A Coarse-Grained Reconfigurable Wavelet Denoiser Exploiting the Multi-Dataflow Composer Tool

N. Carta, C. Sau, F. Palumbo, D. Pani and L. Raffo

DIEE - Dept. of Electrical and Electronic Engineering

University of Cagliari, Italy

Conference on Design and Architectures for Signal and Image Processing

October 8-10, 2013
Cagliari, Italy

Introduction

- Some application-specific fields, such as the biomedical one, push towards **area-** and **power-** effective VLSI implementation of digital signal processing algorithms mainly for the development of wearable or implantable devices.

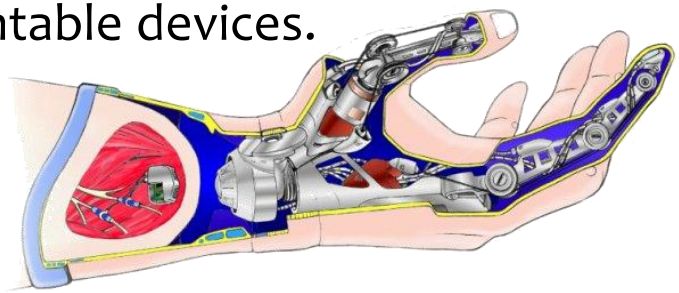

- To this aim, the **Multi-Dataflow Composer** (MDC) tool, originally conceived for image/video processing, can be exploited in other domains leading to optimal hardware implementations.
- The MDC-based hardware platforms are **flexible**, fitting to parametric solutions adjustable at runtime, according to the characteristics of the signal of interest.

Research goals

- Exploit strategies of resources **reusability** through **reconfiguration** to provide area and power minimization.
- Demonstrate the potential **orthogonality** of the MDC approach with respect to the Reconfigurable Video Domain (RVC) domain.
- Test the proposed approach on a runtime reconfigurable **Wavelet Denoiser**, targeted for biomedical applications, prototyped onto an **FPGA** Development Board.
- Evaluate benefits in terms of **area** and **power** in comparison to traditional solutions and verify the system functionality using simulated datasets of neural signals.

Multi-Kernel Datapath Generation

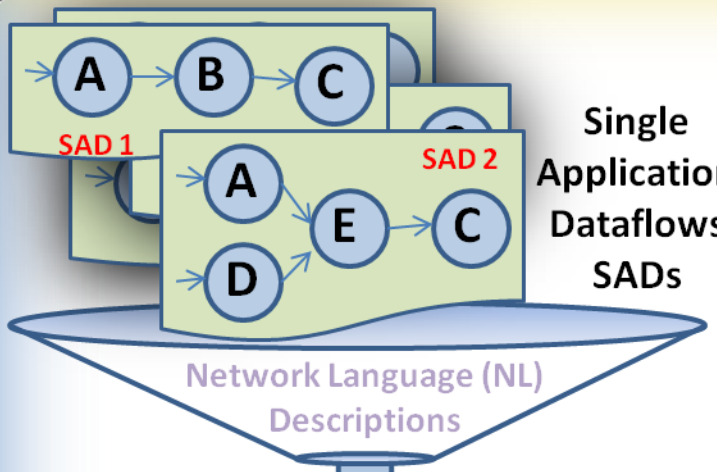

HDL Components library, the *Functional Units* (FUs) implementing the actors:

- IEEE 754-1985 compliant arithmetic modules;
- control flow modules.

All the FUs obey to the same FIFO-based communication protocol.

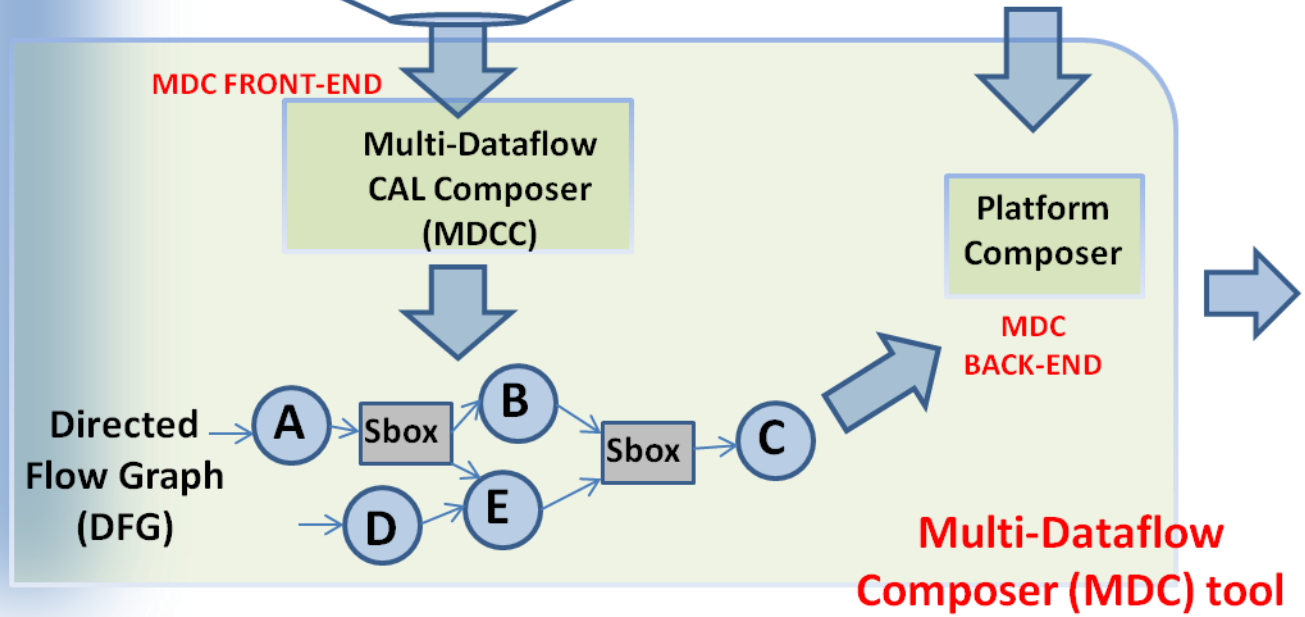

Global Kernel

a Coarse-Grained Reconfigurable Hardware Platform with the minimum number of FUs

Use Case: Wavelet Denoising (WD)

- WD aims at **removing** the **background noise** corrupting the signal of interest, especially when it can be modeled as a Gaussian-distributed random noise.
- Translation Invariant solution: *à-trous* algorithm implementation.

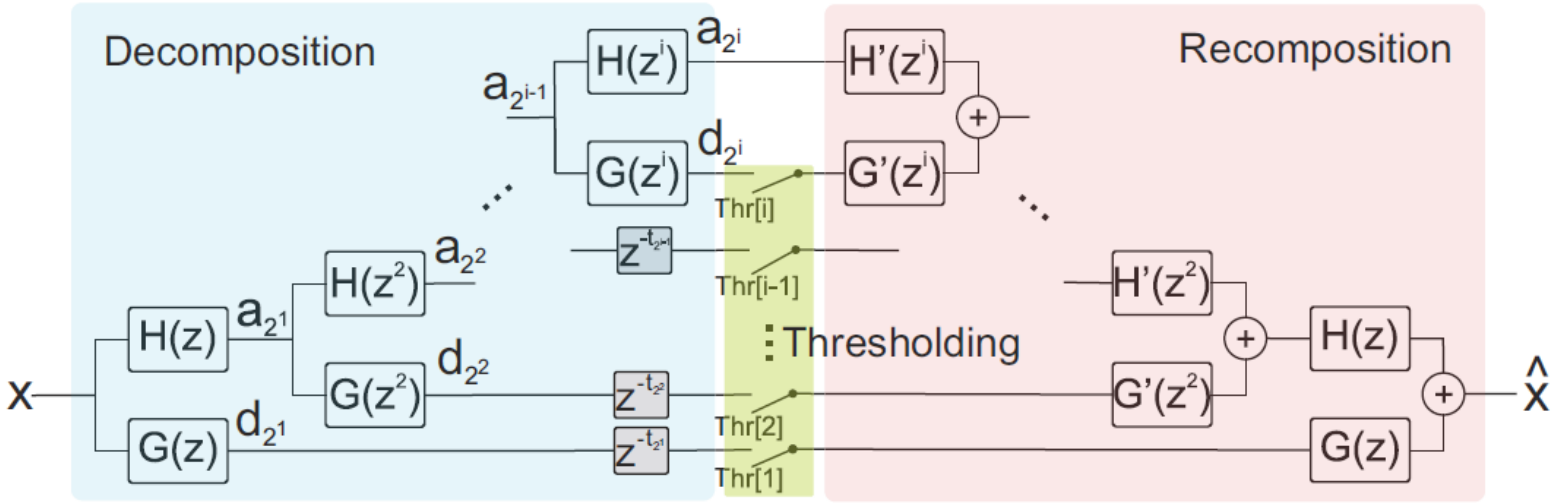

- The number of Decomposition/Recomposition levels depends on the chosen input frequency and on the spectral characteristics of both noise and signal of interest.

Use Case: Wavelet Denoising (WD)

- WD aims at **removing** the **background noise** corrupting the signal of interest, especially when it can be modeled as a Gaussian-distributed random noise.
- Translation Invariant solution: *à-trous* algorithm implementation.

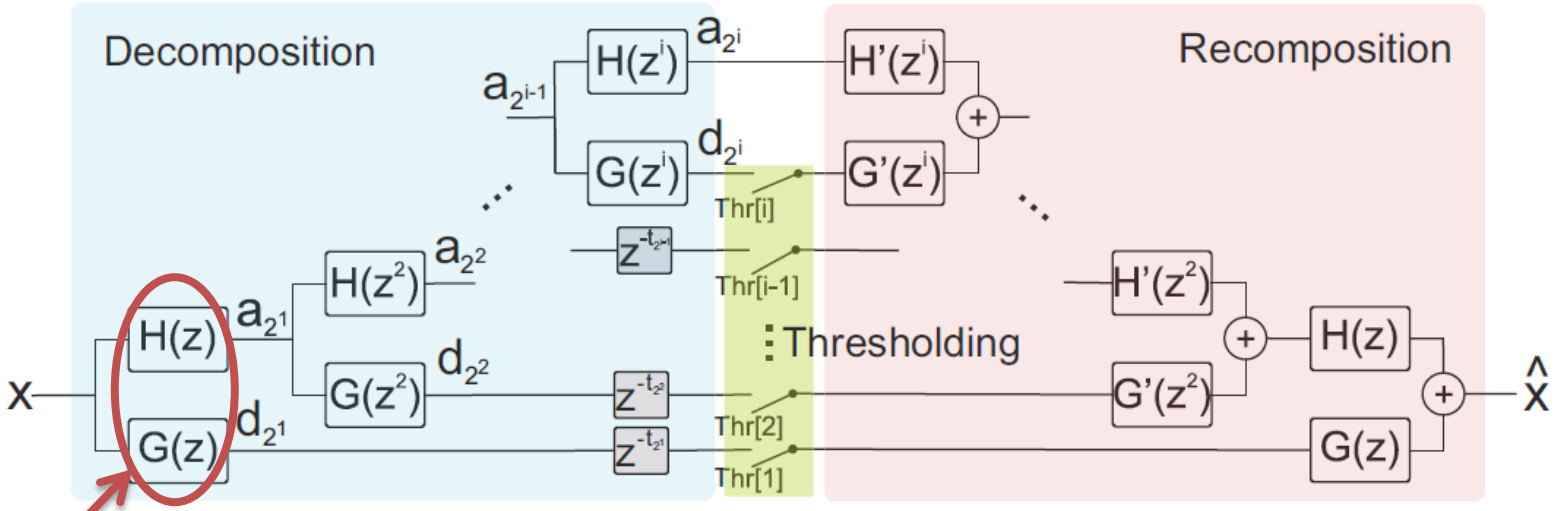

A pair of quadrature mirror Finite Impulse Response (FIR) filters at each level

Use Case: Wavelet Denoising (WD)

- WD aims at **removing** the **background noise** corrupting the signal of interest, especially when it can be modeled as a Gaussian-distributed random noise.
- Translation Invariant solution: *à-trous* algorithm implementation.

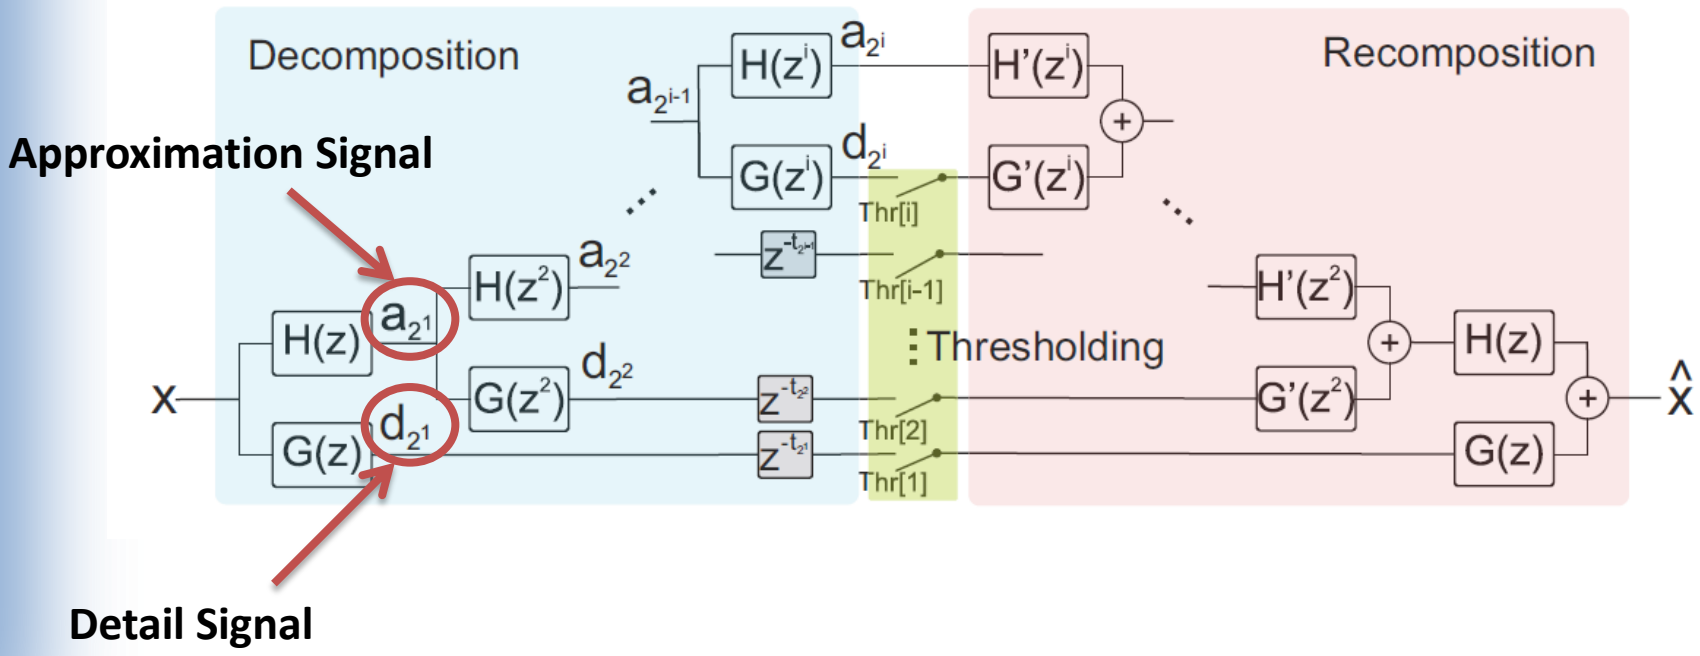

Use Case: Wavelet Denoising (WD)

- WD aims at **removing** the **background noise** corrupting the signal of interest, especially when it can be modeled as a Gaussian-distributed random noise.
- Translation Invariant solution: *à-trous* algorithm implementation.

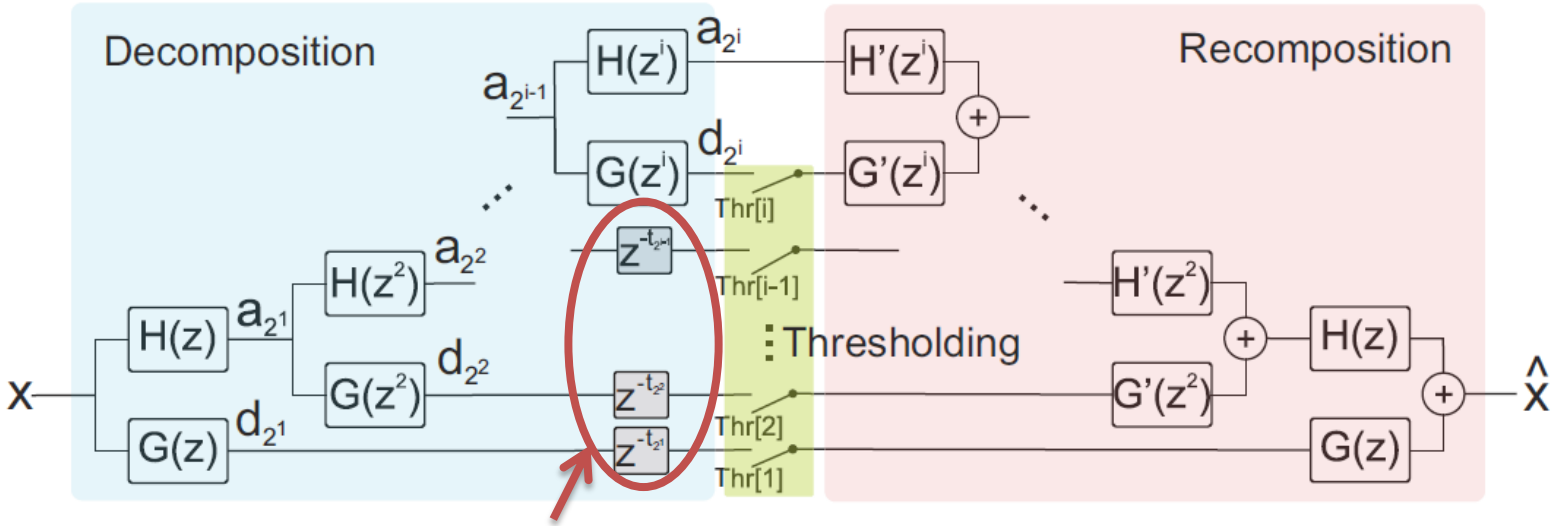

Registers to balance the delay of each path of the trellis

Use Case: Wavelet Denoising (WD)

- WD aims at **removing** the **background noise** corrupting the signal of interest, especially when it can be modeled as a Gaussian-distributed random noise.
- Translation Invariant solution: *à-trous* algorithm implementation.

Band-Pass solution

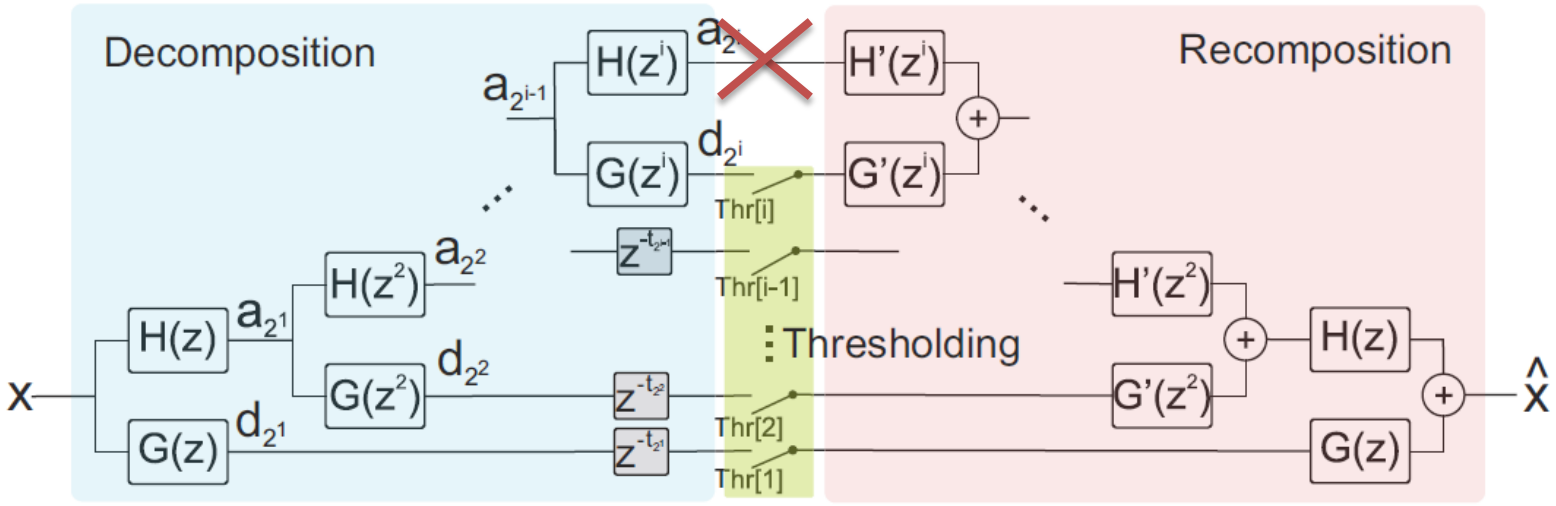

WD: thresholding phase

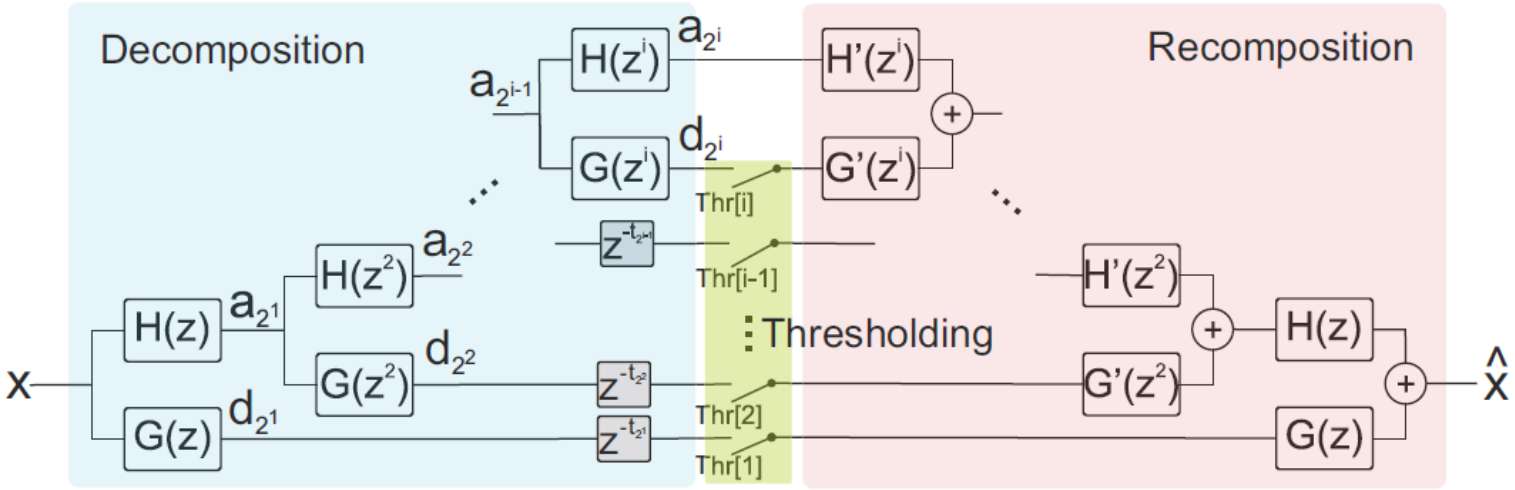

- To minimize the hardware resources, an **hard thresholding** has been implemented: only samples with values greater than θ are passed whereas the others are cleared to zero.
- θ can coincide with a scaled version of the standard deviation σ_n of the additive noise source, computed iteratively analysing consecutive windows of ***b_len* detail samples.**

$$th = SC \cdot \sigma_{2^j} = 3.9\sqrt{\frac{1}{4 * b\_len - 1}\sum_{n=1}^{4} s_{n,2^j}} \qquad s_{n,2^j} = \sum_{k=0}^{b\_len} d^2{}_j[k]$$

Computing Kernels

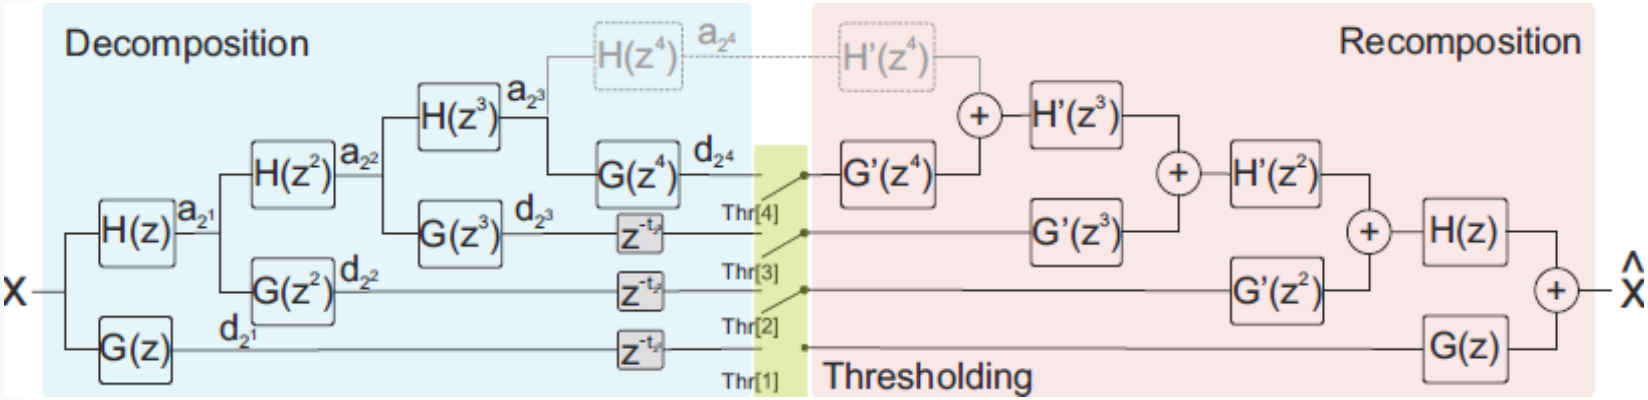

- Considering an input sampling frequency of 12kHz, 4 levels of decomposition/recomposition and removing the approximation signal at the 4th level, the denoised signal has a bandwidth between 375Hz and 6kHz.
- Different computational kernels can be identified for the MDC tool. They should present commonalities at the actor level, enabling **reuse** of **common** hardware FUs.

Computing Kernels

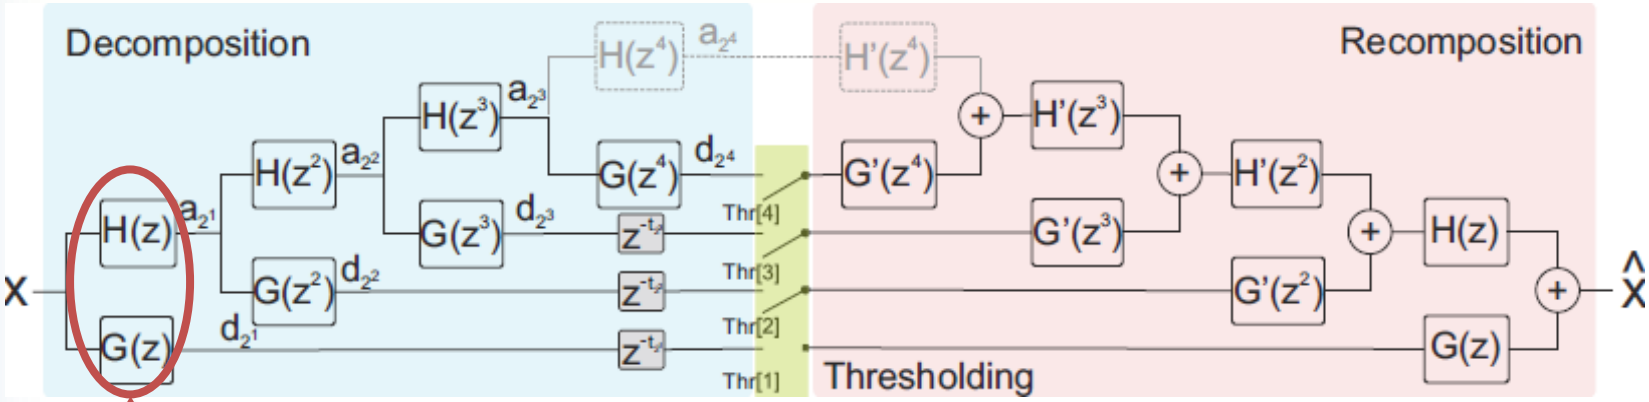

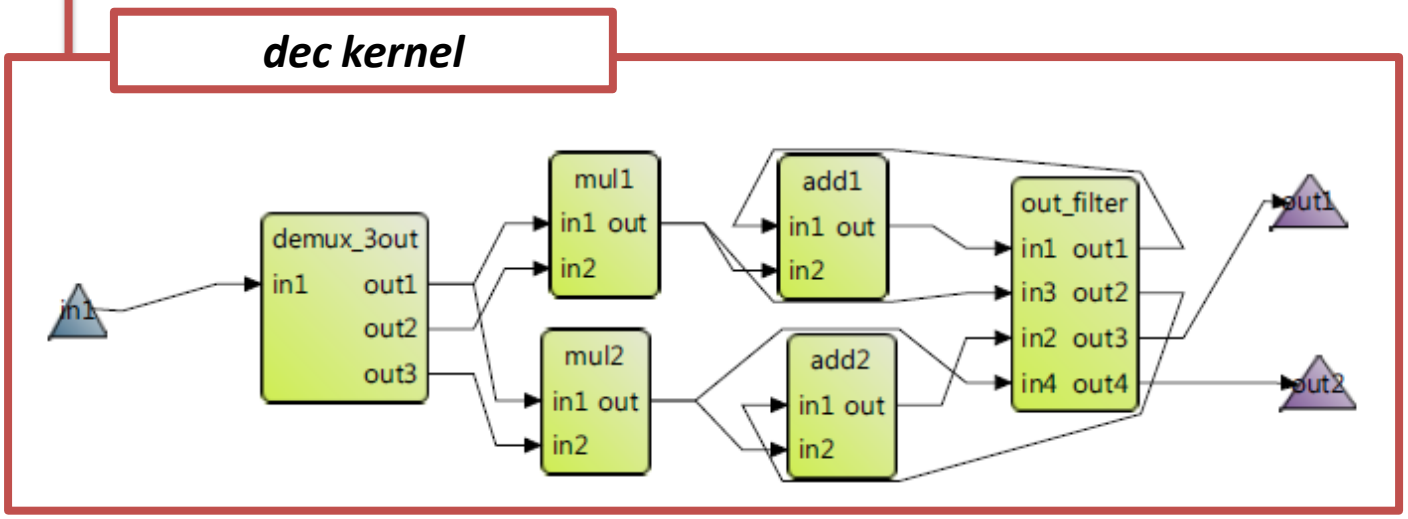

Computing Kernels

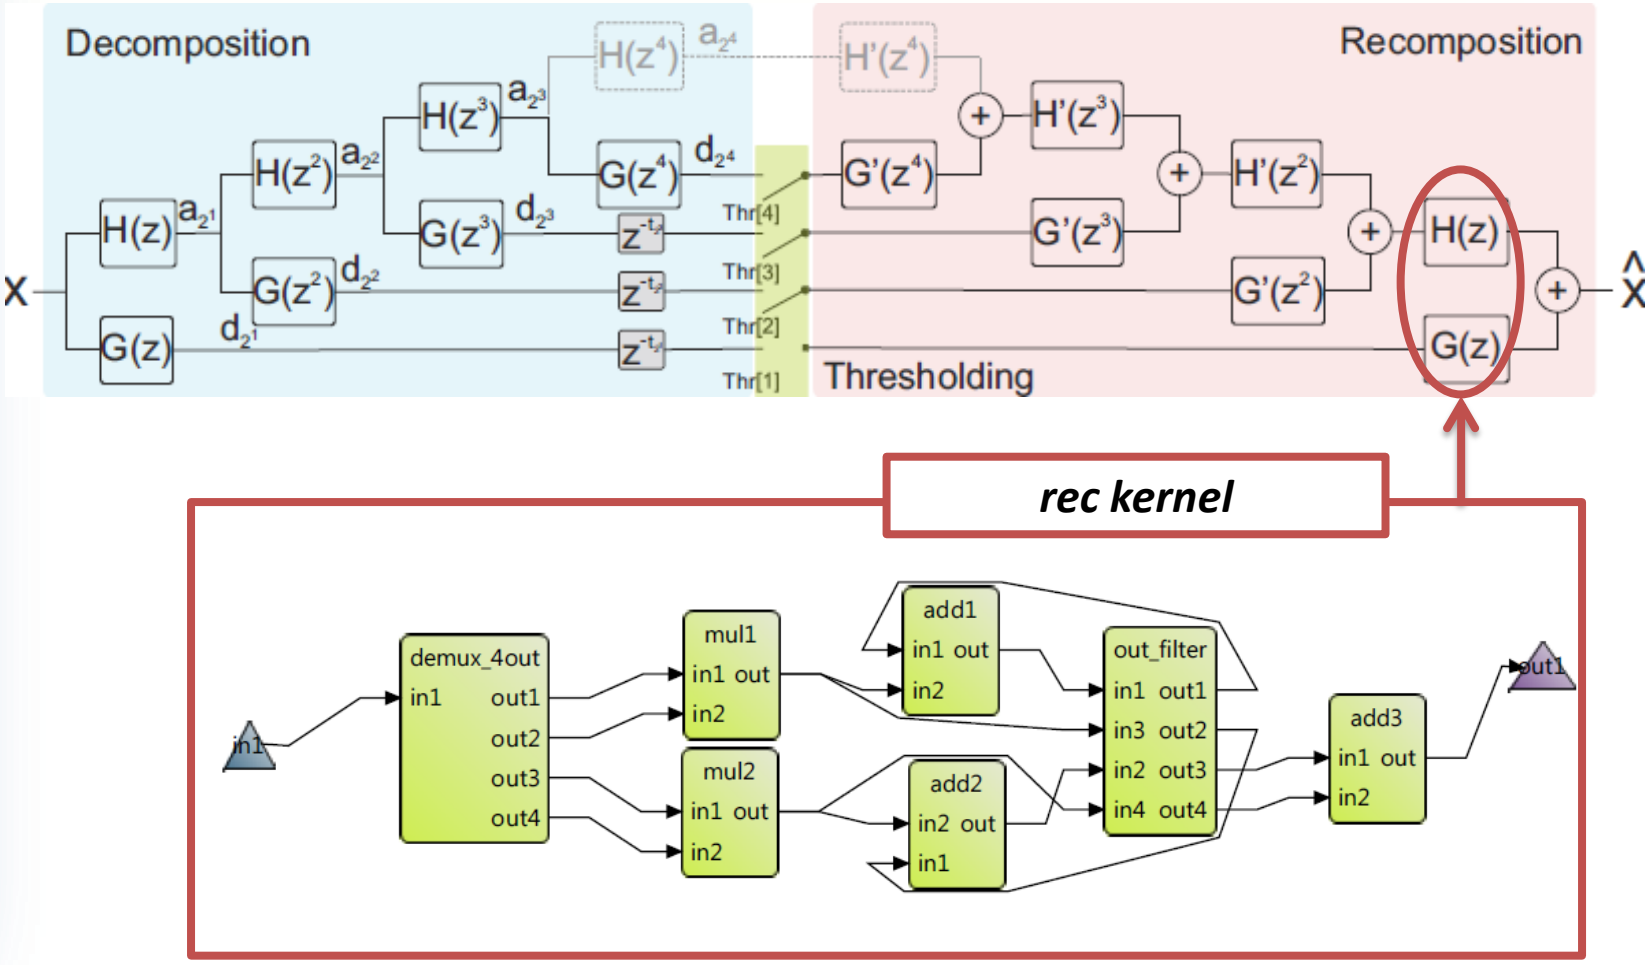

Computing Kernels

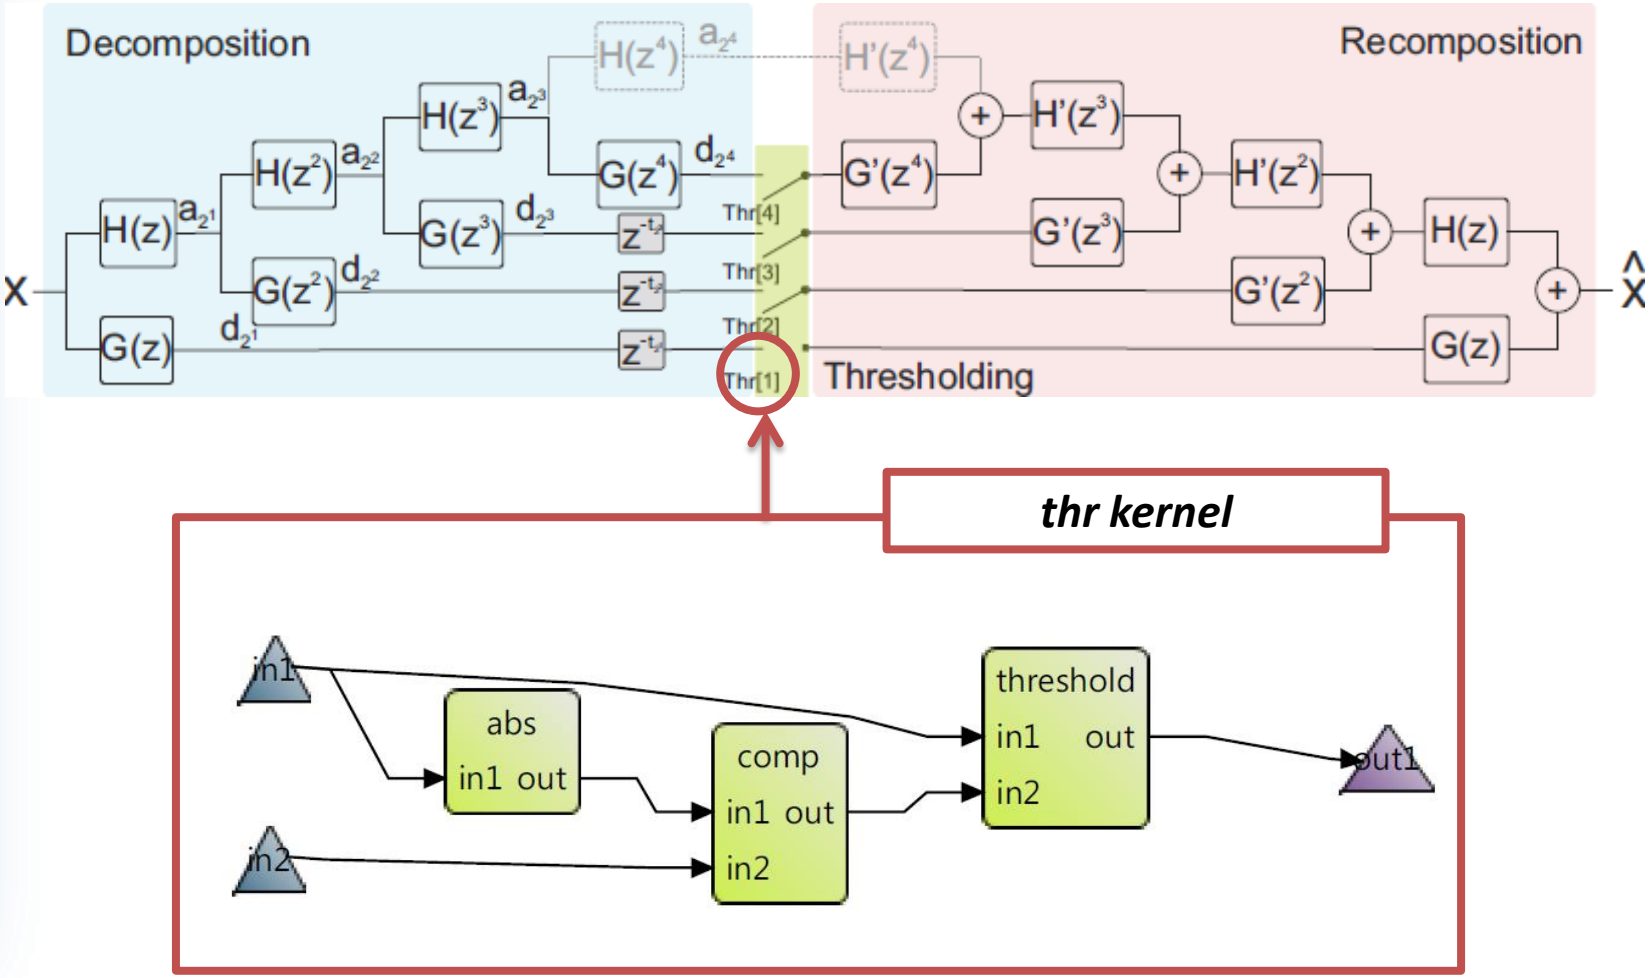

Test Environment

- To evaluate the MDC-based wavelet denoiser functionality, a processor-coprocessor system has been designed and implemented on a Xilinx **FPGA** Spartan-3E 1600 Development Board.

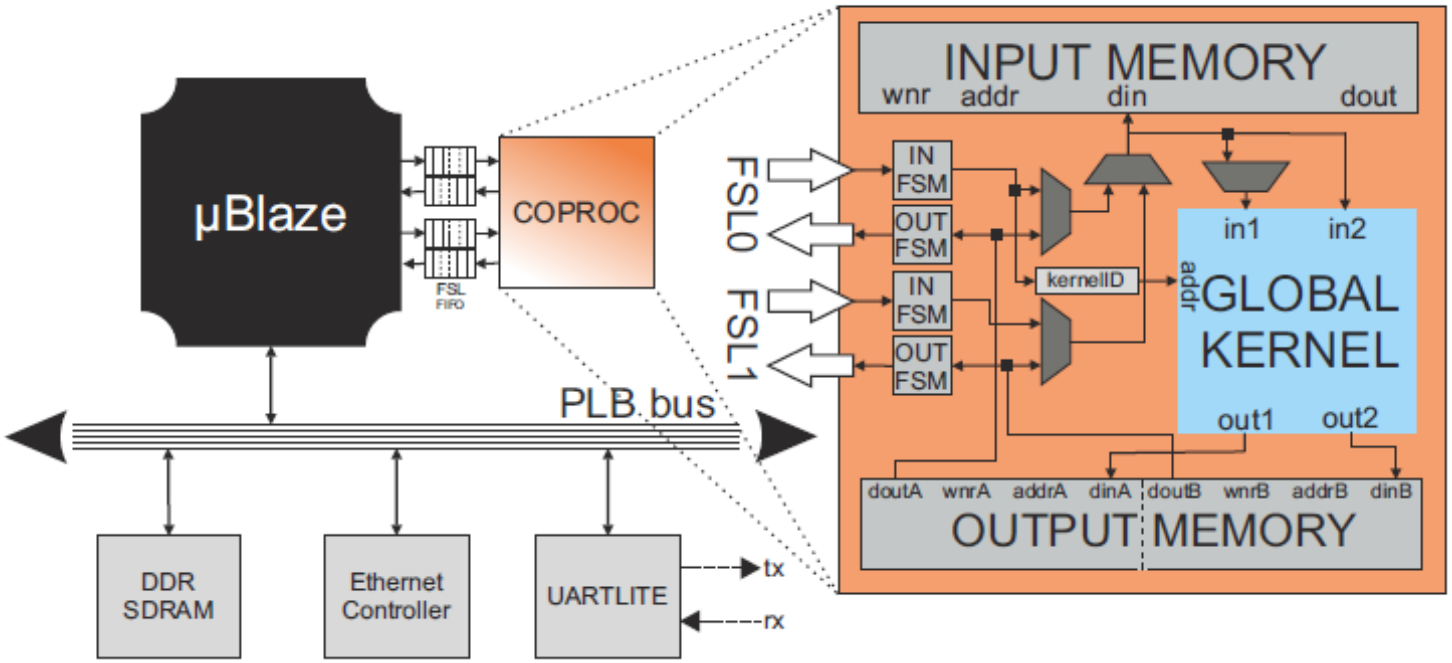

Use-Case Execution Trace

- For every new input buffer frame of *b_len* samples:

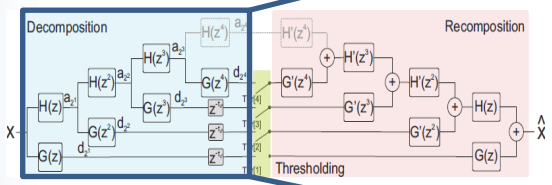

Decomposition
Begin
$d_{2^1}[n] \Leftarrow$ copr_dec_level_1(in_data,$H(z)$,$G(z)$,b_len)
$d_{2^2}[n] \Leftarrow$ copr_dec_level_2_no_coeff(b_len)
$d_{2^3}[n] \Leftarrow$ copr_dec_level_3_no_coeff(b_len)
$d_{2^4}[n] \Leftarrow$ copr_dec_level_4_no_coeff(b_len)
end
Thresholding
Recomposition
Begin
if cut $a_{2^4}[n]$ **then**
copr_rec_level_4(*Zero_coeffs*, $G'(z)$, b_len)
copr_rec_level_3($H'(z)$, $G'(z)$, b_len)
else
copr_rec_level_4($H'(z)$, $G'(z)$,b_len)
copr_rec_level_3_no_coeff(b_len)
end if
copr_rec_level_2_no_coeff(b_len)
den_data_block$\Leftarrow$ copr_rec_level_1_no_coeff(b_len)
end

Use-Case Execution Trace

- For every new input buffer frame of *b_len* samples:

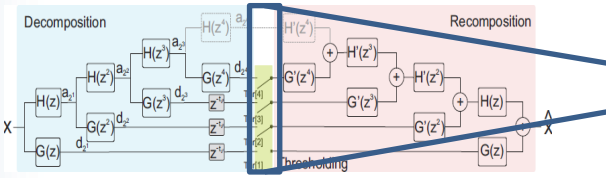

Decomposition
Begin
$d_{2^1}[n] \Leftarrow$ copr_dec_level_1(in_data,$H(z)$,$G(z)$,b_len)
$d_{2^2}[n] \Leftarrow$ copr_dec_level_2_no_coeff(b_len)
$d_{2^3}[n] \Leftarrow$ copr_dec_level_3_no_coeff(b_len)
$d_{2^4}[n] \Leftarrow$ copr_dec_level_4_no_coeff(b_len)
end
Thresholding
Recomposition
Begin
if cut $a_{2^4}[n]$ **then**
copr_rec_level_4(*Zero_coeffs*, $G'(z)$, b_len)
copr_rec_level_3($H'(z)$, $G'(z)$, b_len)
else
copr_rec_level_4($H'(z)$, $G'(z)$,b_len)
copr_rec_level_3_no_coeff(b_len)
end if
copr_rec_level_2_no_coeff(b_len)
den_data_block$\Leftarrow$ copr_rec_level_1_no_coeff(b_len)
end

Use-Case Execution Trace

- For every new input buffer frame of *b_len* samples:

Decomposition
Begin
$d_{2^1}[n] \Leftarrow$ copr_dec_level_1(in_data,$H(z)$,$G(z)$,b_len)
$d_{2^2}[n] \Leftarrow$ copr_dec_level_2_no_coeff(b_len)
$d_{2^3}[n] \Leftarrow$ copr_dec_level_3_no_coeff(b_len)
$d_{2^4}[n] \Leftarrow$ copr_dec_level_4_no_coeff(b_len)
end
Thresholding
Recomposition
Begin
if cut $a_{2^4}[n]$ **then**
copr_rec_level_4(*Zero_coeffs*, $G'(z)$, b_len)
copr_rec_level_3($H'(z)$, $G'(z)$, b_len)
else
copr_rec_level_4($H'(z)$, $G'(z)$,b_len)
copr_rec_level_3_no_coeff(b_len)
end if
copr_rec_level_2_no_coeff(b_len)
den_data_block$\Leftarrow$ copr_rec_level_1_no_coeff(b_len)
end

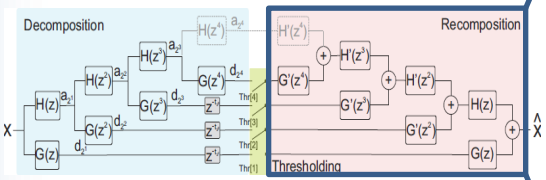

Coprocessor Parameter Set

- The implemented coprocessor is highly **parametric**. The designer can choose at runtime:
  - the number **N** of levels used in decomposition/recomposition;
  - the number and the values of the coefficients used for the FIR filters;
  - the possibility of removing the last *approximation* signal to obtain a band-pass filtering;
  - the threshold values for the *details* at each decomposition level;
  - the **b_len** number of samples for each input frame (allowable maximum: 512).
- The input signal can be applied in input to the testing environment through an UDP communication via **Ethernet** between the FPGA and the host PC.

Experimental results: Test Data

- The performance of the reconfigurable Wavelet Denoiser has been assessed using public available datasets of **simulated neural signals**, constructed starting from a database of real physiological spikes.
- The **background noise** overlapped to the significant signal is obtained considering spikes of different neuronal cells at random times and amplitudes.
- Sampling Frequency: 12kHz Useful Bandwidth: 300Hz – 3kHz

Background noise: 0.1

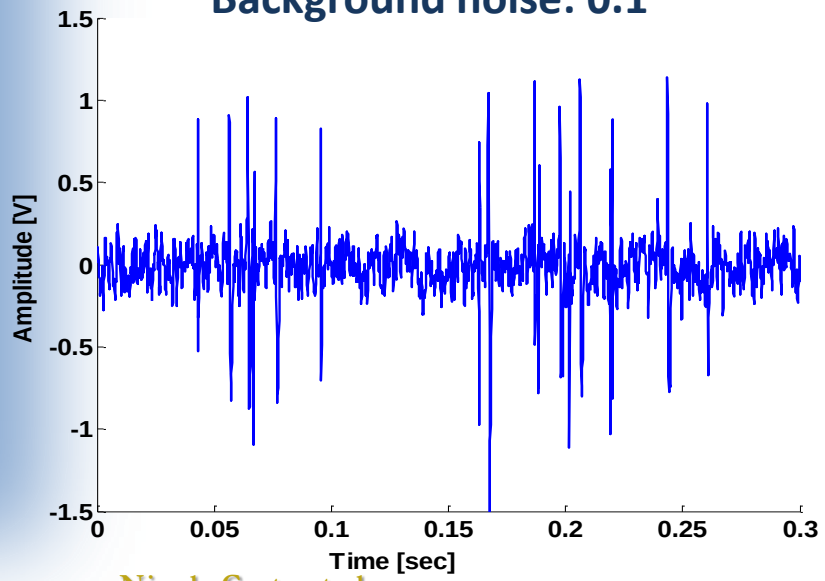

Background noise: 0.3

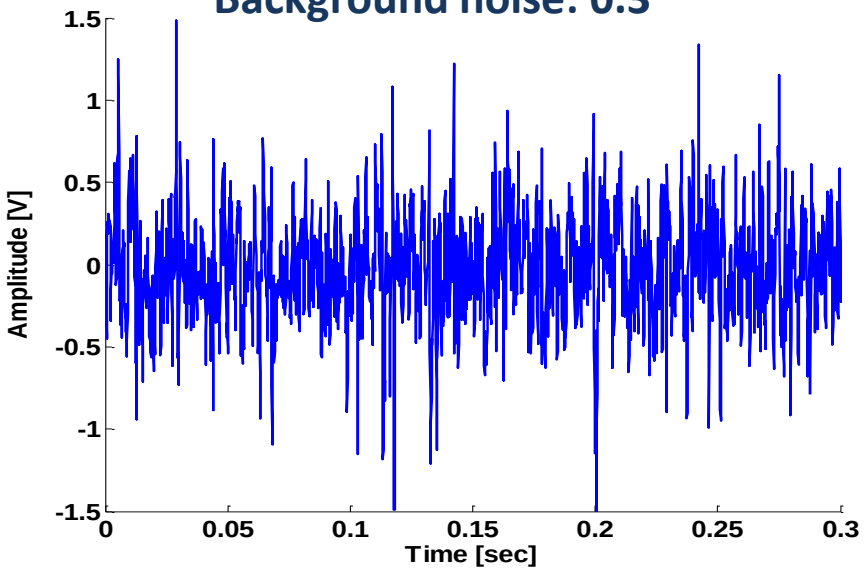

Parameters set-up

- Wavelet Denoiser parameters defined at runtime:
  - **N** = 4 of decomposition/recomposition levels;
  - the removal of the approximation signal at the 4th decomposition level determining an output frequency bandwidth in the range of [375Hz-3kHz].
  - ***b_len*** = 500 samples as length of each input buffer frame;
  - the *Haar/ Daubechies 2* families as **mother wavelets.**

Accuracy: low level noise and *Haar* as mother wavelet

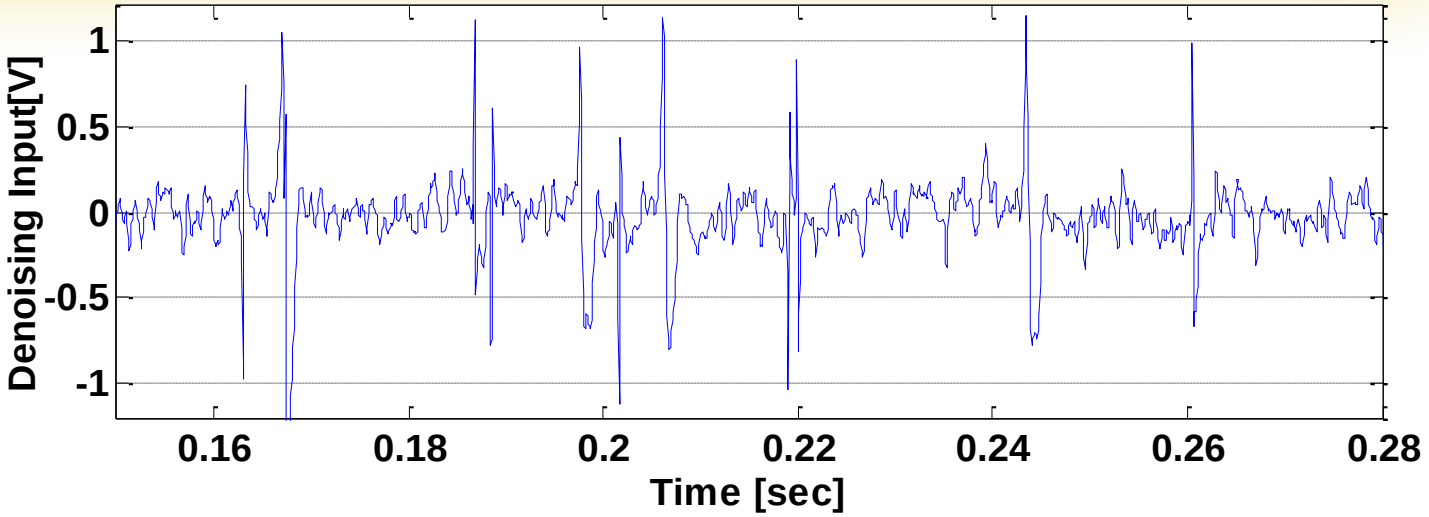

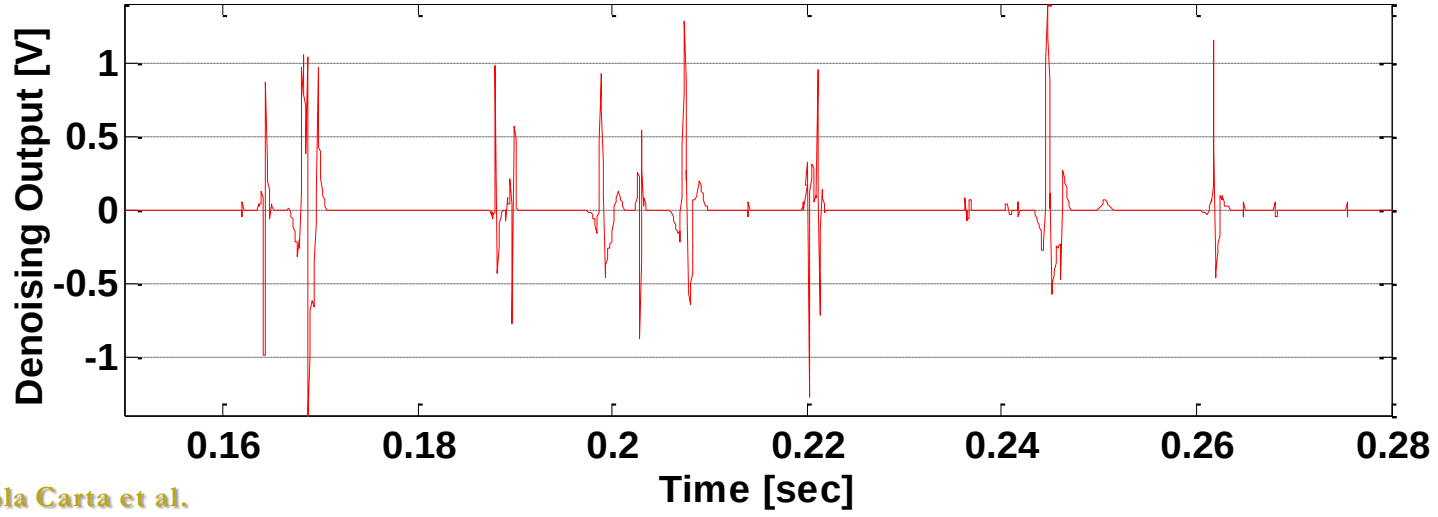

Accuracy: high level noise and *Haar* as mother wavelet

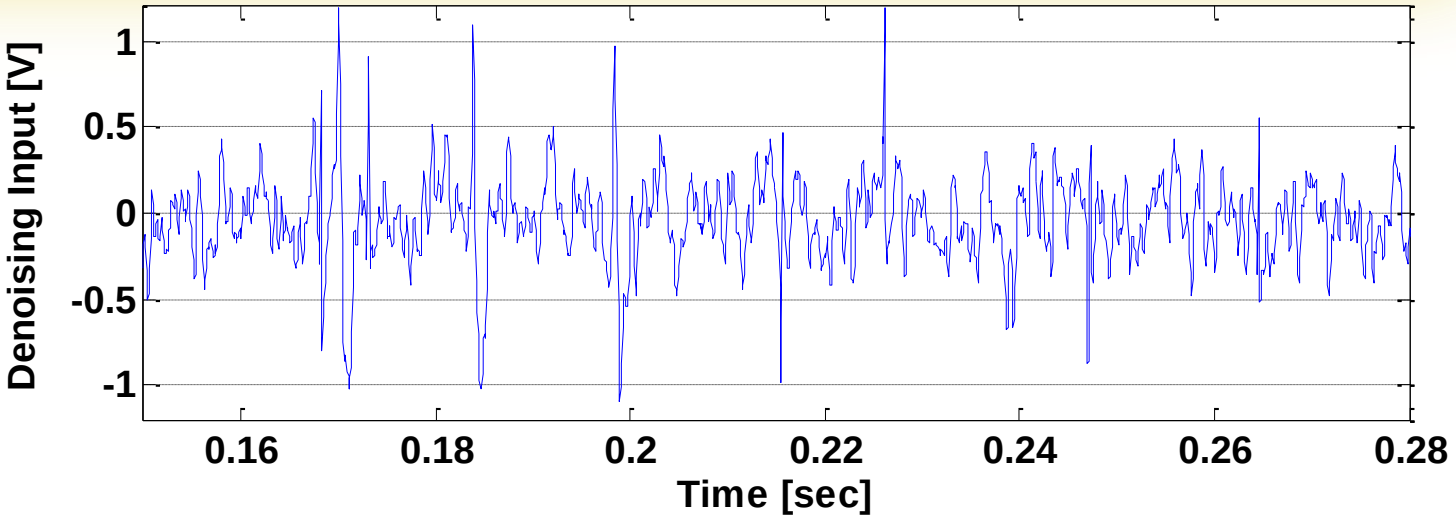

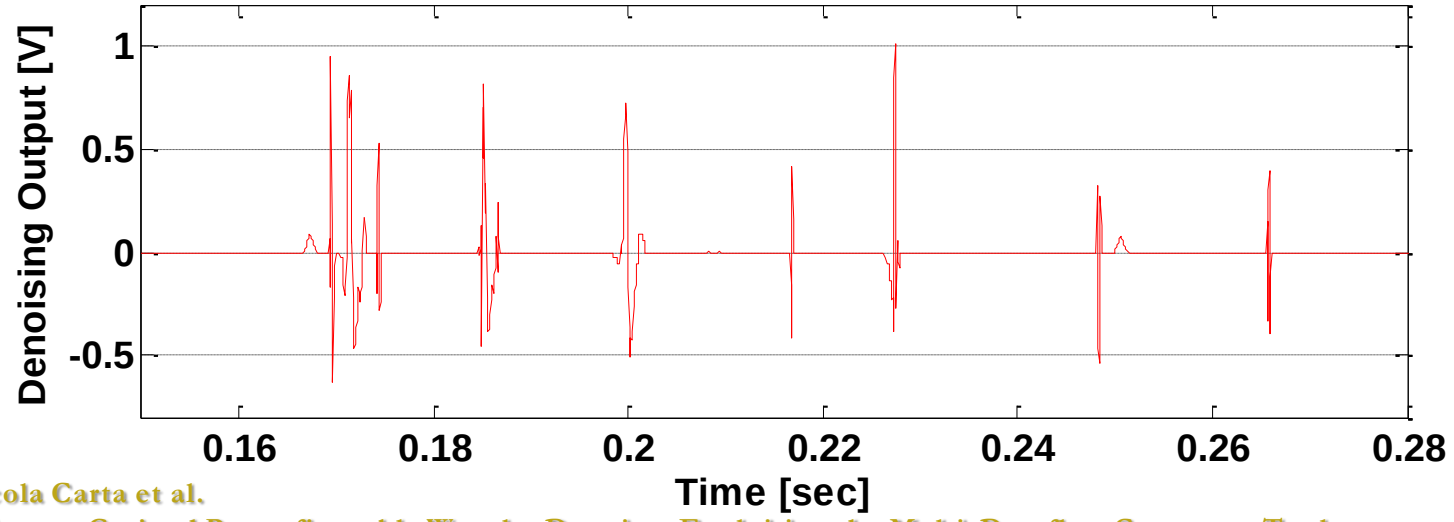

Accuracy: low level noise and *Daubechies 2* as mother wavelet

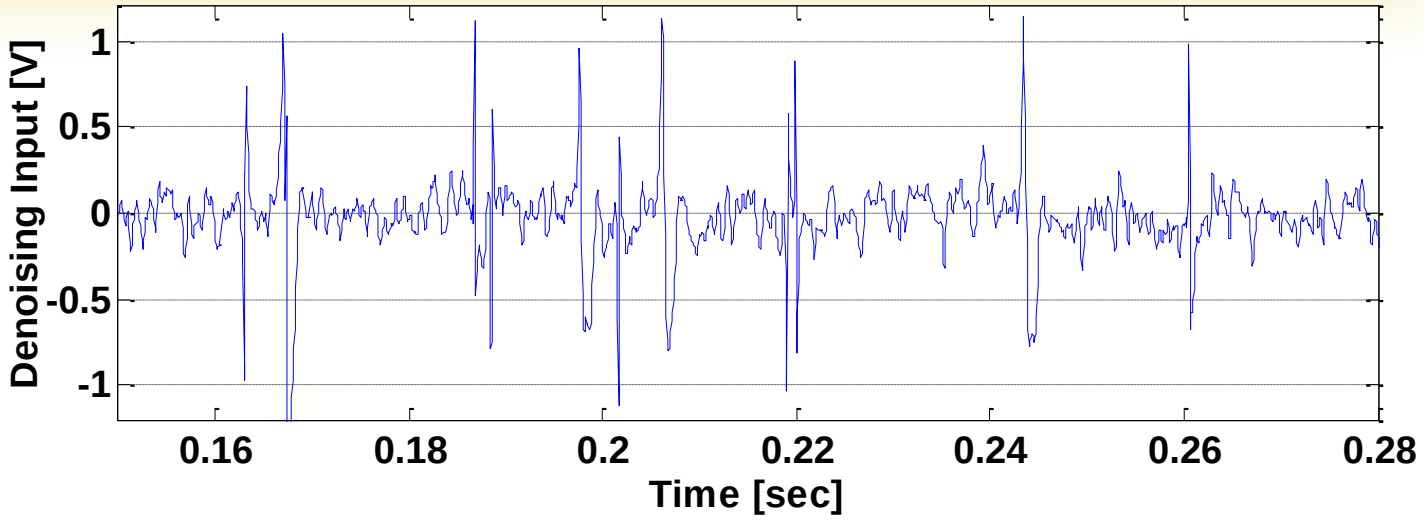

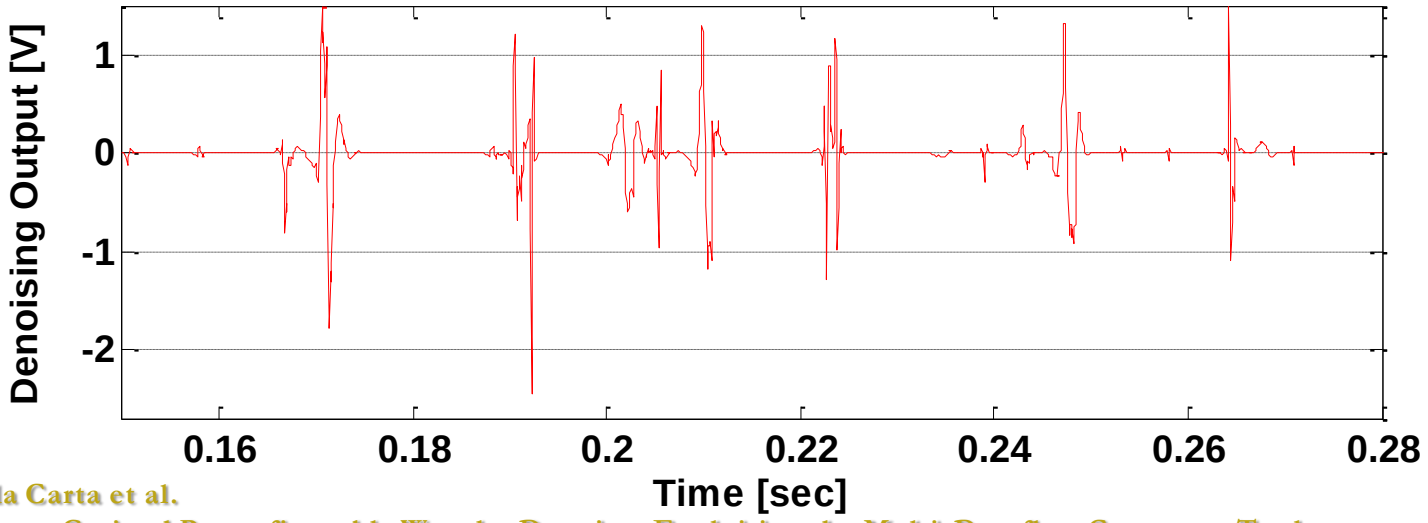

Area occupancy and Power Consumption

- The datapath of the reconfigurable filter has been synthesized in order to evaluate the effectiveness of our approach in terms of **area** and **power saving**.
- We compared the ***Global Kernel*** assembled by the MDC tool with:
  - the datapath obtained by the 3 identified kernels without any resource sharing among them (***Static Global Kernel***);
  - the ***Cascaded Kernel*** implementing the wavelet denoising as a fully-parallel solution without sharing actors among the various decomposition/recomposition levels.
- Results for the ***Cascaded Kernel*** only represent an estimation taking into account the number of required FIR filters and the hardware resources correspondent to the various modules of the HDL components library.

Synthesis Results for *Multiplier* and *Adder* FUs

- "Multiplier" and "Adder" FUs determine the largest contribution in terms of area and power consumption.

Implementation	adders	multipliers
MDC Global Kernel	3	2
Static Global Kernel	6	4
Cascaded WD	20	32

- *FPGA* and *ASIC* (using Cadence RTL Compiler targeting a 90nm CMOS library) synthesis results for the relevant FUs:

FU	FPGA			ASIC		
	Slices	Flip Flops	MULT18x18s	Area [μm^2]	Power [mW]	Frequency [MHz]
adder	341	116	0	8644	1,679	555,56
multiplier	138	131	11	113497	2,031	357,14

FPGA synthesis results

- FPGA synthesis using the Xilinx FPGA Spartan-3E 1600 as target:

Implementation	Slices [%]	Flip Flops [%]	MULT18x18s [%]
MDC Global Kernel	14	5	22
Static Global Kernel	22	8	44
Cascaded WD	76	22	356

- The *MDC Global Kernel* is able to achieve **36%** of saving in terms of slices compared to *Static Global Kernel* and **82%** compared to the *Cascaded WD* solution.
- Considering the Xilinx multiplier primitives, the resource saving percentage is **50%** and **97%** respectively.

ASIC synthesis results

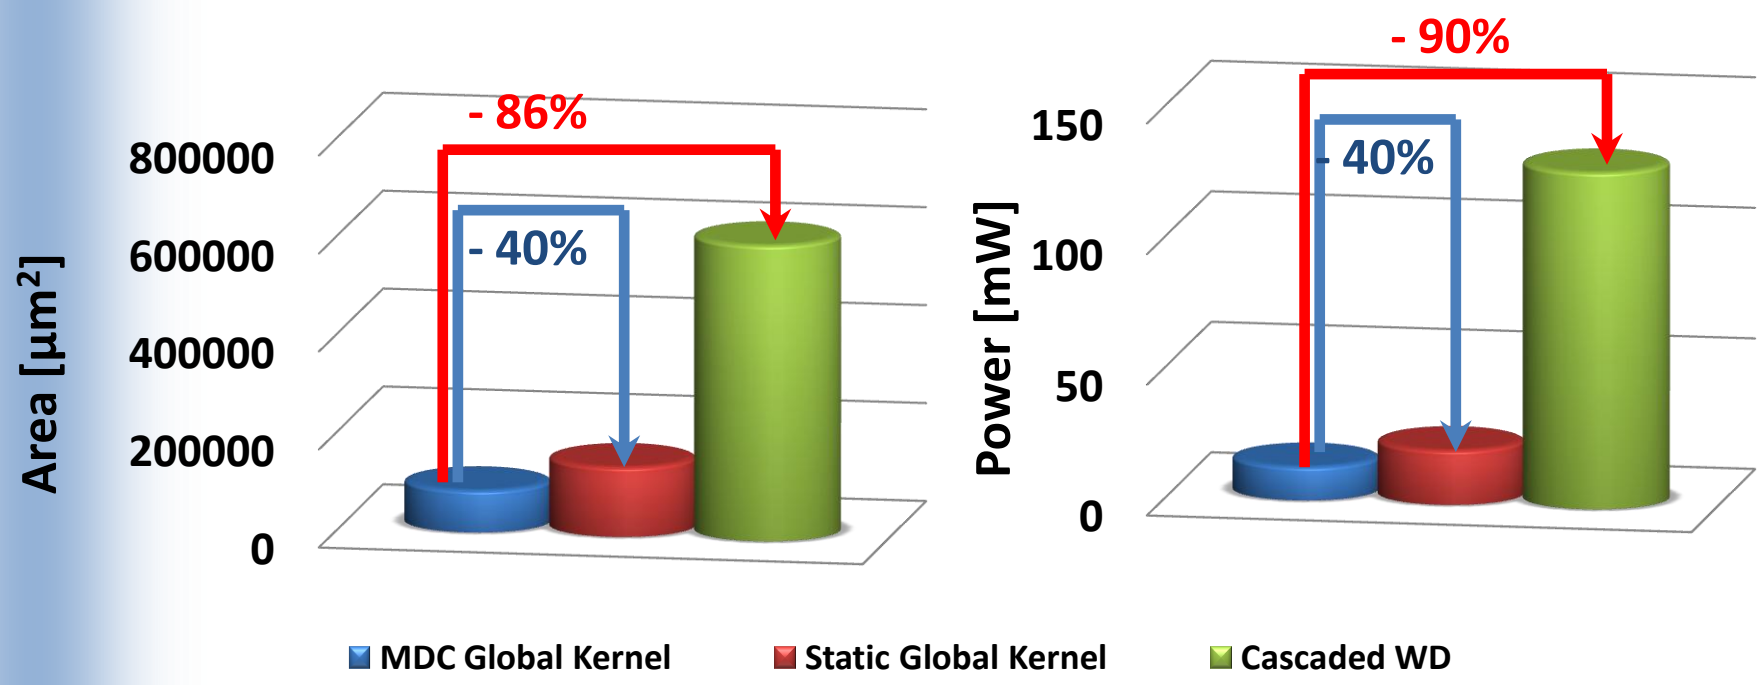

Results evaluation

- Provided that:
  - in the ***MDC Global Kernel*** case, the same resources are reused at each level;
  - in the ***Cascaded Kernel*** case, the number of required resources is linearly proportional to the number of levels;

  the more will be the depth of the denoiser the larger will be the benefits of adopting the MDC-based approach.

- Strict temporal constraints would make impossible the minimization of hardware resources within kernels. In those cases, a compromise design choice must be done to balance resource saving and processing capacity.
- In other cases different kernel models, for example exploiting *pipelining* solutions, would be beneficial.

Conclusions

- In this paper, the implementation of a Wavelet Denoiser allows to demonstrate that:
  - the applicability of the MDC tool on a completely different application field compared to the original one it was conceived for;
  - a non-static implementation of a denoiser is possible leveraging a coarse-grained reconfigurable hardware design platform.
- System correctness and accuracy have been properly verified using a processor-coprocessor test environment and simulated neural signals as input data, varying the values assigned to the various parameters.
- The MDC-based solution is capable of an on-line processing of the incoming samples maximizing the resources reuse and determining area and power minimization.

Acknowledgement

The research leading to these results has received funding from the Region of Sardinia, Fundamental Research Programme, L.R. 7/2007 “Promotion of the scientific research and technological innovation in Sardinia” under grant agreement CRP-18324 RPCT Project (Y. 2010) and CRP-60544 ELoRA Project (Y. 2012). Carlo Sau gratefully acknowledges Sardinia Regional Government for the financial support of his PhD scholarship (P.O.R. Sardegna F.S.E. Operational Programme of the Autonomous Region of Sardinia, European Social Fund 2007-2013 - Axis IV Human Resources, Objective l.3, Line of Activity l.3.1).

A Coarse-Grained Reconfigurable Wavelet Denoiser Exploiting the Multi-Dataflow Composer Tool

N. Carta, C. Sau, F. Palumbo, D. Pani and L. Raffo

DIEE - Dept. of Electrical and Electronic Engineering

University of Cagliari, Italy