

Research Projects

IC and System Design and Test

Coordinator:

Norbert Felber

Reconfigurable Network Processor for LAN/WAN Bridging of Multimedia Networks

Personnel: Thomas Bösch, Eric Roth; Manfred Stadler (BridgeCo)

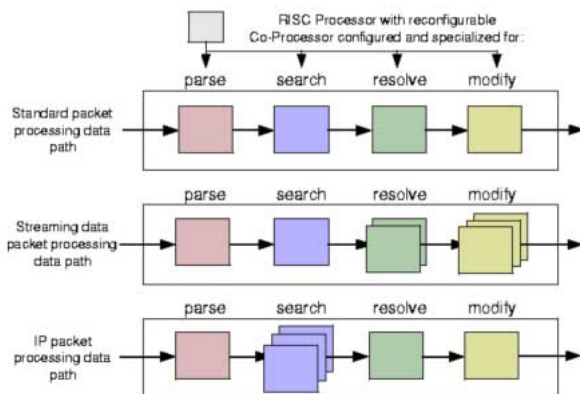
Funding: KTI 4957.2 LANWAN, BridgeCo

Partners: BridgeCo, TIK-ETHZ

Support of timing-critical multimedia content on today's networks becomes an increasingly important issue. However, current network architectures do not satisfy these requirements. Demands are hard Quality of Service (QoS), synchronization and guaranteed high streaming data bandwidth as well as a maximum of flexibility and scalability.

To connect multimedia networks with wide area network technologies, specialized access router devices with high-performance network processors for packet handling are required. Multimedia packet forwarding has to deal with a wide range of different protocols and complex packet manipulations.

Therefore a new network processor architecture is being developed using a combination of RISC cores and reconfigurable co-processors. Data packets are passed through several processing steps like header parse, search, resolve and modify. Each processing step is performed by one or several processor units configured for this task. Each processor unit can be adapted to a specific task by hardware reconfiguration. These hardware structures are rearranged to outsource complex tasks from the RISC core to the hardware. This approach can be compared with modifying and adapting the RISC processors instruction sets on the fly.



Processor configurations for different packet forwarding tasks.

Synchronization of an IEEE-1394 Network for Audio and Video Applications

Personnel: Eric Roth, Thomas Bösch; Georg Dickmann (BridgeCo)

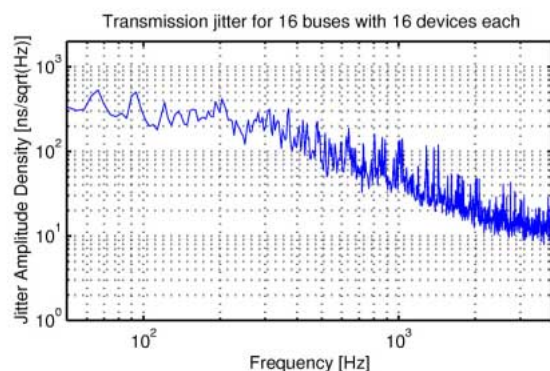

Funding: KTI 4957.2 LANWAN, BridgeCo

Partners: BridgeCo, TIK-ETHZ

Audio and video applications require a notion of time to record and reproduce audio samples and video frames at the correct instants. Phase and frequency of the sample points must be aligned on all devices belonging to the same network. The goal of this project is to develop and implement a synchronization method for a IEEE-1394 (FireWire™) network.

Timing information is broadcast through the IEEE-1394 network such that all devices adjust their sampling rates to a common reference. Bridges that connect 1394 buses replicate the timing information from one bridge portal to the next. Several approaches have been analyzed with respect to time accuracy, jitter performance, and reliability. Much effort was made to influence the standardization of IEEE-1394 bridges to open 1394 technology to professional audio and video applications.

To approach this goal a very flexible recovery circuit which consists of a partially on-chip phase-locked-loop has been implemented that attenuates transmission jitter and generates high-quality audio clocks.



Typical IEEE-1394 network configuration (top) and transmission jitter caused by transmitting timing information across 16 buses with 16 devices each (bottom).

Globally Asynchronous Locally Synchronous (GALS) Design Simplifies Systems on Silicon

Personnel: Thomas Villiger


Funding: Infineon

Partners: Infineon

The continuing technical progress in CMOS technologies, the increase in available die size, and the demand for short time to market ask for integration of large systems on a single chip. The increase in clock frequency and area make proper clock distribution difficult. Complex systems often require a multitude of clock frequencies on a common die, therefore asking for careful use of synchronizers between independent clock domains.

The Globally-Asynchronous Locally-Synchronous (GALS) design methodology facilitates clocking of large systems on silicon since it partitions the circuit into several independently clocked modules that communicate in self-timed fashion. Every synchronous module is surrounded by a so-called asynchronous wrapper that manages all the data transfers. The actual functionality, however, remains located in the locally-synchronous blocks, such that a well established design flow can be used. The local clocks are generated inside the wrappers. When transferring data across clock boundaries a major concern is the danger of metastability. In this approach the problem is addressed by giving the port controller the ability to pause the local clock if data and sampling clock edges occur too close to each other. Metastable data is thus prevented and not resolved as with other known solutions.

Partitioning into modules results in enhanced modularity and reusability and the decoupled timing constraints between modules offer more options for optimization.



Simplified block diagram of an asynchronous wrapper.

SAFER SK-128 Crypto Algorithm as World's First GALS System on Silicon

Personnel: Jens Muttersbach, Thomas Villiger


Funding: KTI-3650.1 DSP, Philips Zurich, Infineon

Partners: Philips Zurich, Infineon

To verify the Globally-Asynchronous Locally-Synchronous (GALS) design methodology a complete chip was implemented in this technique. A secret-key iterated block-cipher algorithm called SAFER was chosen as application. SAFER stands for Secure And Fast Encryption Routine, an algorithm invented by James Massey at ETH. Information is encrypted or decrypted in data blocks of 8 bytes by repeatedly applying the algorithm in up to 12 rounds. The implementation supports the currently strongest version, SK-128, that employs a strengthened key scheduling based on a 128-bit key. The chip supports four cryptographic modes: ECB, CBC, CFB and OFB.

The partitioning chosen mainly reflects the inherent functional groups of the system. The implementation ends up in having five different clock domains, several lookup-tables, an asynchronous RAM, and an asynchronous FIFO that serves as feed-forward buffer and is an example for a completely asynchronous block included in the GALS environment.

The chip called Marilyn was fabricated in a 0.25 μm 5-metal technology and compared to a strictly synchronous counterpart processed on the same wafer. A data throughput up to 780 MBit/s was obtained in the GALS version and an energy saving of 25% compared to the already power-optimized synchronous version was achieved.



Microphotograph of the Marilyn chip.

SAFER SK-128 Cipher Algorithm: a Fully Synchronous Reference Design

Personnel: Marc Oberholzer,
David Studer (students);
Jens Mutersbach, Thomas Villiger

Funding: ETHZ


Partners: Philips Zurich, Infineon

For the comparison of common synchronous designs with the new GALS technique, the same function as in the first GALS implementation has been realized in a student design project. The most important parameters to compare were energy consumption, throughput, and silicon area.

Emphasis has therefore been put on a realization that lies as closely as possible to the GALS version. Circuitry from this implementation has then been re-used as locally-synchronous islands in the GALS design. Gateable clocks have been implemented for fair energy comparisons.

The chip called MERLIN has been fabricated in 0.25 μm technology on the same wafer as the GALS version MARILYN. It features a 64-bit datapass controlled by six finite-state machines and with 128 by 16 bits ram. Full testability is ensured by built-in Ram test and full-scan test access. The 2.89 mm² chip is packaged in a 68-pin chip carrier.

In comparison to this synchronous chip the GALS version saves some 20% energy per encrypted data item at comparable speed. The area consumption of the GALS overhead is 9%.



Chip photograph of the fully synchronous SAFER SK-128 reference design.

Testability of Globally Asynchronous Locally Synchronous Wrapper Modules

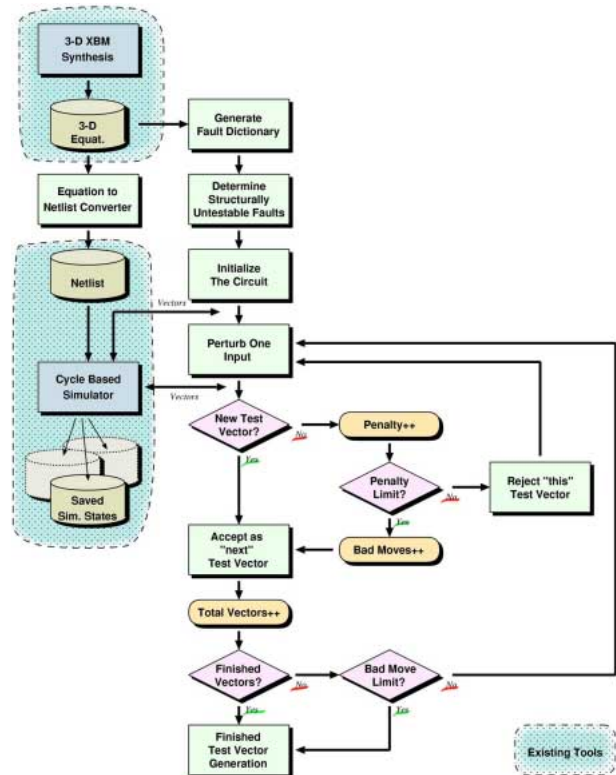
Personnel: Jens Mutersbach, Frank Gürkaynak

Funding: KTI 4897.1 IRRO,
Philips Zurich

Partners: Philips Zurich

Globally-Asynchronous Locally-Synchronous (GALS) architectures have the potential to overcome clock distribution and synchronization problems associated with the design of large Systems on a Chip (SoC). This project addresses the testability problem of GALS-based systems which is not yet solved.

The test methodology basically needs to verify the proper operation of GALS wrapper modules, which contain asynchronous finite state machines (AFSM). Unlike synchronous circuits, asynchronous circuits are susceptible to glitches on internal signals and to path delay faults. In addition, AFSM design methodologies produce circuits that include feedback loops and are not redundancy-free, which further aggravates the testing problem. Current research efforts are concentrated on the development of new software that will allow a seamless integration of testability for GALS systems within industry standard design flows.



The flow chart of the proposed test vector generation algorithm.

IRRQ — a Replacement for IDDQ CMOS Testing


Personnel: Matthias Brändli

Funding: KTI 4897.1 IRRQ,
Philips Zurich

Partners: Philips Zurich

For many years, quiescent current testing (IDDQ) has been a very effective method for detecting fabrication defects in static CMOS circuits. The problem with today's deep-submicron CMOS technologies is that the fault-induced short-circuit currents are exposed to masking by the superimposed MOSFET leakage currents. The phenomenon is bound to become more serious with upcoming process generations as the overall leakage current increases with growing transistor count and diminishing MOSFET structure sizes. As a result, traditional IDDQ testing is no longer possible below a feature size of 0.25 μm or so.

In this project methods for CMOS testing are investigated which can be realized locally, are suitable for built-in self test and other test methods, such as scan path, and provide an effective solution for the future technologies' leakage levels. The investigation on a specific solution called IRRQ is in progress. A patent has been filed on this principle.



Basic fault detection principle of IDDQ testing (white: fabrication defect).

An Integrated Wideband MIMO Receiver for Wireless Multimedia Communication


Personnel: Andy Burg

Funding: Lucent

Partners: Lucent

Wireless communication has become one of the fastest growing fields in the consumer electronics market. New ideas such as multimedia communication are the driving force behind the development of new systems. Especially real-time video and high-quality audio streaming require networks that support very high data rates for a growing number of users. In addition to improvements of the network infrastructure new techniques have to be developed and tested to efficiently and reliably transmit video and audio over these networks.

One possibility to increase the data throughput of communication systems is to use multiple antennas at the transmitter and at the receiver side. Such systems are called MIMO (Multiple Input Multiple Output) systems. Lucent has developed such a technique, called BLAST. The initial goal of this project is the development and implementation of such a system for next-generation wireless networks. The work is carried out in a research collaboration with the Bell-Labs wireless research lab in New Jersey. So far, a prototype of a standard single-antenna UMTS downlink system with transmitter and receiver has been developed making use of a rapid-prototyping methodology. It will be used to investigate advanced receiver algorithms to increase single-antenna system performance and as a starting point for the development of a MIMO-receiver implementation. Based on such a system, multimedia applications can be developed and tested in the future in a real-world environment.



Old and modern telecommunication hardware.

ADAMO — ADaptive Antenna for MOBiles: Adaptive Antenna Training Engine

Personnel: Boris Glass, Bruno Haller

Funding: ESPRIT-27001 ADAMO, BBW


Partners: Ascom Systec, ENSTA, IMST, Thomson

The objective of the European project ADAMO was to validate the potential of small, low-cost adaptive antennas for HIPERLAN I (= ESTI standard for wireless LANs) modems. Improvement of performance and reduction of power consumption are the expected benefits of adaptability.

Instead of using a complex, high-speed (47 Msamples/s) adaptive equalizer (= digital transversal filter) to mitigate the detrimental effects of intersymbol interference caused by multipath propagation, spatial filtering is applied by appropriately weighting and combining the RF signals from three independent antennas.

This "beam-forming network" realized with analog phase shifters and amplifiers consumes considerably reduced power compared to a digital equalizer-based solution.

IIS' contribution was the implementation of the digital signal processing algorithm for the computation of the combiner coefficients of the adaptive antenna. A dedicated ASIC - the AdAnTE chip - has been designed. This high-speed digital chip with 780,000 transistors calculates the set of coefficients on a linear systolic array of pipelined CORDIC (COordinate Rotation DIgital Computer) processor elements in less than one microsecond.



The digital weight controller ASIC AdAnTE proves its full functionality in the demonstrator hardware.

IP Core for Real-Time Solution of Least-Squares Problems using CORDIC and QRD-RLS


Personnel: Boris Glass

Funding: ETHZ

The famous least-squares method was introduced by Gauss in 1795 for the use in survey evaluation and astronomical calculations. Nowadays image processing, robotics or communications are requiring real-time solutions of least-squares problems.

This project arose from the development of a wireless LAN receiver using an adaptive antenna array. The amplitudes and phases of the antenna mixer are determined by comparing received training sequences to the expected undisturbed data. By far the most expensive part of the processing is the least-squares calculation. The combination of QR (triangular) decomposition and recursive least-squares computation is known as the QRD-RLS algorithm (QR-Decomposition Recursive-Least-Square). It is robust, effective, and an implementation using CORDIC processors maps very well to VLSI.

The developed processor array is reusable as an IP core since it is largely scalable. A graphical user interface simplifies to handle all parameters: Transformations from fully parallel to fully sequential processing through several circuit levels can be selected. This enables customizations that closely meet the requested throughput with minimal area consumption. A preview shows the resulting architecture as well as the arrays overall area. Finally, the array is exported as Matlab and VHDL source.



User interface of QRD-RLS IP core overlaid on smart antenna filtering ASIC photograph

Energy-Efficient Processing of Speech Data

Personnel: Jürgen Wassner, Hubert Kaeslin

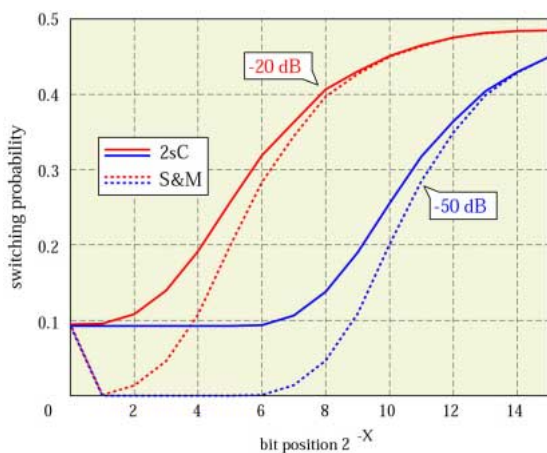
Funding: ETHZ, KTI-5025.1 Bernafon-DSP, Bernafon

Partners: Bernafon

Next to volume, power dissipation is the most stringent constraint for digital signal processing in state-of-the-art hearing aids. This project centers around the question how to take advantage of the statistical characteristics of speech signals to reduce power consumption.

The correlation between various signal properties and overall bit switching activity under different coding schemes has been the subject of an experimental study conducted with real-world application data. The signal parameters investigated included signal variance (AC power), signal-to-noise ratio (SNR), sampling frequency, and quantization resolution. Linear and differential PCM with 2's-complement (2sC) and sign-magnitude (S&M) representations were considered as number formats. Among other things, we have found that S&M provides consistent activity savings over 2sC, especially for the low-amplitude signals that dominate in practice. Also, differential encoding was shown to be beneficial only in conjunction with S&M coding.

In addition, waveform coding techniques known from low-bit-rate telecommunications have been explored for computational tasks. An FIR filter implemented as a dedicated hardware architecture has been chosen to quantify potential power savings subject to application parameters such as bit rate, processing accuracy, and filter order. For the application at hand, our results essentially confirm linear PCM in S&M representation as the optimum choice.



Effect of number representation and signal amplitude on bit-switching probability.

A New Approach for Controlling Series-Connected IGBT Modules

Personnel: Jan Thalheim

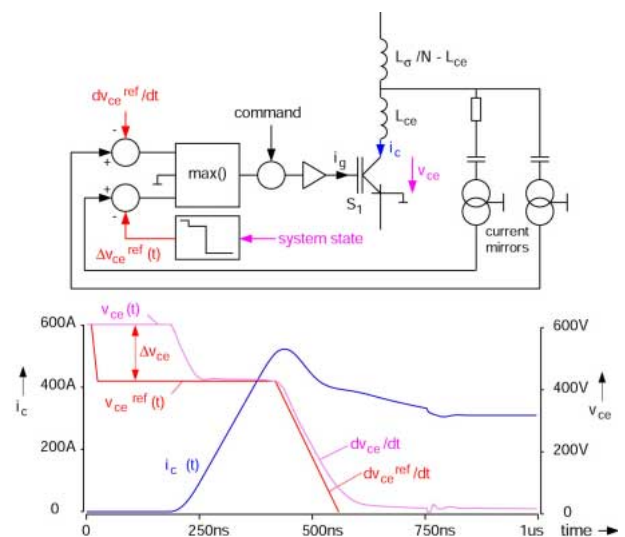
Funding: ETHZ, KTI-3367.1 SIGU

Partners: CONCEPT

Parallel and series connection of medium-power IGBT modules is necessary to achieve the capacity or to exceed the speed of today's SCRs and GTOs. A reconfigurable system using this approach might enhance system reliability compared to systems utilizing single switches.

Hierarchical control of such systems is being investigated. An advanced system-state classification is applied to allow for real-time reconfiguration of the local controller structure. Thereby, a new scheme for controlling the rate of increase of the collector current is introduced. The main object is to control the voltage drop across the system stray inductance while the freewheeling diode is forward biased. The state of diode reverse biasing is derived from the systems state variables and determines the starting point of the main transition of the collector-emitter voltage.

Synchronization of multiple modules is based on an evaluation of the time difference between a globally synchronous and an asynchronous event, which both are derived locally. This avoids the need of high timing resolution in the optical links for the definition of a global sampling time. To attain sufficient accuracy, signal processing is generally based on a small deviation of signals from an operating point taken prior to the switching process.



Concept of turn-on control qualified for series-connected IGBT modules.

Programmable Intellectual Property Modules for System Design by Reuse

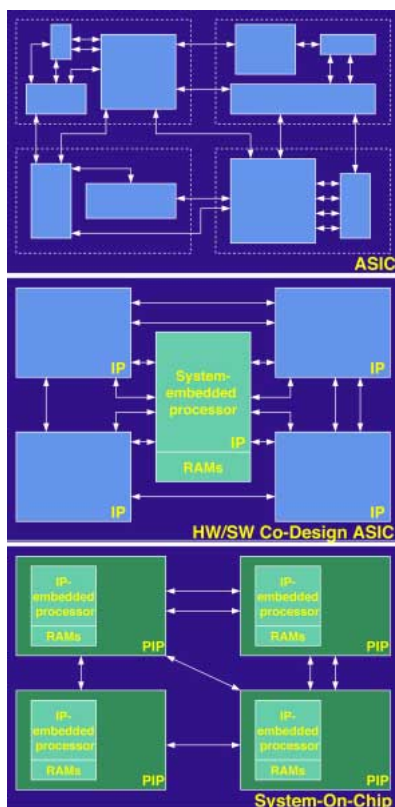
Personnel: Thomas Röwer

Funding: ETHZ, KTI-3297.2.1 ITRASYS,
Siemens Schweiz

Partners: Siemens Schweiz

Hardware/software co-design is one of the most prominent methodologies used to improve the design efficiency and quality. However, system design by reuse is problematic in combination with system-level hardware/software co-design. The tight coupling between a system processor and several hardware blocks does not fit well into a system design flow based on several independent IP blocks.

We propose the concept of programmable intellectual property modules (PIP) to solve this problem. The key innovation is to integrate a processor into each IP module. PIP modules offer the possibility to realize highly reusable IP modules that have superior properties because the embedded processor can be used for several system design tasks. PIP modules fit excellently into a system design flow based on system-level functional partitioning and IP-level hardware/software partitioning.



Evolution of ASIC design from custom-specific blocks to highly reusable programmable intellectual property modules.

Example Implementation of a Programmable Intellectual Property Module

Personnel: Thomas Röwer


Funding: ETHZ, KTI-3297.2.1 ITRASYS,
Siemens Schweiz

Partners: Siemens Schweiz

As an example for the design of programmable intellectual property modules (PIP), a STM-1/STS3 input block was designed. STM-1 and STS3 are the 155 MHz/s hierarchies in the Synchronous Digital Hierarchy (SDH) and Synchronous Optical NETWORK (SONET) standard, respectively. The goal was to design a highly parametrizable soft module that offers maximum flexibility for system integration and customer-specific features. Additionally, the module is usable in both SDH and SONET networks, without any changes to the hardware. By downloading appropriate software, its application for SDH or SONET can be selected after system integration.

Other major benefits of PIP modules that were proven by this implementation are:

- The data transfer protocol of the interfaces that connect PIP modules to their neighboring blocks can be implemented in software running on the embedded processor.
- PIP modules offer enhanced test coverage by running a software-driven BIST using deterministic patterns in addition to a random test.
- The functionality of PIP modules can be altered by a software download. So, even hard-IPs become highly flexible when implemented according to the proposed paradigm.



Chip photograph of the implemented STM-1/STS3 programmable intellectual property module.

Functional Verification of Programmable Intellectual Property Modules

Personnel: Thomas Röwer

Funding: ETHZ, KTI-3297.2.1 ITRASYS,
Siemens Schweiz

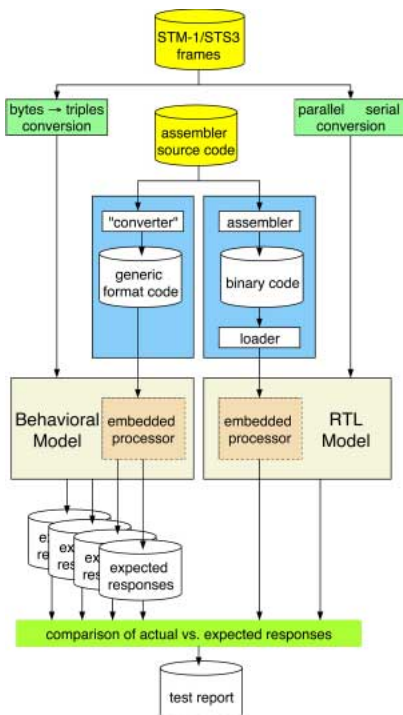
Partners: Siemens Schweiz

Application- or customer-specific scalability introduces severe difficulties into the functional verification flow. The problems are even more stringent for PIP modules since part of the functionality is implemented in software. Designing a test-bench that scales according to the IP solves only part of the problem. The more critical problem is the generation of test vectors that work for every possible IP configuration. To enable this, we propose the following two-step functional verification flow.

The basic idea of this flow is to generate the expected responses for the selected configuration of the IP from configuration-independent stimuli. In the first step the behavioral model is used to generate the test patterns (left side of figure).

The behavioral model is split into processor and data path, just like the RTL model. This enables the generation of test vectors even for features implemented as custom software using a single test bench.

The second step (right side of figure) is the actual verification of the RT-level IP model. In this step the outputs of the RTL model are compared against the generated expected responses from the first step.



Two-step functional verification by simulation.