazione del calcolo, ma anche l'integrazione di funzionalità eterogenee. Sarà quindi discussa un'architettura d'interconnessione denominata Network-on-Chip (NoC). La soluzione proposta è in grado di supportare funzionalità avanzate di networking direttamente in hardware, consentendo tuttavia di raggiungere sempre un compromesso ottimale tra prestazioni in termini di traffico e requisiti di implementazioni a seconda dell'applicazione specifica. Nella discussione di questa tematica, verrà posto l'accento sul problema della configurabilità dei blocchi che compongono una NoC. Quello della configurabilità, è un problema sempre più sentito nella progettazione dei sistemi complessi, nei quali si cerca di sviluppare delle funzionalità, anche molto evolute, ma che siano semplicemente riutilizzabili. A tale scopo sarà introdotta una nuova metodologia, denominata *Metacoding* che consiste nell'astrarre i problemi di configurabilità attraverso linguaggi di programmazione di alto livello. Sulla base del metacoding verrà anche proposto un flusso di design automatico in grado di semplificare la progettazione e la configurazione di una NoC da parte del designer di rete.
Come anticipato, la discussione si sposterà poi a livello di sistema, per affrontare la progettazione di tali sistemi dal punto di vista applicativo, focalizzando l'attenzione in particolare sulle applicazioni di monitoraggio remoto. A tal riguardo saranno studiati nel dettaglio tutti gli aspetti che riguardano la progettazione di un sistema per il monitoraggio di pazienti affetti da scompenso cardiaco cronico. Si partirà dalla definizione dei requisiti, che, come spesso accade a questo livello, derivano principalmente dai bisogni dell'utente finale, nel nostro caso medici e pazienti. Verranno discusse le problematiche di acquisizione, elaborazione e gestione delle misure. Il sistema proposto introduce vari aspetti innovativi tra i quali il concetto di protocollo operativo e l'elevata interoperabilità offerta. In ultima analisi, verranno riportati i risultati relativi alla sperimentazione del sistema implementato.
Infine, il tema del monitoraggio remoto sarà concluso con lo studio delle reti di distribuzione elettrica intelligenti: le Smart Grid, cercando di fare uno studio dello stato dell'arte del settore, proponendo un'architettura di Home Area Network (HAN) e suggerendone una possibile implementazione attraverso Commercial Off the Shelf (COTS).

ABSTRACT

This work deals with embedded computing design, highlighting both: typical challenges coming from technology advances and system requirements concerning the application level. The second ones generally result from end-user needs and market trends.
The discussion takes into account two main points of view: the hardware design and the application or system level design.
Concerning the hardware point of view, the on-chip interconnection will be studied in the detail. This aspect is related to computing parallelization and the integration of heterogeneous functionalities. An interconnection architecture known as Network-on-Chip (NoC) will be discussed. An innovative NoC, offering advanced networking functionalities directly implemented in hardware, will be introduced. The proposed design offers a high flexibility allowing always finding the optimal trade-off between network performances and implementation requirements (derived from specific application). This part will mainly focus on the configurability of the NoC building blocks. Design configurability and reusability is becoming more and more important with the increasing design complexity. In order to effectively address this issue, an innovative methodology called *Metacoding* will be introduced. This methodology consists on the abstraction of the configuration intents using a high-level programming language. Moreover, an automatic design flow based on the metacoding methodology will be proposed. The developed framework simplifies the design and the configuration of a NoC for a network designer.
Concerning the system level point of view, remote monitoring applications will be discussed. This part will focus in particular on the design of a telemonitoring system for patients affected by chronic heart failure. Particular attention will be paid to the system requirements definition coming from end-users: in our case, clinicians and patients. The proposed system is based on an innovative Operating Protocol. The use of well-known communication protocols and standards favours the maximum interoperability with others systems. Vital signals acquisition, data processing and collection will be discussed. Finally, the results of the demonstration phase, involving real patients and clinicians, will be reported. An exploration of energy monitoring for Smart Grids closes the discussion on remote monitoring applications. After a review of the state of the art, a Home Area Network (HAN) architecture will be proposed, taking into account requirements and security issues; a possible implementation based on Commercial Off the Shelf (COTS) is discussed.

CONTENTS

1 INTRODUCTION: COMPUTING CONCEPTS, TREND AND CHALLENGES

"The number of transistors that can be fabricated on a very large-scale integrated (VLSI) chip doubles approximately every two years". Since its introduction in 1965 by the co-founder of Intel, Gordon Moore, and despite having been fixed during years to match market trends, *Moore's Law* still keeps its meaning and validity.
In practical terms, the result is that computing performance doubles every 18 months, and has done so for decades, since early 70's, up to the present.
The ability of the silicon industry to fulfil Moore's law was mainly achieved by the continuous downscaling of the critical dimensions in the integrated circuit: improvements in semiconductor manufacturing processes allow the fabrication of smaller transistors on wider wafers, thus boosting economies of scale. Above all that, advances in design technologies, design paradigms and design methodologies enable engineers to effectively handle more and more complex designs, bridging the design complexity versus the designer productivity gap.
This chapter lays the groundwork for the topics discussed in this dissertation, by presenting current state of the art and future trends of computing architectures and application domains. An overview of the organization of the whole work closes the chapter.

1.1 More Moore and More than Moore

As introduced in the previous section, from a technology perspective, the continuous increase in the integration density proposed by Moore's Law was made possible by a dimensional scaling. The transistors scaling brings two main advantages: on the one hand it allows to increase the density of silicon integration and on the other hand it allows to save power thanks to lower capacitance and supply voltage. High density and low power are very important achievements for computing, in particular in case of battery-powered portable electronic devices, which are among the most requested consumer electronic products.
According to the International Technology Roadmap for Semiconductors (ITRS) classification, this trend in circuit miniaturization by scaling down the transistor and its associated benefits in terms of performances is known as "More Moore" [1].
Continued shrinking of physical feature sizes achieved in the More Moore domain, directly impacts on digital functionalities (i.e. microprocessors, memories, and digital logic). This aspect involves around 70% of the total semiconductor component market.
However, a microelectronic device is usually composed of many digital functionalities and non-digital functionalities. In the past, Application Specific Integrated Circuits (ASICs) implemented digital functionalities of a system, while the non-digital functionalities were developed at Printed Circuit Board (PCB) level, with glue logic or analog circuits providing connectivity with the custom chip.

With technology scaling and present progress in both process technology and design, digital functionalities and non-digital functionalities are being crammed in the same package containing the integrated circuit, or even in the chip area.
The combination of multiple functionalities into a single package is generally known as System in Package (SiP). A SiP may contain for example passives, Micro Electro-Mechanical Systems (MEMS), optical components and other packages and devices. On the other side, the integration of digital and non-digital functionalities on a single chip is generally referred to as System On Chip (SoC). A SoC may contain for example digital, analog, mixed-signal, Radio-Frequency (RF) functions, etc.
Non-digital functionalities do not necessarily scale at the same rate as digital ones because of technology scaling. As a consequence, their incorporation does not contribute to the miniaturization of electronic systems. Nevertheless, they provide an additional value in different ways and allow the migration of non-digital functionalities from the system board-level into the package or onto the chip. The combined need for digital and non-digital functionalities in an integrated system is translated by the ITRS as a dual trend and is labelled as "More-than-Moore" [1] (see Figure 1.1).

1.2 Design challenges

The two trends introduced in the previous section arise a stream of challenges to be overcome when designing an electronic system. As a first stage approximation, we can consider that More Moore trend originates challenges

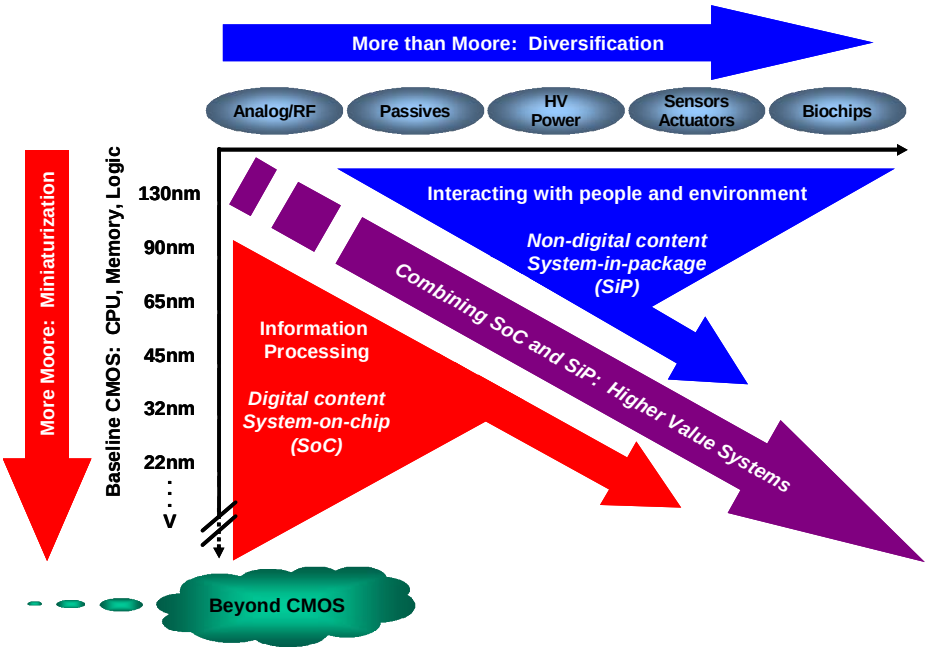

Figure 1.1. More Moore and More-than-Moore trends (source ITRS).

in terms of complexity because of the growing operational demand. At the same time, More-than-More trend introduces a new set of problems quite often concerning various fields.
The most relevant challenges are:

- Design complexity: as circuit complexity grows, design takes much time and it is more expensive. Masks and manufacturing costs become higher and testing costs grow in a comparable way. This issue is dramatically highlighted when considering market effects: design space exploration, mask manufacturing, testing strategies represent an increasing delay for time-to-market, while higher costs can be recovered only with a significant time-market.
- Wiring: as a consequence of transistors scaling and their faster switching activity, wires become the bottleneck for integrated circuits. Wires introduce parasitic capacitances that raise their charging time compromising the overall circuit speed. This effect is particularly highlighted for global wires, like clock wiring nets. The number of required wires generally increases with design complexity and, due to crosstalk effects, this may compromise the integrity of the information carried by wires. Thus, front-end synthesis results are no more reliable since the performance strongly depends on how wires are placed and routed along the circuit. The wiring model provided in front-end design space, is not sufficient to predict circuit performance below 90 nm technologies. This means a continuous feedback between pre-layout and post-layout in the project flow.
- Communication: point-to-point links and traditional shared buses show significant limitations as the number of integrated functionalities increases. The more functionalities are in the chip, the more communication wirings are needed thus compromising scalability, performance and design time. Furthermore, cores need efficient communication infrastructure to exploit intensive parallel communication and to access shared resources.
- Heterogeneity: integration of digital and non-digital functionalities on chip requires dealing with many sub-designs. As discussed, a complex design could foresee the integration of digital, analog, mixed-signal, sensors, radio-frequency functions, etc. in the same chip. This generally implies different expertises to converge in the design group, increasing design space, costs and development time.
- Tools: in the last decades electronic industry productivity has been remarkably pushed forward thanks to CAD (Computer Aided Design) tools. Future SoC development poses significant challenges even for CAD tools, since they must support a larger design space and an increased complexity over functionality. Moreover, an important challenge is also related to verification activities, where simpler techniques are required together with faster and more accurate simulation engines. As the design paradigms evolve even new services of CAD are needed.

- Methodologies: as estimated by SEMATECH, a consortium which performs research and development to advance chip manufacturing, design complexity presents a 58% annual increase in the number of transistors per chip, while the designer productivity grows by only 21% per year (see Figure 1.2). Embedded computing is one partial answer in order to overcome to this gap, since some of the design tasks are moved from hardware to software. However, improved design methodologies are required to deal with the increasing design complexity. Some of the methodology challenges are: to propose effective methods approaching the design by level of abstractions, to simplify the integration of functionalities composing the design and to favour the design reuse.

1.3 Next generation computing

Despite their great performance in terms of computational speed, power consumption and area, ASIC design paradigm undergoes huge limitations approaching billions of transistors designs. ASIC is a strongly application-oriented approach and as a consequence, when increasing designs complexity, it is characterized by lack of flexibility and reusability. Moreover, ASICs are susceptible to the rapid changes of consumer market across technology generations.
On the other hand, Digital Signal Processors (DSPs) and General Purpose Processors (GPPs) can offer relevant benefits in terms of flexibility though showing upper bounds on power supply voltage and clock frequency.
This aspect becomes an important issue especially considering the market trend where electronic devices for consumers represent a large percentage of the electronic request. This sector is mainly dominated by portable electronic devices for entertainment, communication, health & wellness, etc. Some examples are: smart phones, laptops & ultraportable computers, tablets, digital cameras, GPS navigators, multimedia devices, video game consoles, portable devices for fitness, diet & elder care. These devices are generally portable and although most of them offer very high performances in terms of processing, their success on the markets may be related to the battery

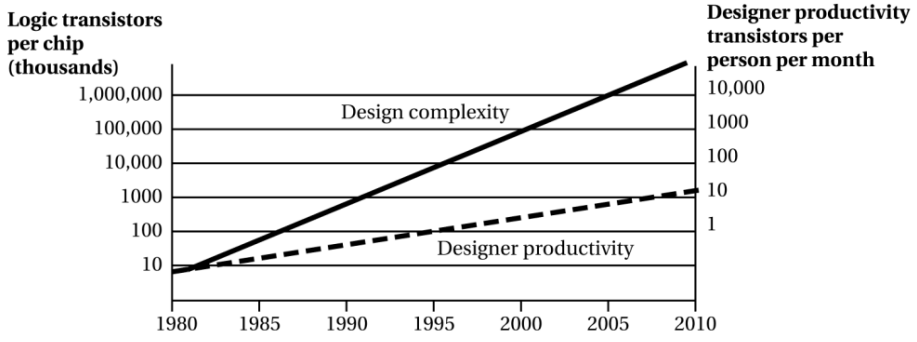

Figure 1.2. Design Complexity vs. Designer Productivity (Source: SEMATECH).

duration and other technical aspects.
Therefore, while power consumption and clock frequency issues cannot be ignored nor overcome, parallelization represents an interesting chance to deal with processing power. Multiple processing cores can be integrated in a single silicon chip offering higher computational power, and at the same time, avoiding the burden of a raising clock frequency, and hence a raising dynamic power consumption. This solution exploits the parallel computation based on many processing units (or cores) running at a lower frequency. This paradigm is also known as Multi Processor SoC (MPSoC). Examples are Intel Core Duo, i3, i5, i7 and Xeon, AMD Phenom II and FX, ARM Cortex-A9 or ARM11-MPCore.
Power consumption and frequency issues do not concern just portable devices, but they are starting to affect other sectors. This is due to an increasing common interest in caring energy utilization, together with emerging changes around computing requirements. Indeed, the wide diffusion of the Internet and its services, in particular the cloud computing, has caused processing loads to be transferred from the end-user to network providers. As a consequence of this scenario, on the one side users can enjoy a large amount of services saving local resources without suffering of portable devices limits. On the other side, service providers are in charge of answering a huge and time-variable request of computing resources.
Thus, recently market interest and research started focussing on low-power servers, thinking at more flexible architectures able to sustain high processing load peaks while reducing power consumption and avoiding power wastefulness during low processing load periods. Also in this scenario MPSoC architecture is definitely the most promising paradigm in terms of flexibility and power efficiency.
Furthermore, as discussed at the beginning of the chapter, according to the More-than-Moore trend, next generation of electronic devices will integrate multiple digital functionalities, non-digital computing, analog & mixed-signal processing, MEMS, RF components etc.
The design issues mentioned in the previous section suggest that the future computing design must be inspired on a flexible and reusable solutions able to face the raising design complexity and to meet time-to-market requirement. New design paradigms and methodologies, as well as advanced CADs are the key of the success for designer engineers in order to face new challenges of the computing design.
Finally, traditional on-chip buses represent the bottleneck on complex designs and in particular for parallel computing, so compromising the performance of the whole system. Network-on-Chip (NoC) is a promising methodology which allows designing of interconnecting architectures independently from the connected cores. The challenge in this case is to minimize the design effort while attempting to cover the widest application space in terms of traffic requirements (i.e. throughput, latency, quality of service, etc.) and implementation requirements (area, clocking scheme, and power consumption).

1.4 Application Domains

This section aims to figure out most promising application domains, trying to highlight opportunities and challenges. In particular, it is focused on healthcare domain. Many works can be found in the literature concerning the subjects dealt with in this section. In particular, the information discussed here mainly come from recent ITRS roadmaps and related reports [1]-[3].
Basically, the success of an application is based on a deep understanding of the societal needs and their relative trends, mixed to the ability to face them with simple, effective and attractive solutions.
According to the ITRS indications, social trends can be mainly grouped into five categories:

- energy & environment
- transport & mobility
- health & wellness
- security & safety
- communication & digital lifestyle

where the latter covers also consumer electronics for infotainment.
Each of these trends represents significant opportunities for future applications in the consumer electronics, automotive electronics, medical applications, communication, etc.
Figure 1.3 shows some examples of application linking societal trends and

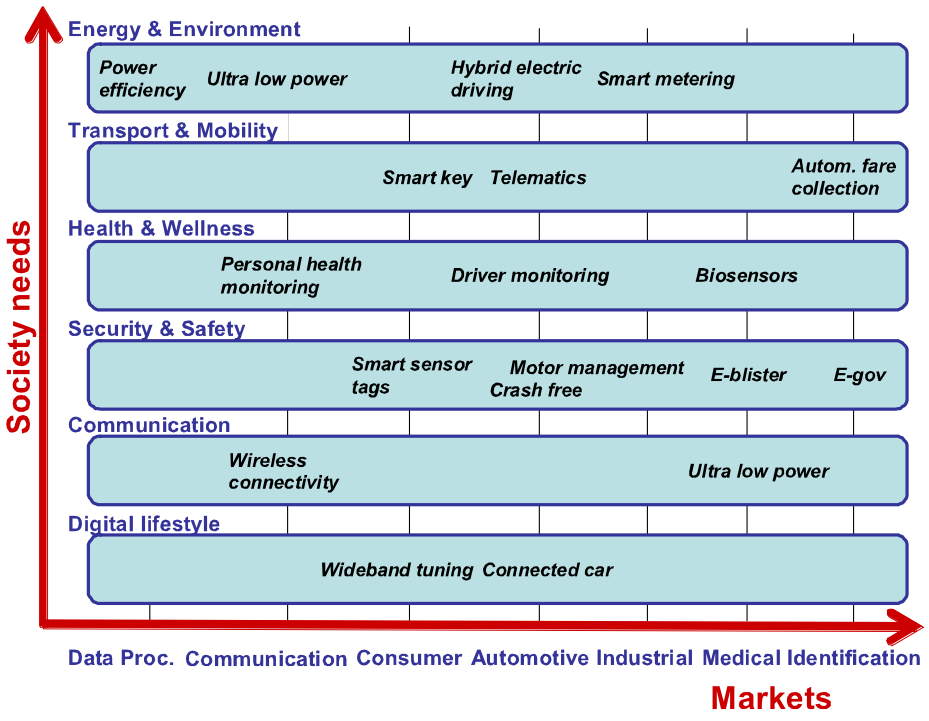

Figure 1.3. Application matrix (source: ITRS).

markets.
Once focused on societal needs and trends, the following step is answering them with an innovative solution. Typically there are two main alternatives to corner the market.
The first one lies in the improvement of an existing solution. As for a portable consumer device, for example, this means improving its performance in terms of computational power, portability or battery duration. In some cases, the improvement is the result of crossing domains.
The second possibility consists in the introduction of a new solution opening a new market. For example, a bio-medical chip that may well revolutionize healthcare.

1.4.1 Healthcare

There is a wide consensus about the necessity of health monitoring application and Wireless Sensor Networks (WSN) as means to deliver higher quality health monitoring at lower cost. This general opinion is supported by many data, related to the high prevalence of certain chronic diseases and the ageing population in western countries. In fact, a large percentage of the total expenditure on health costs in Europe is spent on a few chronic diseases. Moreover, it is well known that the expenditure on healthcare will grow in the near future, because of an ageing population.
The use of health monitoring systems could result in more prevention and more on-line monitoring and this could rapidly reduce healthcare costs.
Health monitoring devices can be divided into four main categories according to their application domains as follows:

- Monitoring: it consists on a long-term collection of relevant data in order to allow the medical professional to make diagnosis.
- Prevention: continuous monitoring of a known disorder in order to provide an alarm to detect and/or prevent the onset of the disorder. Prevention can also be used to assist in preventing the development of a disorder such as obesity, by monitoring the risk factors.
- Closed loop: automatic detection of an adverse event, and automatic supplying outputs to overcome the adverse event.
- Spot-check: whenever desired, the user can measure a body function.

Most relevant elements composing a generic health monitoring system are the following:

- sensors: appropriate sensors for accurately monitoring the relevant signals;
- digital signal processing: providing intelligence for reducing artifacts and recognizing particular events;
- integration technologies: combining electronics components in order to favour a suitably continuous monitoring;
- memory: some applications use local storage of data, and therefore rely on the availability of memory;
- power management: for autonomous devices, particularly relevant for portable and/or wearable devices;

- radio technologies: for real-time monitoring and alerting, in order to transmit messages.
- body/electronics interface: to capture the relevant body signals (i.e. electrodes with or without contact gel), and their corresponding interfaces.

Application domain	Sensors	Artifacts reduction	Event detection	Integration	Memory	Power	Radio
Monitoring	high	high	low	high	high	med	med
Prevention	med	high	med	high	med	high	high
Closed-loop	med	high	high	high	low	high	low
Spot-check	med	low	high	low	med	low	high

Table 1.1. Requirements for health monitoring (source: ITRS).

Table 1.1 gives a general indication about the influence of the monitoring health system components for each application domain.
Recent efforts on medical device development have focused on providing solutions for wearable health monitoring systems. The goal is allowing physiological monitoring and interpretation in daily life environments. However, some challenges are still to be faced. To provide these solutions, devices must be smaller and more comfortable to wear, be robust against motion artifacts, be power efficient, intelligent, and be able to communicate with the user.
Finally, today's interoperability and interconnectivity of these devices is ignored, resulting in limited monitoring functionality and limited intelligence achieved. Development of platform technologies towards wearable health monitors, may lead today's limitations to be overcame.

1.5 Organization of this work

This work aims to give an overview on embedded computing design activities, taking into account two main viewpoints: the hardware level design and the application or system-level design. As discussed in this introduction, challenges and trends appear quite different for the two cases. While in the first case the research is mostly technology-driven, in the second case social needs are the main matter. However, the starting point in order to reach the greater success in future computing is a close collaboration between research domains, designers, industries and markets.
Chapter 2 represents an example of hardware level design. It introduces the design of a NoC interconnecting architecture for MPSoCs, based on high configurable building blocks, supporting advanced networking features directly implemented in hardware. In particular, the discussion focuses on a typical issue affecting complex designs: the configurability of a complex Intellectual Property (IP). This work proposes the *Metacoding* methodology and an

automated design flow to overcome configuration issues. Without loss of generality, the methodology will be discussed referring to the particular case of the NoC building blocks.

Chapter 3 explores a couple of application designs dealing with remote monitoring platforms. At first, it introduces the design of a health monitoring system for patients affected by Chronic Heart Failure (CHF). This part discusses the main outcomes of the Health@Home (H@H) project which has been supported by the Ambient Assisted Living (AAL) programme. This project was successfully concluded with a demonstration over real CHF patients in three European hospitals. The chapter closes with a detailed state of the art for energy monitoring in Smart Grids and proposes an energy Home Area Network (HAN) based on Commercial Off the Shelf (COTS).

2 INTERCONNECTION ARCHITECTURES FOR HETEROGENEOUS MPSOC: NETWORK ON-CHIP

A key element in Multi-Processor Systems-on-Chip (MPSoCs) is the global on-chip communication infrastructure, because its throughput, latency and power consumption set the limit to the overall performance.
The traditional shared bus approach exhibits its limits as the number of integrated IPs increases. Indeed, while gate delay scales with each new technology nodes, global wire delay increases and can be kept constant only by inserting repeaters [4]. For this reason shared bus communications standards are being substituted by multi-layer interconnects, now commonly referred to as Network-on-Chip (NoC). NoC is a scalable packet-switched communication infrastructures, connecting hundreds of IP cells, in MPSoC [5]-[28].
NoCs provide a methodology for designing an interconnect architecture independently from the connected cores that can be general purpose processors, Application Specific Instruction-set Processors (ASIP), Digital Signal Processors (DSP), memories or peripherals. Design flow parallelization, scalability and reusability all benefit from this approach [10]-[14], [29]-[49]. NoCs will be a key component also for the success of future 3D SoC [18], [19].
The NoC paradigm leverages the networking and parallel computing domain experience into the SoC world, implementing packet-switched micro-networks with an ISO/OSI like protocol stack. Examples of NoC architectures include Spidergon [13], [30], Mango [31], Aethereal [16], Arteris [32], Sonics [33], SoCbus [34] and xPipes [35], [36], [63].

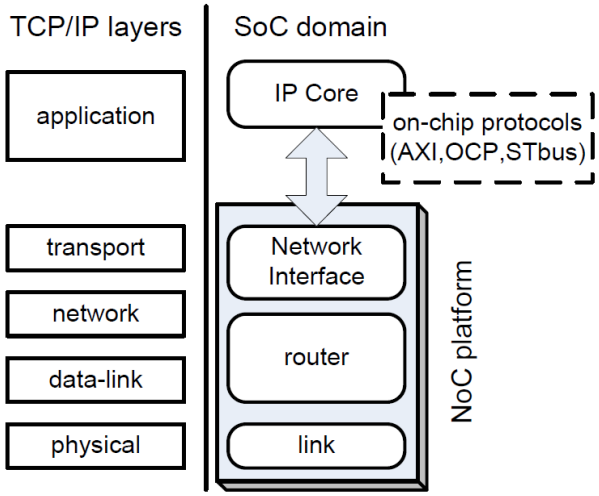

Figure 2.1. Typical ISO-OSI layers for Internet applications and their mapping onto NoC components.

Figure 2.1 illustrates the building blocks of a NoC and the corresponding layers in the ISO/OSI protocol stack. The Network Interface (NI) connects the cores to the NoC domain. The NI is made up of two separate components: shell and kernel. The shell encapsulates the transport layer and transforms local core transactions into NoC packets. The kernel implements the network layer and provides features such as data bus size and frequency conversion between the core and the NoC domain. By splitting transport and network layers into separate sub-components, plug & play design style is extremely simplified. The network is composed of a number of routers that pass packets between nodes. The router implements network and data-link layers. The physical link is responsible for actual signal propagation among routers and/or network interfaces.
For example, Figure 2.2 shows a NoC, based on Spidergon topology (a Ring one with an additional across link for each node to reduce network crossing latency), highlighting its hardware building blocks: connected IP cores, NIs, links and Routers (R) [20], [21], [25], [50].
Next sections are organized as follows. The Router architecture and features are described in Section 2.1. The core NI architecture is detailed in Section 2.2 while Section 2.3 illustrates the advanced NI networking features implemented. Section 2.4 introduces the *Metacoding* methodology and the automated design flow implemented to overcome configuration issues. Finally Section 2.5 collects some CMOS implementation results for different Router and NI configurations and a comparison with the state of the art.

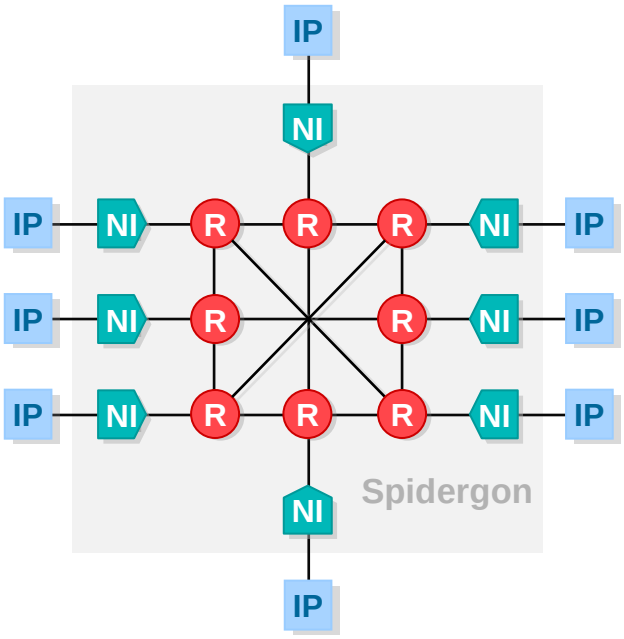

Figure 2.2. Spidergon NoC platform.

2.1 Router Features and Architecture

The router proposed in this work features wormhole packet-switched routing with credit-based flow control and an arbitrary number of virtual channels. All router features are configurable at synthesis time, from arbitration policies to clocking schemes, for optimal trade-offs between network performance and physical design closure.

2.1.1 Topology and Routing

The router supports the Spidergon [13],[30] and ring topologies, plus a family of degree three topologies derived from Spidergon. The router is connected through two unidirectional links with three other routers in the network into directions Right, Left and Across, plus the fourth connection to the local Network Interface. When the Hierarchy link is instantiated, the router acts as a gateway between two NoC sub-networks (see Figure 2.3).

The router adopts the wormhole packet-switching technique, where a packet is subdivided into flits having their unique flow control. Once the first flit of a packet is routed, remaining flits follow the same path reserved to the header. This approach drastically reduces the amount of buffering (queues dimensioned with the granularity of a flit instead of a packet) and allows deeply pipelined data paths when compared to virtual cut-through and store & forward [44].

The routing algorithm is deterministic, so that always the same path is chosen between a source and a destination node, even if multiple paths exist. This choice avoids costly flit reordering at packet reception. The idea is to move clockwise or counterclockwise along the ring to reach destination nodes which are near the source node, and to use the Across link as first or last hop to jump to a part of the network that is far away from the source node.

The router uses a simple source-based routing: the entire path is encoded in the packet header, which has a fixed size due to the symmetry of the topology. This enables fast routing decision at each router, because the

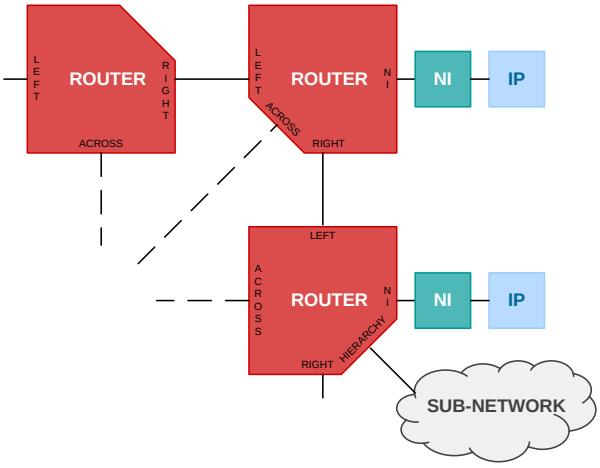

Figure 2.3. Router connections in a NoC.

information is simply extracted from the header, no computation or routing lookup table is required.
The router avoids end to end deadlocks by deploying Virtual Channels. Virtual Channels provide logical links over the same shared physical channels, by establishing a number of independently allocated flit buffers in the corresponding transmitter/receiver nodes. Typically, request and response Virtual Networks (VNs) are implemented on top of two disjoint Virtual Channels for sharing the physical link bandwidth and maximizing wire efficiency. The parametric number of Virtual Channels supported by the router can lead to advanced routing schemes or independent QoS traffic classes for real time and low latency flows.

2.1.2 Architecture

The router general architecture consists of the following main blocks, instantiated as many times as the number of links (see Figure 2.4):

- Downstream Interface (DS ITF), which is connected to the input link and buffers incoming flits (separate Input Buffer for each Virtual Channel)

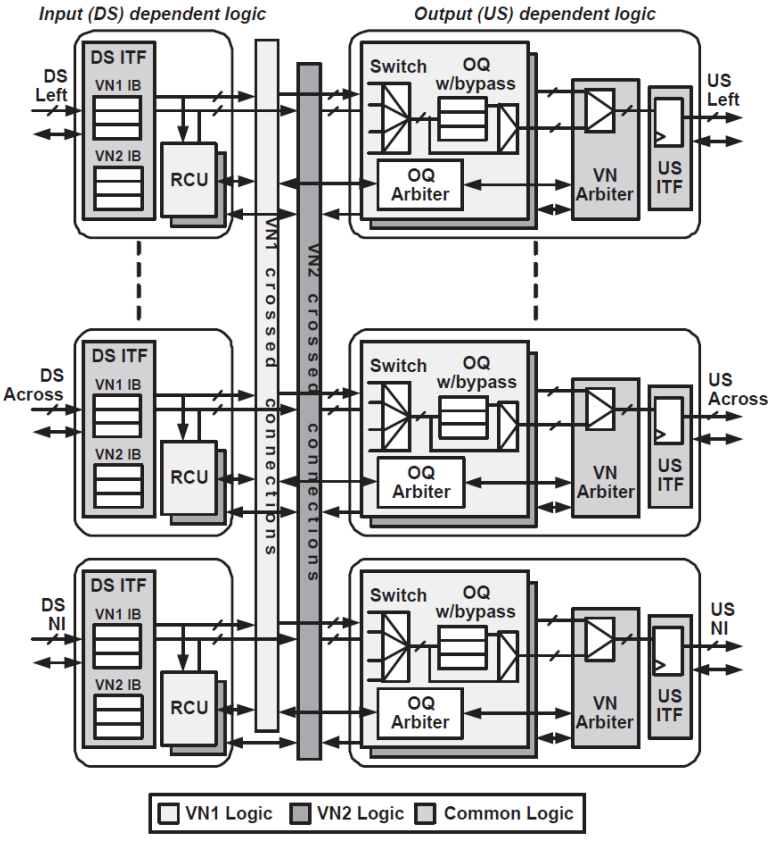

Figure 2.4. Router architecture with two virtual channels.

- Routing Computation Unit (RCU), which extracts routing information from packets' headers
- Switch, which routes inputs to outputs
- Optional Output Queue (OQ) with bypass capability, which stores outgoing flits
- OQ Arbiter, which arbitrates among different inputs requesting the same Output Queue. It is in charge of the per-packet allocation of the OQ, and so of the Virtual Channel
- VN Arbiter, which arbitrates between the two Virtual Networks requesting link access. This is a per-flit arbitration, because the two packets of the two VNs can be interleaved on the link, depending on credits availability
- Upstream Interface (US ITF), which is in charge of managing the credit-based output data flow

Downstream and Upstream Interfaces are dedicated components that implement the physical layer. The different links – synchronous or mesochronous [10] – correspond to different DS/US Interfaces instantiated in the router and in other NoC modules.
The above mentioned blocks can be grouped into DS-dependent blocks (DS Interface and RCU) and US-dependent blocks (Switch, Output Queue, OQ Arbiter, VN Arbiter and US Interface), as shown in Figure 2.4. DS-dependent blocks are instantiated for a link only if the corresponding DS interface exists; similarly, US-dependent blocks are instantiated for a link only if the corresponding US interface exists. The DS-dependent block of each link is connected to the US-dependent blocks of all the other links (represented in Figure 2.4 as crossed connections).

2.1.3 Buffering and Latency

The router adopts the credit-based flow control, which works on a flit per flit basis. Therefore, to guarantee the maximum link throughput, the Input Buffer is sized according to the credit round-trip delay. This delay can be defined as the minimum time between two consecutive credits for the same buffer location and it depends on the credit pipeline between two consecutive routers. The value of this delay is configurable, thus different design trade-offs in terms of working frequency and buffer resources are possible.
The router uses Output Queues for enhanced performance, avoiding head-of-line blocking. Queues are shared among input flows to limit costly time/space speed up factors and they have a bypass feature to reduce the router crossing latency in case of low traffic conditions. Output Queues are optional and are usually not instantiated in low cost (small footprint) implementations.
It is optionally possible to implement a separate Output Queue for each input flow targeting a given output. This configuration improves the offloading capability of a router and it is particularly useful when multiple incoming heavy traffic streams target the same output port (typically the NI port).
The router has a configurable crossing latency, from one up to two clock cycles. This is obtained by means of a flexible pipeline in the data path: one stage is represented by the Input Buffer and the other stage is the optional output retiming register. The Output Queue, when instantiated, can be bypassed in case of low traffic conditions, thus not affecting the overall router

crossing latency. Registering input and output port signals allows orthogonality between the link and the router delay.

2.1.4 Quality of Service

QoS is based on the Fair Bandwidth Allocation mechanism. It allows for a flexible, scalable and low cost management of the allocation of the available bandwidth. The requested bandwidth value is programmed at the injection point (Network Interface) and is not explicitly linked to the path of a data flow through the router as it is done in other NoC architectures [31], [16]. The arbitration logic in the router is thus relatively simple, since it only requires a two-steps arbitration process based on the information available in the network header. When all data flows have the same bandwidth reservation, the arbitration degenerates into a Round Robin, Least Recently Used or fixed priority scheme. Speculation-like techniques allow the two arbitrations to be performed in parallel, thus reducing the critical path delay.

2.2 Core Network Interface Design

A key element of a NoC is the NI which allows IP macrocells to be connected to the on-chip communication backbone in a Plug-and-Play fashion. The NIs are the peripheral building blocks of the NoC, decoupling computation from communication. Basically, the NI is in charge of traffic packetization/depacketization to/from the NoC: it provides protocol abstraction by encoding in the packet's header all data to guarantee successful end to end data delivery between IP cores (transport layer) and all Quality of Service (QoS) information needed by the router at network layer. A NoC packet includes a header and a data payload which are physically split in units called flits. All flits of a packet are routed through the same path across the network. The header field is composed of both a Network Layer Header (NLH), whose content is determined by the NI according to the nodemap network configuration, and a Transport Layer Header (TLH) containing information used by the NIs for end-to-end transaction management.
Some NI designs proposed in literature also implement the conversion of data size, frequency and protocol between the original IP bus and the NoC. The IP bus can be a standardized one such as AMBA AXI (Advanced eXtensible Interface) [51] or OCP (Open Core Protocol) [52] or a custom bus, such STBus [53].
The latest research frontier on NI architecture design aims at integrating more features to directly support in hardware advanced networking functionalities. The challenge in doing so is keeping NI area, power and latency overheads as low as possible with respect to the connected IP cores.
In recent literature some NIs have been presented that add to the basic IP-NoC interface functionalities some features such as handling of out-of-order transactions in [51], [54], [55], detection of error transactions in [24], [56], secure memory access control in [57], QoS management and NI programmability in [58].
However, the literature does not present a design integrating all the abovementioned advanced features in the same NI with limited complexity overhead. With respect to the above NI functionalities other services may be

useful to support in hardware such as end-to-end interoperability between different IP bus types (e.g. end to end connection between an AXI IP core such as an ARM processor and a custom bus IP cell such as an ASIP or DSP coprocessor), management of pending transactions when powering down/up the IP to increase energy efficiency, remapping for master IP cores of their addressable NoC space. Particularly the interoperability feature is important since MPSoCs are often realized as the interconnection of heterogeneous IP cores provided by different vendors.
To overcome the limits of the state of the art this work presents the design and characterization in deep submicron CMOS technology of a NI architecture directly implementing in hardware advanced networking features such as: store&forward transmission, error management, power management, ordering handling, security, QoS management, programmability, interoperability, remapping. Such NI has been conceived as a scalable architecture: the advanced features can be added on top of a basic NI core implementing data packetization and conversion of protocols, frequency and data size between the connected IP core and the NoC. The NI can be configured to reach the desired trade-off between supported services and circuit complexity.

2.2.1 NI Top-level Architecture

IP cores in a NoC infrastructure are commonly classified into Master and Slave IPs: the former (e.g. a processing element) generates request transactions and receives responses, the latter (e.g. a memory) receives and elaborates the requests and then sends back proper responses. Initiator NIs are connected to Master IPs, and convert IP request transactions into NoC traffic, and translate the packets received from NoC into IP response transactions. Dually, Target NIs also exist, associated to Slave IPs. Target NIs present a mirrored architecture: requests are decoded from NoC; responses are encoded. In both NI types, two main domains can be identified (see Figure 2.5a referring to the top view of an Initiator NI and Figure 2.5b referring to the top view of a Target NI): the Shell, IP specific, and the Kernel, NoC specific, each one having its own peculiar functionality and interface.
Figure 2.5 highlight also some advanced networking features such as programming, security, error and power management detailed in Section 0. The aim of the Shell/Kernel separation is to abstract IP specific properties (such as bus protocol and data size) from NoC side properties. This way the NoC becomes an IP-protocol agnostic interconnect, that is whichever protocol, bus size, clock frequency the Master or Slave IP is using, all modules in the system may communicate with each other.

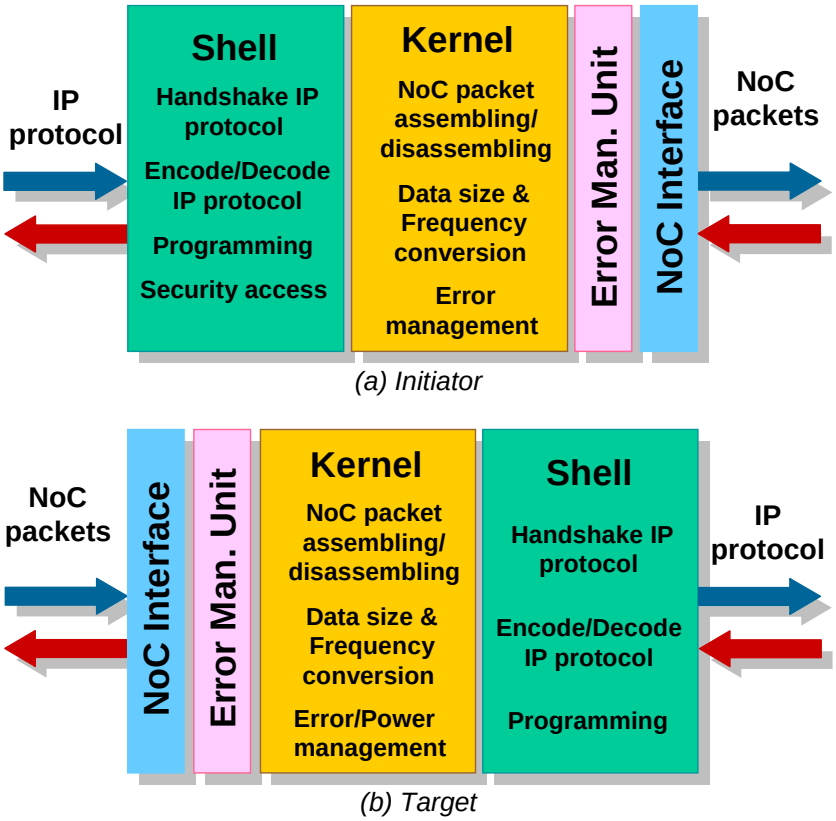

Figure 2.5. Top view of the NI design: (a) Initiator and (b) Target.

Conversion features must be implemented in the two directions, called *request path* (from Master to Slave IPs, blue paths in Figure 2.5) and *response path* (from Slave to Master IPs, red paths in Figure 2.5) respectively. While the Kernel, and the associated NoC interface, is IP protocol independent and its design is common to all possible NIs, the Shell needs to be defined on a per-protocol basis. A specific Shell architectural design is needed for each IP protocol that must be connected to the NoC. The proposed NI supports the following IP bus protocols: AMBA AXI, a de-facto standard in embedded systems; STBus TYPE 3 [53] used by STMicroelectronics as the backbone of its SoC designs; DNP, a distributed network processor interface developed within the SHAPES European Project [5], [6], [10], including STMicroelectronics and ATMEL as main industrial partners, to build a multi-tile MPSoC architecture. In case the programming interface feature is activated, the STBus TYPE 1 or the AMBA APB bus are used as programming bus.

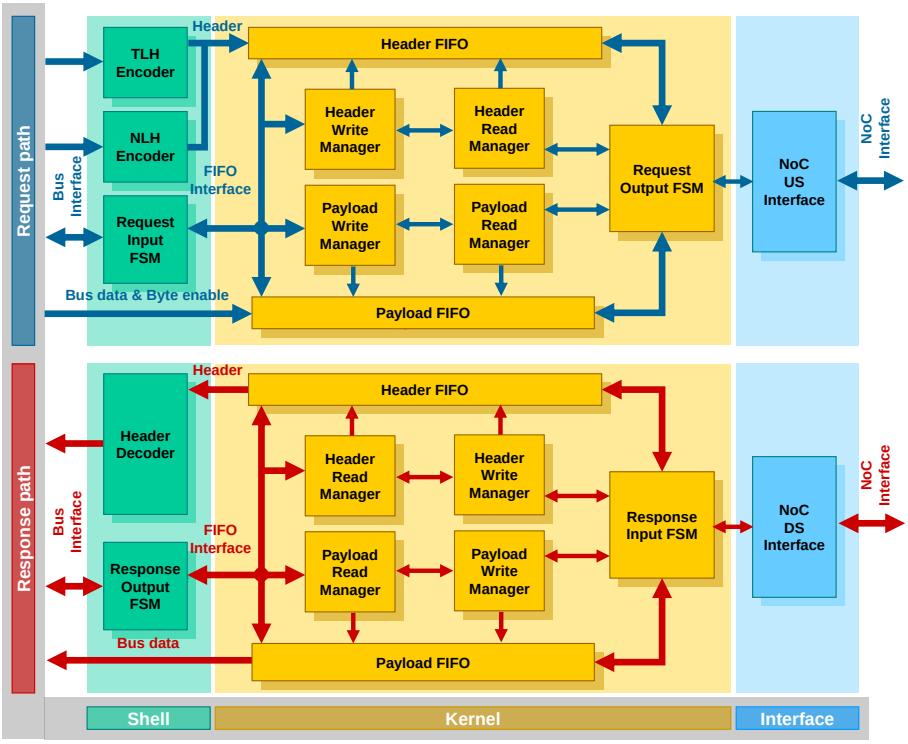

Figure 2.6. Main blocks in the NI micro-architecture.

Figure 2.6 provides a more detailed insight of the core NI Initiator architecture with a clear distinction between *request* and *response* paths. Advanced functionalities, whose implementation is described in Section 2.3, will be added on top of this core NI architecture. From left to right, Figure 2.6 highlights the Shell, the Kernel and the NoC interface respectively. Moreover, the top of the figure refers to the request path, while the bottom part refers to the *response* path. The NoC interface presents an *Upstream* (US) section, to send packets to the NoC, and a *Downstream* (DS) section, receiving packets from the NoC.

The NI Shell part deals directly with the bus protocol, implementing bus specific handshaking rules by means of dedicated Finite State Machines (FSMs). Before passing data on to the Kernel, the Shell also builds the Network and Transport Layer headers, needed by subsequent NoC components (i.e., routers and target NIs) for forwarding the packet and decoding it at destination. The NI Kernel part manages buffering and other services (described in Section 2.2.2) in an IP-protocol independent way.

2.2.2 Kernel-Shell Interface by Bi-synchronous FIFOs

The Kernel is interfaced to the Shell by means of a FIFO-like interface. As reported in Figure 2.6, encoded data coming from the Shell are stored in two FIFOs, an header FIFO (holding transport layer and network layer headers)

and a payload FIFO (holding bus raw data). Each FIFO has its own read and write managers which update FIFO pointers and status, and provide frequency conversion mechanisms. The Kernel is connected to the NoC interface stage through two additional FSMs. In the request path, an output FSM (OFSM) reads headers and payloads and converts them into packets according to the NoC protocol. In the response path, an input FSM (IFSM) collects packets and splits header and payload flits into their respective FIFOs. To be noted that the NI encodes both the TLH and the NLH, while in the decoding action only the TLH is taken into account because the packet has reached its destination and routing data are not needed.
By using bi-synchronous FIFOs in the NI scheme of Figure 2.6 frequency conversion is accomplished between NoC and each connected IP. Each read (write) FIFO manager re-synchronizes in its own clock domain the pointer of the write (read) manager in the other clock domain. Hence, the empty/full status of the FIFO is known by comparison of synchronized pointers, and the header or payload FIFOs can be correctly managed. To increase the robustness of the synchronization the pointers adopt a Gray encoding. Normally, this would limit the possible FIFO sizes to powers of 2, but thanks to a user-transparent pointers initialization any even number of locations can be supported. To be noted that vs. other works in literature [59] that synchronize different clocks but with the limit of an integer ratio between the frequencies, our bi-synchronous FIFO can handle arbitrary ratio clocks.
Figure 2.7 highlights the different four clock domains that can be supported in the proposed NI architecture.
As shown in Figure 2.8, since read (RD) and write (WR) managers can access a FIFO by different basic storage units, also data size conversion between IP and NoC domain is possible. The conversion is managed by exploiting FIFO rows and columns concepts. A FIFO location, or column, is sized according to the larger data size between *data in* and *data out*; a FIFO row is sized according to the smaller data size between *data in* and *data out*. Up-size conversion is accomplished by writing by rows and reading by

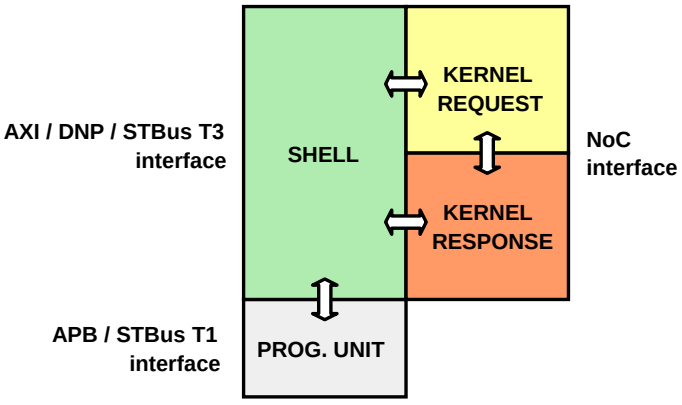

Figure 2.7. Clock domains in the proposed NI architecture.

columns; down-size conversion is exactly the opposite. For example, consider large opcode store operations (i.e., with large amount of payload data) generated by an IP with data bus size of 32 bits and connected to a NoC with flit size of 128 bits (up-size conversion): four 32-bit data write accesses by the IP are necessary to fill a 128-bit payload FIFO location and make it available to the NoC to read it (see Figure 2.8, focusing on a single 6-location FIFO).
In particular conditions when no size conversion nor Shell/Kernel frequency conversion is needed, nor Store & Forward support is required, it is possible to remove the bisynchronous FIFOs, by setting their size to zero, thus saving area and power consumption. This feature is known as Zero-FIFO Kernel.
As far as the NI crossing latency is concerned, its minimum value depends on the pipeline stages used. Typically, at least one retiming is performed due to the presence of the FIFO in the Kernel (unless the Zero-FIFO Kernel feature is enabled). To increase the maximum operating clock frequency, optional pipeline stages can be added at the IP and NoC interfaces. Thus, a maximum of three retiming stages can be implemented. Obviously, the minimum NI crossing latency is equal to the number of retiming stages instantiated, but its actual value may be increased by other factors: for example, in case of frequency conversion the synchronization delay has to be added; or if the current IP traffic has a low priority the NoC QoS support may slow down its access to the interconnect, or the store&forward mechanism (see Section 2.3.1) might be enabled, thus increasing the traffic latency.

2.2.3 NI Physical Link and NoC Interface

As far as the physical link is concerned, at the NoC interface side there are the following hardwired lines for each response or request path, see Figure 2.9:

- N-bit *flits* used to transfer NoC packets (headers and payloads), with N configurable at synthesis time;
- 4-bit *flit_id*, whose 2 LSBs identify first, intermediate and last flits of a packet; the flit can also be a single one. Other optional bits of this signal are used to mark the end of bus packets within compound

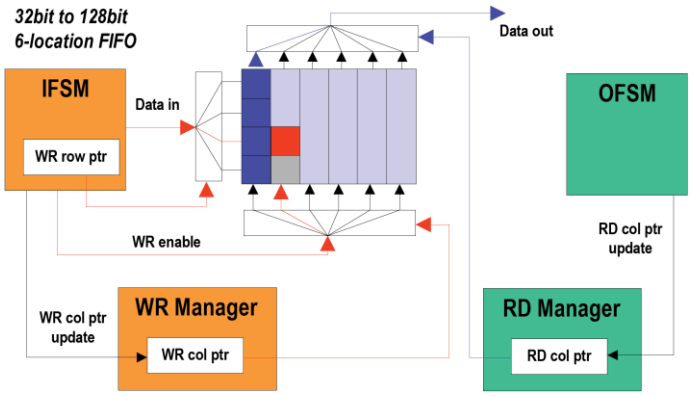

Figure 2.8. Upsize conversion, focus on a single FIFO.

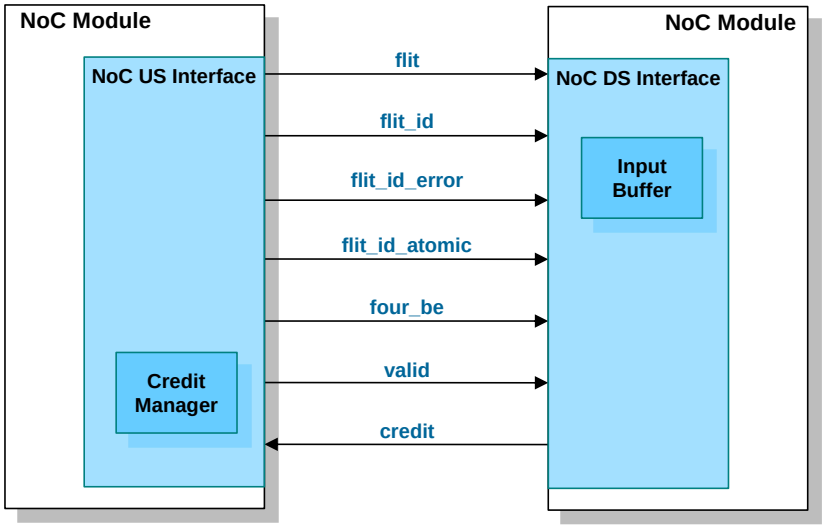

Figure 2.9. NI interface on the NoC side.

transactions (i.e. composed of a number of packets) that are translated into a single NoC packet, and to identify payload flits that cannot be aggregated in case of upsize conversion (necessary in some cases of interoperability);

- the optional K-bit *four_be* (K=N/32) to mark meaningful 32-bit pieces of data within a flit and used in end-to-end size conversion.
- the optional 2-bit *flit_id_error* for signalling a slave side error or an interconnect error;
- the optional *flit_id_atomic* which enables support of atomic operations: a NI can lock paths towards a Slave IP so that a Master IP can perform a generic number of consecutive operations without any interference from other masters;
- *credit* and *valid* signals for credit-based flow control. A flit is sent only when there is room enough to receive it: neither retransmission nor flit dropping are allowed. This is done automatically by setting an initial number of credits in the US interface (in its Credit Manager), equal to the size of the Input Buffer in the DS interface it communicates with. Since the US interface sends flits only if the connected DS interface can accept them, there are no pending flits on the link wires. This approach allows virtual channel flit-level interleaving, so that separate virtual networks can share the same physical link. See Section 2.3.7 for more details on Virtual Networks.

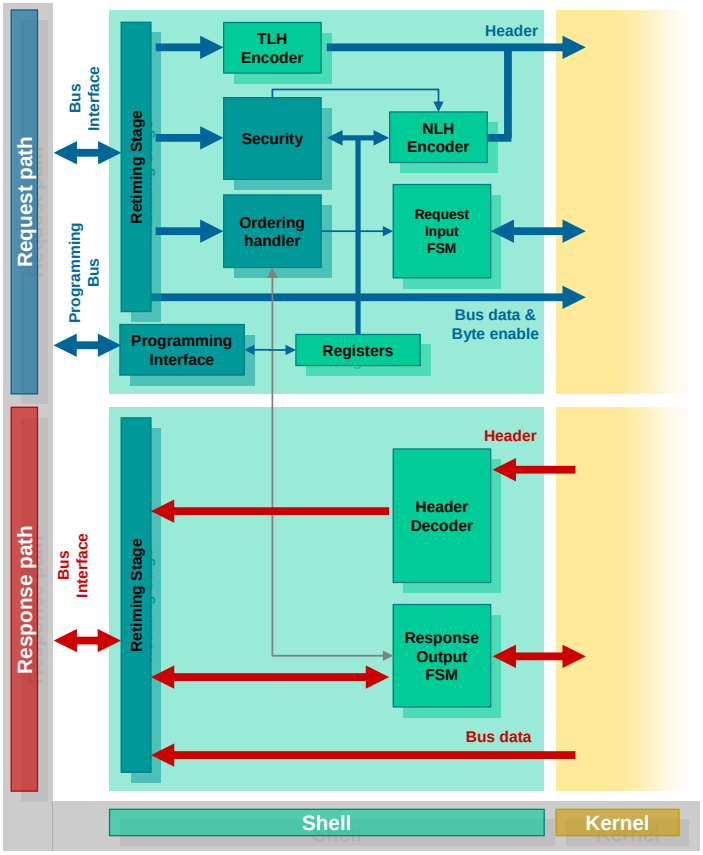

Figure 2.10. Advanced features in a NI Initiator.

2.3 Advanced Network Interface Features

The next sub-sections describe the configurable services available as special features in the novel NI design, and patent filed in US and Europe. Figure 2.10 highlights some of the additional configurable features on a NI initiator.
The proposed design is the first in literature directly implementing in hardware all advanced networking features: store & forward transmissions, error management, power management, security, ordering handling, QoS management, programmability, interoperability, remapping.

2.3.1 Store & Forward (S&F)

Kernel FIFOs in both Request and Response paths contain flits, either received from the NoC and to be decoded towards the IP bus, or encoded from a bus transaction and to be transmitted over the interconnect. Default NI behaviour is that a flit is extracted from the FIFO as soon as it is available. Hence, if the original traffic at an interface (NoC or bus) has an irregular nature, such a shape is reflected also into the other interface (bus or NoC).

When Store & Forward is enabled, flits are kept into the internal kernel FIFOs until the whole packet is encoded/received and then they are transmitted/decoded all together. This way, an irregular traffic is changed to a bubble-free traffic thus improving overall system performance. For example, the interconnect can benefit from the S&F mechanism, since the link is engaged only when the entire transaction is available for transmission.
S&F may be enabled on the request path or response path independently. At NoC-to-bus level, it is possible to enable the per-packet S&F, while different S&F options can be selected at bus-to-NoC level: storing a whole bus packet, or storing an entire compounded transaction, that is a collection of sequential packets tied together by setting the appropriate bus fields (this second option depends on the bus type). The mechanism for the per-packet S&F implementation is quite simple. After completion of a packet, the FSM controlling the FIFOs reading is in a state where only the header FIFO is checked, to extract the beginning of a new packet. The idea to handle per-packet S&F, in both directions, is to keep the packet header (i.e. the flit in the header FIFO) hidden to the reading logic by simply not updating the header FIFO write pointer. When the entire packet is stored in the FIFOs (both header and payload), the header is unmasked and made visible by updating the header FIFO pointer, and the reading logic detects the presence of a new packet.
The management of the bus-to-NoC S&F per compound transactions is a bit more complex, since a compound transaction is composed of a number of packets, that is a number of headers. The header write pointer must be updated upon arrival of any new packet in the compound transaction, to avoid overwriting the previous one, therefore the headers become visible to the reading logic. The trick here is to exploit a field in the header flit to mark the packets' headers as "hidden" (the first packets of the compound transaction) or "visible" (only the last packet of the compound transaction): the reading logic evaluates this field in each header and starts extracting the FIFOs content only when the last packet of the compound transaction is detected in the FIFO. Obviously, if FIFOs go full, then the flits are extracted even though the packet/transaction is not entirely stored.

2.3.2 Error Management Unit (EMU)

The EMU is an optional stage that can be instantiated between the Kernel and the interface to the NoC. The EMU behaviour is different in Initiator and Target NIs. In an Initiator NI, EMU can handle bad address errors or security violations (this second type of errors only if the Security support is enabled; see Section 2.3.4). When the address of the Master IP transaction is not in the range of the assigned memory map, or when the transaction is trying to access a protected memory zone without having the rights, the packet is flagged in its header as an error packet. The EMU then filters the packet directed to NoC US interface to avoid it to enter the network, and builds a response packet re-mapping the request header on a new response header, and if needed adding dummy payload. The response packet is sent back to the Master IP in order to be compliant with protocol rules.
The EMU of a Target NI, instead, encodes the *flit_id_error* value to the US NoC interface in case an error response is produced by the Slave IP. When

the Power Manager is enabled (see Section 2.3.3), the EMU is also in charge of properly managing incoming traffic at DS NoC interface during power down mode. All the traffic received in request during power down mode is flushed by the EMU, so that it never reaches the Slave IP. The EMU itself generates an error response to the Master originating the request.
The EMU is composed by 3 blocks as in Figure 2.11 showing EMU and Power Manager blocks in a Target NI:

- *Error Detector*, which flushes all error traffic. In Initiator NIs, the outgoing error traffic is identified by a flag in the header, while in Target NIs all incoming packets are flushed if the connected Slave IP is in power down mode;
- *Error Encoder*, which assembles a new NoC packet to be channelled in the response path;
- *Error Write Manager*, which is basically a traffic-light to avoid simultaneous traffic to the US NoC Interface from Kernel Response and from the EMU in a Target NI, while in an Initiator NI it avoids interference between the DS NoC Interface and the EMU both, trying to access the Kernel Response.

If a request packet does not contain an error, the EMU behaves transparently and does not add any clock cycles of latency.

2.3.3 Power Manager (PM)

This feature is available only for Target NIs connected to Slaves which may be turned off to save power. The PM is always coupled to an EMU block which rejects incoming NoC packets trying to access the Target NI when the connected Slave IP is in power down mode. The mechanism for building error response packets is the one explained in Section 2.3.2. A simple req/ack protocol controls the power up/down state of the NI, by means of a dedicated

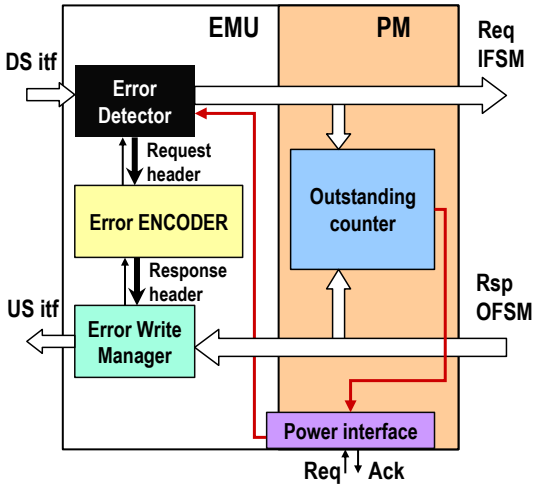

Figure 2.11. EMU and Power Manager in Target NI.

interface: each request (req set to 1) acknowledged by the PM unit (ack set to 1) makes the NI power state switch from UP to DOWN and vice versa. It may happen that a request for power down is sent to the PM while the Slave IP is still elaborating a number of pending transactions.
In this case the Target NI stops accepting packets from the network and waits for all pending transactions to be processed (see the counter of outstanding transactions in Figure 2.11) before acknowledging the request and switching to power down mode. The power manager is a completely new feature introduced by the proposed NI.

2.3.4 Security

The security service, available only in NI Initiators, acts as a hardware firewall mechanism, see Figure 2.12 and Table 2.1. It introduces a set of rules that transactions coming from the Master IP must satisfy to gain access to the network. The security rules involve:

- lists of memory intervals under access control;
- lists of Master IPs that may have access to a certain memory region;
- lists of access types (i.e., Read, Write and Execution permissions,

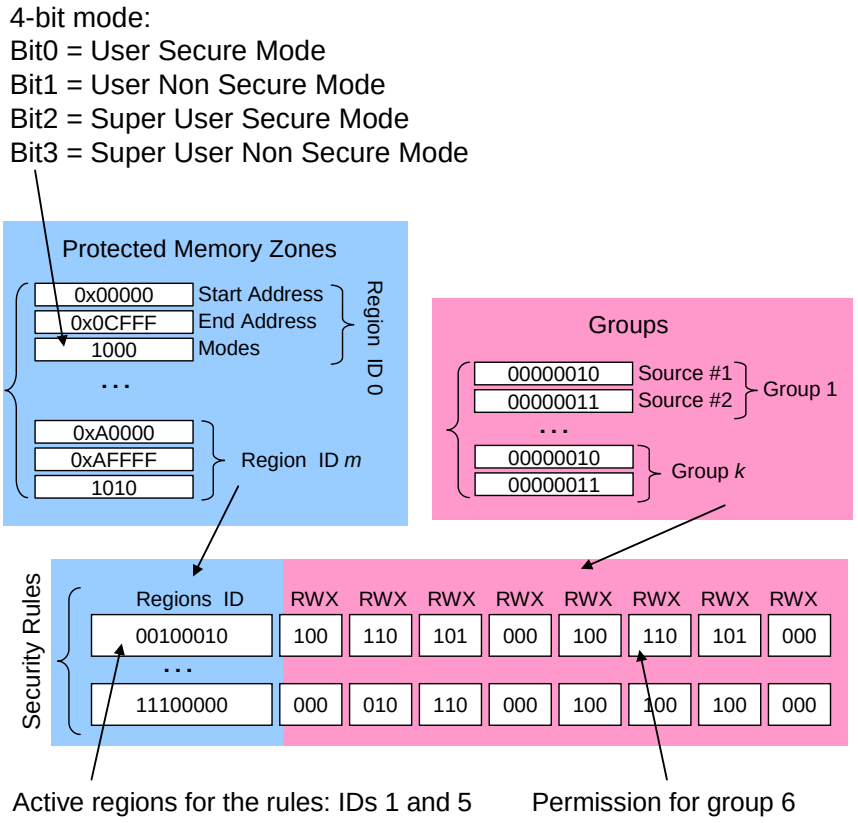

Figure 2.12. Security firewall data structure.

Address in Protected Memory Zone	Address + Opcode Size in Protected Memory Zone	Transaction Source in Groups	Action
yes	yes	yes	rule check
yes	yes	no	access denied
no	yes	yes	access denied
no	yes	no	access denied
yes	no	yes	access denied
yes	no	no	access denied
no	no	yes	access allowed
no	no	no	access allowed

Table 2.1. Security firewall behaviour.

RWX in Figure 2.12) for a certain IP on a certain memory region.
Security rules are applied in the Security block (see Figure 2.10) during packet encoding. If a test fails the security check, the corresponding transaction is marked as an error in the NLH and it is detected by the EMU, which must deactivated as well to properly manage security violations. The illegal packet is then discarded and does not consume network bandwidth, and the error response to the Master IP is directly generated by the EMU itself. The rules that allow a transaction to access the network are described by means of a security map. In this map a number of memory regions are defined, and associated to Region IDs (Protected Memory Zones in Figure 2.12). The same map defines how these regions can be accessed (modes in Figure 2.12). Access to these zones can be allowed only to Master IPs belonging to specified Groups (see Figure 2.12). Finally, a table of read/write/execution permissions is given for each Source and for each protected region (see Security Rules in Figure 2.12).
Depending on the address and the Source, the memory access can be immediately allowed, or immediately denied, or go through a security rule check, as illustrated in Table 2.1. Naturally, there is also a control to block malicious memory accesses to a non protected zone that transfer a number of bytes such that the operation overflows in a protected region.
A region can be any sub range of the address space in the whole system interconnected by the network. The security map may be statically defined at design time or changed dynamically through the programming interface of Section 2.3.8; in the latter case the NI programming interface must be configured to instantiate the registers related to the security category. The implemented security mechanism supports up to 8 protected memory regions, up to 8 access rules with Read/Write/Execution permissions, up to 16 Source groups to classify Masters.

2.3.5 Ordering Handler

Typically, bus protocol rules impose that transactions generated by a single Master IP get their responses with the same order of the requests. In NoC platforms, it may happen that some responses are reordered by the

interconnect because of the existence of alternative paths or paths of different length between the Master and its reachable Slaves. Each transaction generated by a Master is characterized by a destination address and an identification number. The destination address identifies a specific Slave to access. The identification number characterizes the Master itself: this information is used by Target NIs to encode response packets to be routed back. For example, in an AMBA AXI bus, the Master identification or Source field is the ID field. Once a transaction is forwarded to the network nothing can be said about the time the response will be sent back. In general, each Slave has its own list of requests (from several Masters) to respond to, and it may happen that a Slave receiving request ape n from a Master is slower to reply than the Slave receiving request *n+1* from the same Master, due to its longer requests list to handle. As a consequence, there is no guarantee that responses will get back to the corresponding Initiator NIs in the correct order. This is a problem when transactions with the same ID (same Master) but different destinations (different Slaves) hang around the network. When an Initiator NI receives the response transactions it cannot distinguish from which specific Slave it comes because the ID field is the same (and, generally, the address information is not available in the response path).
To avoid the risks of this situation, or the necessity to re-order the responses, the proposed NI may be configured to support the Ordering Handler feature. The Ordering Handler block is placed in the Initiator NI Shell (see Figure 2.10) and is responsible for applying filtering rules that avoid out-of-order transactions. A Master IP can access a generic number of Slave IPs via NoC. The ordering filter just prevents transactions with the same ID and directed to different destinations from accessing the network simultaneously. A transaction with some previously used ID is accepted only if the intended destination is the same of still pending transactions or if there are no pending transactions with that ID. Request transactions with the same ID going to the same Slave can be forwarded since, in this case, it is a slave's responsibility to send back responses in the correct order. This is a common bus protocol rule to manage multi-slave system scenarios.
The Ordering Handler filtering mechanism exploits a buffer to store the history of the Master pending transactions. A single buffer entry is represented in Figure 2.13. The entry is allocated upon reception of a request transaction with a new ID/destination pair. Any new transaction with the same pair increments the outstanding transactions counter in the associated entry (ISSCAP field in Figure 2.13, whose size is configurable). The counter is decremented upon delivery to the Master of a corresponding response packet (characterized by the same ID). When the counter is zero, there are no more pending transactions, and the entry is again available for other ID/destination pairs. The filtering procedure is such that any request transaction with the same ID of a valid buffer entry but different destination is stalled on the bus. Obviously, also any transaction with a new ID/destination pair is stalled if the buffer does not have empty entries to allocate: for this reason the buffer size is configurable, to adapt to different applications requirements. A simplified filtering scheme is also possible, where a Master can only access a single Slave at any particular time: in this case a single buffer entry is enough, to filter incoming transactions on the basis of their destination only.

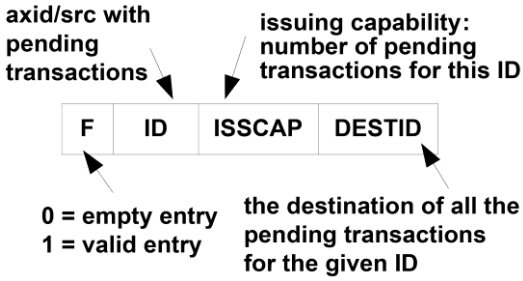

Figure 2.13. Ordering Handler: pending transactions buffer entry.

2.3.6 Remap

This feature, available only for Initiator NIs, allows the definition of more than one master-to-slave address configuration map or nodemap. Consider, for example, a system where a Master needs to communicate with different sets of Slave IP cores or memory regions according to some conditions; or, for example, a platform where masters need to configure particular devices during bootstrap initialization. The remap feature aids in associating to a single Initiator NI multiple sets of addressable NoC regions. The user may select the different maps, at run-time, by means of an external signal. The proposed Initiator NI implementation supports up to 14 different sets of addressable slaves.

2.3.7 QoS Scheme

Since different traffic classes can interoperate on the same interconnect, a QoS is necessary to avoid them to interfere. In addition, in an interconnect there is a need for an easy real time reconfigurability of bandwidth allocation. The proposed solution is based on Virtual Channels for separating traffic classes plus the FBA (Fair Bandwidth Allocation) scheme for SW-controlled real-time bandwidth allocation.
Virtual Channels are a widely known concept. They create a single physical network and efficiently share it through virtualization, thus creating virtual networks. This ensures that there is no interference among different traffic classes, and the Virtual Channel flow control allows long packets to be overtaken by high priority packets. At the same time, this approach reduces interconnect wires, since the different VNs share the same physical link.
The *FBA QoS* scheme allows the support of different bandwidth allocations for different targets. Moreover, it is software programmable and independent from the interconnect topology.

The basic principle of FBA arbitration is to share the Slave available bandwidth among the Masters during peak request period. Since the arbitration is distributed among the different Router crossed by the packets, traditional weighted Round Robin arbitration algorithms cannot be used: the solution is to apply a faction tag at packet injection (i.e. in the NI), and to keep together, in the same faction round, packets with the same tag in the interconnect (i.e. in the Routers, where the distributed arbitration is performed). This scheme should not be confused with the TDMA approach: the faction round duration is variable; if only packets belonging to a new tag are received, they win the arbitration, so that there is not wasted bandwidth. For example, while in Faction round *i* in Figure 2.14 all IPs are using their reserved bandwidth, in Faction round *i*+1 IP1 is not producing traffic and the total bandwidth is redistributed among the other IPs (see the two pie charts referring to the Faction rounds).
The FBA QoS scheme is summarized hereafter. The NI tags the packets with a faction identifier and if needed with their priority; each injected flow specifies the requested bandwidth. The requested bandwidth is the global amount of data (computed in bytes from the opcode size) transferred by the considered NI flow in a given faction round at the specific target (see IP Faction Thresholds in Figure 2.14). The round at the specific target is a given number of available accesses; the number of bytes read or written in that round represents the percentage of available bandwidth (bandwidth in the round) demanded by this initiator flow.
The requested bandwidth corresponds to a threshold that must be reached by a counter to switch the faction identifier bit. The counter (inside an Initiator NI shell) computes (from the opcode) the number of bytes that flow to each target and enables the faction bit switching when the threshold is reached.
Two different schemes can be enabled: one offering to configure a separate threshold (or bandwidth) for each target, and a simplified scheme with a unique threshold for all targets.

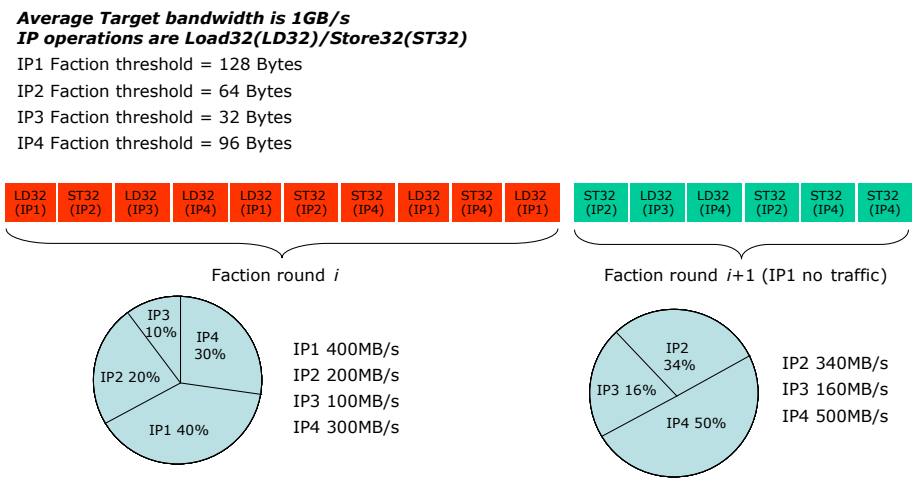

Figure 2.14. FBA QoS scheme - Example.

The proposed FBA scheme offers a number of benefits:

- need to program only the NIs with the requested bandwidth value (number of transferred bytes), by means of a tagging mechanism based on a simple counter;
- the QoS is not explicitly linked to the path in the network, but only to the injection point (the NI). Therefore, for instance, routing can change without any effort to recompute the path followed by the flow and the consequent QoS parameters along this new path;
- routers do not need to be programmed. Their behaviour is definite, and they implement simple arbiters, without any need of slow complex logic;
- the same scheme is used for any VNs, and if all packets injected in the network have the same faction the FBA degenerates into the basic Round Robin, Least Recently Used or packet priority schemes.

2.3.8 Programming Interface

A 32-bit programming slave interface may be enabled on the NI to dynamically configure from the external world the information used to encode packets. This interface exploits a simplified bus protocol that can be chosen between STBus TYPE 1 and AMBA APB.

The programming interface gives access to NI internal registers, which can be grouped into three different categories:

- *routing*: these registers contain information about routing fields to be encoded in the NLH and used by routers to deliver the packet to its destination;
- *QoS*: they control packets priority and bandwidth, encoded in the NLH according to the Fair Bandwidth Allocation QoS scheme;
- *security*: if the security manager is enabled, these registers may be instantiated to control access rules.

Registers, 32 bits wide, are available in both Initiator and Target NIs with the difference that Target NIs does not have security (and so neither the associated registers). If a register category is not enabled then that kind of information will be hardcoded in the NI hardware using statically defined values. Since the configuration registers affect the behaviour of the NLH encoder, the programming unit is integrated into the Shell (in the Request path for Initiator NIs, in the Response path for Target NIs). The registers addressing space is defined *a priori* by the architecture: it starts from 0 and all the enabled registers are consecutive in a pre-defined order.

2.3.9 Interoperability and End-to-end Size Conversion

A novel important feature provided by the proposed NI vs. the state of the art is the interoperability across the NoC between IPs using different bus sizes and even different kinds of bus, without the need to add specific bus-to-bus bridges: the NIs are capable to handle the protocol, size and frequency conversion not only at IP-to-NoC level (and vice versa), but also at end-to-end level, obviously only for the supported IP protocols. This way, from an end-to-end point of view, the NIs perform the protocol, frequency and data size conversion between Master and Slave IPs. The NoC traffic packetization/depacketization, that is the real conversion performed by the

NIs, is transparent to the connected IPs: each NI collects IP traffic from the core it is connected to and then converts such traffic into NoC packets sending them to the network of routers; upon arrival at the destination NI, the NoC packets are translated into the correct IP transactions according to protocol, frequency and bus size of the destination IP.
This is achieved by enabling different levels of end-to-end size conversion supports, together with interoperability support if the Master and Slaves do not use the same bus protocol.
With proper limitations on the managed transaction opcodes the NIs can handle end-to-end size conversion without any additional logic. With no restriction on opcodes but the guarantee of addresses aligned to the Slave data bus size it is possible to enable a simplified end-to-end size conversion hardware based on a Byte Lane Matrix for reshuffling correctly the 32-bit pieces of payload in the transfer, depending on address and opcode (see Figure 2.15a). In other cases when the limitations cannot be applied, a specific support may be required for address realignment coupled to payload cells reshuffling through specific Byte Lane Matrix and Keep/Pass logic (Figure 2.15b): while the Byte Lane Matrix changes the Byte Lane position within the same transfer (vertical reshuffling), the Keep/Pass logic changes the Byte Lane position between two transfers (horizontal reshuffling), which might be needed for some wrap operations. When the network connects Master and Slaves using different protocol, interoperability support may be necessary in some cases. Here, again, it is possible to have an incremental level of interoperability, depending on the transaction types to be handled and on address alignment.

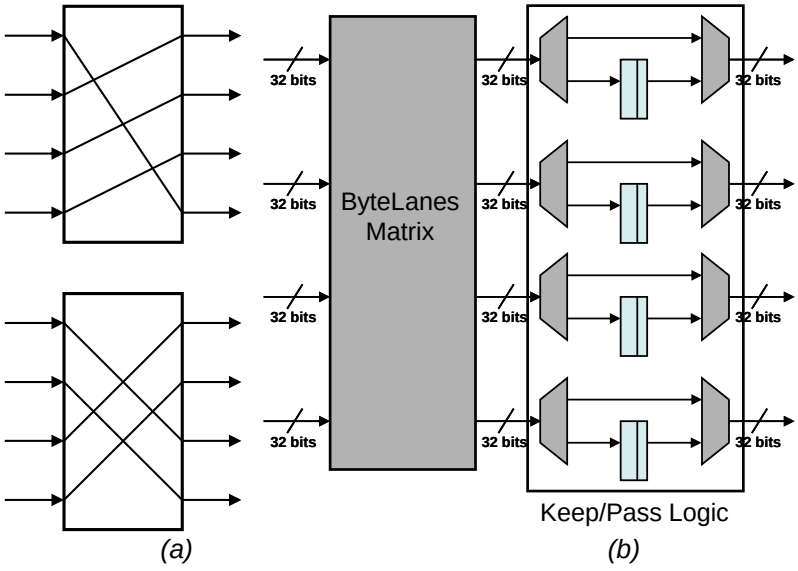

Figure 2.15. (a) Examples of reshuffling in the Byte Lanes Matrix and (b) Byte Lane Matrix coupled to the Keep/Pass logic.

2.4 Metacoding design methodology and automatic design flow

One of the major challenges when designing a communication platform is to minimize the design effort while attempting to cover the widest application space in terms of traffic requirements (high and/or guaranteed bandwidth, low latency, etc.) and implementation requirements (area, clocking scheme, power consumption, etc.). A number of algorithms [37],[38] support high-level decisions like network topology, routing schemes, and partitioning of clock domains. However, the actual implementation would not be feasible without NoC building blocks (router, NI, link) that provide the configurability necessary to match these high-level requirements. In order to meet these requirements a new methodology, named *metacoding*, has been adopted. This methodology supports the design of such configurable components by providing a proper abstraction of the coding process. Metacoding overcomes the limit of current HDLs in capturing configuration intents and reflecting them into an optimized and easy to use RTL codebase, which satisfies the following requisites:

- **Coded consistently**: unnecessary code is never generated, the internal components are the smallest required to provide a given functionality, unconnected signals and ports are removed, unused control signals are always driven with proper values;
- **Neutral to tools:** the codebase is read as-is by *any* standard front-end toolchain (LINT checking, functional simulation, RTL synthesis);
- **Verification friendly**: high coverage scores (both code and functional) are achieved with a reasonable number of configurations.

Differently from other works in literature that present top-to-bottom system level design flows [39], [40], hereafter is only considered the abstraction of the RTL coding process. The input of the flow is not a high-level specification, but RTL code templates and a set of properly defined rules to assemble them.
The rationale of this choice is to streamline the generation of those design aspects that are repeatable, structured and prone to errors, while retaining the full advantages of manually coding RTL blocks, like the ability of achieving timing closure with extremely tight constraints.
Similar ideas have been applied in the software design domain; one of the most remarkable is the FFTW library [41] and its inner code generator *genfft* [42]. Using the terminology of the new IP-XACT standard [43], i.e. the XML format used to package reusable IP cores.
Without loss of generality, the methodology will be discussed referring to the particular case of the NoC building blocks.

2.4.1 Router Configuration Space

The switching matrix of the router implemented is fully configurable: the existence of *each* stream direction (up and down) in *any* virtual network can be independently specified for *each* link (L, R, A, NI, H). This subset of the configuration space is referred to as *router topology* or *backbone*. The ability to fully control the router topology enables to strip away unnecessary paths in the network, thus saving hardware resources (e.g. FIFO buffers, arbiters, muxes, etc.).

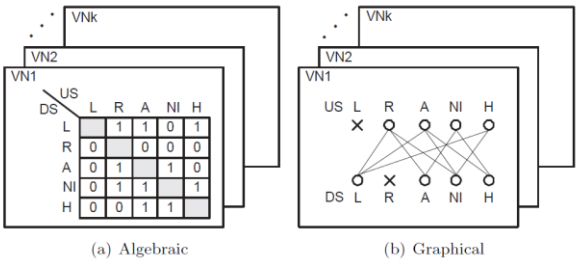

Figure 2.16. Representation of router configurable topology.

A convenient representation for the router topology is shown in Figure 2.16(a), where a switching matrix is defined for each virtual channel. Rows represent DS ports, while columns are associated to US ports. A non-null element at position (*i*, *j*) indicates that traffic entering port *i* can be routed toward port *j*. A null row/column represents a missing DS/US port, as illustrated in the equivalent graph representation of Figure 2.16(b). Note that DS and US port sets overlap if and only if the interconnection matrix is symmetric.
The router configuration space is summarized in Table 2.2. Features are classified according to their functional meaning. The *configuration layers* identify the coding technique required to implement that specific feature in a configurable fashion.
Features in Layer1 represent design parts that can either exist or not, thus reflecting the router topology.
Features in Layer2 need different code layouts for each legal configuration of Layer1. Indeed, the type of sub-components to instantiate and the signals connecting them to the surrounding logic depends on the router topology. These features are described with conditional instance statements (e.g. IF..GENERATE in VHDL/Verilog2001) and a consistent scenario for any possible configurations has to be considered during the coding phase. Moreover, conditional statements do not allow modification of port maps, thus

Feature	**Family**	**Configuration Layer**
Link	topology	Layer1
Stream		
VN		
Reserved queues	Buffering & latency	Layer2
Queue size		Layer3
Crossing latency		Layer2
Credit round-trip delay		
Arbitration schemes	QoS	
Clocking scheme	link	
Flit size		Layer3

Table 2.2. Router configuration space.

additional pre-processor directives are required.
For example, when configuring the crossing latency, registers are either inserted or removed across the data-path and mutual connections need to be changed accordingly. If different arbitration schemes are selected, then different arbiter components should be instantiated and properly connected to the surrounding logic.
Layer3 includes parametric features such as size of buses and memories. Parametric features are easily implemented by means of array-like data types supported in any HDL.

2.4.2 NI Configuration Space

The proposed NI is designed to support a wide configuration space, see Table 2.3. Not only the advanced features can be enabled or disabled, but also the basic characteristics can be configured like flit size, IP bus data size, payload and header FIFO size, frequency and size conversion support, crossing latency. By changing the configuration set different trade-offs between performance and complexity are achieved.

Feature	**Value**
Flit size (NoC) & bus size	Parametric
Bus protocol	STBus, AXI, DNP
Bus/NoC Size conversion	Enabled / disabled
Frequency conversion	Enabled / disabled
Store&Forward	Enabled (per packet; per compound transaction) / disabled
Ordering handler	Enabled/disabled
Security	Enabled/disabled
Programming Unit	Enabled (programming bus, STBus or APB) / disabled with per-feature granularity
EMU (Initiator only)	Enabled / disabled
EMU & PM (target only)	Enabled / disabled
Memory Remap	Enabled (parametric) / disabled
QoS (periodic faction switch)	Enabled (single or multiple thresholds) / disabled
FIFO size (headers)	US path (parametric), DS path (parametric)
FIFO size (payload)	US path (parametric), DS path (parametric)
Retiming stages	Up to 3 depending if FIFO and retiming registers are enabled or not
End to end size conversion	Different types can be enabled / disabled
Interoperability	Different types can be enabled / disabled

Table 2.3. NI configurability.

Feature	Configuration Layer
Flit size & bus size	Layer3
FIFO size	
Bus protocol	Layer2
Bus/NoC Size conversion	Layer1
Frequency conversion	
Store & Forward	
Ordering handler	
Security	
EMU & PM	
Retiming stages	
Memory Remap	Layer1&3
QoS	
Programming Unit	Layer1&2
End to end size conversion	
Interoperability	

Table 2.4. NI configuration space.

The NI configuration space is summarized in Table 2.4. NI features are classified in configuration layers in order to highlight the configuration technique required to implement them in a configurable fashion as introduced in the previous section.

2.4.3 The Metacoding Approach

As explained in previous section, the RTL coding requires an effort which grows exponentially with the dimension of Layer1 space. For example, considering a router with two virtual channels, five links and independent US and DS paths for each link, the number of configurations is greater than 2^{16}.
Hand-writing the whole codebase is not only time-consuming, but also prone to errors, because of the high similarity among code portions related to the same sub-component. Lastly, since each configuration comes from the combination of different portions of code, full code coverage can be achieved only by verifying any legal router topology.
The greatest limitation coming from a traditional coding style based on conditional instance statements and pre-processor directives is the *locality* of such statements. Portions of code activated by the same condition, but located in different sub-components have to be replicated, because each statement queries the configuration space independently from each other.
This work refers to the abstraction process as *metacoding*, while *metadesign* is used to denote the code generator (or *metacompiler*). Consistently with the terminology adopted in [42], the term *codelet* designates a group of HDL statements describing some part of the design.

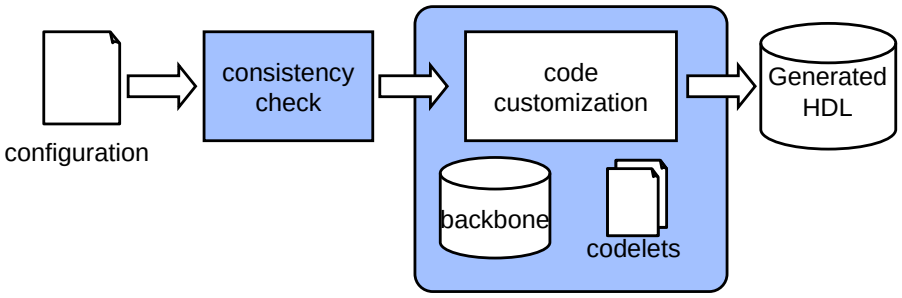

Figure 2.17. Main steps in the metacoding flow.

The metadesign is an equivalent representation of the design composed of three modules which interact as illustrated in Figure 2.17:

- **Checker**: constraints the configuration space by filtering any illegal set of parameters;
- **Backbone**: stores the configuration and provides the codelets with methods for querying parameters' value;
- **Codelet**: a basic class equipped with methods for HDL code customization. Since different codelets may require different customization procedures, the codelet class can be specialized for each of them.

A commercial tool for IP packaging (Synopsys coreTools) embodies the checker. The configuration parameters are annotated with proper expressions that the tool uses to determine whether a configuration is legal or not.
The backbone and the codelets are custom software components which are aware of the design structure and are hooked to the above mentioned tool. The backbone captures the configuration intents in such a way that the codelets can query to customize themselves. Codelet objects are in turn composed of:

- An HDL template describing some part of the design without explicit references to design items (components, signals, ports, etc.);
- A set of rules for substituting instance specific items within the template;
- A signature, used as a key for querying the backbone. The signature is built concatenating the coordinates of the codelet in the Layer1 space (i.e. Layer1 parameters determining the existence of that codelet in the HDL).

As already highlighted in Section 2.4, the metacoding approach raises the level of abstraction of the design description, not of the design itself. In other words, the codelets use the same hardware abstraction layer of the design (RTL in this case).

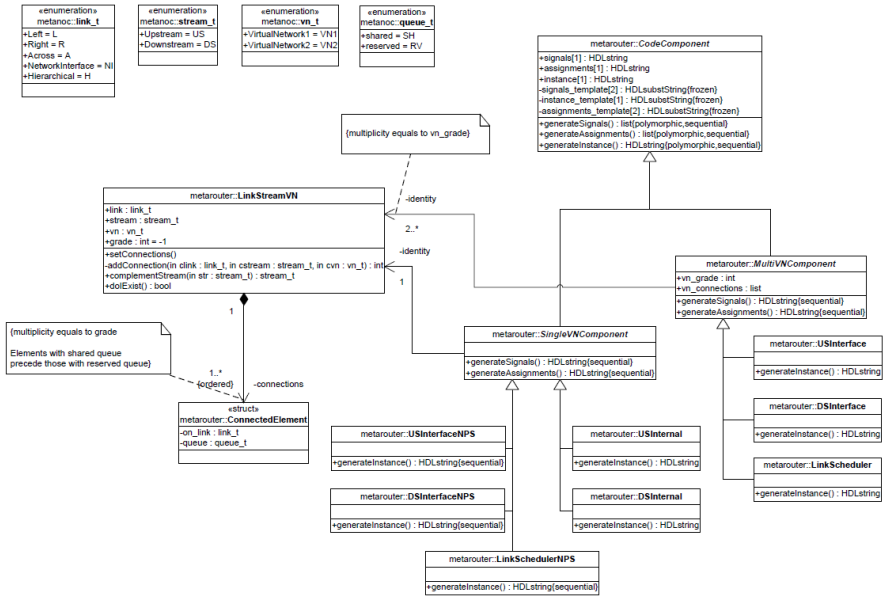

Figure 2.18. UML class diagram of the metarouter plug-in.

Figure 2.18 illustrates the internals of the router metadesign (*metarouter*) in the form of a UML class diagram. The LinkStreamVN class impersonates the backbone. The codelet class is further specialized according to the components highlighted in Figure 2.4. An intermediate level of specialization differentiates between single and multiple VN components. All classes are implemented using XOTcl, an aspect oriented scripting language based on Tcl [45].

Table 2.5 compares the metacoding approach with a standard coding approach using the router codebase as a benchmark. The router was originally implemented using plain VHDL and pre-processor directives. The migration to the metacoding approach occurred before a number of features (like additional ports) were added. The statistics reported in the '*hand-coded*" column are a projection of the original codebase considering the configurability implemented in the metarouter.

The metadesign achieves a reduction in the codebase size greater than 60%. The number of distinct topologies to verify is decreased of three orders of magnitudes. This reduction is possible since router topologies can now be

	hand coded	**metadesign**	**Saving**
HDL (lines)	11158	4093	63%
Configuration Intents (lines)	3211 (preprocessor macro)	994 (Tcl/XOTcl)	69%
Topologies to verify	> 216	56	10-3

Table 2.5. Router configuration space.

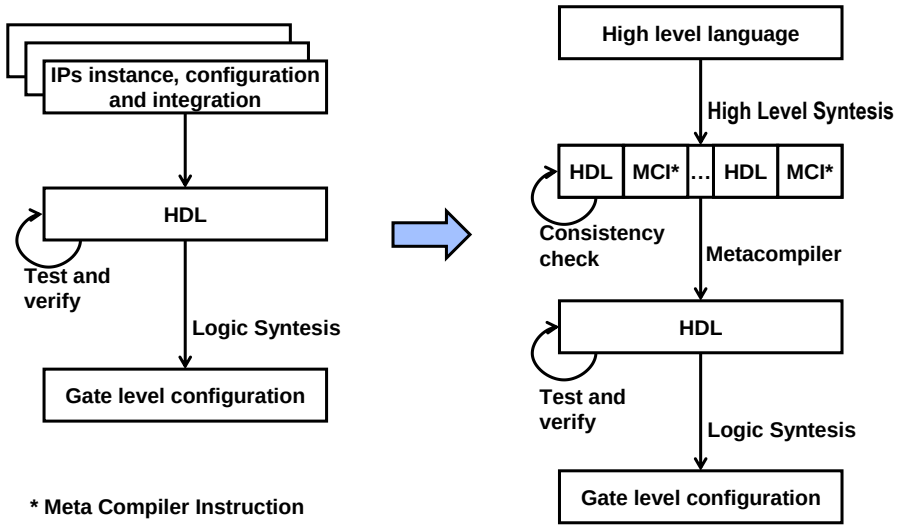

Figure 2.19. Automated Design Flow.

grouped into equivalence classes and only one representative for each class needs to be verified.
Similar techniques have been applied to the design of the network interface. *Metaheader* and *Nodemap* abstract features such as network and transport packet format, routing tables and network access policy. By abstracting these aspects, it was possible to write the HDL code of this component focusing on performance optimization without sacrificing code readability, hence maintenance. This approach also ensured consistency among all the instances of the network components.

2.4.4 Automated Design Flow

Configuration complexity described in the previous Sections is not just related to the design and the verification of the single blocks composing a NoC (i.e. NI or router). A hard configuration process can represent a real deterrent for NoC adoption if not effectively addressed. As already discussed, next to AMBA bus, ST bus or other traditional buses, NoC is becoming a fundamental commodity for semiconductor manufacturers. Nevertheless, it is crucial to have a simple and quick configuration mechanism to meet any specific application requirement without introducing an excessive overhead coming from the configuration of the NoC blocks. In order to clarify this aspect it is possible to consider a simple example: an advanced NI (like the one described in this work) may have around 140 parameters while a router offering the flexibility discussed may have around 450 parameters. In this scenario, the number of parameters a NoC architect should be able to manage, in order to configure an 8 nodes NoC is higher than 4,7 K-parameters.
To overcome this problem some steps of the traditional design flow were revised, as sketched in Figure 2.19.

The typical design flow was extended defining a high level language composed of Hardware Descriptor Language (HDL) instructions and Metacompiler Instructions (MCI). MCIs are processed by the metacompiler which is able to generate the configured HDL code on the basis of the configuration parameters. Once the whole HDL code is generated it continues the typical design flow through the logic synthesis and the following steps.
The architecture of the developed framework is depicted in Figure 2.20. The INoC is a GUI allowing assembling the NoC platform and customizing their building blocks through a friendly and semi-automatic interface.
Through the INoC interface, the NoC architect can easily define the network architecture and tune the configuration parameters considering the application requirements. Any operation, as instances and interconnections, can be made in a graphical fashion. Parameters are generally calculated automatically and propagated to neighbouring blocks. For example: when configuring the US interface of a router, the DS of the connected router is configured coherently. Otherwise, when an automatic value cannot be calculated, the default value is set.
The INoC interface is developed as an Eclipse plugin.
Moreover, the metacompiler (i.e. metarouter, metaheader and nodemap), which is composed by checkers, backbones, and codelets, is in charge to check the configuration and to generate the configured HDL as described in Section 2.4.3. Finally, the metacompiler is able to customize the test environment for the whole generated interconnection.

2.5 CMOS Implementation Results

The correct functionality of the proposed NoC design in multiple configurations has been verified at different abstraction levels. First a constrained-random functional verification environment has been created and applied to multiple. To this aim the *e*-language for functional verification of digital designs has been exploited. The creation of the constrained-random

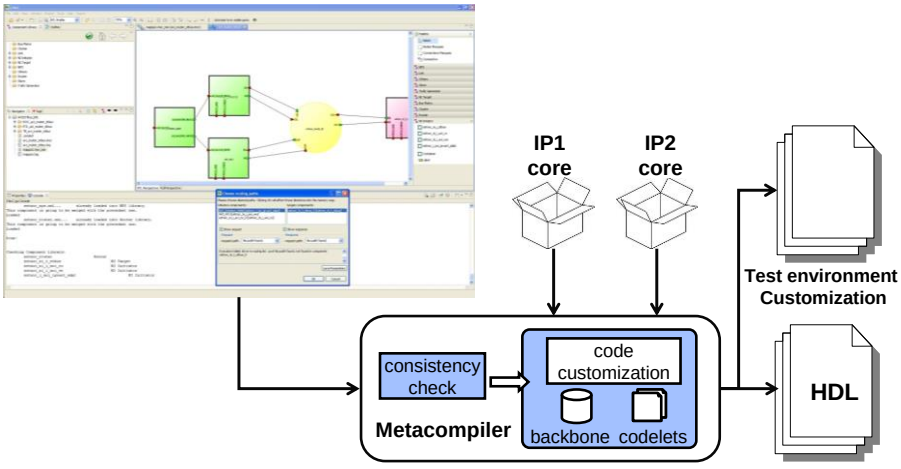

Figure 2.20. Automated Design flow Elements.

functional verification environment and its application to NoC building blocks such as NI and Router is discussed in detail in [15], [50], [61].
The validated data base has been then synthesized on 45nm and 65nm CMOS standard-cells technology from STMicroelectronics and on FPGA devices from Xilinx and Altera. The developed test benches have been reapplied on several synthesized instances allowing also for timing verification and validation in all corner cases (best, typical and worst considering process-temperature-voltage variations). Beside simulations also real-time emulation on FPGA-based prototyping platforms has been accomplished. The data base, verified and validated at different levels by both simulations and emulations, has been then successfully used for the implementation of a real-world 8-tile multi-core system in 45 nm CMOS technology inside the project SHAPES, in collaborations between University of Pisa, ATMEL and STMicroelectronics [5], [10]. The developed NoC technologies have been also integrated in several STMicroelectronics projects.

2.5.1 CMOS Synthesis Results

As discussed in the previous Section the verified database has been characterized in submicron CMOS technology for different configurations evaluating area occupation, power consumption, crossing latency, throughput. This Section reports the achieved results in 65nm CMOS standard-cells technology, using 1.1V supply.
The router has been characterized on STMicroelectronics 65 nm CMOS standard-cells technology, always achieving optimal trade-off between performance and complexity for the different router configurations. As example, a 5-port Router configuration with a data size of 128 bits, with 1 and 2 VNs, using input buffers but without output queues and retiming stage has been synthesized and results are reported in Table 2.6. It shows a circuit complexity of less than 33 kgates considering one VN and less than 61 kgates for the configuration implementing two VNs and it achieves a clock frequency of 400MHz. For the first configuration the static power consumption is around 40uW. The static power consumption (leakage) in typical conditions of the first configuration is around 40uW and around 73uW for the second configuration.
Obviously, performance in terms of area, maximum frequency and leakage power is strictly tied to router configuration at design time, mostly depending on topology (i.e. number of ports and VNs) and internal buffers size. As example, a 3-port Router with 3-bit data size and a single VN, with no OQs instantiated, has a circuit complexity lower than 9 kgates.

VN	Area	Power
1	32,8 kgates	10,1mW dyn + 40,1uW leak
2	60,6 kgates	18,7mW dyn + 73,2uW leak

Table 2.6. Router complexity and power consumption in 65nm (at 400 MHz).

Different NI configurations have been considered, and some of them are reported in Table 2.7, which can be suitable for different scenarios. The NI instance labelled 'A' is an advanced configuration directly supporting in hardware all basic and advanced networking features and with large FIFO buffers and large size for flits and for the IP bus data. Such configuration is suitable for MPSoC designs requiring high on-chip communication bandwidth and the hardware support of complex networking functionalities. The instances labelled from 'B' to 'D' refer to typical NI configurations with all main features enabled and with different sizing for the FIFO buffers, flits and IP bus data. Finally the configuration labelled 'E' is a simple NI configuration, implementing basic functionalities as in most of state-of-the-art designs [9], [10], [17], [59], [62], [63].
By comparing the results for configurations 'A' and 'B', having the same sizing for FIFO and data, the overhead of the advanced networking functionalities can be easily evaluated.
An important difference among the configurations from 'A' to 'E' in Table 2.7 is also the increasing bandwidth, by a factor of 4 for the IP bus interface, from the Basic 'E' configuration to the Advanced 'A' configuration due to different data sizing. As bandwidth increases also internal buffering with FIFOs increases (note the higher FIFO size in the Advanced NI).
Another important remark in discussing the synthesis results is the size of the address map for the NI. In fact, the NLH encoder contains address comparators whose number and size depend on the network nodemap configuration. In the present case, all NI configurations use a map specifying 8 slaves, each one having 4 different memory regions defined by threshold 10

Configuration	**A**	**B**	**C**	**D**	**E**
Flit size	64	64	64	32	32
STBus size	128	128	64	32	32
Size conversion	√	√	×	×	×
Frequency conv.	√	√	√	√	√
Retiming Stage	√	√	√	√	√
Store&Forward	√	×	×	×	×
Ordering handler	√	×	×	×	×
Security	√	×	×	×	×
Programming Unit	√	×	×	×	×
EMU	√	×	×	×	×
FIFO req header	64	64	16	8	8
FIFO req payload	128	128	64	32	8
FIFO rsp header	64	64	16	8	8
FIFO rsp payload	128	128	64	32	8
Interoperability	√	×	×	×	×
Memory Remap	√	×	×	×	×
QoS with FBA	√	×	×	×	×

Flit, Bus and FIFO sizes are expressed in bytes.

Table 2.7. Configurations used for NI characterization.

bits wide.
It should be noted that in the configurations of Table 2.7 the Shell implements an STBus Type 3 interface. Similar results are obtained for other Shell configurations (e.g. AXI Shell) since the main area contribution is due to the Kernel (protocol bus independent) which contains the header and payload FIFOs.
With reference to the configurations in Table 2.7, Table 2.8 shows the achieved results in terms of circuit complexity when considering a target frequency of 500 MHz for the NoC. The results of Table 2.8 refer to a NI initiator but similar results are achieved for a target NI (with similar configuration) which is a mirrored version of the Initiator one.
In Table 2.8 the 'E' NI instance with just the essential logic to exploit NI basic services is around 8 kgates. Its complexity is comparable to state-of-the-art NI designs such as [59], implementing minimal features, synthesized in the same 65 nm CMOS technology.
To further reduce the circuit complexity vs. the 'E' configuration, to few kgates, the NI can be configured with zero-FIFO Kernel as discussed in Section 2.2.2.
The proposed NI architecture is scalable: if more services and a larger flits size and FIFO depth are needed then more advanced NI macrocells can be configured and generated. As example, the 'A' NI configuration, with all the main advanced features enabled as indicated in Table 2.7, with 128/64-bit IP/NoC size has a complexity of 41.5 kgates. The range from few kgates to 41.5 kgates determine the complexity variation range for all the other configurations in Table 2.7. By comparing the results for the different configurations in Table 2.8 it can be noticed that the NI complexity is strongly affected by the storage buffers implemented (FIFOs, retiming stages, programming registers), see as example the configurations from 'B' to 'E'. Instead, the cost of the advanced services in terms of complexity overhead is limited: by comparing 'B' and 'A' configurations, having the same FIFO size and the same STBus/NoC data size, it can be noted that the overhead of all advanced networking features (interoperability, QoS FBA, memory remap, EMU, Programming Unit, Security, Ordering handler, Store&Forward) is limited at 8 kgates.

Configuration	**A**	**B**	**C**	**D**	**E**
kgates	41.54	33.35	18.4	12.08	8.2
Area m	86399	69379	38266	25124	17058
IP bus throughput	64 Gb/s	64 Gb/s	32 Gb/s	16 Gb/s	16 Gb/s

Table 2.8. Complexity and throughput for the different Initiator NI configurations of Table 2.7 in 65nm (at 500 MHz).

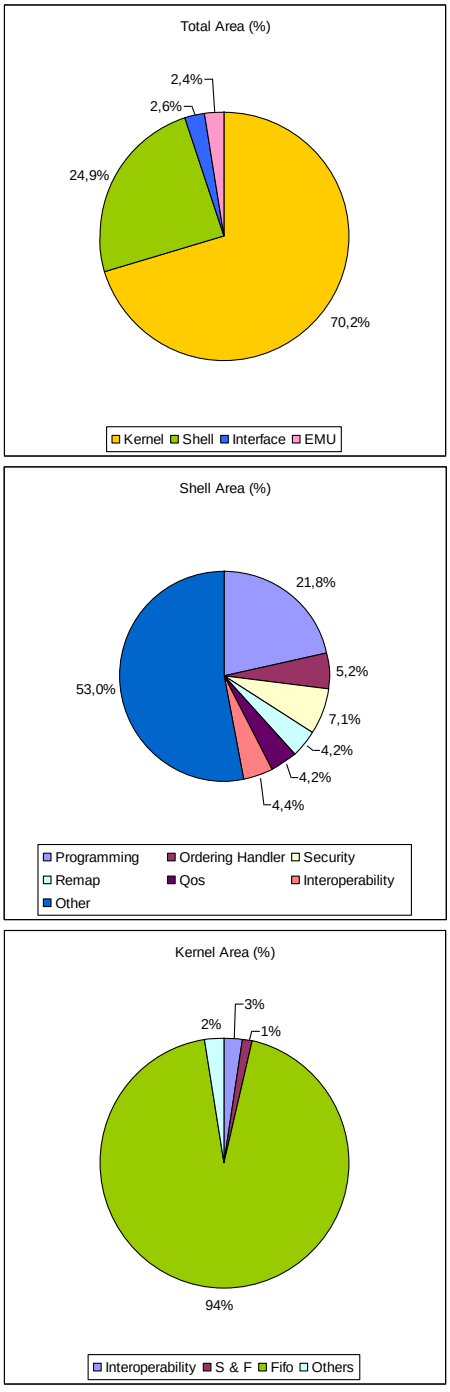

Figure 2.21. Complexity of an advanced NI due to the different sub-blocks.

With reference to NI instances with advanced 'A' configuration Figure 2.21 shows how the macrocell complexity is shared among the different sub-blocks. From such figure it is clear that the NI complexity is dominated by the kernel (70%) and particularly by the FIFO size.
Table 2.9 compares the proposed NI results to other NI IP cores found in literature in terms of clock frequency and circuit complexity (measured in terms of equivalent Nand2 logic gates since they are realized in different technology nodes). In Table 2.9 the advanced features supported by each state-of-the-art NI are also highlighted. Our proposed NI is the only one implementing all the advanced networking features while most of the other NI designs implement just basic functionalities. Therefore for a fair comparison in Table 2.9 for our NI design we report two different configurations: the advanced 'A' configuration and the basic 'E' configuration.
From the analysis of Table 2.9 it emerges that: (i) the implementation complexity of our design in basic configuration is comparable to other designs offering similar services [59], [62], [63]; (ii) the advanced 'A' configuration offers an optimal trade-off between the features supported and the complexity overhead vs. the state of the art: e.g. with respect to [56] the advanced NI has a lower complexity and directly support in hardware more features.
As far as the NI crossing latency is concerned the proposed architecture well behaves vs. the state of the art. As discussed before, the minimum crossing latency in the proposed NI is equal to the number of the inserted retiming stages that is configurable in the range 1-3, while [51] declares 3-4 cycles of latency, [54] takes 4-10 cycles and [23], a NI optimized for low latency needs 3 cycles in the request path and 4 in the response path. The number of retiming stages also affects the maximum achievable clock frequency: the latter together with the size of flits and IP bus data determine the supported throughput on the NoC side and on the IP side. With reference to a NI in 65nm CMOS 1.1V technology a clock frequency up to 1 GHz can be achieved with 3 retiming stages and 500 MHz with 1 retiming stage (as proved by the configurations of Table 2.7 whose implementation results are reported in Table 2.8).

IP	kgates	Tech.	Freq.	Advanced features
'E'	8.2	65nm	500 MHz	Basic
'A'	41.54	65nm	500 MHz	See Table 2.7
[55]	18.9	90 nm	N/A	Re-ordering
[51]	35	130nm	150 MHz	Re-ordering
[57]	20.75	130nm	500 MHz	Security memory access
[54]	21.2	130nm	500 MHz	Programmability, Re-ordering, QoS
[56]	53.75	130 nm	312.5 MHz	Error management (CRC detection, retransmission)
[62]	10.75	180 nm	126 MHz	Basic
[59]	7.4	65 nm	500 MHz	Basic
[63]	7.6	130 nm	500 MHz	Basic

Table 2.9. Comparison of the proposed NI to state of the art.

The blocks implementing advanced features are designed so that in normal conditions they do not introduce extra latency cycles. Obviously, if size or frequency conversions are enabled or the store&forward feature changes the traffic shape to eliminate bubbles, then extra delay cycles occur, as explained in Section 2.2.2.

2.5.2 Complexity of the NoC vs. the connected IP in real MPSoC implementations

It is also important to evaluate the overhead of the NoC blocks with respect to the complexity of the connected IP cells in real MPSoC implementations. The NI has been successfully integrated in a real-world 8-tile multi-core system in 45 nm CMOS technology in the framework of the project SHAPES, in a collaboration between University of Pisa, ATMEL and STMicroelectronics. The SHAPES MPSoC connects 8 tiles each composed by a VLIW floating-point DSP based on the 64-bit mAgicV architecture by ATMEL, 1 Mbits of program memory and 640 kbits of RAM, a RISC processor based on the ARM926 core, a distributed network processor (DNP) interface for extra tile communication which is interfaced to the NoC through a NI, a set of peripherals for off-chip communication. The area of each tile is 7 mm^2 and the number of logic gates is 4.46 Millions. The area of the whole MPSoC platform is 8 x7.1 mm^2 and the number of gates is about 36 Millions. Considering for the platform a target frequency of 250 MHz the power consumption for a tile is in the order of 350 mW using a voltage supply of 1.1 V. The static power consumption (leakage) of the tile is 8 mW at 1.1V. The total MPSoC dynamic core power consumption has been estimated in 2.8 W in typical conditions (3.7 W in worst case). The static (leakage) power consumption is 65 mW in typical conditions.

The area and power overhead of the NoC in the SHAPES MPSoC resulted negligible vs. the connected computing tiles. The occupied area of the synthesized NoC interconnect (NIs plus routers and links), after place and route, is 0.123 mm^2: 1/3 due to the 8 NIs and 2/3 due to 8 4-port Routers. The overall power consumption (dynamic power plus leakage power) for the NoC in the SHAPES MPSoC platform is less than 4 mW, 40% due to the 8 NIs.

The power contribution of the NoC could be further reduced adopting proper end-to-end data coding scheme as proposed in [22].

The above implementation results refer to the following NI configuration:

- NI with data bus size and flit size of 32 bits with DNP-compliant Shell;
- IP and NoC running at 250 MHz;
- header and payload FIFOs in the Kernel (Request and Response paths) have 2 locations of 32 bits;
- no advanced NI services support (security, order handling, EMU, frequency/data size conversion,..)

And the following router configuration:

- Router flit size is 32 bits;
- running at 250 MHz;
- 4 ports (left, right, across and NI);
- One Virtual Network;
- Output queue is 2 locations;

- Input buffer is 2 locations.

3 EFFICIENT APPLICATION DESIGN EXPLORATION

This chapter explores computing application designs from the system-level point of view. In particular it deals with remote monitoring platforms focussing on system requirements, efficient architectures, data acquisition & processing and security aspects.
Section 3.1 introduces the design of a health monitoring system for patients affected by Chronic Heart Failure (CHF). After briefly introduction, Section 3.1.1 describes the system architecture. The features of all involved sensors are discussed in Section 3.1.2 while Section 3.1.3 is focused on the developed front-end IC for cardiac sensor. The sensor data processing strategy is described in Section 3.1.4. Section 3.1.5 and Section 3.1.6 present the results of the demonstration phase and the comparison with the state-of-the-art.
Section 3.2 deals with energy monitoring for Smart Grid, discussing aims (in Section 3.2.1), architecture (in Section 3.2.2) and security issues (in Section 3.2.3). Section 3.2.4 proposes a possible realization of a Home Area Network (HAN) for Smart Grid. Possible hardware-software architectures and implementations using COTS components are presented respectively in Section 3.2.4.1 and Section 3.2.4.2. Security issues are discussed in Section 3.2.4.3.

3.1 Health monitoring system for CHF patients

Chronic heart failure represents one of the most relevant chronic disease in all industrialized countries, affecting approximately 15 million people in Europe and more than 5 million in the US, with a prevalence ranging from 1% to 2% and an incidence of 3.6 million new cases each year in Europe and 550.000 cases in US [64], [65], [66]. It is the leading cause of hospital admission especially for older adults reaching a prevalence of 1.3%, 1.5%, and 8.4% in 55–64 years old, 65–74 years and 75 years or older segments respectively [65]. Admission to hospital with heart failure has more than doubled in the last 20 years [64] and it is expected that CHF patients will double in 2030. Hospital admissions caused by CHF result in a large societal and economical issue, accounting for 2% of all hospitalizations [67]. The CHF management accounts for 2% of the total healthcare expenditure [68], [69] and hospitalizations represent more than two thirds of such expenditure [66].
The current healthcare model is mostly in-hospital based and consists of periodic visits. Previous studies pointed out that in patients with a discharge diagnosis of heart failure the probability of a readmission in the following 30 days is about 0.25, with the readmission rate that approaches 45% within 6 months [70]. It is acknowledged that changes in vital signs often precede symptom worsening and clinical destabilization: indeed, a daily monitoring of some biological parameters would ensure an early recognition of heart failure de-compensation signs, allowing appropriate and timely interventions, likely leading to a reduction in the number of re-hospitalizations. Due to lack of resources at medical facilities to support this kind of follow-up, the use of ICT (Information and Communication Technologies) has been identified by physicians and administrator as a possible valid support to overcome this

limit. There is in literature some evidence that a multidisciplinary management program [71],[72] including a home-based follow-up strategy can improve outcome of heart failure patients, including a reduction in mortality, hospital readmissions, and lengths of hospital stays, and increase patient satisfaction [73]-[75].

The platform presented in this chapter has been developed within the Health at Home project (H@H) of the Ambient Assisted Living Programme (AAL). It takes into account the recent AAL Roadmap guidelines [76], the future challenges in telecare [77] and some recent studies conducted on AAL solutions [78], [79].

The H@H platform aims at connecting in-hospital care of the acute syndrome with out-of-hospital follow-up by patient/family caregiver, being directly integrated with the usual cardiology departmental HIS. Patients' signs, symptoms and raised alarms can be received by healthcare providers, and aggravations can be quickly detected and acted upon. Thanks to the collection of vital parameters at home, the sensor data signal processing and the automatic data transmission to the medical centre, a more frequent (usually daily) assessment of clinical status than in conventional practice is permitted [80].

3.1.1 H@H Telecare System Overview

The H@H development Consortium is composed by industrial and research partners with qualified competences in sensing and data processing as well as very important healthcare providers (Hospitales Virgen del Rocio, Spain; Dom Koper Hospital in Slovenia and the research clinical centre Fondazione Gabriele Monasterio in Italy). The system requirements come directly from the

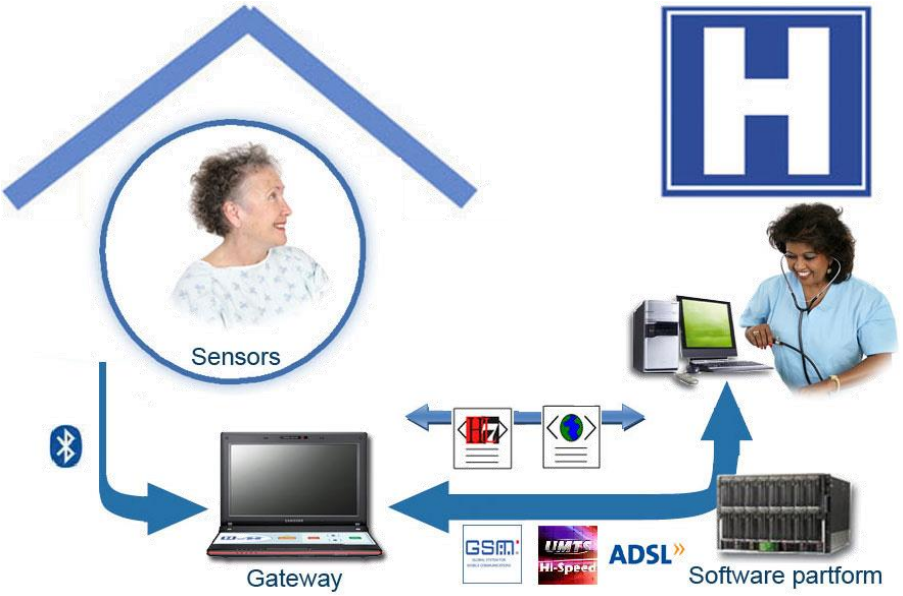

Figure 3.1. H@H System Architecture.

long experience in the CHF field of the involved physicians. The resulting platform takes into consideration both medical expectations, patients' features (elderly, with comorbidity and cognitive deficit) and the progressive nature of the disease. For these reasons we propose an intuitive home monitoring system based on a configurable follow-up operating protocol (OP), integrated

Parameters	SAMPLING	**Basic**	**Advanced**
3 lead ECG	500 S/s/lead (12bit/S)	√	√
SpO2	3 S/s (10 bit/S)	√	√
Blood pressure	1 S/type (32bit int)	√	√
Weight	1 S (32bit float)	√	√
Chest impedance	25 S/s (10 bit/S)		√
Respiration	25 S/s (10 bit/S)		√
Posture	3 axes x 1 S/s/axis (8 bit/S)		√

Table 3.1. Sensing requirements and versioning.

with the HIS of the cardiology department through a server software platform.

The complete H@H system has the client/server architecture shown in Figure 3.1. The clients are typically located at patient's home and consist of a set of wireless sensors to measure the main vital signs (see Table 3.1) and an additional device, the home gateway (see Table 3.2), that centralizes all computation and communication resources. These domestic subsystems have in turn a client/server structure where sensors are the clients of the collection and transmission point represented by the home gateway. The server platform, installed at health service facilities, accepts data from gateways making them available in the HIS and finally allows the management of all patients' data since their enrolment.
The home gateway [81] receives sensors data via point-to-point Bluetooth connections initiated by the sensing modules. Upon received, it processes all data to detect dangerous alterations and then forwards them to the hospital server through ADSL or mobile Broadband, to be further analyzed and flowed

Intel ATOM N450 1.66 GHz
Memory RAM 1 GB DDR2
Hard Disk Sata 250GB
Display 10.1"
Ethernet 10 / 100 LAN
Wireless 802.11bg/n
Bluetooth 2.0
GSM/GPRS/EDGE and UMTS/HSPA
Built-in audio

Ad-hoc 5 keys keypad

Table 3.2. Home gateway hardware resources.

into the HIS. In case of an alarm situation, caregivers or relatives are contacted via SMS (i.e. reporting the abnormal values that lead to the alarm) and all pending data are sent to the server. Section 3.1.4 will explain in depth the sensors data processing and the alarm detection task. The gateway normally operates connected to the power-line, but the internal battery ensures about 5 hours of autonomy in case of power failure. The peculiar features of the target patients require the design of an intuitive and simple home gateway user interface. This is able to display reminder messages, guide animations and sounds when a planned activity time is reached according to the OP. Patient can read the last measured values and the status of sensors battery charge. Green, yellow and red background colours are used for information, warning and error messages respectively. To simplify the use of the gateway a 5-keys membrane keypad is provided (i.e. Yes, No, Alarm sending and Up/Down scroll buttons).

The server platform is based on the web services paradigm for data reception and presentation and also for the interaction with the cardiology department HIS and the patient's information management. The user interface allows the clinicians to interact with the system, also in mobility, using the web-browser.

The follow-up OP consists of a formal XML specification of the list of daily actions to be performed during the monitoring period (i.e. types and frequencies of measurements and drugs assumptions) and it defines the behaviour of the home gateway in terms of alarm thresholds for each scheduled measures, transmission policy (i.e. daily, weekly, immediate send), and selectable symptoms. It is configurable according to the patient's needs and remotely updatable if required via the server platform. This novel concept embeds medical prescription for the given patient in the home gateway.

The availability of multiple communication paths ensures a good adaptability of the system in overall operating areas and improves the fault tolerance. As the coverage of GSM is close to 99% the system reaches a very high degree of connectivity. In the worst case the GPRS upload data rate of 20Kbps is sufficient to transmit all data in few minutes. Surely better performance becomes available with EDGE and UMTS. Furthermore the gateway is able to exploit the GSM network to send SMSs to the physicians, patient's relatives and caregivers in case of alert situation.

Authentication, integrity and confidentiality of the communication are guaranteed by the HTTPS protocol. The use of international standard for data communication, ANSI HL7-RIM Clinical Document Architecture (CDA) v2 [82]-[84] and XML, improves the interoperability of the system as well as the integration with existing HIS. All numeric and waveform observations use SNOME CT [85] or LOINC [86] standards codes.

The proposed system is conceived to allow a better assessment of vital signs identified by clinicians as the most significant in CHF through one or few daily measurements, being in contrast to those systems that offer a continuous monitoring for limited period. It does not introduce any remarkable overhead with respect to regular activities of the medical staff. Indeed all signs are flowed as row data into the patient's electronic health record (EHR) and, thanks to the provided automatic signal processing capability, clinicians and caregivers are timely informed in case of alarm detection. Moreover H@H minimizes the impact on the patient. The wireless biomedical sensors avoid

connection-cable encumbrance. The number of sensing modules is minimized and the signal quality is not excessively dependent on transducer positioning. The domestic gateway reminds the scheduled activities, provides a graphical assisted procedure that shows how to use the sensors and acquires data without requiring any preventive action to the user.

3.1.2 H@H Sensor Devices

In general biomedical sensors address the wearability/portability, non-invasivity, wireless communication and battery duration concerns in order to be easily used autonomously by the patients at home. Moreover, the system minimizes the number of devices and sensors/electrodes to be positioned on the patient's body (e.g. 3 recording sites ECG instead of a more complex 12-lead ECG is adopted to limit the effects of electrodes misplacement [87]). The measurement experience consists of wearing/using the sensors periodically, once or twice a day, only for the duration of the acquisition without any long-period application of the sensors as in different solutions.
According to the analysis carried out by the clinicians, interesting vital parameters to monitor in a CHF patient are ECG, SpO2, weight, blood pressure, chest impedance, respiration and posture (see Table 3.1 for sampling details). To achieve an asset for the implementation and to increase scalability of the system, the sensing modules have been clustered into two possible configurations: basic and advanced. Basic partitioning is intended as the minimum set of requirements to ensure a complete and useful telecare system in CHF. As advanced we refer to additional features in order to widen the kind of CHF patients to be possibly enrolled into the telemonitoring and to cope with other chronic diseases (i.e. Chronic Obstructive Pulmonary

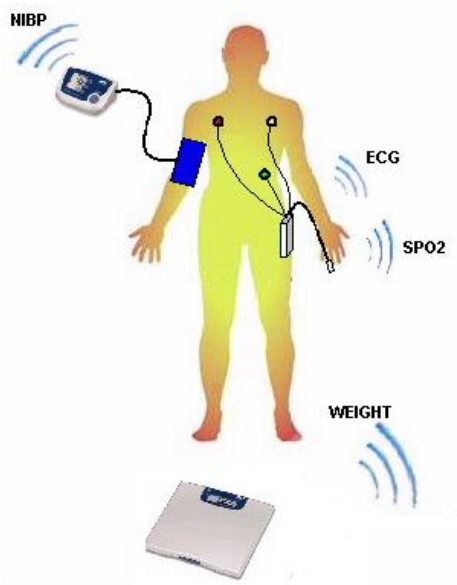

Figure 3.2. Sensors Positioning.

Disease, Diabetes). Since the basic configuration provides to the HIS the needed data set for accurate CHF telemonitoring, then the basic set in Table 3.1 is the one implemented for the medical technology test reported in Section 0.

To achieve our goals, according to the basic configuration, the H@H system is formed by a bunch of three Bluetooth and battery-powered sensing devices, the commercial standalone modules UA-767BT arm cuff device for blood pressure readings and UA-321PBT digital scale by A&D Medical and the new integrated ECG-SpO2 module ad-hoc developed in the framework of this project. The latter is based on a new multi-channel front-end IC developed for cardiac sensor interfacing. The overall set of devices is shown in Figure 3.2, where the sensors positioning and the limited impact on the patient are also visible.

All sensing devices exploit the Bluetooth 2.0 wireless technology for the communication with the home gateway. They send only raw data and any form of signal processing is completely demanded to the gateway. Class I transceivers, with near 7dBm of transmission power, ensure an appropriate coverage in the domestic environment.

The modules implement the Service Discovery Protocol (SDP) and the Serial Port Profile (SPP) to discover and wirelessly communicate with the access point collection service, which is identified in the home gateway, with optional link-level security (128bit encryption). Each sensor acts as initiator of connections according to the following rules:

1. if it knows the default remote peer address (DRPA), it begins the page procedure to establishes a point-to-point connection (GAP Connectable Mode enabled), sends all data and closes the connection;
2. if no DRPA exists, the module initiates the inquiry procedure for 10.24 sec to discover remote peers in range (GAP General Discoverable Mode enabled) offering the SPP collection service (i.e. gateways);
3. all peers identified in step 2 are sequentially asked for the authentication PIN (pairing mechanism);
4. the device that matches the PIN receives data from the sensor and in case of successful transmission it becomes the new DRPA;
5. if in the above step 1 the page procedure fails, steps 2 to 4 are performed to identify the new DRPA.

The sensors are provided to the patients already configured to work with the given home gateway, hence at runtime only the above step 1 is performed while the whole procedure is required once at configuration time. In this condition a measurement simply consists of turning on the sensing device and waiting until the end of the measurement process without any preventive interactions with the gateway (e.g. device searches, PIN typing, connection opening), introducing also benefits in battery saving. Actual gateway implementation deals with one active connection at a time, as required by the physicians which defined a time-division strategy for bio-signal acquisition. Simultaneous connection attempts lead to discard attempts other than the first. Anyway this meets the user requirements as the system guides the patient to follow the OP activities requesting actions one by one.

With respect to alternatives such as Zigbee or point-to-point Wi-Fi links, Bluetooth 2.0 provides the desired trade-off among available bandwidth, security and reliability of the connection, cost and power consumption of the node, link distance. It operates at 2.4GHz, does not require line of sight positioning of the units and the frequency-hopping and fast acknowledgment scheme improves the robustness of the link in noisy environments. Indeed the nominal data rate is around 2Mbps against ~18Kbps required by the ECG-SpO2 device. Zigbee, although it's very interesting for its power consumption, is still a new product with respect to the cheaper and more diffused Bluetooth technology.

3.1.2.1 ECG-SpO2 module

The ECG-SpO2 module is a new sensing device, developed ad hoc in the H@H project. Specifically, the module provides electrocardiographic, pulse oximetry and plethysmographic measurements by means of proper sensors and electrodes placed on patient's body, implementing non-invasive techniques. It is hosted in a robust and small size ABS case (92 x 150 x 28 mm, 200g) powered by an integrated rechargeable 3.7V and 1700mAh Li-Ion battery able to ensure 18 hours of continuous operability (i.e. more than 3 months of acquisitions considering 5 minutes track twice a day).

The module uses an ECG patient trunk cable with 4 lead wires: RA, LA, LL and RL (neutral) to provide the electrocardiographic signal. The sensor for pulse oximetry and plethysmographic measurements is a classical finger clip reader type to be applied at patient's first finger. In accordance with the physicians the Einthoven's 3 leads ECG configuration is considered sufficient for our purposes (e.g. detection of heart rate and rhythm) and not excessively dependant on the transducer positioning. All such signals are conditioned, digitalized and then packetized and finally transmitted via Bluetooth protocol. The ECG-SpO2 device outputs digitalized waveforms of two standard limb leads, the oxygen saturation in the blood, the plethysmographic waveform and the battery level. The final appearance of the sensing module and its schematic view are shown in Figure 3.3.

Integrating ECG and SpO2 functionalities in a novel single device reduces the number of sensing devices in the final system and also enables to acquire synchronized ECG and SpO2 or plethysmogram traces. This allows a larger and more specific amount of information for advanced analysis and multisensors data fusion (e.g. the Pulse Transit Time estimation).

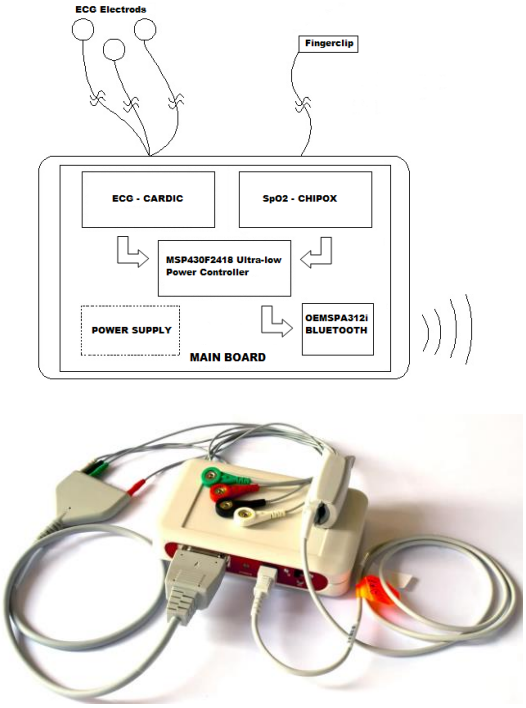

Figure 3.3. Architecture and final appearance of the ECG-SpO2 module.

As reported in Figure 3.3, the ECG-SpO2 module is realized assembling its building blocks on a single printed circuit board. All communications within the board take place under the coordination of a dedicated firmware running on the MSP430F2418 Ultra-Low Power Mixed Signal Controller, responsible for mixing raw data before passing them to the Bluetooth interface.
The core of the ECG block is an ad-hoc developed ASIC, very compact and high configurable, able to integrate the ECG functionality into portable and

Function	**Schedule**
Multi-channel ECG block features	8 input differential ECG channels fully configurable + configurable gain and offset per channel +pace maker detector + RL and Shield driver to reduce 50 Hz noise
Multi-parameter channels	Body temperature channel, blood pressure channel, general purpose auxiliary channel
ADC section	Resolution: 12 bit @83 KS/s, INL: ±1 LSB typ, DNL: 0.75 LSB max, conversion time: 14.1µs
Interface	Serial, 1.3 MHz, pin 3.3V/5V compliant
Power	5V/ 20 mA (max), 10 mA (typ) consumption
Technology/package	CMOS/TQFP128 14 x 14 x1.4 mm

Table 3.3. ASIC CARDIC technical specification.

Figure 3.4. UA-767BT blood pressure monitor.

wearable devices. The technical specifications of the chip, called CARD*IC*, are summarized in Table 3.3, while a detailed description of the chip architecture and performance is provided in Section 3.1.3.
The ECG-SpO2 module also hosts the ChipOx OEM by Envitec for pulsioximetry measuring, being fully configurable, highly dimension-contained and with low power consumption. Together with SpO2 readings, the device gives also Heart Rate (HR) digital data and a digitalized Plethysmographic Waveform (PPG). The range of measure is from 45% to 100%, with an accuracy of 1.5 – 2% for oxygen saturation and 0-255 LSB at 100 Hz, with accuracy > 6ppm / LSB for plethysmography. It consumes up to 25mA at 3.3 Volt of power supply and it communicates over an UART-TTL with a baud rate of 9600. Its dimensions are 31 x 14 x 5 mm.
The Bluetooth communication leverages the OEMSPA312i by Connect Blue, which is a small size module based on the Phillips BGB203 SiP. The BGB203 has on chip SRAM and FLASH stacked in the same package. The chip is compliant with the 2.0 standard and the Class I specification.

3.1.2.2 UA-767BT blood pressure monitor

Regarding non invasive blood pressure measuring, the market offer and the cost effectiveness led to the choice of a commercial stand-alone solution by A&D Medical (see Figure 3.4). This arm cuffed automatic meter, based on the oscillometric method, is already used in various telemedicine systems as it is certified for medical use and can be employed also at home. It is equipped with Bluetooth class I communication capabilities in order to send the acquired systolic and diastolic values to a base station and it is able to store up to 30 measurements waiting for transmission. It outputs also heart rate frequency. All values have 8bit precision; the sensibility range and the accuracy are 20-280 mmHg and +/- 3mmHg respectively. Dimensions are 147 x 64 x 110 mm with about 0.3Kg of weight. The power supplied by 4 x 1.5 Volt AA batteries ensures near 6 months of operability in normal OP (i.e. two or three measures per day).

3.1.2.3 UA-321PBT digital scale

The market of digital scales offers interesting Bluetooth enabled products allowing for wireless transmission of multiple readings, thus avoiding cable

encumbrance. These devices are user-friendly, with a low-profile design and are not expensive, commonly used in fitness and telecare systems.
The A&D scale shown in Figure 3.5 measures only weight and belongs to the same series as the UA-767BT, so they have the same Bluetooth data communication specification. The current reading is visualized on a on-board LCD display and sent to the base station. All values are expressed with 5bytes and a simple conversion formula is needed. It has a maximum capacity of 200Kg and an accuracy of +/- 0.1Kg. The dimensions and weight are 300 x 300 x 30 mm and 1.0Kg respectively. The battery life of the 4 AA batteries ensures approximately 1000 measurements (i.e. more than 1 year in classical OP).

3.1.3 IC Front-End for Cardiac Sensors

CARD*IC* is an ASIC designed and manufactured in order to be very compact and fully configurable for multi-sensor integration in wearable and portable medical devices. It ensures multiparameters medical acquisition, offering a cost effective solution able to meet requirements coming from both monitoring and diagnostic demand. Indeed as in the automotive case [88] the medical monitoring field has a potential large volume market justifying the development of a mixed-signal IC for high performance sensor front-end.
The chip is developed in AMS CMOS 0.8 µm CXZ 2 Metal Layers 2 Poly technology and is fully assembled in a 14 x 14 x 1.4 mm TQFP 128 package. The core architecture has been implemented taking into account the typical constraints of biomedical signals monitoring: like standard ECG, blood pressure, body temperature. Typical power supply voltage is 5V. Digital In and Output pins are compatible with 3.3V power supply level. The operating temperature is from 0 °C to 70 °C. The overall IC characteristics are reported in Table 3.3.
Referring to the IC functional block in Figure 3.6, the following main parts can be pointed out as a description of the chip features and functionalities:

- a fully configurable multi-channel ECG block including 8 input differential channels for signal conditioning (amplification, filtering, offset regulation) of 3-5-12-leads ECG systems, an adder generating the Central Terminal Point (CTP) for prechordial leads, a right leg (RL) driver and a SHIELD driver in order to reduce the Common Mode 50Hz noise, a Pace Maker detector section for pace pulses

Figure 3.5. UA-321PBT digital scale.

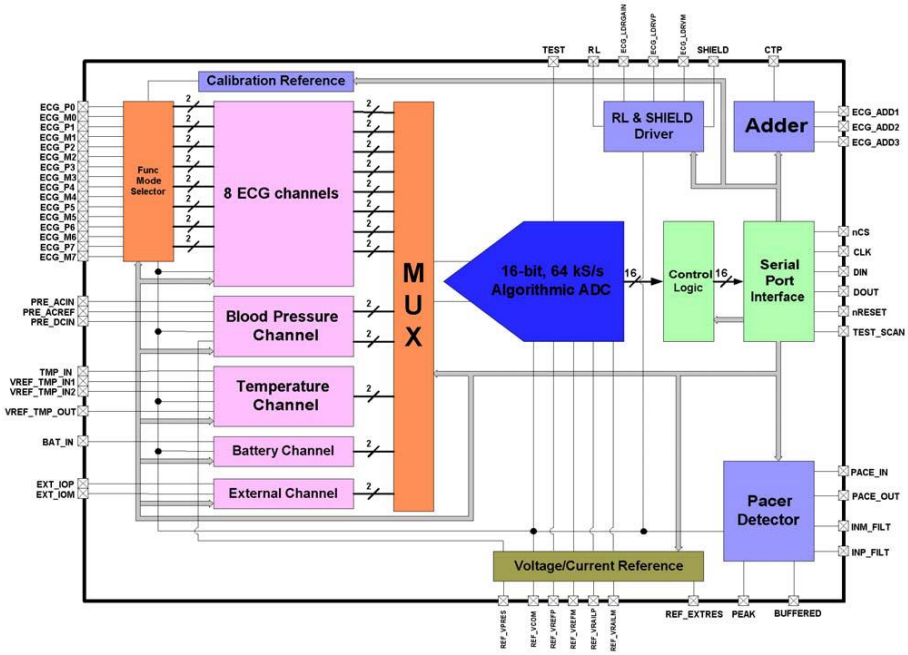

Figure 3.6. CARDIC chip schematic view.

detection on a dedicated pin. All is in conformity to IEC 60601- 2-51 2003-02 standard;

- one analog channel designed for decoupling and amplifying a blood pressure signal coming from an optional external pressure sensor;
- one analog channel able to process a temperature signal provided by an external temperature sensor;
- one programmable analog multiplexer for switching among the channels to be converted by the ADC;
- one 16 bit (12 bit linearity) cyclic algorithmic ADC to convert the voltages from the analog channels;
- a Serial Peripheral Interface (SPI) to configure the chip settings and for data readout;
- one battery channel for monitoring the battery status.

The structure of a single ECG channel is shown in Figure 3.7. It is mainly composed by an Instrumentation Amplifier (IA) with high Common Mode Rejection Ratio (CMRR), a Programmable Gain Amplifier (PGA), a BUFFER and an Offset Regulator block. The IA has a high CMRR, 100 dB typical, 92 dB minimum, in order to reduce environmental electric interferences, like the 50Hz noise from the industrial network, always present in both electrodes connected: with human body and ground. The power supply rejection ratio (PSRR) is 100 dB typical, 96 dB minimum. The first order High Pass filter, obtained through an external capacitor, removes low frequency baseline wandering that is so common in ECG circuits (usually due to electrodes). An anti-aliasing filter is obtained with external RC network. The PGA stage can

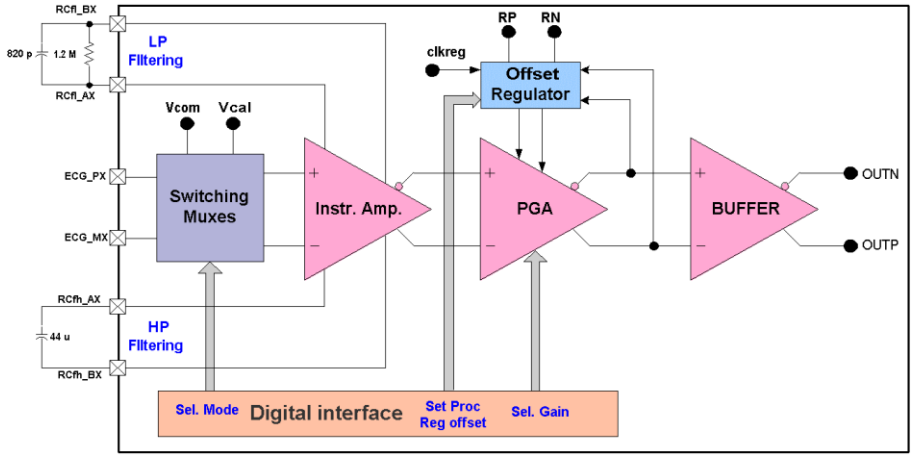

Figure 3.7. ECG signal channel for analog pre-processing.

provide four gains: 18, 24, 36 and 48, for standard ECG measurement. The IA has a gain of 15.6 set with an external resistor of 1.2 MΩ. Therefore the total gain can be configured up to roughly 700. The third stage is a buffer section with fast settling which sends out the signal to the ADC. Each ECG channel has some switching multiplexers placed in the input section and it is configurable via SPI command. The architecture of the ECG channel allows to implement the baseline fast restoring, useful any times variations in the baseline of the signals due to artefacts cause a temporary saturation of the IA. Offset regulator and gains regulator procedures are also available for each channel. If the ECG measurement is chosen, a cyclic procedure is activated and the analog multiplexer (MUX) connects this section to the A/D converter. Within 100 µs (worst case) all the channels are processed. Each ECG processing channel has been also characterized in terms of noise: the input referred noise measured in the range 0.1 Hz to 150 Hz is within 10 µVrms.
The Adder block is used to obtain the CTP signal reference to be used as inverting input for each channel that receives a pre-cordial signal. A driven right leg circuit (RL Driver block) helps to set the common mode and it is safer than connecting the right leg to voltage reference. The circuitry able to drive in active mode the shield of the ECG cables (SHIELD Driver block) helps to reduce the Common Mode 50 Hz noise.

The pacemaker detector block (see Figure 3.8) provides band pass filtering on the ECG signal, full wave rectification for detecting pacemaker pulses of either polarities, peak detection on the filtered and rectified signals, and discrimination relative to a programmable threshold level. Once a pacemaker pulse is detected, it provides on the dedicated output digital pin a pulse whose duration is set by an external RC network.

The blood pressure section includes two channels. An analog channel that converts a DC single ended, ground referred voltage to a differential voltage suitable to be converted by the on chip ADC, and it amplifies this signal. Input signal ranges between 0.5V (0mmHg) and 1.7V (300mmHg). The offset voltage for single ended/differential conversion is provided by the internal references generator block. Another analog channel amplifies the AC component of the input pressure signal and translates the DC component around a reference voltage. Extraction of AC component is assured by two external capacitors and one external resistor. If the blood pressure measurement is chosen, a cyclic procedure is activated and the MUX connects this section to the A/D converter.

The skin temperature measurement channel uses a NTC resistor as input sensor: the NTC resistor has a non linear Voltage-Temperature characteristic so an additional resistive network, showed in Figure 3.9, is used to improve the linearity of the response according to the following equation:

$$R_{LIN}(T) = (R_{NTC}(T) + R_3) \| R_1 \quad (1)$$

The chosen values in Figure 3.9 (R1=18 kΩ, R2= 3.75 kΩ, R3= 1 kΩ, NTC with 10 kΩ at 25 ^{0}C and beta of 3976, VREF_TMP_IN1&2 of 3.5 V and 1.5 V respectively) optimize the response linearity of the acquisition system in the range 32^{0}C to 46 ^{0}C, including the human body temperatures.

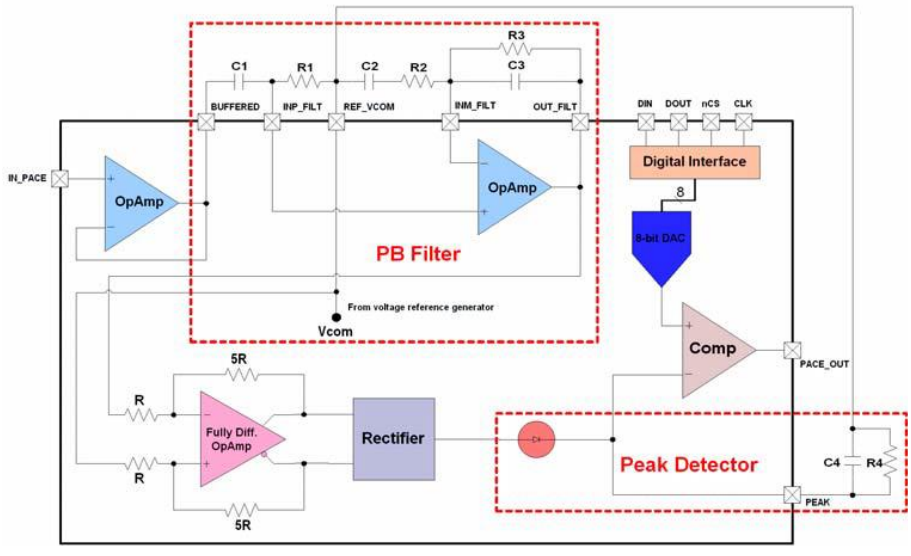

Figure 3.8. Pace maker detector block diagram.

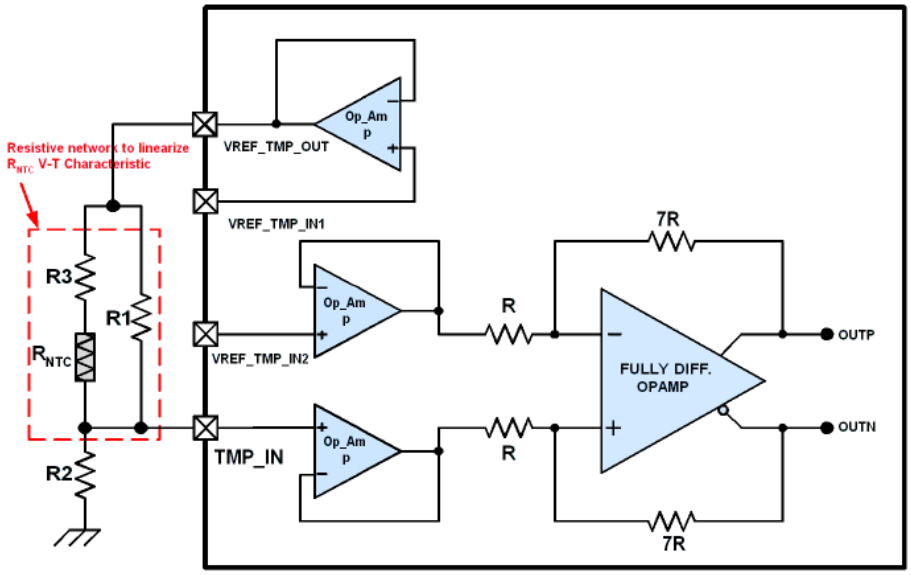

Figure 3.9. Temperature channel block diagram.

The ADC structure used in this design is a Cyclic/Algorithmic architecture with 1.5bit per cycle, made by three main blocks: a Sample and Hold stage, comparators stage and a Residue Multiplying DAC stage. ADC needs 16 clock cycles (1MHz) to produce a 16 bits output code. ADC code output is serially transmitted on output pin. The ADC has an INL of ± 1 LSB and a DNL of 0.75 LSB (worst case). The offset and error gain of the ADC are below 0.5% and 1.5% of the full scale range, respectively. Finally the battery channel in Figure 3.6 is used to monitor the battery status. To be noted that integrated bandgap circuits are used to internally generate reference levels from 1.1 V to 3.5 V.

In the H@H telemonitoring platform the CARD*IC* IC has been configured so that the ECG is acquired involving only 3 channels and the RL driver while the optional blood pressure channel is not used (blood pressure is monitored as discussed in Section 3.1.2.2). Although temperature has not been identified by the clinicians as a key parameter for CHF, the presence of this channel allows also for body temperature monitoring useful for upgrades of the H@H platform to other disease in order to create a multi-purpose configurable monitoring system.

3.1.4 H@H Sensor Signal Processing

Sensor data signal processing, implemented in the home gateway, assumes an important role within remote health monitoring system, being in charge of detecting as soon as possible alterations in the vital parameters that are potentially dangerous for the patient. The aim is to extract the meaningful information in the fastest way. In this way the physicians are timely alerted if

anything out of the ordinary is found and they have in the HIS all patient's data to plan the following actions.
The processing core is an Intel ATOM N450 at 1.66 GHz, providing a computational capability of about 5000 MIPS, with 512KB L2 on-chip cache memory and off-chip 1GB DDR2 memory. Its computational and storage capability are enough for a real-time processing of data collected from the sensors.
Data processing involves three main steps:

- pre-processing: raw data provided by the sensors are filtered to remove noise and main interferences (i.e. power line or movements artefacts). Extraction and gathering of derived information is also accomplished (i.e. heart rate from ECG track). Resulting low quality signals arising from sensors misplacement or corrupted by heavy motion artefacts are detected and the system asks for the measurement repetition;
- analysis: the results of the previous block are compared with the thresholds that establish the admissibility range for punctual values or trends over a medium period. All values must lay within the safe zone defined by the clinicians. Maximum, minimum and average values are checked by an expert system;
- false positive avoidance: to reduce the number of alarms generated for example due to temporary stress, in case of abnormal value the same measurement is shortly deferred and only if confirmed the alarm is raised. Both observations are stored and tagged accordingly.

The signal processing chain is implemented in C language leveraging respectively VSIPL library, libbluetooth, libSSL, libXML, libgtk to support filters implementation, Bluetooth management, security, XML parsing and graphical development. The main concerns of the system, data acquisition, processing algorithms, data transmission, data presentation and scheduling of activities, were implemented in separated threads to ensure a high modularity and the maximum flexibility. A permanent storage of pending data waiting for transmission avoids that power supply failures will result in data loss. The gateway provides also the possibility to store in a local database all collected vital signs occupying less than 1 Gb for one year observation.

3.1.4.1 ECG

Electrocardiogram (ECG) signal is, obviously, one of the most significant and reliable sources of information for CHF patients monitoring. Expert cardiologists within the project have considered that special attention should be given to four signs of heart deterioration:

- Abnormal Heart Frequency, above 120 or under 50 beats per minutes
- Emergence of atrial fibrillation episodes
- QRS Complex of more than 120 milliseconds with complete left bundle branch block morphology
- Signs of myocardial ischemia

Early detection of these symptoms is decisive in the prevention of cardiac threats. Home monitoring, such as the one developed in this project, allows physicians to periodically check the ECG of the patients avoiding unnecessary visits neither to the patient's home nor to the clinic. But this can still be

improved. Making a prior study of the ECG at the patient's house makes it possible to repeat those measures that obtained unexpected results and to raise early alarms for the physicians to check certain patient's ECG sooner. At this project, the ECG analysis is separated in two parts. The first one extracts basic features for heart rate calculation or respiration analysis. The second part uses the extracted information to evaluate the presence of atrial fibrillation episodes.

3.1.4.1.1 Basic Feature Extraction

The starting point for any ECG analysis is feature extraction, especially the QRS Complex detection that is nearly always the reference point within the ECG signal. A first derivative based algorithm [89] combined with a rule based system [90] is used for the QRS complex detection. Figure 3.10 shows the main points of the ECG and the steps of the processing algorithm. The signal is filtered using an optimum Kaiser filter band pass 85th order symmetric FIR in order to keep just the central frequencies (8 to 30 Hz) where the QRS complex information lays. This kind of filter has linear phase, constant delay and it is reasonably sized. Afterwards, the signal is differentiated, squared and averaged by using a rectangular sliding window. Comparing the averaged signal against a threshold creates a set of windows that allow recognizing the R peaks in the filtered signal (maximum positive within the window). The R peak has still to pass through a rule based system that evaluates whether the detected QRS is a valid QRS complex or not basing on the distance in time between consecutive peaks: i.e. if two peaks are closer than the refractory

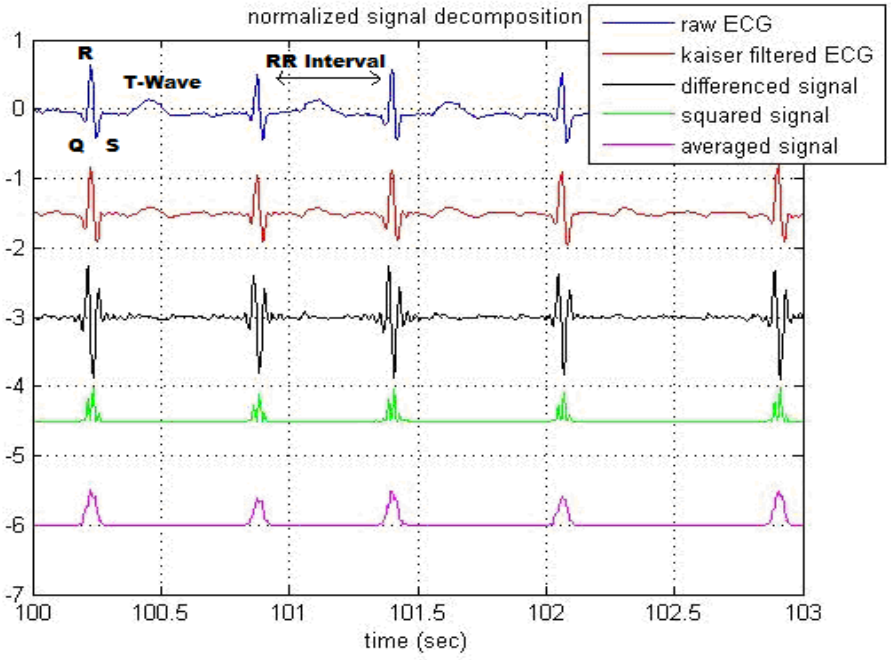

Figure 3.10. ECG main points and signal processing for QRS detection.

period of the myocardium (200ms), one of them is discarded. Then the T-wave discrimination is performed, being stricter about those peaks situated 200 – 360ms later than an accepted R peak. The rule system is also used to check for possible missed peaks when the current RR interval is 1.5 times the previous RR interval. In addition a 50Hz notch filter is applied to the signal. This filtering does not affect the detection but it improves the legibility of the track for further medical review. Figure 3.11a and Figure 3.11b show the raw and filtered signal. There is also a cubic spline data interpolation algorithm that extracts the envelope of the R peaks as an indication of the respiratory activity. Figure 3.11c shows the envelope signal of the peaks. Maximum, minimum and average heart rates along the track are calculated (RR interval) and analyzed using a 30 seconds window that is shifted along the time axis on 5 seconds length steps. The VSIPL library linear FIR function was used to implement the Kaiser filtering, and additional IIR filtering and cubic spline interpolation functions have been developed based on the basic functions of the library.

3.1.4.1.2 Abnormal Heart Frequency

At the basic configuration, Heart Frequency is calculated using a 30 seconds window shifted along the time axis on 5 seconds length steps. For each step, the number of beats within the window is counted and that value is extrapolated to a 60 seconds window to the beats per minute value. Maximum, minimum and average heart rates along the track are calculated. If

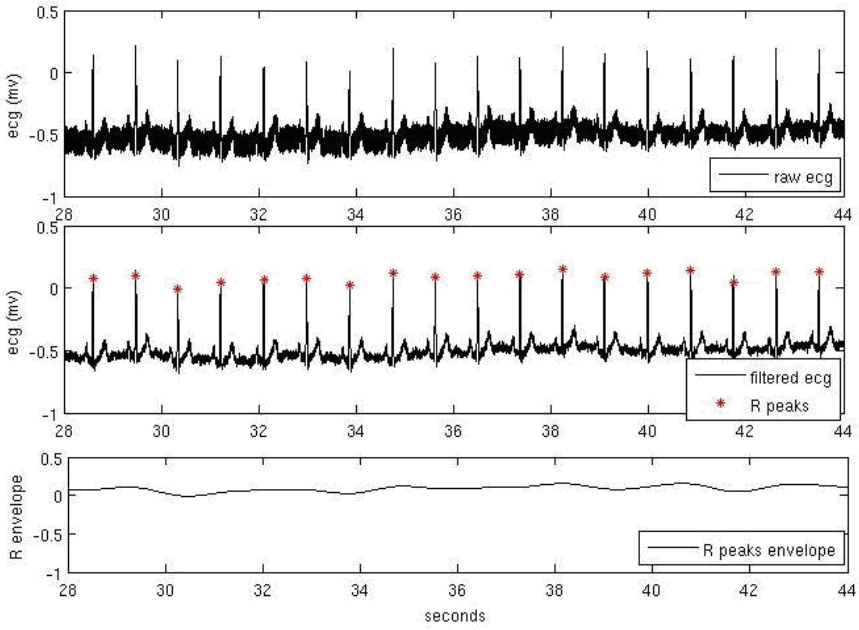

Figure 3.11. a) ECG signal; b) filtered signal plus R peaks; c) R envelope signal.

these values goes higher or lower than the limits established by the physicians, an alarm is raised.

3.1.4.1.3 Atrial Fibrillation

Atrial fibrillation (AFIB) is one of the most common arrhythmias, especially between elder people. It is the cause of approximately one third of hospitalizations for cardiac rhythm disturbances. An estimated 2.3 million people in North America and 4.5 million people in the European Union suffer from AFIB and with the aging of the population the number of patients suffering of this condition is definitely going to increase. AFIB is a heart condition where the activation in the atria is fully irregular or chaotic. This might be caused by an irregular focal triggering mechanism involving automatic cells or multiple reentrant wavelets that randomly excite tissue that has previously just been activated by the same or another wavelet. Irregular atrial activity is reflected in the ECG as the absence of P waves before the QRS complex and fluctuating waveforms in the baseline. Organized ventricle activity allows QRS complex to maintain its usual shape, but with irregular rates so RR interval will show sharp variations, see Figure 3.12.

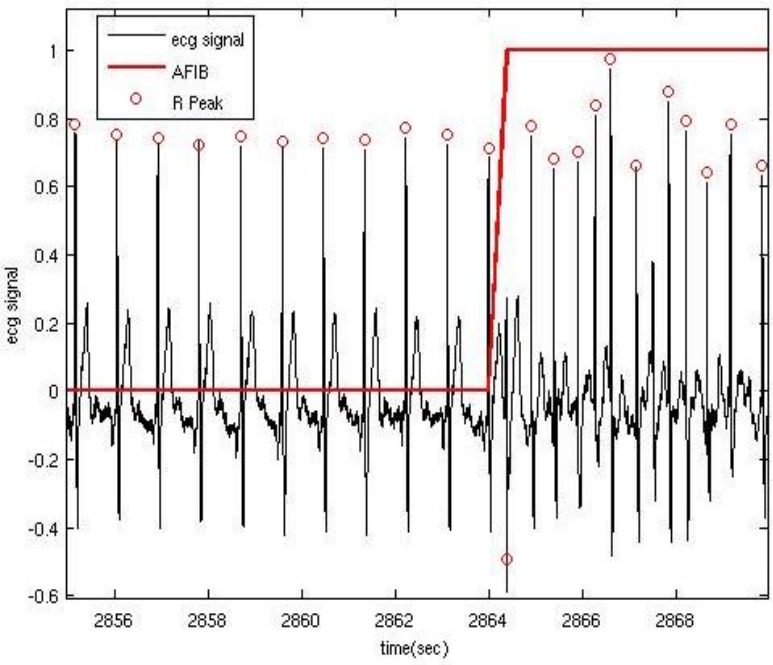

Figure 3.12. Atrial Fibrillation detection.

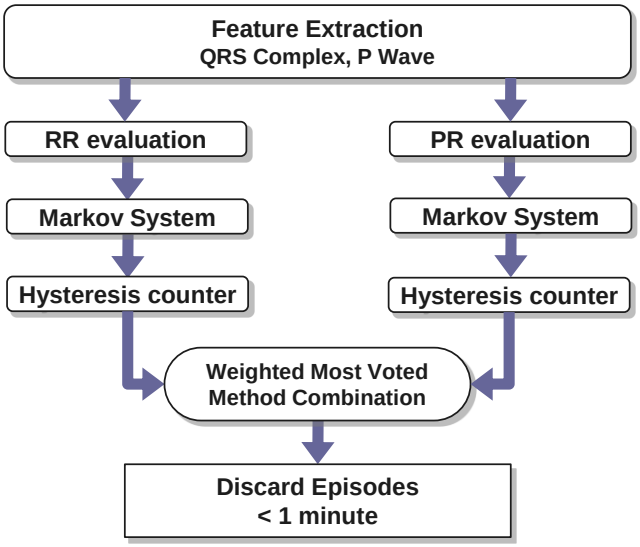

Figure 3.13. AFIB detection Flow chart.

There are multiple approaches to AFIB detection. Many of them are based on RR interval variability and its statistics: in [91] the authors compare the standard density histogram of the RR interval length (RR) and RR difference (ΔRR) with previously compiled standard density histograms segments during AFIB. In [92] the previously calculated histograms of the normalized ΔRR for AFIB and non-AFIB episodes are used to calculate the probability that a sequence of consecutive RR is part of an AFIB episode. They compare the histograms of the normalized ΔRR for both AFIB and non-AFIB episodes and demonstrate that they can be approximated by Laplace and Gaussian distributions. In [93] a linear discriminant classifier fed with RR intervals statistics is used. A method based on the random characteristics of the atrial fibrillation intervals and the Shannon Entropy is discussed in [94]. Both [91] and [93] need atrial fibrillation templates or training data while [94] is a purely statistic method that does not need any training data.

Another common approach is to separate the atrial activity from the ventricular activity in order to analyze the atrial behaviour. It is necessary to take into account that atrial and ventricular activity occurs in the same frequencies and sometimes, as during AFIB, at the same time. There are multiple documented methods to do this. Some are based on detecting and subtracting the QRS-complex. Others are based on representing the ECG signal in a distributed way (PCA, wavelets, etc) that allows the atrial activity to be separated. [95] presents a solution based on Principal Component Analysis (Karhunen-Lòeve Transform). After calculating first 12 coefficients of KLT of the V1 lead it concludes that coefficients 1-2 contain ventricular information, 3-8 contain atrial activity information and 9-12 contain noise. [96] reconstructs the atrial activity using components 3-8 and after makes a study in frequency of that activity. [96] presents a similar solution based on

AFIB	S	R	L
S	0.26	0.18	0.25
R	0.50	0.64	0.57
L	0.23	0.18	0.20

Non AFIB	S	R	L
S	0.22	0.04	0.17
R	0.26	0.95	0.62
L	0.50	0.01	0.20

Table 3.4. Transition probability matrixes (TransMat) for RR interval.

wavelets. Atrial activity is mainly contained in 4-9 Hz frequencies and can be isolated using certain coefficients of the stationary wavelet transform (SWT). Other authors mix classic RR interval analysis with additional indicators to decrease the number of false positives detected with the previous methods. [97] proposes a method based on RR intervals mixing it with P-wave detection. It also includes a hysteresis counter to determine the beginning of an episode only if several consecutive sets of parameters have been classified as AFIB. It uses Markov process modelling to analyze RR interval regularity. P wave detection considers P wave location (PR interval variation) and morphology.

For this project, a variation of the approach in [97] has been adopted. RR interval variation and PR interval variation are used to detect the presence of AFIB episodes. As shown in Fig. 12 each of them is modelled as a separate three state Markov process. Each interval is evaluated to decide whether it is short (S), regular (R) or long (L) with respect to the mean of the interval length (IntLen); interval means (IntMeam) will be determined recursively following the relation:

$$IntMean = (mw) \cdot IntMean(i-1) + (1-mw) \cdot IntLen(i) \quad (2)$$

The Mean Weight (mw) parameter determines how sensible should the mean be to the variations of the actual interval. Any Markov process is characterized by two transition probability matrixes, see Table 3.4, where the probability of transition from one state to another is represented for each of the studied cases (AFIB or non AFIB beat).

The best feature of Markov processing is its capability to evaluate a sequence of events before reaching a conclusion. In this case it recursively evaluates the probability of a beat belonging to an AFIB episode considering the state transitions of the previous *chainLength* beats.

$$Initialcondition: proAFIB(i) = 1; proNonAFIB(i) = 1;$$

$$recursivecondition: for\ k \in (2, chainLength)$$

$$proAFIB(i) = proAFIB(i) \cdot transMatAF(RRcode(k), RRcode(k-1))$$

$$proNonAFIB(i) = proNonAFIB(i) \cdot transMatNonAF(RRcode(k), RRcode(k-1))$$

Finally, the Markov score for a single beat will be:

$$score(i) = \log\left(\frac{proAFIB(i)}{proNonAF(i)}\right) \quad (3)$$

Beat to beat AFIB detection has a huge problem with very short detected episodes. A three beat fibrillation episode can hardly be of any interest. In order to avoid false positive detections, two symmetric independent hysteresis counters have been implemented. They evaluate separately the possibility of a beat being AFIB giving a score from -10 to 10. Table 3.5 summarizes the results obtained with an external evaluator (Epicmp function included in physioTools [102]) excluding from the statistics all episodes shorter than 1 minute, as they are discarded in the application. Sensitivity gives an idea of the accuracy of the system, while Positive Predicted Value gives an idea of the robustness of the system against false alarms (false positives). The obtained results show the robustness of the system. It is important to remark that this algorithm is only used to give additional information and warnings to the medical personal of the project, not for diagnose.

3.1.4.2 SpO2 and Plethysmographic wave

Oxygen saturation indicates the percentage of haemoglobin molecules in the arterial blood which are saturated with oxygen. Concerning SpO2 signal analysis, normal values in healthy adults are between 94% and 100%. A level outside this range could be assumed as warning situation. A frequent monitoring is necessary to avoid these changes to go undetected. Typically SpO2 is acquired once or twice per day and each acquisition lasts some minutes. The analysis step for SpO2 has in the worst case a complexity linear with the number of samples both in computation and memory. It involves maximum and minimum level over the track and also the average value, extracted by an on-line digital low-pass FIR filter. An example of trend of the average SpO2 value along a week is shown in Figure 3.14. In this case all values are included in the safety zone and no alarms are raised.

The plethysmographic wave is a simple way to detect blood volume changes in the microvascular bed of tissue. The waveform is composed by two main parts: a pulsatile physiological waveform due to cardiac synchronous changes in the blood volume with each heart beat, and a slowly varying (‘DC’) baseline with various lower frequency components attributed to respiration [98]. The availability of this signal in the system represents another source for the heart rate calculation and moreover leaves open the possibility to extract information about the respiration frequency in non-invasive fashion. There is also in literature some evidence that the blood pressure could be extracted

		RR Int	RR int + PR int
Gross Duration	Se	94.75 %	95.47%
	ppv	94.92%	95.44%
Gross episode	Se	93.30%	92.72%
	ppv	80.11%	80.73%
Average duration	Se	96.17%	96.06%
	ppv	69.57%	71.78%
Average episode	Se	91.87%	92.39%
	ppv	77.04%	78.26%

Table 3.5. Statistic results (Se: sensitivity = TP/(TP + FN), PPV: Positive predictive Value = TP/(TP + FP)).

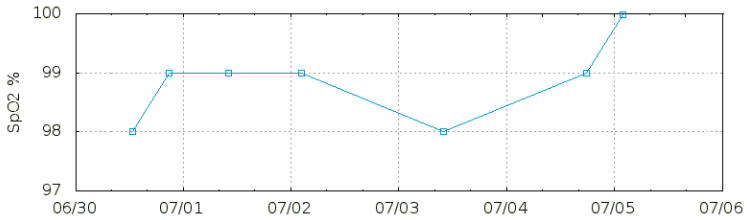

Figure 3.14. Example of SpO2 average level trend (29th of June to 6th of July).

from this signal [99].

3.1.4.3 Blood pressure

Blood pressure (BP) is the pressure exerted by circulating blood upon the walls of blood vessels. During each heartbeat, BP varies between a maximum (systolic) and a minimum (diastolic) level. In this segment of population the abnormality of punctual values of BP along with its variability in a short period are the main manifestations of cardiac instability. For these reasons more measurements per day are suggested. The systolic and diastolic punctual values provided by the sensor are analyzed to find under or over threshold situations and the general trends of both parameters are verified looking for suspicious variability. The complexity in terms of memory and computation is linear with the number of values considered. An example of observation of BP, one month long, is shown in Figure 3.15 along with the safety thresholds. It is possible to observe an under threshold and an over threshold situation respectively for the diastolic and systolic parameters.

3.1.4.4 Weight

This parameter is related to the body mass. As far as weight signal processing, it is very easy to measure as well as very effective in the CHF

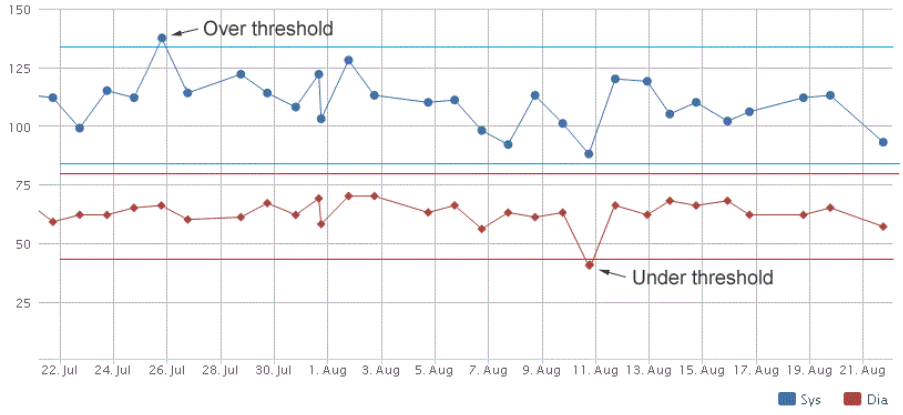

Figure 3.15. Example of blood pressure trends, comparison thresholds defined in the OP are also visible (1 month).

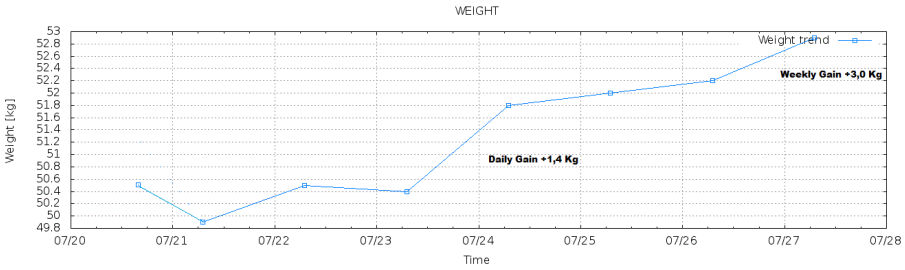

Figure 3.16. Example of weight trend from 20th to 28th July.

management. A rapid gain of this parameter means fluid retention in the patient, and it is one of the cardinal and dangerous manifestation of CHF. Frequent monitoring of the weight trend allows for simple detection of a gain. Increases of 1 Kg in a day or 3 Kg in a week is generally considered as alarm situation. The system is able to store a sufficient series of punctual values sent by the scale in order to extract the trend. To avoid impossible values of weight, i.e. other people uses the scale, a difference of more than 3 kg respect to the last observation leads to discard it. Figure 3.16 shows an example of weight trend where the signal processing step finds out two dangerous situations: a gain of 1.4 Kg within 24 hours and the second one due to 3.0 Kg gain in the last 7 days.

3.1.5 Testing and Results

Exhaustive testing of HW/SW platforms is a key issue for their adoption in telemedicine systems [100],[101]. Each developed component passed the unit test. From the hardware point of view, the novel ECG-SpO2 module was subjected to the electrical certification and the home gateway was tested after the replacement of the native keyboard with the custom membrane keypad. The operational correctness of the software elements has been verified using conventional software validation procedures. Both the sensor firmware and the multithreading application software executed on the home gateway have been heavily tested, with particular attention to the concurrency and the responsivity. Moreover all the signal processing algorithms were tested making use of signal simulators (i.e. ECG simulator) and manually generated values. Additionally, the QRS detection algorithm has been tested using PhysioToolkit and PhysioNet databases [102], obtaining a 99.7% of Sensitivity and 99.8% of Positive Predictive Value on MIT_BIH_Arrythmia database [103] and also on MIT_BIH Atrial Fibrillation [103] obtaining a 99.54% of Sensitivity and 99.36% of Positive Predictive Value. The following integration test verified the end-to-end communication in the system, and the overall system-level interactions of the different HW and SW parts, also with the execution of ad-hoc prepared scenarios and the involvement of healthy users to test if the behaviour of the systems meets the expected one. Another important phase before the final technical demonstration was conducted enrolling for a month two patients, minimally aware of ICT and younger than the CHF average age.

The first impressions of physicians and patients coming from these tests were taken into account to tune the final version of the system.
Once completed the unit and system prototype, a technical validation of the system has been implemented involving 30 patients with CHF disease in NYHA class III and IV, with an average age of 62 years and recently hospitalized for HF. The size of the set of CHF affected patients is similar to those of other works published in literature about ICT systems for CHF monitoring [104]. The minimum period of monitoring was one month. All devices were provided in a single bag, with the dimension of those carrying 15" laptop computer, and an overall weight of 2,8 kg to be transportable everywhere. Patients were enrolled in the study at time of discharge from the Hospital where they were admitted for acute heart failure or during a routine ambulatory visit. Main inclusion criteria were diagnosis of heart failure, class NYHA ≥III, at least one hospitalization for acute heart failure in the previous 6 months, agreement to take part in the study. Acute coronary syndrome within 3 months before the enrolment was the only exclusion criterion. Physicians that took part to the study were all cardiologist, involved in the management of patients with heart failure. They also checked out information arrival in HIS, evaluating the quality and coherence of data collected and the relevance of the alarms. The quality and the reliability of acquired signals and generate alarms, the robustness of data transmission and the system effectiveness from the medical perspective were evaluated as a key point of system functionalities. On the other hand, the ergonomic of patient's interfaces was evaluated, as well as the general end-user usability of the sensing devices and the home gateway. A specific testing protocol and a questionnaire have been developed to gather patients, caregivers and physicians feedbacks and validate the system.
The results show a very limited number of activity misses (<3%), mostly in the first days of monitoring, confirming also the property of such system to improve the therapy compliance. Moreover, the number of false positive alarms is less than 5%. No connectivity and transmission problems, including data lost, occurred. All end-users reported a positive feedback and good satisfaction level in the final questionnaire. The score reached for each macro-parameter is shown in Table 3.6 and Figure 3.17 respectively for medical staff and patients. Physicians reported that the use of this platform does not load up in a significant way their regular activity, but represents a

Macro-parameter	**Score**
Simply decision and increase effectiveness of diagnosis and treatment of patient based on better evidence	9.125 / 10
In general terms, easy to use with clear and understandable interactions	9.5 / 10
Flexibility of the system and compatibility with other systems already in use	9.75 / 10
Quality of the provided signal	9.1 / 10
Sensibility of the alarm detection function	9.15 / 10
In favour of the adoption of the H@H system	9.2 / 10

Table 3.6. Physicians' aggregation feedbacks.

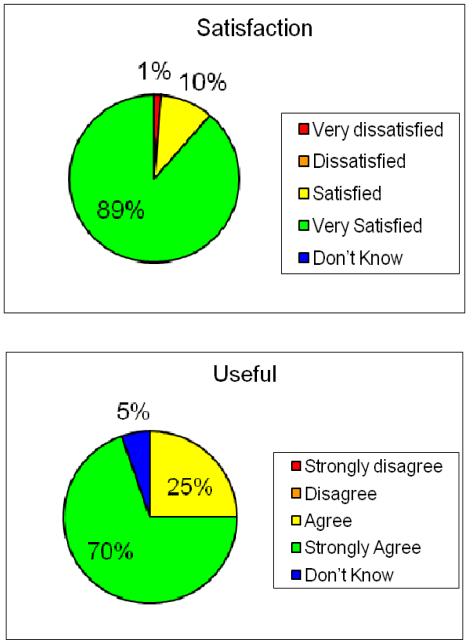

Figure 3.17. Feedbacks aggregation result of patients.

valid means to control at distance the evolution of the followed patients thanks to the high quality of acquired signals and alarms detection capability. All physicians involved in the demonstration are definitely in favour of the adoption of the H@H system. The 89% of the patients report a very high satisfaction level, highlighting the friendliness of the solution and the easiness to follow the daily therapy.
Due to the success of the H@H technology test under medical control, a clinical validation, including an economical evaluation (with OPEX/CAPEX capital and operational expenditure analysis), with more than 500 patients has been already planned in the Italian Regional Tuscany Health System.

3.1.6 Comparison with the State of the Art

This section underlines the distinctive features of the H@H platform with respect to telemedicine systems specific for CHF or suitable for this kind of disease and with some dedicated prototypes and projects. Systems based on telephone calls or web-portal [105] to report symptoms and measures have to be discarded for scalability and usability issues.
Similar commercial systems in terms of supported vital signs [106]-[108] use proprietary data formats and dedicated databases limiting the degree of interoperability and integration with existing departmental HIS, while they rely on dedicated call centres. The user interfaces and the collection gateways (i.e. PDA or smartphone) are in some cases quite complicated with respect to the patient's ability [107],[108]. Others involve 12 lead ECG prone to mispositioning issues [109]. In literature there are systems where often one

main parameter, pulse oximetry as in [110] or ECG only as in [111] or [112], is acquired (not enough for CHF). Despite the importance of the activity level, home monitoring using motion sensors [113],[114] is useful more for presence or fall detection rather than for early diagnosis of CHF. In [115] weight and blood pressure are monitored and the user has to manually type the values in the gateway for sending, while [116] is based on a T-shirt equipped with sensing areas to measure the ECG, temperature and activity level. In different ways they result scarcely usable for elderly patients and provide less vital signs that identified in the medical requirements. Our architecture differs from similar ones such as [117]. In that work only ECG and SpO2 are included. As discussed previously, this does not match with the minimum requirement defined by clinicians for clinical monitoring of patient with CHF. Moreover, the user interface running on PC is not suitable for elderly people. All previous solutions do not include advanced signal processing functions in the home gateway for early warning of the remote system. Instead they demand all processing to the remote host, applying only noise removal on the captured signals.

The chipsets for ECG front-end and parts of the signal processing (i.e. Texas Instruments [118][119] and IMEC [120]-[123]) can not be compared to H@H system since they represent only a subsystem as the CardIC sub-unit. CardIC offers a multichannel analog front-end IC for bio sensor interfacing. Texas Instruments has a family of IC ADS119x and ADS129x and ADs1258/ADs1278 with up to 8 channels (CMRR up to 100 dB, noise levels in the order of few µV), with programmable gain amplifier, 16 bit or 24 bit ADC integrating the circuitry needed to interface 1 or multiple ECG or EEG electrodes. However some external components to build up a complete monitoring system (i.e. channels for external temperature sensor or blood pressure sensors, pace maker detector) are missing. CardIC instead integrates all missing elements (see Table 3.3) reaching a deeper integration level than the TI 16-bit ADC version. IMEC has developed several versions of ICs for ECG and EEG with integrated signal processing in the analog or digital domain. Such systems are optimized in terms of low power consumption (less than 1mW with designs realized in 0.5 µm or 180 nm while CardIC is in 0.8 µm) but have only 1 channel and just dedicated to ECG while CardIC has 12 channels for interfacing ECG electrodes or other sensors. Hence IMEC design is more suited for ultra low-power wireless heart rate monitoring system for wellness or general healthcare. CardIC is more suited for clinical applications where more channels and sensors are required. H@H systems do not require ultra low power and the advanced signal processing is demanded to the home gateway, which is realized using a COTS programmable core, the Atom processor, rather than designing application specific DSP system integrating single-chip microcontroller and hardware accelerator IP cells [124],[125]. Overall the H@H system addresses the usability concerns through an easy-user interface designed to allow the patient to follow autonomously the therapy at home. It allows multi-parametric monitoring, according to the clinicians recommendations, based on easy-to-use sensing devices (i.e. 3-lead ECG, blood pressure, weight in basic configuration) and flexibility, through the operating protocol, to meet the individuality of the end users.

3.2 Energy monitoring for Smart grid

Smart Grid is the evolution of the current power grid, into a new smarter network [126],[127]. It is a modernization, a re-engineering of the electricity delivery system, through the exploitation of information and communication technologies for power system engineering. The result should be an intelligent network that can monitor, protect and optimize the operation of all its nodes, from the central and distributed generator layer to the end users [128]-[131]. The primary purpose of this innovation is to increase energy efficiency, reliability and sustainability to address the growing electricity demand, and to mitigate the climate changes reducing gas emissions. Thanks to continuous monitoring of all power grid nodes and the interconnection with classic ICT networks, Smart Grid may be used to increase the energy-awareness of the society suggesting and stimulating "green behaviors".

3.2.1 Limits of existing power grid and challenges of Smart Grid

The typical structure of the existing power grid is shown in Figure 3.18. Utility companies all around the world designed their power grid imposing clear demarcation between its main subsystems: generation, transmission and distribution systems. This approach has brought different levels of automation in the various subsystems, and each subsystem has separately experienced different evolutions and transformations. Moreover, the hierarchical structure of the grid can cause domino effect failures.

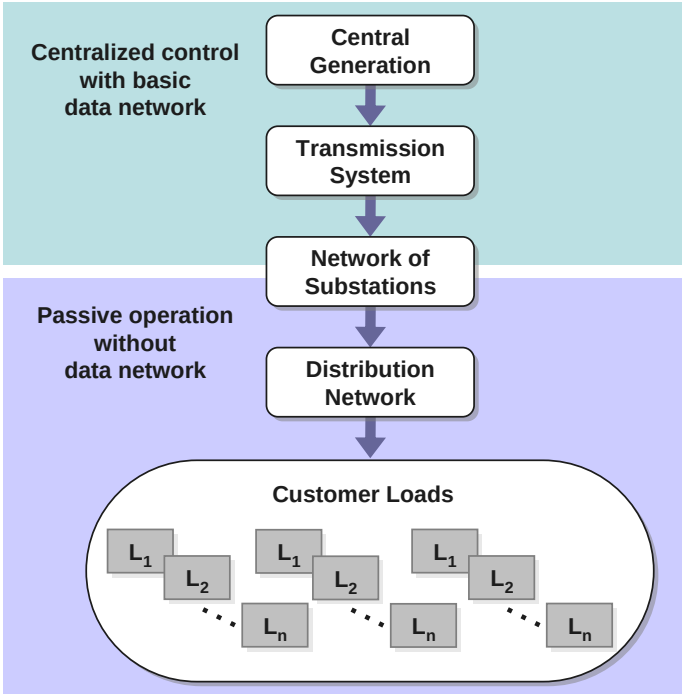

Figure 3.18. Architecture of existing power grid.

In term of efficiency, along the existing grid there is waste of energy in various forms: only one-third of fuel energy is converted into electricity (and waste heat is not recovered), 8% of the produced energy is lost along transmission lines and 20% of the generation capacity exists only to support a potential peak demand [126]. This last point is very important. The existing grids are actually over-engineered to stand maximum peak demand, which is very infrequent, limiting therefore the whole system efficiency.
Moreover, present electricity grids are mainly unidirectional: generators produce energy and distribute it to the lower level, with very few information about grid status and end-users energy consumption. Typically, the electric power source has no real-time information about service parameters of termination points, and cannot control energy production according to the real request of the grid. The new challenges for present grids can be summarized in five main points:

- Introduction of new forms of power generation, in particular those using renewable energy sources such as wind, sun and biomass. These types of generators have intermittent and small outputs, and need therefore a different management from traditional generators.
- Need of uninterrupted electricity supply.
- Need to decrease peak demands during the day, and to reduce energy waste to ensure adequate energy reserves.
- Diffusion of new digitally controlled devices able to change the behavior of the electrical load (e.g. switching itself on or off), smart power meters, and energy control units implementing energy management strategy and improving energy-awareness of users.
- Security threats, that involve not only the electricity supply but also cyber attacks [133],[134].

The evolution towards the Smart Grid begins with innovations in the existing grid by incorporating new ICT technologies in many point of the infrastructure. Table 3.7 shows the main differences between existing grids and Smart Grid. The starting point of this revolution is the bottom layer of the system: the electrical distribution subsystem. The first step is the insertion of distributed and networked monitoring and control systems in the electrical grid. Such systems can assist utility companies in grid monitoring and can identify potential risks, taking corrective actions in time. Secondly a complete overhaul of the ICT infrastructure is required. Communication and data management will provide a layer of intelligence over both the existing grid and the future infrastructure, allowing the introduction of new applications. This organic growth of the grid allows companies to gradually shift old grid's function into the new grid and consequently to improve their critical services.

Existing Grid	Smart Grid
Electromechanical	Digital
One-way communication	Two-Way Communication
Centralized Generation	Distributed Generation
Hierarchical	Network of Networks
Few Sensors	Sensors Throughout
Blind	Self-Monitoring
Manual Repairing	Self-Healing
Failures and blackouts	Adaptive and Islanding
Manual check/test	Remote Check/Test
Limited Control	Pervasive Control
Few Customer Choices	Many Customer Choices

Table 3.7. Smart Grid innovations vs. existing power grid.

The change of the unidirectional approach of the classic power grid with the bidirectional one introduced by the Smart Grid concept, can favor the diffusion of distributed generators or co-generators, along the existing grid. Indeed, Smart Grid can provide an easier integration of alternative sources of energy (i.e. sun, wind, etc.) characterized by time-varying energy production level with storage systems, in order to fill the gap between when/where the energy is produced and when/where the energy is required. Smart Grid can aid utility companies to make a more efficient use of the existing infrastructures, introducing step by step some key features as: demand response, peak shaving and service quality control [126].
The evolution of the grid requires the coexistence between Smart Grid and existing grid [135] . This permits the gradual growth of the grid, increasing step by step its capacity and adding new capabilities. A way to perform this evolution is the introduction of microgrids [136] which are networks of distributed energy systems, loads and generators, that can work connected to the grid or not. They can be, for example, houses or factories, having their own local energy source, that want to optimize their energy consumption.
An example of a microgrid has been developed at University of Pisa in the framework of the Nanocatgeo project [137] in collaboration with industrial partners such as Acta Energy and Edi Progetti, see Figure 3.19. The idea was to develop a micro Smart Grid for wind-based energy autonomous homes located in windy zones. In the system a wind energy source (non constant energy production) plus an AC/DC converter provides energy on a DC bus wherein are also connected: (i) an inverter DC/AC system to power the home or sell any excess production to the energy utility company; (ii) an AC/DC converter to supply an Hydrogen Electrolyzer [138] to store in the form of hydrogen any excess when the wind electrical energy production is higher than the users' needs and to obtain energy back from hydrogen when the wind energy production is lower than users' consumption. All the subsystems of this Smart Grid are interconnected (wired) through a control box, implementing energy management strategies; the control box is based on an ATmega microcontroller core plus RS485 and CAN interfaces.

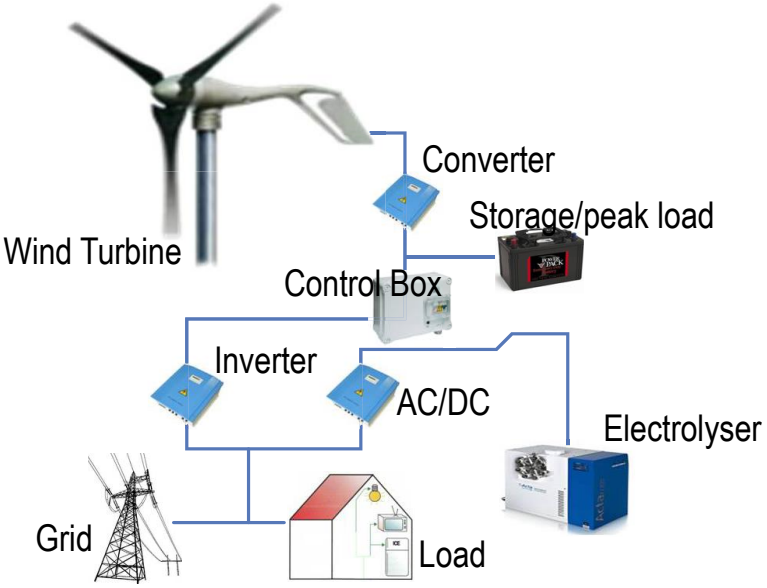

Figure 3.19. Nanocatgeo microgrid developed at University of Pisa.

3.2.2 Network architecture

Smart Grid can be viewed as a network of networks, see Figure 3.20. Starting from the customer side, the HAN is the network of communicating loads, sensors and appliances within the customer's premises. Customers are connected to the energy distribution level through a LAN. LAN identifies the network of smart power meters, gateways and elements in the distribution system. Last, we find the WAN. This is the network of upstream utility assets that include power plants, substations, distributed storage and so on. Substation gateways interface WAN and LAN networks.

3.2.3 Security and privacy problems

Since the existing grid is moving from a centralized network to a dynamic

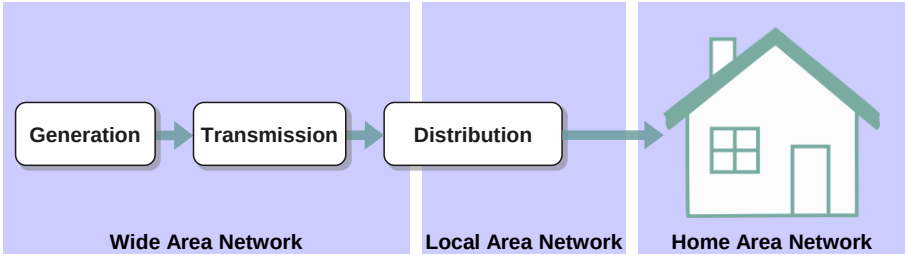

Figure 3.20. Smart Grid network hierarchy.

peer-to-peer network, with a growing complexity, it is also becoming more vulnerable to local and global disruptions. Smart endpoints introduced into the network become portals for intrusion and malicious attacks. Moreover, Smart Grid is growing over systems not designed with security criteria, thus with significant security holes [133],[139],[140]. Security problems do not involve only cyber security aspects, but it concern also failures in the grid and protection against natural disasters. Following, a list of potential risks for a Smart Grid:

- the complexity of the grid increases accidental errors and possible points of intrusion;
- the deployment of new technologies can introduce new issues in the network;
- the presence of many network links increases potential cascading failures, and gives more opportunities to compromise the system;
- smart nodes can be vulnerable entry points for Denial of Service (DoS) attacks.

Particularly, the focus of Smart Grid security is on the HAN: indeed WAN and LAN in Smart Grid are known computer networks whose security issues are widely discussed in literature. The HAN network is deployed into the customer domain, and its security is a critical point strictly related with customer's privacy.

A typical HAN is composed of four elements:

- A gateway that connects the HAN network to the outside information services, in the LAN or WAN network.
- The access points or network nodes composing the HAN network.
- A network operating system and a network management software.
- Smart endpoints, such as: smart meters, displays, refrigerators, appliances and thermostats.

So far, many technologies have been considered in order to implement the HAN by different groups and organizations. The most significant standards are ZigBee [141], Z-Wave, INSTEON and Wavenis. They are all standards for wireless networks. The main features of these standards are presented in Table 3.8. Talking about security, if present, similar encryption algorithms are used by them: AES and in some cases 3DES. AES is the most reliable

	ZigBee	**Z-Wave**	**Insteon**	**Wavenis**
RF band (MHz)	868/915/ 2400	868/908/ 2400	904	433/868/ 915/2400
Range (m)	10-100	30-100	45	200-1000
Bit rate (Kbps)	20/40/250	9.6/40/200	38.4	4.8/19.2/ 100
Message size (bytes)	127	64	14-28	NA
Security algorithm	128b AES	128b AES	NA	3DES/128b AES

Table 3.8. HAN standards and security algorithm, main characteristics.

encryption algorithm between them [134] and its implementation (hardware and software) offers better performances than 3DES. Moreover, AES encryption can be performed using ad hoc AES hardware, i.e. an AES coprocessor, which is present in many of the solutions proposed for implementing HAN networks. In the example network, ZigBee standard will be used. ZigBee security issues will be discussed in detail in Section 3.2.4.3.

From the customer point of view, a fundamental requirement is the protection of the information exchanged between the utility company and the smart power meters installed at the customers' premises. Far from old electro-mechanical measuring systems, the new generation of power meters is fully electronic [142]-[147], they provides advanced power measurement and management capabilities thanks to power ASIC provided by semiconductor suppliers like STMicroelectronics [148]. The new generation of smart meters integrates a two-way communication system. In particular, power consumption data are transmitted over low voltage power lines, using packet-switched digital Power Line Communication standards [149], from the customer's premises towards data concentrators, based on Echelon technology [150]. On the other side, from each data concentrator point, information are sent to the servers of the utility company using the Internet network. Vice versa, the utility company can easily operate on remote smart meters by accessing through internet the data concentrators and from them, though Power Line communication over low voltage residential power lines, the smart meters at the customers' premises. This way the utility operator can turn power on/off to customers, read usage information, change customer's billing plan, and also detect service outages or unauthorized electricity use.

Power line communication is based on the following idea. AC power is transmitted over high-voltage transmission wires at 50-60 Hz, so it is possible to impress a higher frequency signal carrying digital information in both directions (from customers' premises towards the utility company and vice versa). The carrier used for data transmission in power line communication has generally a frequency of about 100-200 kHz, for data rates of few Kbps, so that data signals can be easily separated from power ones. More details on power line communication in Smart Grid, and the relevant packet formats and standards can be found in [150].

However, consumption records obtained through the smart meters can reveal a lot of information about customer's activities, thus it is important to satisfy some requirements in terms of *confidentiality*, *integrity* and *availability*.

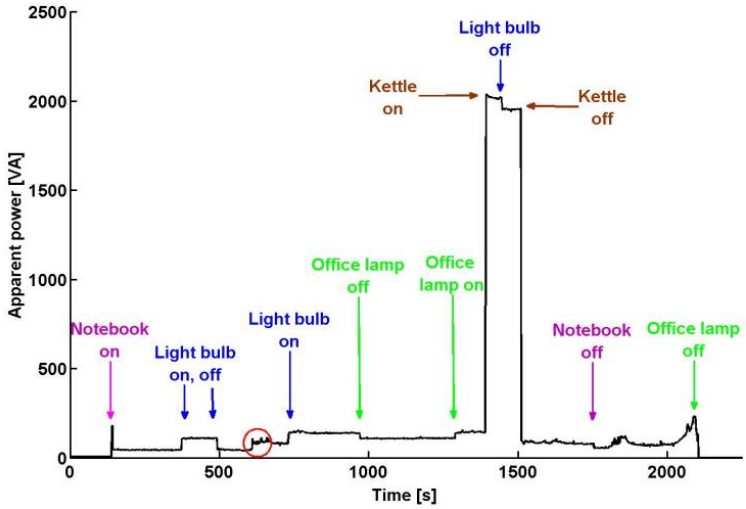

*Figure 3.21. House electricity demand and information extracted: apparent power (Volt*Ampere) for notebook, lighting sources and a kettle [152].*

Confidentiality deals with information protection from unauthorized access. It is a customer side requirement. In Smart Grid, the focus is on data stored in the utility companies servers and transmitted from customer's smart meter. These data contain energy usage information and billing data. To protect them, it is important to properly implement the least privilege principle: a user has no more privileges than necessary to perform its function. A detailed guide for implementing security in the organization data server, that follows this principle, can be found in [151]. The protection of these confidential data assures customer's privacy. In fact, energy usage information reveals user activities during the day, allowing to deduce what kind of device or appliance was in use at a given time. In literature there are a lot of load signature algorithms, results of NALM (Non-intrusive Appliance Load Monitoring) research branch [147],[152], that can extract detailed information from electricity usage records. An example of the results of these algorithms, over real consumption records, is shown in Figure 3.21 and Figure 3.22. Signature of the usage of a specific power appliance can be extracted considering typical power consumption of the load, frequency of use, transient response since different appliances have different load types: non linear resistors for heaters and bulb lamps, inductive for electric motors, reactive for microwave oven, diode-like for led lights.

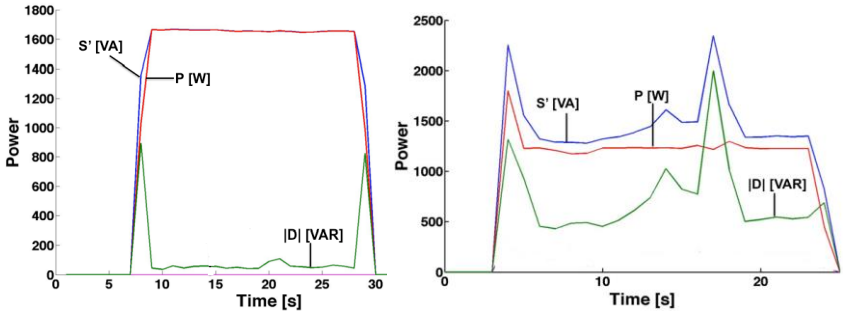

Figure 3.22. House electricity demand and information extracted: real power P, apparent power S and reactive power D of an iron (left) and of a washing machine (right) [152].

Integrity ensures the correctness of information protecting data against modification attacks. A countermeasure to prevent this type of attack is based on the access control. With this, only authorized users can modify the information.
Availability ensures that services are always available to users. Security must prevent out-of-service due to human factor, or DOS attack against utility companies that can compromise power distribution. Redundancy is a good practice to prevent environmental threats.

3.2.3.1 Practice to secure the HAN

In information technology, there are a lot of codes and rules in order to achieve the security requirements emphasized. An example is ISO/IEC 27000 series, a set of standards of certified best practices for information security [153]-[155]. Another security program, not certified, is the Information Security Forum (ISF) [156]. ISF is a no-profit organization that distributes the Information Security Forum's Standard Of Good Practice free of charge.
These guides can be applied to ensure information security for every kind of systems, including Smart Grid, and particularly for the HAN. A code of technical practice for security in the HAN of a Smart Grid can be summarized in the following twelve points.

1. *Threat Modeling.* Possible threats must be identified, for preparing proper countermeasures. This study can be conducted analyzing use case vs. abuse case.
2. *Segmentation.* To minimize the impact of attacks, segmentation can be adopted limiting, for example, data traffic in specific area through a firewall: attack damage would be confined to such area.
3. *Firewall rules.* Proper rules must be used for firewall, proxy server and content filtering.
4. *Signing.* Software codes running in the grid have to implement digital signing. This allows the execution only of trusted applications, and ensures integrity of the information exchanged within the Smart Grid.
5. *Honeypots.* The deployment of honeypots, traps for hackers, permits identification of a new type of attack and alerts organizations in time. Thus, honeypots show weakness and security hole of the system.

These elements can be placed in the Smart Grid environment and in its peripheral areas.

6. *Encryption.* Through the encryption algorithm, sensitive information are protected from unauthorized disclosure. Encryption must be adopted on the transport layer, on the archived data and in the control network.
7. *Vulnerability analysis.* Utility companies have to create control centers to analyze network traffic and systems to identify any exposures that increase vulnerability to attacks.
8. *Penetration testing.* Beside vulnerability analysis, simulating an attack usually done by a malicious hacker on the Smart Grid must be performed periodically; this way a snapshot of the effectiveness of the Smart Grid security can be obtained.
9. *Source code review.* Smart Grid applications must present no vulnerabilities. Thus, their source code must be reviewed carefully in order to meet high quality requirements. After the identification of vulnerabilities in the code, these can be fixed.
10. *Configuration hardening.* All the elements of the Smart Grid, especially smart end points, have to be tested before their deployment. This can be done with vulnerabilities scanners and benchmarking tests.
11. *Strong authentication.* There are three main types of authentication methods: something the user knows (e.g. password), something the user has (e.g. hardware key), and something the user is (e.g. biometric id). At least two of these methods must be used.
12. *Logging and monitoring*. They are powerful tools for providing information for attack identification and for reconstructing events in case of natural disasters. Starting from stored data, data mining techniques and signal processing analysis give important information about attacks and grid behavior during certain events. However, if not correctly managed, data logging could represent a further backdoor into the system. For this reason it is important to define an accurate log planning process including log management planning, policies, and procedures taking into account security issues [157].

This set of practices can be also used as a backbone for the development of future Smart Grid standards.

3.2.3.2 Security standards and proposed solutions

Several associations and groups in different countries have developed many standards for security in Smart Grid. The IEC 62351 standard, developed by the International Electrotechnical Commission (IEC), is one of them. This standard concerns power system management and associated information exchanged, and is divided in eight core standards, reported in Table 3.9. The scope of the IEC 62351 standard is information security for power system control operation [158]. In Table 3.9 other two IEC standards are mentioned: IEC 60870 and IEC 61850. The first one, the IEC 60870 standard defines system used for telecontrol. Part 5 of this standard deals with communication between nodes directly connected. IEC 61850 is an electrical substation automation standard for modeling data, reporting schemes, fast transfer of events, setting group, sample data transfer, commands and data storage.

Core standard	Topics
IEC 62351-1	Communication network and system security - introduction
IEC 62351-2	Glossary of terms
IEC 62351-3	Profiles including TCP/IP
IEC 62351-4	Profiles including manufacturing message specifications (MMS)
IEC 62351-5	Security for IEC 60870-5 and derivatives
IEC 62351-6	Security for IEC 61850
IEC 62351-7	Network and system management (NSM) data object models
IEC 62351-8	Role-based access control

Table 3.9. IEC 62351 core standards.

NIST interagency report 7628 for cyber security in Smart Grid is another important document [159]. This report contains a framework for cyber security risk management, and a list of requirements for power meter security, as well as a discussion about privacy and Smart Grid. Moreover, this document contains power system use cases for security requirements and bottom up security analysis of Smart Grid.
Besides these standards, in literature there are some solutions and models proposed for implementing security in Smart Grid. One of the most interesting solutions is based on public key infrastructure (PKI). This is based on the fact that security and privacy technologies use a key to encrypt and protect data, in order to meet the desired security requirements. The problem, in a large network as a Smart Grid, is the key management system. The PKI infrastructure proposed is composed of five main elements:

- PKI standards
- Smart Grid PKI tools
- Device attestation
- Trust anchor security
- Certificate attributes

PKI standards would be used to determine requirements on the PKI operations of energy service provider. PKI however is notoriously hard to deploy and to use, due to the fact that PKI standards provide only high level framework, and leave to companies the detailed implementation. Smart Grid PKI tools give users an easy way to manage the infrastructure, and enable the development of future applications, which meet PKI security requirements. An important feature of these tools is to eliminate the need of symmetric key configuration, which is an insecure and expensive process. In a secure system, each component must be a trusted component. Device attestation techniques are used to identify devices, and to find out if the device has been tampered. Within a network based on PKI infrastructure, an important aspect is the management of devices' certificates. These certificates can be organized in trees, and the root is called Trust Anchor (TA). It is important to secure operations on TA: loading and storing, identification, management of local policy database (a set of rules defining how a device should use its

certificates, and what type of certificates it should accept), etc. It is essential in Smart Grid that any device in the network can determine the authorization status of another device, and authenticate it. This can be done using the attributes present into the certificate, and contacting a security server. Therefore, it is important to distribute local security servers in various part of the network, and not to rely only on a back-end server (single point of failure problem).
The solution proposed is only a high level description of how security and privacy can be achieved in Smart Grid, and many problems may come out during the implementation of a PKI infrastructure. Some of these problems were discussed previously (i.e. need of distributed authentication servers, implementation of PKI standards, secure management of devices certificates, etc.).
Alternatively, the PAKE (Password-Authenticated Key Exchange) research [160],[161], explores an approach to protect passwords without relying on PKI at all. PAKE aims to achieve two goals. First, it allows zero-knowledge proof of the password. One can prove the knowledge of the password without revealing it to the other party. Second, it performs authenticated key exchange. If the password is correct, both parties will be able to establish a common session key that no one else can to compute.
As far as privacy of the customers is concerned, a possible solution is based on the anonymization of smart meters data. The idea is to distinguish smart meters data on the basis of their generation frequencies:

- High-frequency data, sent by smart meters to utility data concentrators in order to control power generation and distribution network, and to enable a real-time response to power quality. These data do not need to be attributable to a particular customer, and are sent, for example, every minute.
- Low-frequency data, sent to utility company, for billing and account management. These data must be attributable to a customer or an account, and are sent every day / week / month.

Only high frequency data are "anonymized", because of their sending rate. A smart meter, using this technique, has two ID for its message: an HFID for high frequency data, and a LFID for low frequency data. The method proposed ensures the anonymity of the HFID, thus of high frequency messages. The utility company and customers know only their LFID, the HFID is known only by the manufacturer of the smart meter, that so it is the only one that has the correspondence between LFID and HFID. The HFID for example can be hardware encoded. This solution however considers only data sent from smart meters. The limits in terms of security of an approach similar to the proposed one (HFID and LFID correspondence is known by the smart meter manufactures and stored in its archives) are discussed in [162].

3.2.4 Home Energy Network Possible Implementation

This section presents a possible implementation of a home area network for smart energy management, discusses security issues and analyzes some commercial hardware/software solutions for its implementation. The network proposed is derived from the experience gained in Smart Grid projects proposal in Italy such as the Energy@Home project [163], carried out by

industrial partners such as Electrolux, Enel, Indesit and Telecom Italia, and the SEAS proposal, by Italian academic partners. The aim is to develop a communication infrastructure, for exchanging information related to energy usage, consumption and tariffs in the home area network.
The general architecture of the smart Energy HAN is presented in Figure 3.23. The HAN network contains a smart information box called Home Energy Angel, realized as an electronic control unit with on-board memory, computing capabilities (32-bit microprocessor with non-volatile, SRAM and SDRAM memories) and digital networking interfaces.
The Home Energy Angel box implements these main functions:

- collecting data from the power meters and from the smart endpoints in the home domain, monitoring the energy sources (from the electricity provider, or from local renewable energy sources such as photovoltaic panels or wind-based systems), the energy loads (recharge point of electric vehicles if any, lighting, air conditioning, household appliances and infotainment devices) and the energy buffers (Li-ion batteries or H2-based energy storage [138]);
- collecting data through the HAN from environmental sensors (temperature, light, humidity);
- forecasting of users' needs, based on data provided by sensors and

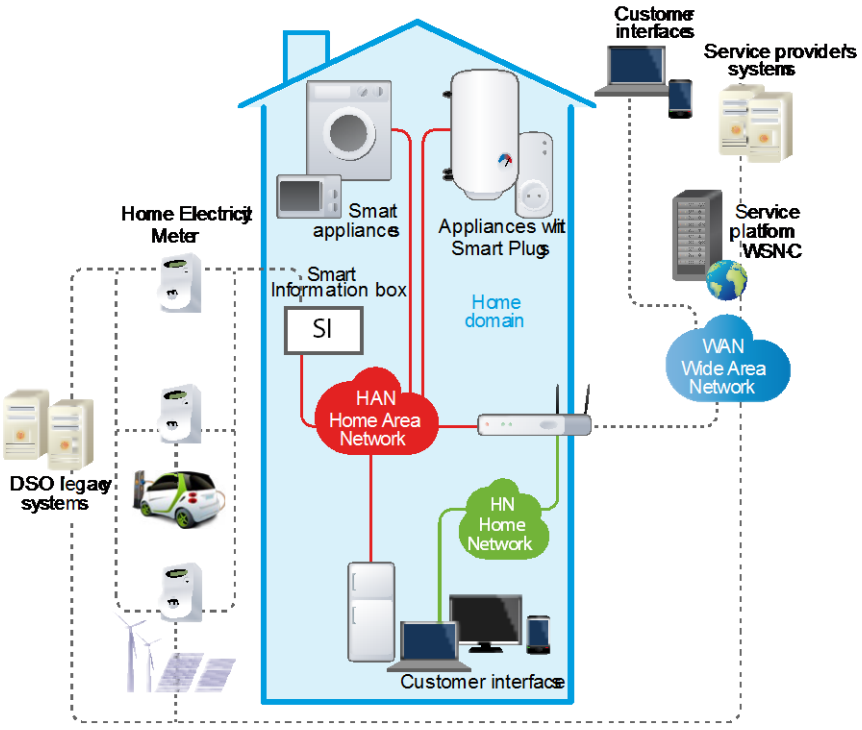

Figure 3.23. Smart Energy HAN general architecture.

by profiling methods;

- sending commands to smart appliances according to pre-programmed strategies to implement power saving strategies (e.g. turn-off/on lights adaptively on the environment conditions, proper time programming of washing machines or oven to avoid peak consumption,..),
- providing information to the users about their energy behavior through their tablet PC or smart phones .

Beyond the Home Energy Angel box, a home gateway is also connected to the HAN. This provides internet access for users through a Wi-Fi network. The home gateway is the interface between the HAN and the WAN network (Internet in this architecture). Users can obtain information about home consumption contacting the Home Energy Angel information box through the home gateway, using a simple internet connection, or locally using the connected interfaces on their tablet, laptop or smartphone. The Home Energy Angel smart information box provides energy services to make customers aware about their energy consumption. These services are automatic load management, energy efficiency, active demand service and networking with smart appliances.

A Graphical User Interface (GUI) will be developed for the Home Energy Angel, enabling a better user experience of the whole system. The GUI will be designed for two main purposes. Firstly, users will be able to easily provide information on their preferences in using the energy at home (i.e. on the appliance that they are willing to use, on the time window to start/stop each device). Secondly, it will be used to visualize the optimal energy plan calculated by the Home Energy Angel and to access additional information such as the load consumption profile, costs or statistical data, thereby settling the general lack of awareness people have of their energy consumption.

The proposed Home Energy Architecture will provide benefits in terms of:

- Provide an easy-to-use support to optimize the production and consumption of electricity, reduce electricity cost and minimize electricity waste.
- Increase user awareness on energy consumption/saving.
- Improve the grid efficiency by leveling peaks in the demand.

Energy consumption information, collected by meters, can be sent directly to the Distribution System Operator servers (DSO legacy systems in Figure 3.23) exploiting Power Line Communication protocol instead of using the HAN connection. In the architecture detailed in Figure 3.23 there is another particular element: the smart plug. Smart plugs (or home plugs) are systems able to add intelligence to old generation devices. They are simple socket points with a wireless connection (for example ZigBee) providing consumption monitoring of the connected devices. Smart plugs can also control the status of the connected devices (powering them on/off) and share the power among them.

The system architecture of the network is shown in Figure 3.24. A possible implementation of this system uses a ZigBee network for realizing the HAN. Within the network there is a smart information box, connected to the HAN through a ZigBee transceiver, and equipped with a Wi-Fi interface for contacting the home gateway. The home gateway is a simple Wi-Fi router. The role of the smart information box is to collect data from the HAN, to compute them using information coming from the Internet network (customer's tariff, billing account information) and to present them by means of a user-friendly interface.
The home gateway acts as an interface between the WAN (Internet) and the HAN. A Wi-Fi router can play this role: the smart information box can be accessed remotely from users, and can easily contact the utility service servers. In the example of Figure 3.24 there are four smart appliances: an oven, a refrigerator, a washing machine and a smart plug; together with the smart information box and the smart meter transceiver, they form the HAN.
ZigBee protocol assigns a role to each node into the network. There are three possible roles:

- ZigBee Coordinator (ZC): it is the smartest device in the network. The coordinator node is the root of the network and can also act as a bridge between different networks. It can contain information about the network as well as store the security keys. In each ZigBee network there is only one coordinator. The smart information box is the ZC of the example network of Figure 3.24.
- ZigBee Router (ZR): it acts as router in the network, exchanging data between nodes (not present in the example network of Figure 3.24).
- ZigBee End Device (ZEDs): they are the simplest nodes of the network, and they can communicate only with the coordinator or routers. ZEDs require little amount of memory. The devices in the network of Figure 3.24 are all ZEDs, and they communicate only with the smart information box.

It is worth noting that ZigBee is not a protocol for peer-to-peer networks (i.e. networks composed by nodes that have all the same role and where there is

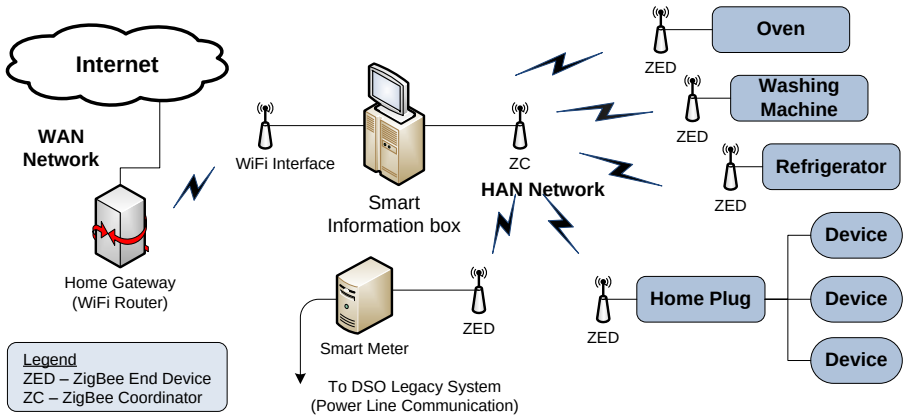

Figure 3.24. Smart Energy HAN architecture implementation example.

no distinction between them). ZigBee instead assigns a role to each node, and ZEDs cannot communicate directly, but only through a router or coordinator. Using routers within ZigBee network allows the deployment of a network architecture similar to mesh network.
Moreover the Home Energy Angel in our vision is a smart device that runs applications specific for energy management. It is a point of presence for every smart device within the home domain, and for third party's domotic solutions. For our purpose, the Home Energy Angel smart information box can integrate also the Wi-Fi router to provide an Internet access to users. In such a case all the applications for network and energy management run on the Home Energy Angel smart information box that acts also as a gateway.
Implementing the proposed energy HAN will allow energy saving and cost saving benefits for the end users.
As discussed at the last SustainIT2012 conference in several papers [164]-[166] in Europe the household contribution to the overall electricity consumption is about 29% corresponding for a country like Italy [166] in 70 TWh per year, 12 billions of Euros of cost, 2.5 MWh/user per year. Simulations carried out considering the power cost of typical house appliances and the consumption profiles of typical users allow the following estimation: the introduction of the energy HAN, supporting energy awareness of users and implementing automatic energy management policy, can reduce the user consumption per year by 20% from actual 2.5 MWh to 2 MWh.
A further cost saving can be achieved, thanks to the HAN, by enabling users to automatically exploit the high variability of energy cost which, as reported in Figure 3.25, can vary by a factor 3 from peak hours (F1 tariff in red in Figure 3.25) to off-peak hours (F3 tariff in green in Figure 3.25).

3.2.4.1 HW Architectures of Building Nodes: Smart Plugs, Home Energy Angel Box and Smart Power Meters

Figure 3.26 shows the block diagrams of smart devices present in the example network. Every device has a microcontroller (MCU) core, e.g. a 32-b RISC Cortex managing system activities, and an interactive user interface. The connectivity module enables them to join the network and to be remotely controlled. This subsystem can be a simple transceiver integrated with the MCU in the same PCB board or a single-chip wireless microcontroller can be used.

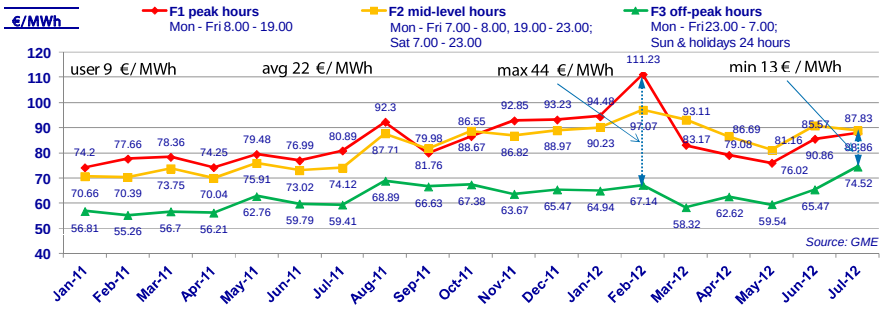

Figure 3.25. Cost of energy in different time ranges in Italy.

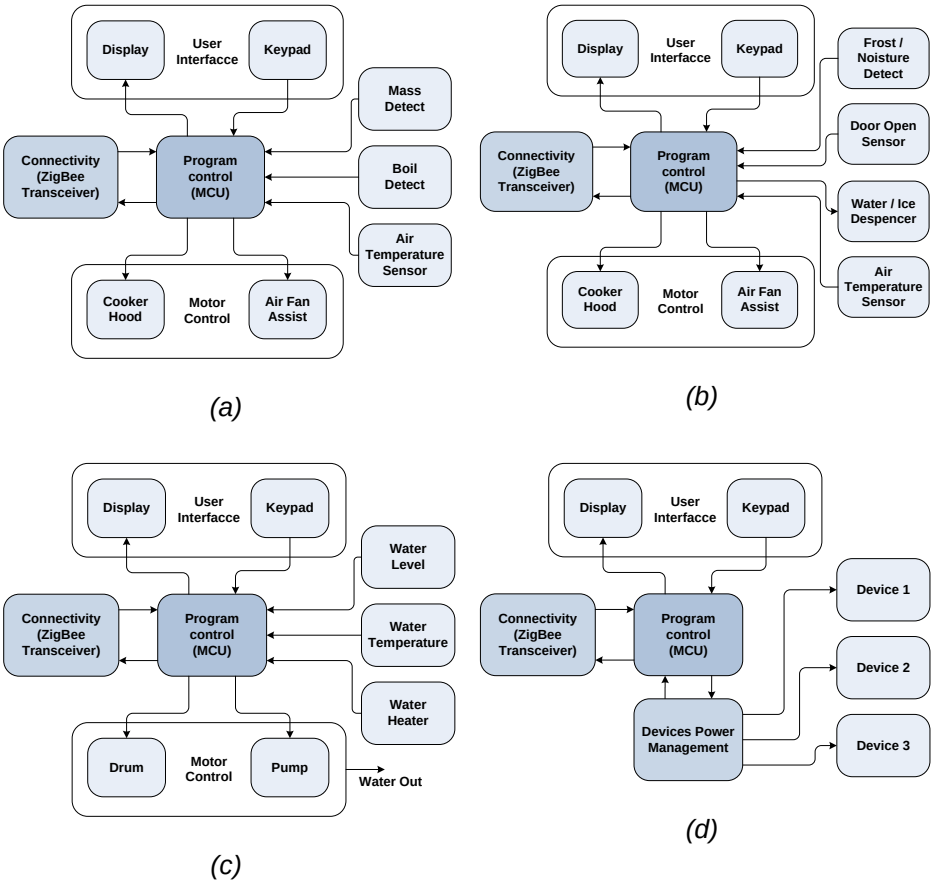

Figure 3.26. Oven (a), refrigerator (b), washing machine (c) and home plug (d) system diagram.

An example of stand-alone ZigBee transceiver is the Texas Instruments CC2520. In the solution that integrates the wireless microcontroller and the transceiver in the same chip, the antenna can be directly printed on the PCB board achieving enough gain with a limited size, as demonstrated by recent works done at University of Pisa where multi-loop multi-frequency antennas have been realized as PCB printed antenna for sub GHz applications [167],[168].

Both solutions can be used to upgrade the existing device's design enabling them to join home area networks. If we use a transceiver, the existing microcontroller could be connected to it using GPIO (General Purpose Input/output) and SPI lines. However, this introduces an overhead to the device MCU, that now has to control the system and to implement the communication protocol. On the other hand, using a wireless microcontroller avoids this problem. Actually, a wireless microcontroller can be used as co-processor, placing it side by side with the device MCU in charge of the system

control. The wireless microcontroller implements the communication protocol and manages the transmission and the reception of packets while the other MCU continues implementing the control algorithm, and when it needs to communicate with other devices in the network it sends a request to the

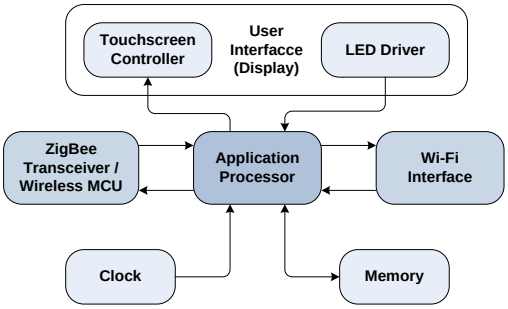

Figure 3.27. Architecture of the Home Energy Angel smart box.

wireless microcontroller. The overhead introduced by this scenario is limited.
The architecture of the Home Energy Angel smart information box is shown in Figure 3.27. A touch-screen display (in-home display) provides an easy way for users to interact with the smart box. Through this display the user can manage the energy settings of the devices connected to the HAN network, and can check the energy consumption records. An external permanent memory is needed to store past records, files for software running on the smart information box and any significant information about the network. The smart information box needs a connection to Internet in order to retrieve information about customer energy account, that are used by energy management applications. For that reason there is a Wi-Fi interface connected to the device MCU.
Figure 3.28 presents the block diagram architecture for the smart meter to be installed by utility companies (e.g. ENEL in Italy). The core of the smart meter is represented by the electronic meter (E-meter), able to calculate the energy consumption by sensing current and voltage from the electric network through an analog front-end. This information is digitized and elaborated by an MCU equipped with digital signal processing capabilities (e.g. a 32-bit Cortex processor), and then presented through a display. Also a power line communication modem is connected to the meter in order to send usage information to data concentrators. The extension required with respect to the current state of the art is the ZigBee transceiver. This component enables communication with the smart information box, so users can check their real-time consumption. The smart meter can also store consumption data into its memory for later use.
The described smart meter is the result of an evolution started from AMR (Automated Meter Reading) systems. These meters allow utility companies to read consumption records, status or alarms occurred. AMRs provide only

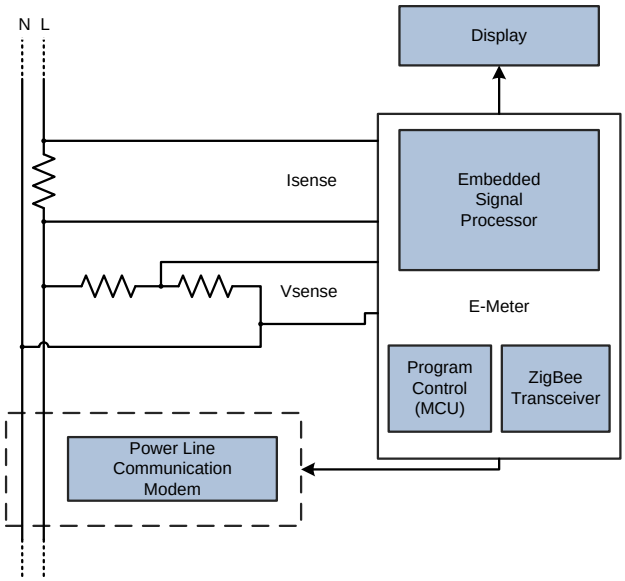

Figure 3.28. Architecture of the smart meter.

one-way communication: utility companies cannot take corrective actions on the customer grid.

The evolution of AMR is the AMI (Advanced Metering Information) meter whose hardware architecture has been detailed in Figure 3.28. AMI is characterized by a built-in two-way communication system, allowing the modification of customers' service level parameters. In this way customers can control energy cost choosing between different billing plans. However, as recognized in the state of the art, deploying a high number of smart meters in the environment is cumbersome. To reduce the number of deployed power metering devices without affecting the information reliability new approaches have been investigated, i.e. load disaggregation for extracting individual appliance power consumption information from single-point circuit-level measures. Such approach is based on the observation that each appliance has its own power consumption profile over time, as proved in Figure 3.21 and Figure 3.22, which can be isolated from the single-point measure. Disaggregation aims to extract the signatures of the different appliances from the aggregate measures. With respect to the state of the art, where artificial neural networks [169] or a Bayesian approach [170] have been used to perform the appliance recognition, we aim at avoiding the training phase. To this aim we propose adopting a coarse description of the appliance 'energy signatures' (i.e. how each type of appliance works and therefore its power consumption in any possible state) and to recognize most of the appliance types used in a residential building by indentifying: which unit is active, how long and how much is consuming. To this aim, the advanced hardware architecture proposed in Figure 3.28 is needed, since both a smart analog front end for the power meter measure (E-meter) and a powerful processing

Device	ATZB-24-A2/B0 [173]	JN5139	JN5148	STM32W108C8	CC2530
CPU	8b RISC ATMega128	32b RISC	32b RISC	32b RISC Cortex3	8b RISC 8051
Radio Freq.	2.4 - 2.485 GHz	2.4 GHz	2.4 GHz	2.4 GHz	2.394 - 2.507 GHz
Flash/ ROM	128 kB Flash	192 kB ROM	128 kB ROM	64 kB Flash	32, 64, 128, 356 kB Flash
RAM	8 kB	8 kB	128 kB	8 kB	8 kB
data-rate	250 kbps	250 kbps	250, 500, 667 kbps	250 kbps	250 kbps
Vsupply	1.8V- 3.6V	2.2 V-3.6 V	2 V-3.6V	2.1 V- 3.6V	2 V-3.6V
RX current	19 mA	34 mA	17.5 mA	27 mA	24 mA
Tx current	18 mA	34 mA	15 mA	31 mA	29 mA
Standby current	6 µA	1.3 µA	1.25 µA	0.8 µA	0.4 mA
Wakeup time	N.A.	N.A.	840 µs	110 µs	600 µs
RX sensitivity	-101 dBm	-97 dBm	-95 dBm	-99 dBm	-97 dBm
TX power	3 dBm	3 dBm	3 dBm	3 dBm	4.5 dBm

Table 3.10. COTS components to implement the Energy HAN nodes.

unit (the embedded signal processor) are required to implement the disaggregation DSP algorithms. Moreover the smart meter supports communication by wireless connection with the home area network and by Power Line Communication with the main energy grid.
For anti-tapering reason a 3-axis tilt sensor is also integrated in the smart meter architecture and is connected to the ZigBee and the Power line communication interfaces (i.e. if the smart meter is tampered the end user and/or the utility is notified), see [171]. A detailed review of anti-tampering techniques for smart meters is reported in [172].

3.2.4.2 COTS Components Selection to Build the Energy HAN

Table 3.10 presents a selection of COTS components suitable for developing devices capable of forming a ZigBee network, according to the architectures presented in the previous Sections. They are SoCs integrating at least a ZigBee transceiver, on chip non-volatile memory, RAM memory and a CPU with a security AES co-processor (see on-chip architecture in Figure 3.29).

As far as the RF part is concerned all devices in Table 3.10 implement the ZigBee physical layer at 2.4 GHz with a transmitted power in the order of few mW (3-4.5 dBm being 0 dBm=1 mW) and a sensitivity from -95 to -101 dBm. Hence considering a full TX-RX link the proposed devices can face path losses up to -100 dB which is enough to build reliable Home Area Networks according to the topology discussed in Section 3.2.4. Since 2.4 GHz is a worldwide unlicensed frequency, this allows the use of these transceivers in every country, without difficulties. By using an integrated MURATA antenna, printed on the PCB board, 3 dBm of TX power allows reaching 30 meters indoor/urban or 100 meters and more outdoor line-of-sight. As discussed in [174], several strategies are foreseen in ZigBee in order to solve frequency co-existence problems with other communication technologies (e.g. Wi-Fi, Bluetooth) in the crowded 2.4 GHz frequency spectrum.
As far as the power consumption is concerned it is in the order of several tens of mW in RX or TX active mode; implementing power cycling strategies the power consumption can be kept as low as few µW using the STM32W108C8 device which moreover has a short wake up time of 110 µs. To implement smart metering or energy management function devices with a 32 bit CPU have to be preferred and with both RAM and Flash (rather than ROM) on-chip capabilities. To this aim we have selected for the implementation the STM32W108C8 which has 8 kbytes and 64 kbytes or RAM and Flash respectively and a 32 bit Cortex M3 processor which has a computation efficiency of 1.25 Dhrystone MIPS/MHz, enhanced instructions such as Hardware Divide, Single-Cycle (32x32) Multiply, Saturated Math Support, 149 µW/MHz when realized in 180 nm CMOS technology. Some of the typical MAC operations required during communication (such as ACK management, automatic back-off delay and packet filtering) are implemented via hardware, in order to meet the strict timing requirements imposed by IEEE 802.15.4-2003 standard.
When optimizing the network, a specific customization can be done according to the specific device under control. For example to control simple power appliances like oven, refrigerators, boiler, lights or air conditioning, it is not needed a continuous control. They have to send their consumptions, alerts in case of troubles (oven overheat during cooking, failures in refrigerator's

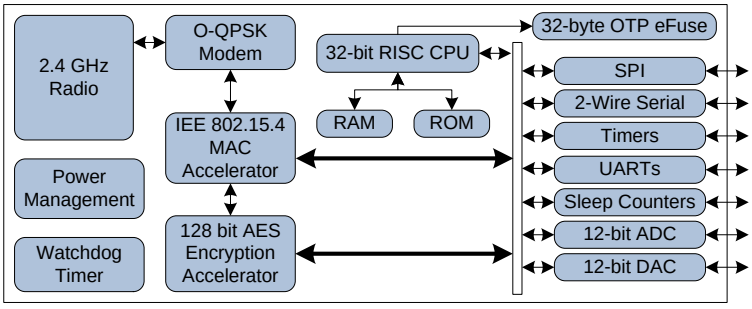

Figure 3.29. SoC architecture for wireless sensor networks in energy HAN.

components, etc.) and the capability of turning on or off them remotely is needed. For these devices a simple transceiver can be added, and hence the overhead introduced is minimum. These devices have only to transmit and to receive messages at low rate: their status, including energy usage, is checked few times in a day (4/5 times in a day) and alerts do not occur frequently. So a RF transceiver can be added to the MCU already present in such machines or their MCU can be replaced with a wireless SoC such as the STM32W108C8. In case of smart plug (e.g. turn-on/off control of lights) where a microcontroller is not present (since no logic control is required) a simple transceiver is added without any CPU core.

When dealing with other appliances such as washing machines, or rechargeable systems for electric vehicles, a continuous control can be useful to program their work and hence to find an optimal trade-off between user needs, time-based energy tariff and production peaks of renewable home energy generators (wind, photovoltaic) if any. A smart washing machine can be programmed to work during low cost time slots. To do this, once the device is programmed, it must be in standby mode until the job can be performed. Small standby consumption is required. The smart information box of Figure 3.27 can contact the washing machine when the low cost slot begins. Then the device can start its work. A good solution to implement the smart washing machine could be the STMicroelectronics STM32W108C8, mainly thanks to its low consumption of current during standby mode (0.8 µA). It could be envisaged the possibility to use this wireless microcontroller also to manage washing machine operations, for those models that are not too complex. Finally for the Home Energy Angel smart information box and for the smart meter with local processing capabilities (data disaggregation), powerful architectures are needed. As example in the smart meter, the STM32W108C8 SoC could just implement a part (the MCU and the wireless transceiver) of the

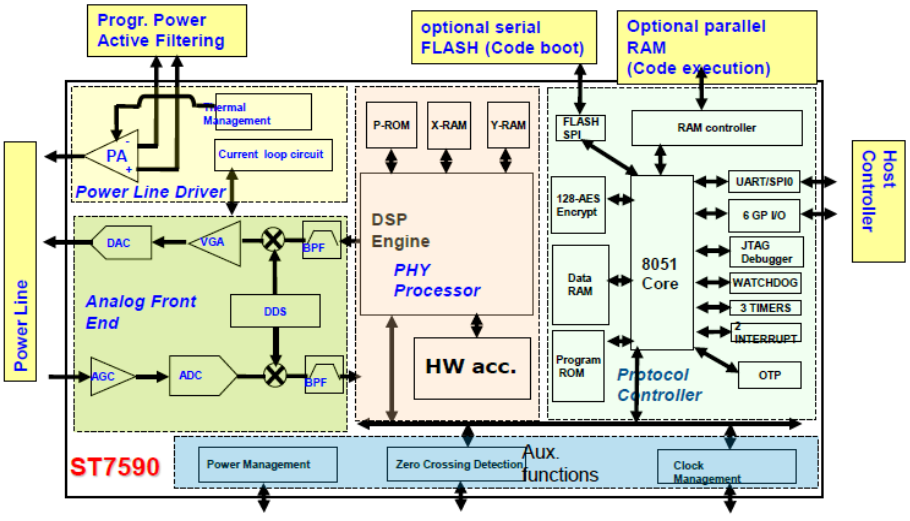

Figure 3.30. Architecture of the ST7590 PLC modem [148].

architecture of Figure 3.28. The board of the smart meter should be equipped also with: a display driver, an E-meter ASIC, a power line communication modem chipset and a touch screen/display controller. In the case of the Home Energy Angel also a Wi-Fi communication controller is needed. The ST7590 IC can be used as the power line communication modem. This device is able to operate at 28.8 kbps, its architecture is reported in Figure 3.30.
For the E-meter, the ASIC reported in Figure 3.31 by STMicroelectronics can be used. It implements measures of active, reactive and apparent energy by acquiring voltage or current value through dedicated acquisition channels. The accuracy is 0.1% of full scale value.

3.2.4.3 Security in the Proposed ZigBee/802.15.4 HAN

All the devices discussed in the previous section contain an AES dedicated processor to implement ZigBee/IEEE 802.15.4 secure communications. AES is the encryption algorithm used in the proposed network. On-chip one time programmable memory can be used to store 64-bit MAC ID and 128-bit AES security key. As reported in Figure 3.32, with respect to the ISO/OSI protocol stack, ZigBee implements the first three layers (Application, Transport and Network layer), while IEEE 802.15.4 provides protocols for data link layer (that is divided in Logical Link Control and Media Access Control). This standard has several versions, named by year. The most important are the 2003 and 2006 versions. IEEE 802.15.4 uses 27 channels divided in three main bands; the most interesting are the 16 channels available in the worldwide available 2.4 GHz unlicensed band. To avoid the simultaneous transmission of several nodes, the standard uses CSMA-CA (Carrier Sensing Multiple Access - Collision Avoidance) or GTS (Guarantee Time Slots) protocol. A node using CSMA-CA, before sending packets in the network listen the medium: if it is free the node will send its packets, otherwise it will wait for a certain period of

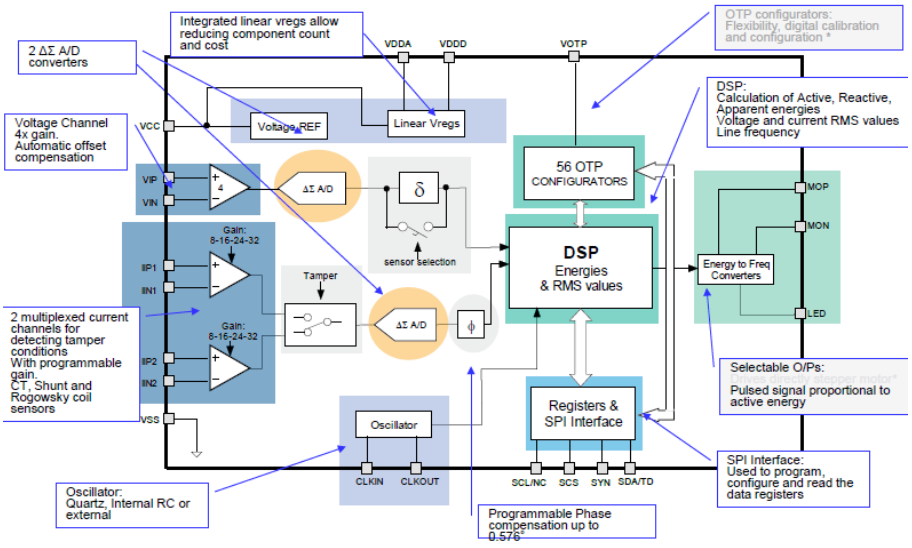

Figure 3.31. Architecture of the E-meter ASIC [148].

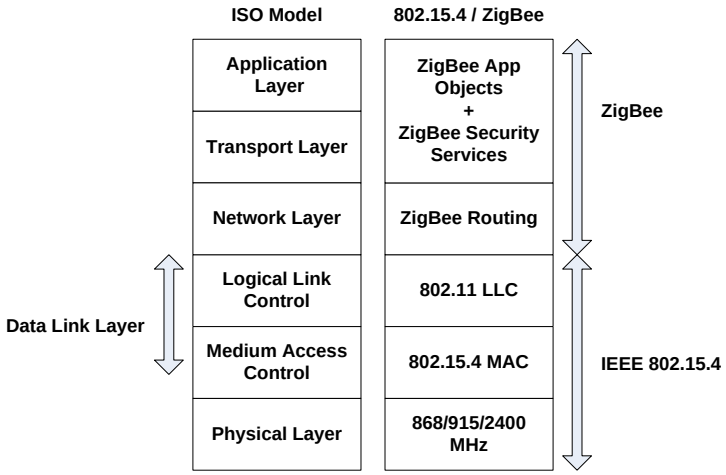

Figure 3.32. IEEE 802.15.4 and ZigBee role in the ISO/OSI stack.

time, computed with specific back-off algorithms (for example, exponential back-off algorithm). The GTS protocol, instead, uses a coordinator node which gives to the other nodes time slots, so that anyone knows when it has to transmit its data. An interesting feature of IEEE 802.15.4 is the Channel Energy Scan. Before using a channel, network sees how much energy (other network activity, noise, interference) there is.
This mechanism saves energy, choosing free channels when setting the network. IEEE 802.15.4 is a low consumption protocol. Nodes that use this protocol can keep their transceiver sleeping most of the time (up to 99%), and receiving and sending tasks can be set to take small part of the devices' energy. ZigBee [141] is a standard for high level communication based on IEEE 802.15.4 [175]-[178] data link standard. ZigBee is widely used in short range wireless communications requiring low date-rates, low energy consumption and a secure channel. The standard offers four main services:

- **Extra encryption service**: the application and network layers use 128-bit AES encryption.
- **Association and authentication**: only trusted nodes can join the network.
- **Routing protocol**: the Ad hoc On-Demand Distance Vector (AODV) routing protocol specifies how nodes communicate with each other.
- **Application service**: ZigBee introduces the concept of "cluster". Each node belongs to a cluster and can perform only actions allowed for the cluster. For example "House light system" cluster has only two possible actions: "turn lights on" and "turn lights off".

ZigBee nodes have a 16-bit network address, assigned by the coordinator during the association process. This address is used for routing information. Nodes within the network play different roles: coordinator, router, end device. Coordinator and routers cannot sleep. They must be always awake in order to manage the network and to send packets along the network.

It is important to remind that the ZigBee network has not a peer-to-peer architecture but a hierarchy one in which end devices can only communicate with routers and coordinators.
IEEE 802.15.4 supports only the encryption algorithm 128-bit AES. The reason is mainly due to the possibility to easily find on the market specific devices able to carry out encryption and decryption at the hardware level. The selected SoC has the AES processor embedded directly into transceivers and requires low resources. This standard does not specify how the keys have to be managed or the authentication policies to be applied. These details are leaved to the high level standards. AES is used for data security (payload encryption) and for data integrity. In particular, the integrity is achieved using Message Integrity Code (MIC). MIC is obtained encrypting part of the MAC (Medium Access Control) frame, using the network key, and its length is usually 128 bits.
Figure 3.33 shows the IEEE 802.15.4 MAC frame. There are three important fields for security issues: Frame Control, Auxiliary Security Header and Data payload.
Auxiliary Security Header field is meaningless if the Security Enable Bit (within the Frame Control field) is unset. Otherwise, this field is divided into three sub-fields described hereafter.
Security Control: this field is used to select what kind of protection is used for the frame (i.e. security policies adopted): what is encrypted and how long is the key. The first 3 bits specify the Security Level, related codes are listed in Table 3.11.
Frame Counter. To prevent replying attacks, every frame has an unique id.
Key Identifier. This field contains information about the type of key used in the communication with the other node. Keys can be implicit (known by nodes that are communicating) or explicit. In this last case, Key Index and Key Source sub-fields give indication about the key used.
Payload fields change according to Security Control field bits.

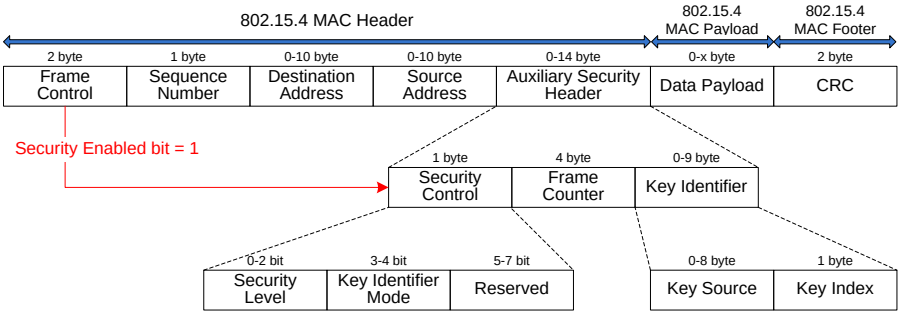

Figure 3.33. IEEE 802.15.4 MAC frame and security issues.

Code	Security type	Authentication	Security services
0x00	No security	-	No security
0x01	AES-CBC-MAC-32	MIC-32	Data integrity
0x02	AES-CBC-MAC-64	MIC-64	
0x03	AES-CBC-MAC-128	MIC-128	
0x04	AES-CTR	-	Data security
0x05	AES-CCM-32	AES-CCM-32	Data integrity and security
0x06	AES-CCM-64	AES-CCM-64	
0x07	AES-CCM-128	AES-CCM-128	

Table 3.11. Security Control codes.

Every node within the network has an Access Control List (ACL), a list of "trusted brothers". Each node before sending data to another node checks if the receiver is a trusted brother using ACL table. If the receiver does not appear into the list, the node can take two possible actions, according to the security policy adopted for the network: reject the message or begin an authentication process. ACL fields are specified in Table 3.12.
With respect to the 802.15.4 layers, ZigBee adds two additional security layers: the network and the application layer. As all security mechanisms use 128-bit AES encryption, devices designed for IEEE 802.15.4 standard can be used without any modification. ZigBee standard uses three type of keys. These keys are actually used or not, according to the policy chosen for the network. ZigBee keys are:

- Master key: It is used for keeping link keys confidential and checking their correspondence.
- Link keys: These keys are unique between pair of nodes. The use of link keys introduces a significant overhead for the node, requesting more memory resources, due to the fact that all data exchanged between two nodes must be encrypted with this key. Link keys are used only in commercial mode policy.
- Network key: it is an unique 128-bit key shared between the devices composing the network. Network key is generated by the Trust Center and it is regenerated after specific time interval. The old key is used to encrypt the new key, that is sent to nodes.

Field	Description
Address	Address of the destination node
Security suite	Security policy used
Key	128-bit key used in AES algorithm
Last Initial Vector (IV)	Used by the source to avoid reply attacks
Replay Counter	Replay counter is equal to IV but is used by the destination node

Table 3.12. Access Control List fields.

Master and link keys are used by the application layer, network key is used both by the ZigBee and the MAC layers. The Trust Center is a special device that is trusted by the other nodes within the network. Generally, the coordinator is the trust center, even if this role can be played also by another node.
To ensure security, the ZigBee network can use both master and link keys, or if a simple connection is needed only the network key.
In the first case, ideally, every device has the trust center address and an initial master key pre-installed. Otherwise, master keys can be distributed by trust center, during initial network setup using an insecure channel. After all nodes have the master key, link key can be obtained using agreement or transport process. Link keys can be also pre-installed. An example of this use of keys is the commercial mode policy (shown in Figure 3.34).
When the ZigBee network uses only the network key there is an initial distribution of this key, which is done by the trust center through an insecure channel. Only after the network key is acquired by all nodes the communications between nodes become secure.
Security policies decide which keys are used to make safe the network:

- Commercial mode where both master key and link keys are used. In this case more memory resources are required.
- Residential mode where data exchange within the network are encrypted using only network key. This mode is the easiest to implement but is less secure.

To ensure security and privacy protection in the example Energy HAN, both residential and commercial modes (see Figure 3.34 and Figure 3.35) can be

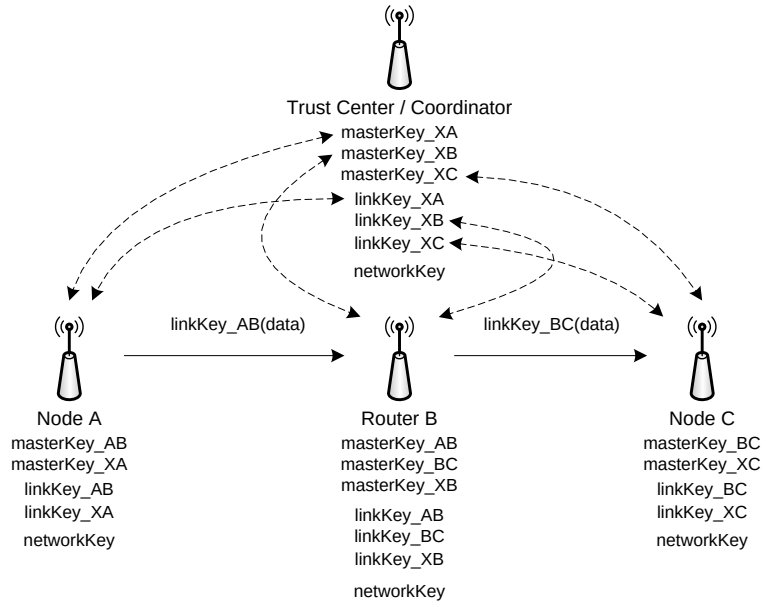

Figure 3.34. ZigBee Commercial mode.

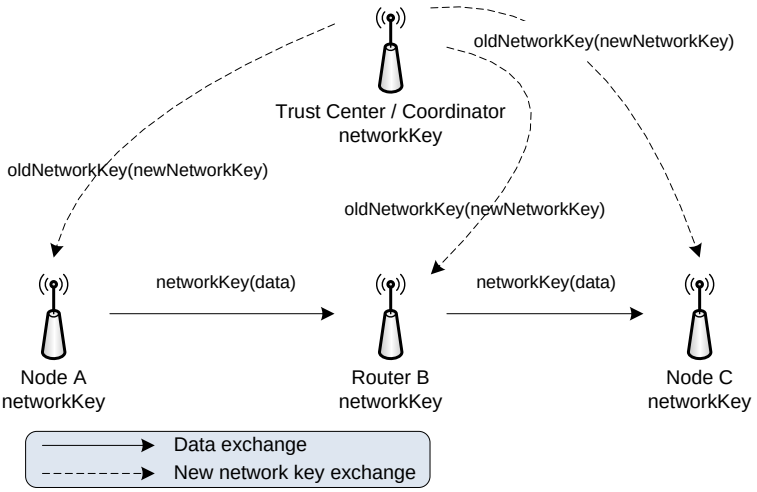

Figure 3.35. ZigBee residential mode.

used. In both modes, the trust center role can be played by the smart information box, that is also the coordinator of the network. This eliminates the need of a special node performing only security operations. Security parameters can be easily set using the interfaces of the smart information box, and information about network behavior (logs) can be stored and later accessed. Moreover, the home Energy Angel smart information box is also accessible remotely, so users can control security status of the network also outdoor, using an Internet connection.

To implement the residential mode the network will need only a network key. The Home Energy Angel Smart information box can establish a first key and then distribute it through an insecure connection to other nodes. Otherwise users and operators can "write" it into the devices' memory, this operation is more secure. Actually any key is transmitted through an insecure channel. Summarizing, if residential mode is chosen security problems can occur during the initial setup of the network.

Commercial mode provides stronger security than residential mode, but requires more resources (memory and CPU time). Actually each connection between two nodes uses different keys, and if security is broken in one link this will not affect the whole network. Also in this mode, the initial key setting is a critical point. A secure method is to assign to each node a first master key "manually". Then, this will be changed by the trust center using a secure connection. If a first master key is not assigned, this task must be done by the trust center using an insecure channel. A successful conclusion of the initial setup of the network assures the confidentiality and the integrity of the network.

Information exchanged between nodes are always encrypted and message integrity can always be checked, if the highest security level (AES-CCM-128) was selected. This choice does not affect devices performance since ZigBee transceivers have dedicated AES processors for encryption and decryption.

4 CONCLUSION

This dissertation discussed embedded computing design activities taking into account two main points of view: the hardware design and the application or system level design. Sometimes challenges and trends appear quite different for the two cases. While in the first case the research is mostly technology-driven, in the second case social needs are the main matter. However, the starting point in order to reach the greater success in future embedded computing is a close collaboration between research domains, designers, industries and markets.

4.1 NoC interconnection architectures for MPSoC

The design of a novel Network-on-Chip interconnecting architecture has been presented in this work. The architecture and the main features of the key building blocks (i.e. router and Network Interface) have been discussed.
The proposed Network Interface offers a wide set of advanced networking functionalities: store&forward transmission, error management, power management, ordering handling, security, QoS management, programmability, interoperability and remapping. Also the switching matrix of the implemented router is fully configurable, resulting in a very flexible topology. Furthermore the router offers different QoS arbitration schemes and many features (i.e. queues, latency and credit round-trip delay) which are configurable at synthesis time.
The capability to support all these advanced features in the same hardware represents a novelty with respect to the state of the art. Furthermore, a wide and fine-grained configuration space ensures an optimal scalability of the design, to reach the desired trade-off between supported services and circuit complexity.
The implementation of the NoC blocks was supported by a new methodology called metacoding that generates correct-by-construction technology independent RTL codebases. The RTL coding itself is abstracted and modelled with an Object Oriented framework.
Compared with traditional coding styles based on pre-processor directives, the proposed methodology produces smaller codebase (more than 60% reduction) and a considerable reduction of the verification space by decoupling the verification of the codelets and of the backbone.
Thanks to the metacoding methodology, the typical design flow was slightly reviewed in order to implement an automatic design flow. The framework developed offers a simple and quick configuration mechanism to meet any specific application requirement without introducing an excessive overhead coming from the configuration of the blocks composing the NoC.
Several configurations have been characterized on nanoscale CMOS technology, and they have shown complexity and performance figures comparable with architectures found in literature supporting the same features subset. The proposed NoC represents a complete solution that can be customized for different scenarios, from multimedia to real-time applications, just analyzing the system requirements and tailoring features and parameters to obtain an optimized hardware description.

The research activity was carried out in collaboration with the Advanced System Technology group of ST Microelectronics in Grenoble (France). The proposed architecture and methodology were filed in some US and European patents and are currently embedded in an industrial flow for assembling interconnection platforms in commercial chipsets.

4.2 Health monitoring system for CHF patients

The requirements and the realization in terms of sensing devices and sensor data signal processing of a complete and integrated ICT platform to improve the provisioning of healthcare services for patients affected by Chronic Heart Failure patients have been presented. The Health@Home system proposes an innovative home care model in order to support the whole process of the patient treatment in integrated and coordinated fashion, connecting in-hospital care with out-of-hospital follow up. With the remote monitoring, the medical staff can perceive changes of patients' parameters of without frequently visiting them and consequently they can take proper actions to prevent possible aggravations. The benefits extend beyond the early detection of clinical exacerbation to optimizing specialized resources scheduling and to reducing unnecessary travels to hospital. The system definition was completely driven by the end-users resulting in a platform particularly effective and practical with respect to other telemonitoring trials and state-of-the-art products. One of the main system novelties is represented by the home gateway which embeds, through the so called operating protocol, the medical prescription for any given patient. Moreover, it locally performs all the sensor signal processing and alarm detection. The operating protocol is configurable according to the patient’s needs and remotely updatable if required via the server platform. The basic sensors kit established by physicians includes two commercial devices for weight and blood pressure and a new developed module for acquiring synchronized ECG and SpO2 track leveraging a multi-channel front-end IC for cardiac sensors. This again represents a novelty aspect with respect to the state-of-the-art together with the use of international standards for data exchange to favour the integration/interaction of the platform with other systems or biomedical sensors. The overall cost of a kit, considering the first prototype and test phase of H@H with 30 kits produced, is close to 1000 euros. Adopting the H@H telemonitoring CHF system in a large scale, e.g. in national health system, the kit cost can be reduced in the range of hundreds of euros. First technology assessment in a real medical scenario with tens of patients affected by CHF disease NYHA class III and IV, under remote control for some months by cardiologists using their usual Hospital Information System, proves the effectiveness of the telemonitoring system from both patients and caregivers point of view. Key points underlined by the end-users are the goodness of the sensor data analysis implemented, enabling early and accurate detection of destabilizations in vital signs and naturally the usability of the system. Summarising the H@H system meets end-user expectations and, in particular, physicians have already planned a clinical validation, including an economical evaluation, with more than 500 patients in the Italian Regional Tuscany Health System.

The H@H project obtained second prize in the element14 Medical Design Award on October 2011 at the Medical Technology Event, organized by Selezione di Elettronica (magazine by Il Sole 24 ORE Group); moreover, it was among the four finalists at the AAL PROJECT AWARD 2012, which recognizes the most successful projects of the Ambient Assisted Living Joint Programme demonstrating great promise in terms of innovation, human-centric approaches to development and market potential.

4.3 Energy Home Area Network for Smart Grid

The energy monitoring application for Smart Grid has been presented in this work. Network architecture, aims and requirements are discussed in the detail highlighting actual needs and challenges. An energy Home Area Network, a key element of Smart Grid, is presented, dealing with its security and privacy aspects and showing some solutions to realize a wireless network, based on ZigBee. Implementation challenges from the hardware and software point of view and possible architectures and implementation using COTS components are proposed for key nodes of the smart energy HAN: smart power meters, smart plugs and a Home Energy Angel information box essential for energy management/saving policy and for energy awareness. Architecture and solutions proposed in this work represent some effective answers for Smart Grids implementations; these represent the initial step for future research activities already scheduled.

BIBLIOGRAPHY

[1] International Technology Roadmap for Semiconductors, "More-than-Moore," White paper, 2010.

[2] International Technology Roadmap for Semiconductors, "2012 Update Overview," 2012.

[3] CATRENE Scientific Committee "Towards a More-than-Moore roadmap," CATRENE Scientific Committee report, 2011

[4] R. Ho, K. Mai, and M. Horowitz, “The future of wires," Proc. IEEE, vol. 89, no. 4, pp. 490 - 504, April 2001.

[5] P. S. Paolucci, F. LoCicero, A. Lonardo, M. Perra, D. Rossetti, C. Sidore, P. Vicini, M. Coppola, L. Raffo, G. Mereu, F. Palumbo, L. Fanucci, S. Saponara, and F. Vitullo, “Introduction to the tiled HW architecture of SHAPES,” in Proc. Design, Automation and Test in Europe, 2007, pp. 77–82.

[6] P. S. Paolucci, A. A. Jerraya, R. Leupers, L. Thiele, and P. Vicini, “SHAPES: a tiled scalable software hardware architecture platform for embedded systems,” in Proc. 4th Int. Conf. Hardware/Software Codesign and System Synthesis CODES+ISSS ’06, 2006, pp. 167–172.

[7] N. Parakh, A. Mittal, and R. Niyogi, “Optimization of MPEG-2 encoder on Cell B. E. processor,” in Proc. IEEE Int. Advance Computing Conf. IACC 2009, 2009, pp. 423–427.

[8] J. Nickolls and W. J. Dally, “The GPU computing era,” vol. 30, no. 2, pp. 56–69, 2010.

[9] S. Saponara, M. Martina, M. Casula, L. Fanucci, and G. Masera, “Motion estimation and CABAC VLSI co-processors for real-time high-quality H.264/AVC video coding,” Microprocess. Microsyst., vol. 34, pp. 316–328, November 2010.

[10] F. Vitullo, N. E. L’Insalata, E. Petri, S. Saponara, L. Fanucci, M. Casula, R. Locatelli, and M. Coppola, “Low-complexity link microarchitecture for mesochronous communication in networks-on-chip,” vol. 57, no. 9, pp. 1196–1201, 2008.

[11] M. A. U. Rahman, I. Ahmed, F. Rodriguez, and N. Islam, “Efficient 2DMesh Network on Chip (NoC) considering GALS approach,” in Proc. Fourth Int. Conf. Computer Sciences and Convergence Information Technology ICCIT ’09, 2009, pp. 841–846.

[12] H. G. Lee, N. Chang, U. Y. Ogras, and R. Marculescu, “On-chip communication architecture exploration: A quantitative evaluation of point-to-point, bus, and network-on-chip approaches,” ACM Trans. Des. Autom. Electron. Syst., vol. 12, pp. 23:1–23:20, May 2007.

[13] M. Coppola, M. Grammatikakis, and R. Locatelli, Design of cost-efficient interconnect processing units: Spidergon STNoC, ser. System-on-chip design and technologies. CRC Press, 2008.

[14] L. Benini and G. De Micheli, “Networks on chips: a new SoC paradigm,” Computer, vol. 35, no. 1, pp. 70–78, 2002.

[15] F. Vitullo, S. Saponara, E. Petri, M. Casula, L. Fanucci, G. Maruccia, R. Locatelli, and M. Coppola, “A reusable coverage-driven verification environment for network-on-chip communication in embedded system

platforms," in Proc. Seventh Workshop Intelligent solutions in Embedded Systems, 2009, pp. 71–77.

[16] K. Goossens, J. Dielissen, and A. Radulescu, "AEthereal network on chip: concepts, architectures, and implementations," IEEE Design & Test of Computers, vol. 22, no. 5, pp. 414–421, 2005.

[17] S. Saponara, L. Fanucci, and E. Petri, "A multi-processor NoC-based architecture for real-time image/video enhancement," Journal of Real-Time Image Processing, pp. 1–15, 2011, 10.1007/s11554-011-0215-8.

[18] W. Zhong, S. Chen, F. Ma, T. Yoshimura, and S. Goto, "Floorplanning driven network-on-chip synthesis for 3-D SoCs," in Proc. IEEE Int Circuits and Systems (ISCAS) Symp, 2011, pp. 1203–1206.

[19] S. Murali, C. Seiculescu, L. Benini, and G. De Micheli, "Synthesis of networks on chips for 3D systems on chips," in Proc. Asia and South Pacific Design Automation Conf. ASP-DAC 2009, 2009, pp. 242–247.

[20] R. S. Ramanujam, V. Soteriou, B. Lin, and L.-S. Peh, "Design of a high-throughput distributed shared-buffer NoC router," in Proc. Fourth ACM/IEEE Int Networks-on-Chip (NOCS) Symp, 2010, pp. 69–78.

[21] C.-H. Chan, K.-L. Tsai, F. Lai, and S.-H. Tsai, "A priority based output arbiter for NoC router," in Proc. IEEE Int Circuits and Systems (ISCAS) Symp, 2011, pp. 1928–1931.

[22] M. Palesi, G. Ascia, F. Fazzino, and V. Catania, "Data encoding schemes in networks on chip," vol. 30, no. 5, pp. 774–786, 2011.

[23] B. Attia, W. Chouchene, A. Zitouni, A. Nourdin, and R. Tourki, "Design and implementation of low latency network interface for network on chip," in Proc. 5th Int. Design and Test Workshop (IDT), 2010, pp. 37–42.

[24] H. Kariniemi and J. Nurmi, "NoC Interface for fault-tolerant Message-Passing communication on Multiprocessor SoC platform," in Proc. NORCHIP, 2009, pp. 1–6.

[25] G. Leary, K. Mehta, and K. S. Chatha, "Performance and resource optimization of NoC router architecture for master and slave IP cores," in Proc. 5th IEEE/ACM/IFIP Int Hardware/Software Codesign and System Synthesis (CODES+ISSS) Conf, 2007, pp. 155–160.

[26] E. Rotem, R. Ginosar, A. Mendelson, and U. Weiser, "Multiple clock and voltage domains for chip multi processors," in Proc. MICRO-42 Microarchitecture 42nd Annual IEEE/ACM Int. Symp, 2009, pp. 459–468.

[27] F. Baronti, E. Petri, S. Saponara, L. Fanucci, R. Roncella, R. Saletti, P. D'Abramo, and R. Serventi, "Design and verification of hardware building blocks for high-speed and fault-tolerant in-vehicle networks," IEEE Transactions on Industrial Electronics, vol. 58, no. 3, pp. 792–801, 2011.

[28] N. E. L'Insalata, S. Saponara, L. Fanucci, and P. Terreni, "Automatic synthesis of cost effective FFT/IFFT cores for VLSI OFDM systems," IEICE Transactions on Electronics, vol. E-91C/4, pp. 487–496, 2008.

[29] G. Maruccia, R. Locatelli, L. Pieralisi, and M. Coppola, "Buffering architecture for packet injection and extraction in on-chip networks," US Patent Application US 2009/0 147 783 A1, 06 11, 2009.

[30] L. Bononi and N. Concer, “Simulation and analysis of network on chip architectures: ring, spidergon and 2D mesh," in *Proc. Design, Automation and Test in Europe (DATE)*, Mar. 2006, pp. 167 - 172.

[31] T. Bjerregaard and J. Sparso, “A router architecture for connection-oriented service guarantees in the MANGO clockless network-on-chip," in *Proc. Design, Automation and Test in Europe (DATE)*, vol. 2, Mar. 2005, pp. 1226 - 1231.

[32] P. Martin, “Network on chip: The future of SoC power management," in *CDNLive! EMEA Conference Proceedings*, 2006.

[33] SonicsMX SMART Interconnect Data Sheet, Sonics Inc.

[34] D. Wiklund and D. Liu, “SoCBUS: switched network on chip for hard real time embedded systems," in *Proc. International Parallel and Distributed Processing Symposium*, April 2003.

[35] A. Pullini, F. Angiolini, P. Meloni, D. Atienza, S. Murali, L. Raffo, G. De Micheli, and L. Benini, “NoC design and implementation in 65nm technology," in *Proc. First International Symposium on Networks-on-Chip (NOCS)*, May 2007, pp. 273 - 282.

[36] D. Bertozzi and L. Benini, “Xpipes: a network-on-chip architecture for gigascale systems-on-chip," *IEEE Circuits Syst. Mag*., vol. 4, no. 2, pp. 18 - 31, 2004.

[37] S. Murali, M. Coenen, A. Radulescu, K. Goossens, G. De Micheli, "Mapping and configuration methods for multi-use-case networks on chips," *Design Automation, 2006. Asia and South Pacific Conference on*, Jan. 2006

[38] D. Bertozzi, A. Jalabert, Srinivasan Murali, R. Tamhankar, S. Stergiou, L. Benini, G. De Micheli, "NoC synthesis flow for customized domain specific multiprocessor systems-on-chip," *Parallel and Distributed Systems, IEEE Transactions on* , vol.16, no.2, pp. 113- 129, Feb. 2005

[39] A. Sangiovanni-Vincentelli, Guang Yang, S.K. Shukla, D.A. Mathaikutty, J. Sztipanovits, “Metamodeling: An Emerging Representation Paradigm for System-Level Design,” in *Design & Test of Computers*, vol 26, 2009, pp. 54 – 69.

[40] R. Damasevicius; V. Stuikys, “Application of UML for hardware design based on design process model,” in *Design Automation Conference*, 2004.

[41] M. Frigo and S. Johnson, “The design and implementation of FFTW3," *Proc. IEEE*, vol. 93, no. 2, pp. 216 - 231, Feb. 2005.

[42] M. Frigo, “A Fast Fourier Transform compiler," in *Proc. ACM SIGPLAN Conf. on Programming Language Design and Implementation*, vol. 34. ACM, May 1999, pp. 169 - 180.

[43] "IEEE Standard for IP-XACT, Standard Structure for Packaging, Integrating, and Reusing IP within Tools Flows," *IEEE Std 1685-2009*, pp.C1-360, Feb. 18 2010

[44] W. Dally and B. Towles, Principles and Practices of Interconnection Networks. Morgan Kaufmann, 2003.

[45] G. Neumann and U. Zdun, “XOTcl, an object-oriented scripting language," in *Proceedings of Tcl2k: The 7th USENIX Tcl/Tk Conference*, Feb. 2000.

[46] G. Maruccia, R. Locatelli, L. Pieralisi, M. Coppola, M. Casula, L. Fanucci, and S. Saponara, “Method for transferring a stream of at least one data packet between first and second electric devices and corresponding device,” US Patent Application US 2009/0 129 390 A1, 05 21, 2009.

[47] P. Teninge, R. Locatelli, M. Coppola, L. Pieralisi, and G. Maruccia, “System for transmitting data between transmitter and receiver modules on a channel provided with a flow control link,” US Patent Application US 2008/0 155 142 A1, 06 26, 2008.

[48] G. Maruccia, R. Locatelli, L. Pieralisi, and M. Coppola, “Method for transferring data from a source target to a destination target, and corresponding network interface,” US Patent Application US 2008/0 320 161 A1, 12 25, 2008.

[49] V. Catalano, M. Coppola, R. Locatelli, C. Silvano, G. Palermo, and L. Fiorin, “Programmable data protection device, secure programming manager system and process for controlling access to an interconnect network for an integrated circuit,” US Patent Application US 2009/0 089 861 A1, 04 02, 2009.

[50] S. Saponara, F. Vitullo, E. Petri, L. Fanucci, M. Coppola, and R. Locatelli, “Coverage-driven verification of HDL IP cores,” in Solutions on Embedded Systems, ser. Lecture Notes in Electrical Engineering, M. Conti, S. Orcioni, N. M. Martinez Madrid, and R. E. E. D. Seepold, Eds. Springer Netherlands, 2011, vol. 81, pp. 105–119, 10.1007/978-94-007-0638-5 8.

[51] X. Yang, Z. Qing-li, F. Fang-fa, Y. Ming-yan, and L. Cheng, “NISAR: An AXI compliant on-chip NI architecture offering transaction reordering processing,” in Proc. 7th Int. Conf. ASIC ASICON ’07, 2007, pp. 890–893.

[52] B. A. A. Zitouni and R. Tourki, “Design and implementation of network interface compatible OCP for packet based NOC,” in Proc. 5th Int Design and Technology of Integrated Systems in Nanoscale Era (DTIS) Conf, 2010, pp. 1–8.

[53] T. Tayachi and P.-Y. Martinez, “Integration of an STBus Type 3 protocol custom component into a HLS tool,” in Proc. 3rd Int. Conf. Design and Technology of Integrated Systems in Nanoscale Era DTIS 2008, 2008, pp. 1–4.

[54] A. Radulescu, J. Dielissen, S. G. Pestana, O. P. Gangwal, E. Rijpkema, P. Wielage, and K. Goossens, “An efficient on-chip NI offering guaranteed services, shared-memory abstraction, and flexible network configuration,” vol. 24, no. 1, pp. 4–17, 2005.

[55] M. Ebrahimi, M. Daneshtalab, P. Liljeberg, J. Plosila, and H. Tenhunen, “A high-performance network interface architecture for NoCs using reorder buffer sharing,” in Proc. 18th Euromicro Int Parallel, Distributed and Network-Based Processing (PDP) Conf, 2010, pp. 546–550.

[56] Y.-L. Lai, S.-W. Yang, M.-H. Sheu, Y.-T. Hwang, H.-Y. Tang, and P.-Z. Huang, “A high-speed network interface design for packet-based NoC,” in Proc. Conf. Int Communications, Circuits and Systems, vol. 4, 2006, pp. 2667–2671.

[57] L. Fiorin, G. Palermo, S. Lukovic, V. Catalano, and C. Silvano, “Secure memory accesses on Networks-on-Chip,” vol. 57, no. 9, pp. 1216–1229, 2008.

[58] D. Matos, M. Costa, L. Carro, and A. Susin, “Network interface to synchronize multiple packets on NoC-based Systems-on-Chip,” in Proc. 18th IEEE/IFIP VLSI System Chip Conf. (VLSI-SoC), 2010, pp. 31–36.

[59] A. Ferrante, S. Medardoni, and D. Bertozzi, “Network interface sharing techniques for area optimized NoC architectures,” in Proc. 11th EUROMICRO Conf. Digital System Design Architectures, Methods and Tools DSD ’08, 2008, pp. 10–17.

[60] Synopsys Inc, “Synopsys coreTools: IP based design and verification,” pp. 1–3, 2008.

[61] S. Saponara, L. Fanucci, and M. Coppola, “Design and coverage-driven verification of a novel network-interface IP macrocell for network-onchip interconnects,” Microprocessors and Microsystems, vol. 35, no. 6, pp. 579 – 592, 2011.

[62] D. Matos, L. Carro, and A. Susin, “Associating packets of heterogeneous cores using a synchronizer wrapper for NoCs,” in Proc. IEEE Int Circuits and Systems (ISCAS) Symp, 2010, pp. 4177–4180.

[63] S. Stergiou, F. Angiolini, S. Carta, L. Raffo, D. Bertozzi, and G. De Micheli, “×pipes Lite: a synthesis oriented design library for networks on chips,” in Proc. Design, Automation and Test in Europe, 2005, pp. 1188–1193.

[64] SHAPE Survey Results to the General Public, *Annual Congress of the European Society of Cardiology* in Vienna, September 2003

[65] F. Zannad et al., “Heart failure burden and therapy” *Europace*, 11 (5):1-9, 2009

[66] V. Roger et al., "Heart disease and stroke statistics–2011 update: a report from the American Heart Association," *Circulation*, 123:18–209, 2011

[67] F. Alla, F. Zannad, G. Filippatos. “Epidemiology of acute heart failure syndromes”, *Heart Fail Rev*, 12:91–95, April 2007

[68] C. Berry, D. Murdoch, J. McMurray, "Economics of chronic heart failure", *Eur J Heart Fail*, 3(3): 283-291, 2001

[69] A. Bundkirchen, et al. “Epidemiology and economic burden of chronic heart failure” *European Heart Journal Supplements*, 6: 57–60, 2004

[70] J. Ross, et al “Recent National Trends in Readmission Rates after Heart Failure Hospitalization”, *Circulation: Heart Failure*, 3: 97-103, 2009

[71] S. Stewart, "Financial aspects of heart failure programs of care", *European Journal of Heart Failure*, 7 (3):423-428, 2005

[72] F. McAlister, et al. “Multidisciplinary strategies for the management of heart failure patients at high risk for admission: A systematic review of randomized trials”, *J. of the American College of Cardiology*,44(4), 2004

[73] E. Seto, “Cost comparison between telemonitoring and usual care of heart failure: a systematic review”, *Telemedicine and e-Health*,14(7), 2008

[74] C. Klersy, et al. “A Meta-Analysis of Remote Monitoring of Heart Failure Patients”, *Journal of the American College of Cardiology,* 54(18), 2009

[75] SC Inglis, et al. “Structured telephone support or telemonitoring programmes for patients with chronic heart failure” *The Cochrane Library*, 8, 2010
[76] G. van den Broek, et al., “Ambient Assisted Living Roadmap,” March 2010 Available: http://www.aaliance.eu
[77] S.J. Devaraj, K. Ezra, "Current trends and future challenges in wireless telemedicine system," *IEEE ICECT 2011,* pp. 417-421
[78] C. Fabbricatore, et al., “Towards an unified architecture for smart home and Ambient Assisted Living solutions: A focus on elderly people,” *IEEE DEST 2011,* pp.305-311, June 2011
[79] A. Jara et al., “An Architecture for Ambient Assisted Living and Health Environments,” *Lecture Notes in Computer Science*, 5518: 882-889, 2009
[80] J.P. Riley, M.R.Cowie; “Telemonitoring in heart failure”; *Heart and Education in Heart*; 95 (23): 1964-1968, 2009
[81] T. Bacchillone et al., “A flexible home gateway system for telecare of patients affected by chronic heart failure”, *IEEE Int. Symposium on Medical Information and Communication Technology,* pp.139-142, 2011
[82] R. *Dolin et al., “HL7 Clinic*al Document Architecture, Release 2” *JAMIA* 2006; 13:30-39, 2006
[83] Implementation Guide for CDA Release 2.0 Personal Healthcare Monitoring Report (PHMR). Available: http://www.hl7.org
[84] M. Yuksel et al., "Interoperability of Medical Device Information and the Clinical Applications: An HL7 RMIM based on the ISO/IEEE 11073 DIM", *IEEE Trans. Inf. Tech. In Biomedicine*, 15 (4): 557-566, 2011
[85] Systemized Nomenclature of Medicine – Clinical Terms. Available: http://www.nlm.nih.gov
[86] Logical Observation Names and Identifier. Available: http://loinc.org/
[87] B. Schijvenaars et al., "Intraindividual variability in electrocardiograms", *J. Electrocardiol*.,41(3):190-196, 2008
[88] S. Saponara et al., "Sensor modeling, low-complexity fusion algorithms, and mixed-signal IC prototyping for gas measures in low-emission vehicles", *IEEE Trans. Instr. and Measurements*, 60 (2): 372-384, 2011
[89] Pan Jiapu et al., “A Real-Time QRS Detection Algorithm”, *IEEE Trans. on Biomedical Engineering* , 32 (3): 230-236, 1985
[90] N. Arzeno, Z. Deng, C Poon, “Analysis of First-Derivative Based QRS Detection Algorithms”, *IEEE Trans. Biomed. Eng*., 55(2): 478-484, 2008
[91] K. Tateno, L.Glass, “Automatic detection of atrial fibrillation using the coefficient of variation and density histograms of RR and ΔRR intervals”, *Medical & Biological Engineering & Computing*, vol. 39, 2001
[92] A. Ghodrati et al., “Statistical analysis of RR interval irregularities for detection of atrial fibrillation”, *Computers in Cardiology, vol. 35, 2008*
[93] R. Shouldice et al., “Automated detection of paroxysmal atrial fibrillation from inter-heartbeats intervals”, *IEEE EMBC 2007*, pp 686-689, 2007.
[94] S.Dash, K.H.Chon, S.Lu, E.A.Raeder, “Automatic Real Time Detection of Atrial Fibrillation”, *Annals of Biomedical Engineering*, 39 (9), 2009.
[95] I.Romero Legarreta, “Component Selection for Principal Component Analysis-Based Extraction of Atrial Fibrillation”, *Computers in Cardiology*, 33: 137-140, 2006

[96] B. Weng, J.J.Wang, F.Michaud, M.Blanco-Velasco, “Atrial fibrillation detection using stationary wavelet transform analysis”, *IEEE EMBC 2*004
[97] S. Babaeizadeh et al., “Improvements in atrial fibrillation detection for real-time monitoring”, *Journal of Electrocardiology*, 2009.
[98] J. Allen, “Photoplethysmography and its application in clinical physiological measurement”, *Physiological Measurements*., 28(3), 2007
[99] M. Asif-Ul-Hoque, “Measurement of Blood Pressure Using Photoplethysmography”, *13th Int. Conf. on Computer Modelling and Simulation (UKSim), pp. 32.35, 2011*
[100] B-K. Miller, W. MacCaull, "Toward Web-based Careflow Management Systems", *Journal of Emerging Tech. in Web Intel.,* 1 (2): 137-145, 2009
[101] Bin Chen et al., "Analyzing Medical Processes”, *ICSE '08,* 623-632, 2008
[102] A. L. Goldberger, et al., “PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals”, *Circulation*, 101: 215-220, 2000
[103] Harvard-MIT Division of Health Sciences and Technology. MIT-BIH ArrythmiaDatabase, and AtrialFibrillationdataBase, Available: http://www.physionet.org
[104] L.Pecchia et al., "Discrimination power of short-term heart rate variability measures for CHF assessment", *IEEE Tran. Inf. Tech. Bio.,* 15 (1), 2011
[105] Insight Telehealth System. Available: http://www.itsmyhealthyheart.com
[106] Aerotel Medical System. Available: http://www.aerotel.com
[107] Tunstall. Available: http://www.tunstall.it
[108] Telbios. Available: http://www.telbios.com
[109] Parsys teleme*dicine. Available:* http://www.parsyssante.com
[110] P. Crilly et al., "An integrated pulse oximeter system for telemedicine applications", *IEEE Instr. Meas. Tech. Conf. (IM2TC)* 1997, pp. 102-104
[111] C. De Capua et al., "A Smart ECG Measurement System Based on Web-Service-Oriented Architecture for Telemedicine Application", *IEEE Trans. Instr. and Measurements*, 59 (10): 2530-2538, 2010
[112] Ying-Wen Bai et al., "Design and Implementation of an Embedded Remote ECG Measurement System", *IEEE Instr. and Meas. Tech. Conference* 2001, pp. 1401-1406
[113] M. Alhamid et al., "Hamon: An activity recognition framework for health monitoring support at home", *IEEE IM2TC* 2011, pp. 1-5
[114] G. Gupta, " Design of a low-cost physiological parameter measurement and monitoring device", *IEEE Instr. Meas. Tech. Conf (IM2TC*07), pp.1-6
[115] A. Gund, et al., “Desing evaluation of an home-based telecare system for Chronic heart failure patients”*, IEEE Eng Med Biol Soc. 2008*, pp. 5851-5854.
[116] Villalba, et al. “Wearable and Mobile System to Manage Remotely Heart Failure,” *IEEE Trans*. *Information Technology in Biomedicine,* 13(6): 990-996, 2009
[117] Jianchu Yao et al., "A wearable point-of-care system for home use that incorporates plug-and-play and wireless standards", *IEEE Trans. on Information Technology in Biomedicine*, 9 (3): 363-371, Sept. 2005

[118] K. Soundarapandian, M. Berarducci, "Analog Front-End Design for ECG Systems Using Delta-Sigma ADCs", TI report SBAA160A, pp. 1-11, Apr. 2010

[119] Texas Instruments, "Low-Power, 8-Channel, 16-Bit Analog Front-End for Biopotential Measurements", report n. SBAS471C, pp. 1-76, Nov. 2011

[120] R.F. Yazicioglu et al., "A 30 µW Analog Signal Processor ASIC for Portable Biopotential Signal Monitoring", *IEEE Journal of Solid State Circuits*, 40 (1): 209-223, 2011

[121] R. F. Yazicioglu, P. Merken, R. Puers, C. Van Hoof, "A 60 µW 60 nV/Hz readout front-end for portable biopotential acquisition systems," *IEEE J. Solid-State Circuits*, 42 (5): 1100–1110, 2007

[122] H. De Groot et al., "Human++: key challenges and trade-offs in embedded system design for personal health care", *IEEE Euromicro conference on Digital System design*, pp. 4-10, 2011

[123] H. Kim et al., "A configurable and low-power mixed signal SoC for portable ECG monitoring applications", *IEEE Symp. on VLSI circuits,* pp. 142-143, 2011

[124] N.E. L'Insalata et al., "Automatic synthesis of cost effective FFT/FFT cores for VLSI OFDM systems", *IEICE Transactions on Electronics,* E91-C (4): 487-496, 2008

[125] L. Fanucci et al., "Power optimization of an 8051-compliant IP microcontroller," *IEICE Transactions on Electronics,* E88-C (4): 597-600, 2005

[126] H. Farhangi, "The path of the smart grid", *IEEE Power and Energy Magazine*, vol. 8, n. 1, pp. 18–28, 2010.

[127] S. Massoud Amin, B. F. Wollenberg, "Toward a smart grid: power delivery for the 21st century", *IEEE Power and Energy Magazine*, vol. 3, n. 5, pp. 34–41, 2005.

[128] W.-K. Park, C. sic Choi, I. woo Lee, J. Jang, "Energy efficient multi-function home gateway in always-on home environment", *IEEE Transactions on Consumer Electronics*, vol. 56, n. 1, pp. 106–111, 2010.

[129] M. Jahn, M. Jentsch, C. R. Prause, F. Pramudianto, A. Al-Akkad, R. Reiners, "The Energy Aware Smart Home", *Proc. 5th Int Future Information Technology (FutureTech) Conf*, pp. 1–8, 2010.

[130] D. Niyato, L. Xiao, P. Wang, "Machine-to-machine communications for home energy management system in smart grid", *IEEE Communications Magazine*, vol. 49, n. 4, pp. 53–59, 2011.

[131] D. Y. R. Nagesh, J. V. V. Krishna, S. S. Tulasiram, "A real-time architecture for smart energy management", *Proc. Innovative Smart Grid Technologies (ISGT)*, pp. 1–4, 2010.

[132] P. Kulkarni, S. Gormus, Z. Fan, B. Motz, "A mesh-radio-based solution for smart metering networks", *IEEE Communications Magazine*, vol. 50, n. 7, pp. 86–95, 2012.

[133] E. Pallotti, F. Mangiatordi, "Smart grid cyber security requirements", *Proc. 10th Int Environment and Electrical Engineering (EEEIC) Conf*, pp. 1–4, 2011.

[134] A. R. Metke, R. L. Ekl, "Security Technology for Smart Grid Networks", *IEEE Transactions on Smart Grid*, vol. 1, n. 1, pp. 99–107, 2010.
[135] EPRI, "Report to NIST on Smart Grid interoperability standards roadmap", August 2010.
[136] R. H. Lasseter, P. Paigi, "Microgrid: a conceptual solution", *Proc. IEEE 35th Annual Power Electronics Specialists Conf. PESC 04*, pp. 4285–4290, 2004.
[137] NanoCatGeo, "NanoCatGeo Project" , https://sites.google.com/site/nanocatgeo.
[138] Acta, "Hydrogen generators and fuel cells systems" , http://www.actagroup.it.
[139] E. L. Quinn, "Privacy and the New Energy Infrastructure", *Social Science Research Network (SSRN)*, vol. , pp. 43, 2009.
[140] T. Flick, J. Morehouse, "Securing the Smart Grid: Next Generation Power Grid Security", Syngress Publishing, pp. 320, 2010.
[141] ZigBee, "The ZigBee Specification version 1.0", ZigBee Alliance, (Q4/2007).
[142] M. Zeifman, K. Roth, "Nonintrusive appliance load monitoring: Review and outlook", *IEEE Transactions on Consumer Electronics*, vol. 57, n. 1, pp. 76–84, 2011.
[143] T.-S. Choi, K.-R. Ko, S.-C. Park, Y.-S. Jang, Y.-T. Yoon, S.-K. Im, "Analysis of energy savings using smart metering system and IHD (in-home display)", *Proc. Transmission & Distribution Conf. & Exposition: Asia and Pacific*, pp. 1–4, 2009.
[144] Z. Wang, G. Zheng, "Residential Appliances Identification and Monitoring by a Nonintrusive Method", *IEEE Transactions on Smart Grid*, vol. 3, n. 1, pp. 80–92, 2012.
[145] M. Venables, "Smart meters make smart consumers [Analysis]", *Engineering & Technology*, vol. 2, n. 4, 2007.
[146] F. Benzi, N. Anglani, E. Bassi, L. Frosini, "Electricity Smart Meters Interfacing the Households", *IEEE Transactions on Industrial Electronics*, vol. 58, n. 10, pp. 4487–4494, 2011.
[147] C. Efthymiou, G. Kalogridis, "Smart Grid Privacy via Anonymization of Smart Metering Data", *Proc. First IEEE Int Smart Grid Communications (SmartGridComm) Conf*, pp. 238–243, 2010.
[148] Y. Gourdou, "Smart Grid/Metering Solution", EMCU, 2011.
[149] M. Nassar, J. Lin, Y. Mortazavi, A. Dabak, I. H. Kim, B. L. Evans, "Local Utility Power Line Communications in the 3--500 kHz Band: Channel Impairments, Noise, and Standards", *IEEE Signal Processing Magazine*, vol. 29, n. 5, pp. 116–127, 2012.
[150] Echelon, "DCN 1000 Series Data Concentrator", Echelon, 2012.
[151] IEEE, "IEEE Guide for Smart Grid Interoperability of Energy Technology and Information Technology Operation with the Electric Power System (EPS), End-Use Applications, and Loads", n. 2030-2011, IEEE, pp. 1–126, 2011.
[152] M. Weiss, A. Helfenstein, F. Mattern, T. Staake, "Leveraging smart meter data to recognize home appliances", *Proc. IEEE Int Pervasive Computing and Communications (PerCom) Conf*, pp. 190–197, 2012.

[153] ISO/IEC, "ISO/IEC 27000:2009. Information technology -- Security techniques -- Information security management systems -- Overview and vocabulary", ISO/IEC.
[154] ISO/IEC, "ISO/IEC 27001:2005. Information technology -- Security techniques -- Information security management systems -- Requirements", ISO/IEC.
[155] ISO/IEC, "ISO/IEC 27002:2005. Information technology -- Security techniques -- Code of practice for information security management", ISO/IEC.
[156] ISF, "Information Security Forum's Standard of Good Practice" , http://www.isfsecuritystandard.com.
[157] K. Kent, M. Souppaya, "Guide to Computer Security Log Management"; *NIST Special Publication 800-92*, 72 pages, 2006
[158] IEC/TS, "IEC/TS 62351. Power systems management and associated information exchange - Data and communications security", IEC/TS.
[159] A. Lee, T. Brewer, "Smart Grid Cyber Security Strategy and Requirements"; *NISTIR draft n. 7628*, 236 pages, 2009.
[160] F. Hao, P. Y. A. Ryan, "Password authenticated key exchange by juggling", *Proceedings of the 16th International conference on Security protocols*, Springer-Verlag, Berlin, Heidelberg, pp. 159–171, 2008.
[161] F. Hao, P. Ryan, "J-PAKE: Authenticated Key Exchange without PKI"Gavrilova, M., Tan, C. and Moreno, E., eds., *Transactions on Computational Science XI*, Springer Berlin Heidelberg, pp. 192-206, 2010.
[162] M. Blaze, "Protocol failure in the escrowed encryption standard", *Proceedings of the 2nd ACM Conference on Computer and communications security*, ACM, New York, NY, USA, pp. 59–67, 1994.
[163] C. Borean, "Energy@home: a “User-Centric” Energy Management System", *5th European ZigBee Developers' Conference*, Munich, 2011.
[164] T. Jakobi, T. Schwartz, "Putting the user in charge: end user development for eco-feedback technologies", *IFIP SustainIT2012 conference*, Pisa, 2012.
[165] N. Goddard, J. Moore, C. Sutton, W. J., H. Lovell, "Machine Learning and Multimedia Content Generation for Energy Demand Reduction", *IFIP SustainIT2012 conference*, Pisa, 2012.
[166] F. Bellifemine, "Smart Consumption: the Energy@home approach", *IFIP SustainIT2012 conference*, Pisa, 2012.
[167] S. Genovesi, S. Saponara, A. Monorchio, "Parametric Design of Compact Dual-Frequency Antennas for Wireless Sensor Networks", *IEEE Transactions on Antennas and Propagatio*, vol. 59, n. 7, pp. 2619–2627, 2011.
[168] S. Genovesi, A. Monorchio, S. Saponara, "Compact Triple-Frequency Antenna for Sub-GHz Wireless Communications", *IEEE Antennas and Wireless Propagation Letters*, vol. 11, pp. 14–17, 2012.
[169] A. G. Ruzzelli, C. Nicolas, A. Schoofs, G. M. P. O'Hare, "Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor", *Proc. 7th Annual IEEE Communications Society Conf. Sensor Mesh and Ad Hoc Communications and Networks (SECON)*, pp. 1–9, 2010.

[170] A. Marchiori, D. Hakkarinen, Q. Han, L. Earle, "Circuit-Level Load Monitoring for Household Energy Management", *IEEE Pervasive Computing*, vol. 10, n. 1, pp. 40–48, 2011.
[171] Freescale, "Electronic Tamper Detection Smart Meter Reference Design", 2012.
[172] J. McCullough, "Deterrent and detection of smart grid meter tampering and theft of electricity, water, or gas", Elster, 2010.
[173] Atmel, "ZigBit™ 2.4 GHz Wireless Modules - ATZB-24-A2/B0", 2009.
[174] ZigBee, "ZigBee and Wireless Radio Frequency Coexistence", ZigBee Alliance, 2007.
[175] G. Anastasi, M. Conti, M. Di Francesco, "A Comprehensive Analysis of the MAC Unreliability Problem in IEEE 802.15.4 Wireless Sensor Networks", *IEEE Transactions on Industrial Informatics*, vol. 7, n. 1, pp. 52–65, 2011.
[176] C. Gomez, J. Paradells, "Wireless home automation networks: A survey of architectures and technologies", *IEEE Communications Magazine*, vol. 48, n. 6, pp. 92–101, 2010.
[177] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, E. Cayirci, "Wireless sensor networks: a survey", *Computer Networks*, vol. 38, n. 4, Elsevier North-Holland, Inc., New York, NY, USA, pp. 393–422, 2002.
[178] D.-M. Han, J.-H. Lim, "Smart home energy management system using IEEE 802.15.4 and zigbee", *IEEE Transactions on Consumer Electronics*, vol. 56, n. 3, pp. 1403–1410, 2010.

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

AAL : Ambient Assisted Living

ACL : Access Control List

AFIB : Atrial Fibrillation

AMI : Advanced Metering Information

AMR : Automated Meter Reading

AODV : On-Demand Distance Vector

ASIC : Application Specific Integrated Circuit

ASIP : Application Specific Instruction-set Processor

AXI : Advanced eXtensible Interface

BP : Blood pressure

CAD : Computer Aided Design

CDA : Clinical Document Architecture

CHF : Chronic Heart Failure

CMRR : Common Mode Rejection Ratio

COTS : Commercial Off the Shelf

CSMA-CA : Carrier Sensing Multiple Access - Collision Avoidance

CTP : Central Terminal Point

DNP : Distributed Network Processor

DoS : Denial of Service

DRPA : Default Remote Peer Address

DS : Downstream

DSO : Distribution System Operator servers

DSP : Digital Signal Processor

ECG : Electrocardiogram

EHR : Electronic Health Record

EMU : Error Management Unit

FBA : Fair Bandwidth Allocation

FSM : Finite State Machine

GPIO : General Purpose Input/Output

GPP : General Purpose Processor

GTS : Guarantee Time Slots

GUI : Graphical User Interface

H@H : Health@Home

HAN : Home Area Network

HDL : Hardware Descriptor Language

HR : Heart Rate

IA : Instrumentation Amplifier

ICT : Information and Communication Technologies

IEC : International Electrotechnical Commission

IP : Intellectual Property

ISF : Information Security Forum

ITF : Interface

ITRS : International Technology Roadmap for Semiconductors

MAC : Medium Access Control

MCI : Metacompiler Instruction

MEMS : Micro Electro-Mechanical Systems

MIC : Message Integrity Code

MMS : Manufacturing Message Specifications

MPSoC : Multi Processor SoC

NALM : Non-intrusive Appliance Load Monitoring

NI : Network Interface

NLH : Network Layer Header

NoC : Network-on-Chip

NSM : Network and System Management

OCP : Open Core Protocol

OP : Operating Protocol

OQ : Output Queue

PAKE : Password-Authenticated Key Exchange

PCB : Printed Circuit Board

PGA : Programmable Gain Amplifier

PKI : Public Key Infrastructure

PM : Power Manager

PPG : Plethysmographic Waveform

PSRR : Power Supply Rejection Ratio

QoS : Quality of Service

RCU : Routing Computation Unit

RD : Read

RF : Radio-Frequency

S&F : Store & Forward

SDP : Service Discovery Protocol

SiP : System in Package

SoC : System On Chip

SPI : Serial Peripheral Interface

SpO2 : Oxygen Saturation

SPP : Serial Port Profile

SWT : Stationary Wavelet Transform

TA : Trust Anchor

TLH : Transport Layer Header

US : Upstream

VLSI : Very Large-Scale Integration

VN : Virtual Network

WR : Write

WSN : Wireless Sensor Networks

ZC : ZigBee Coordinator

ZED : ZigBee End Device

ZR : ZigBee Router

LIST OF PUBLICATIONS

International peer-reviewed journals

[J1] Donati Massimiliano, Bacchillone Tony Salvatore, Fanucci Luca, Saponara Sergio, and Costalli Filippo, “Operating Protocol and Networking Issues of a Telemedicine Platform Integrating from Wireless Home Sensors to the Hospital Information System,” *Journal of Computer Networks and Communications*, vol. 2013, Article ID 781620, 12 pages, 2013. doi:10.1155/2013/781620

[J2] Fanucci Luca, Saponara Sergio, Bacchillone Tony Salvatore, Donati Massimiliano, Barba Pierluigi, Sánchez-Tato Isabel, Carmona Cristina, "Sensing Devices and Sensor Signal Processing for Remote Monitoring of Vital Signs in CHF Patients", *IEEE Transactions on Instrumentation and Measurements*, pp. 1-17, 2012, doi:10.1109/TIM.2012.2218681.

[J3] Saponara Sergio, Bacchillone Tony Salvatore, Petri Esa, Fanucci Luca, Locatelli Riccardo, Coppola Marcello, "Design of a NoC Interface Macrocell with Hardware Support of Advanced Networking Functionalities", *IEEE Transactions on Computers*, vol. 61/11,pp 1-13, 2012, doi:10.1109/TC.2012.70.

[J4] Sergio Saponara, Tony Bacchillone, "Network Architecture, Security Issues, and Hardware Implementation of a Home Area Network for Smart Grid", *Journal of Computer Networks and Communications*, vol. 2012, Article ID 534512, pp 1-19, 2012. doi:10.1155/2012/534512.

Conference proceedings

[C5] Saponara Sergio, Donati Massimiliano, Bacchillone Tony Salvatore, Fanucci Luca, Sanchez Isabel, Carmona Cristina, Barba Pierluigi, "Remote monitoring of vital signs in patients with chronic heart failure, sensor devices and data analysis perspective", IEEE *Sensors Applications Symposium (SAS-2012)*, pp.1-6, Brescia, Italia, 2012, doi:10.1109/SAS.2012.6166310.

[C6] Bacchillone Tony Salvatore, Donati Massimiliano, Saponara Sergio, Fanucci Luca, "A flexible home monitoring platform for patients affected by chronic heart failure directly integrated with the remote Hospital Information System", *SPIE Microtechnologies 2011*,vol. 8068,pp 8068-26,tot.pag 8,2011, doi:10.1117/12.886465.

[C7] Bacchillone Tony Salvatore, Donati Massimiliano, Saponara Sergio, Fanucci Luca, "A flexible home gateway system for telecare of patients affected by chronic heart failure", *IEEE Medical Information & Communication Technology (ISMICT)* 2011, pp 139-142, Montreaux,2011, doi:10.1109/ISMICT.2011.5759814.

[C8] Fanucci Luca, Bacchillone Tony Salvatore, Donati Massimiliano, Saponara Sergio, Passino Claudio , Costalli Filippo , Petrucci Stefano , Sanchez-tato Isabel , Pascual Francisco , Zlatko Vukovic , Hrvatin Orjana, Health@Home: lesson learnt and future perspective in the home monitoring of patients affected by chronic heart failure, *Ambient Assisted Living Forum*, pp 3-5, Lecce, Italia,vol. 1, 2011.

[C9] Fanucci Luca, Saponara Sergio, Donati Massimiliano, Bacchillone Tony Salvatore, A flexible home gateway system for telecare of patients affected by chronic heart failure, *Ambient Assisted Living Forum*, Lecce, Italia, 2011.

Filed Patents

[P10] Filed US patent: A.-M. Coppola, R. Locatelli, E. Petri, L. Fanucci, S. Saponara, T. Bacchillone. "Zero-Cycle Router for Networks On-Chip", 2012.

[P11] Filed EU patent: V. Catalano, A.-M. Coppola, R. Locatelli, E. Petri, L. Fanucci, S. Saponara, T. Bacchillone. "Automated metacompiler for designing a configurable hardware IP component", 2012.

[P12] Filed EU patent: A.-M. Coppola, R. Locatelli, E. Petri, L. Fanucci, S. Saponara, T. Bacchillone. "Networks on-chip router", 2011.

Awards

[A13] The H@H project was among the four finalists at the AAL PROJECT AWARD 2012, which recognizes the most successful projects of the Ambient Assisted Living Joint Programme demonstrating great promise in terms of innovation, human-centric approaches to development and market potential.

[A14] The H@H project obtained the second prize in element14 Medical Design Award on October 2011 at the Medical Technology Event, organized by Selezione di Elettronica (magazine by Il Sole 24 ORE Group).