Research Repository UCD

Title	Coarse Grained Reconfigurable Array Based Architecture for Low Power Real Time Seizure Detection
Authors(s)	Patel, Kunjan, Bleakley, Chris J.
Publication date	2015
Publication information	Patel, Kunjan, and Chris J. Bleakley. "Coarse Grained Reconfigurable Array Based Architecture for Low Power Real Time Seizure Detection" 82, no. 1 (2015).
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/7026
Publisher's statement	The final publication is available at www.springerlink.com
Publisher's version (DOI)	10.1007/s11265-015-0981-9

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Coarse Grained Reconfigurable Array based Architecture for Low Power Real Time Seizure Detection

Kunjan Patel and C.J.Bleakley
UCD School of Computer Science and Informatics
University College Dublin, Dublin, Ireland

Abstract

There is increasing research and commercial interest in miniature on-body and implantable devices for continuous real-time biosignal monitoring. A key challenge in realizing this vision is in implementation of biosignal processing algorithms with acceptably low energy consumption. In this article, we investigate implementation of the REACT algorithm for real-time epileptic seizure detection on a Coarse Grained Reconfigurable Array (CGRA) based architecture. Computationally expensive biosignal processing tasks are offloaded from a conventional Digital Signal Processor (DSP) to the CGRA. The CGRA is designed to support low power biosignal processing by means of a systolic architecture, flexible interconnect and low resource usage. The CGRA architecture is shown to provide 38% and 60% improvements in energy consumption and in performance, respectively, for the REACT system, without the use of voltage scaling or increased clock frequency.

Index Terms

Coarse grained reconfigurable architecture, biosignal processing, low energy, seizure detection, ambulatory monitoring

I. INTRODUCTION

Epilepsy is a common neurological condition, affecting approximately 1% of the world's population [1]. Diagnosis can be difficult due to the nature of the physiological signals involved. The diagnostic gold standard is long-term in-patient Electroencephalogram (EEG) and video monitoring [2]. This process is very costly in terms of hospital resources, including equipment, ward occupation and staff time. Not only must the patient be cared for, but the physiological data generated must be analysed by qualified medical staff. Furthermore, the patient is removed from their natural environment which may affect their condition. This can lead to long and numerous monitoring periods.

Currently, there is significant research and commercial interest in developing on-body and implantable biomedical devices which can perform continuous real-time biosignal monitoring and analysis, allowing for rapid condition detection and diagnosis of conditions [3]. Ambulatory EEG (AEEG) recording devices have recently been developed and deployed so that patients can have their EEG recorded at home. However, these devices do not perform data analysis. The recorded data must be analysed on return to the hospital. If no events are detected, the patient is sent home for further recording.

Given these problems, there is a move to implement automated real-time seizure detection algorithms on ambulatory devices [4]. The REACT (Real-time EEG Analysis for event deTeCTion) algorithm developed by our collaborators at University College Cork achieves high accuracy in epileptic seizure detection based on EEG analysis. The algorithm operates on EEG data to extract set of features in the time, information theoretic and spectral domains. A Support Vector Machine (SVM) classifier uses these features to identify periods of seizure and non-seizure activity. To date, the REACT algorithm has been used to automate hospital in-patient monitoring. Herein, we investigate the potential for reducing the power consumption of REACT implementation, so as to enable its inclusion in AEEG devices. This would mean that seizure

events could be detected at home and medical staff notified immediately. If no events are detected, the patient does not need to go to hospital.

The battery lifetime of ambulatory devices is important factor. Ideally, a real-time seizure detection and recording device should have a battery lifetime of roughly 1 week. This would be sufficient to obviate the need for recharging by the user. However, in order to maximize user comfort, the battery must be as small and lightweight as possible. Even though EEG is sampled at a low rate (hundreds of Hertz), the multi-channel nature of the biosignals (possibly tens of channels) and the difficulty of the detection and classification problem make the overall REACT system power hungry. The computational complexity of the detection algorithm leads to high power consumption when implemented on a conventional, commercial Digital Signal Processor (DSP). Previous work has focused on reducing the computational complexity of the REACT algorithm by reducing the number of channels and/or the number of features [5]. While effective in reducing computational complexity, this approach has the undesirable side-effect of reducing accuracy.

Herein, we investigate implementation of the REACT algorithm on SYSCORE, a novel Coarse-Grained Reconfigurable Array (CGRA) processor architecture designed for low power on-chip biosignal processing. The overall SYSCORE architecture consists of a conventional programmable DSP processor, that performs irregular control oriented tasks and a CGRA unit that performs regular, computationally intensive signal processing tasks. The CGRA unit consists of a grid of interconnected reconfigurable processing units which can perform logical and arithmetic operations. Unlike conventional DSPs, the CGRA supports systolic signal processing, significantly reducing the number of data and program RAM accesses and lowering power consumption. In addition, the higher degree of parallelism allows for more aggressive voltage scaling. Unlike Field Programming Gate Arrays (FPGAs), the processing units are reconfigurable at the operation level rather than at the bit-level. This significantly reduces power consumption compared to FPGAs while maintaining flexibility and increasing performance [6].

The SYSCORE architecture is specifically designed for low-power biosignal processing and so differs from previous CGRA architectures in a number of ways. Firstly, the architecture is systolic in that input and output data is pumped synchronously in parallel between nearest neighbour processing units arranged in an 2-dimensional pipeline manner. Secondly, the architecture allows mapping of the Fast Fourier Transform (FFT), by means of a novel interconnection scheme, Roundabout Interconnect (RAI), whereby non-nearest neighbour data transfer is supported without dense interconnect. Thirdly, in order to reduce power consumption, the SYSCORE architecture uses a minimal number of resources in each functional unit. Each unit is fixed-point with 2 operational units and 4 data registers.

In this paper, we consider the problem of mapping the REACT real-time seizure detection algorithm to the SYSCORE CGRA architecture. Firstly, the computational complexity of the constituent routines in the REACT algorithm is determined. Secondly, the suitability of mapping each routine to the CGRA is assessed. Thirdly, a comparison is made between the performance and power consumption of the REACT algorithm mapped to a conventional DSP processor architecture and to the SYSCORE architecture. Overall, the SYSCORE approach is shown to provide an energy consumption saving of 38% and a speed up of 60% compared to the conventional DSP solution. To the authors' knowledge, this was the first work to study implementation of a biosignal processing algorithm on a CGRA architecture.

The remainder of the paper is structured as follows. Section II presents a discussion of relevant CGRA architectures. A description of the REACT system and the proposed architecture are given in Sections III and IV, respectively. In Section V implementation details are given. Section VI presents results and the paper is concluded in Section VII.

II. RELATED WORK

A. Automated Seizure Detection Systems

There have been many publications in the area of automated seizure detection [7]. Few of these proposals have progressed to realistic implementation, with notable exceptions being [8], [9] and the REACT system.

The performance of the REACT system has been investigated and validated in a number of papers including [10], [11] and [12].

B. Architectures for Biosignal Processing Applications

Previous work on processor platforms for biosignal processing has focused on multi-core and ASIC architectures. Multicore architectures allow parallel processing of multichannel data. The authors of [13] presented a multiprocessor system-on-chip for real-time human heart monitoring and analysis. An architecture with 12 DSP processors was proposed to process 12 channel ECG data. Since the DSP cores run concurrently, the architecture implements a semaphore and interrupt system for communication and resource sharing. Multi-core architectures provide performance increase over single core architectures but do not typically provide energy consumption reductions, other than by voltage scaling. In fact, the resource sharing and communication overhead are often significant in terms of power consumption and area.

A low energy biomedical signal processing platform was presented in [14]. The platform is based on a 16 bit CPU with hardware accelerators for common DSP operations such as FFT, CORDIC, FIR filter and median filter. The inclusion of hardware accelerators reduces energy consumption but limits flexibility since the accelerators only support specific operations.

ASIC designs can achieve high performance with low power consumption. The authors of [15] presented an ASIC for heart rate variability monitoring and assessment. The ASIC was used in conjunction with a microcontroller. Specific tasks were offloaded to the ASIC in which dedicated blocks executed specific functions. The ASIC reduced power consumption by a factor of 7 compared to a standalone microcontroller. An energy-efficient ASIC for ultra low power wireless sensor nodes was presented in [16]. The ASIC was designed to perform the main functions of the proposed wireless Body Area Network sensor node. The leakage current and overall current consumption of the node was very low. Of course, the main disadvantage of ASICs is lack of flexibility. Also, targeting low volume biomedical applications may not be cost effective since ASIC redesign is a lengthy and costly process.

C. CGRA Architectures

CGRAs have been previously proposed for multimedia, embedded and DSP applications [17], [18], [19], [20], [21], [22], [23], [24]. Most previous work on CGRAs has focused on increasing performance rather than on explicitly reducing energy consumption. The authors of [21] reported the power consumption of a CGRA but didn't propose, discuss or prove the effectiveness of particular power saving techniques.

The authors of [25] proposed a scalable reconfigurable architecture for wireless communication. The algorithms were mapped by solving a set of concurrent matrix computations. A compiler designed by the authors was used for algorithm mapping. A dynamic context compression technique was proposed in [26] to reduce the power consumption of re-configuration.

CGRA architectures with Single Instruction Multiple Data (SIMD) processing model (Morphosys, ADRES, CGRA Express) are efficient for algorithms which can be vectorized but can be inefficient because processing steps must be concatenated, as in biosignal processing. Multimedia CGRA architectures (Morphosys, ADRES, CGRA Express) primarily provide support for vectorize-able two-dimensional image and video processing algorithms.

For accurate EEG analysis, a fixed bit-width of at least 12-bit for IO and 24-bit for internal processing is required [3]. Certainly, architectures with less than 16-bits are insufficient (PipeRench). Complex interconnections are required to support algorithms that have irregular data access patterns, such as the FFT. This support is provided in some CGRAs, such as SmartCell. However, this typically requires dense interconnections, increasing chip area and power. For example, in SmartCell, the rich interconnections and instruction memory consume 53% of the total power.

Two previous CGRAs claim to operate in a systolic fashion - PolySA and SmartCell. Neither are designed for low power. PolySA is floating-point and SmartCell is rich in interconnections. Some previously proposed architectures (e.g. RaPiD, PPA) are capable of passing data to nearest neighbour processing units

but cannot pass input data and output data in parallel with operation execution. Most previous CGRA functional units are large in terms of area. For example, REMARC uses floating point number format and AMBER contains 64 registers. Based on publicly available information, we estimate that the functional units of all previous architectures are more than double that of SYSCORE except for PACT XPP and PolySA which are 20% and 50% larger, respectively.

In SYSCORE, minimum energy is achieved by means of a fixed-point architecture, low area units, flexible interconnect and input-output systolic data passing.

III. REACT ALGORITHM

REACT (Real-Time EEG Analysis for event deteCTion) is an automated seizure detection system for adults [4][10] [12][11][27]. EEG signals contain a mixture of activity from all over the brain, artifacts from the body, such as muscle movement, and artifacts from external sources such as lights and electrical equipment. Repetitive patterns evident in the EEG may be due to seizure, but might also be due to mains electricity or the patient tapping their finger or simply breathing. Examples of non-seizure and seizure EEG traces are shown in Figure 1 and 2.

Fig. 1: Non-seizure EEG

The steps in the REACT seizure detection algorithm are shown in Figure 3. The steps are downsampling, feature extraction, classification and post-processing. The EEG signal is 6 channels and recorded at 250Hz to meet clinician requirements [28]. The EEG signal is then down-sampled from 250Hz to 32Hz using an anti-aliasing filter with cut off at 16Hz [4]. Prior to feature extraction, the EEG is split into 8.192 second epochs (2048 samples) with 50% overlap to avoid end effects and to reduce the impact of data at the end of an epoch. REACT can operate concurrently on multiple channels of EEG, treating each channel individually until a final classification decision is made for each epoch. Feature extraction is performed on the downsampled epoch, using a total of 55 features which provide an evaluation of the spectral content, energy and structure of the EEG epoch. The list of extracted features is given in Table I. These features are described in detail in [12].

The extracted features are passed to the SVM classifier which uses a Gaussian kernel. The SVM has been pre-trained using seizure and non-seizure EEG data. The output of the SVM is converted to a

Fig. 2: Seizure EEG

Fig. 3: Basic block diagram of the REACT system.

probability that the epoch contains seizure activity. Probabilities from successive epochs are smoothed using a central moving average filter and transformed into binary seizure/non-seizure decisions. Finally, individual decisions for each EEG channel are combined using a binary OR to produce the final result.

The REACT system has produced excellent performance on both neonatal and adult data, with Receiver Operating Characteristic (ROC) curve areas of 0.96 and 0.94 respectively [11] as shown in Figure 4.

The feature extraction stage of the algorithm requires the greatest number of instructions to complete, consuming the most computational power. The authors of [29] measured the number of clock cycles required by each of the features in the REACT system on a Blackfin processor. The resulting computational complexity distribution of the REACT system is shown in Figure 5. The computational complexity distributions of the feature extraction tasks are shown in Figures 6 and 7. The results show that the SVD entropy and Fisher Information features make up nearly 74% of the feature extraction clock cycle count, due to the fact that they consist of a large number of repeated loops to estimate decomposition

TABLE I: List of features extracted from the EEG data

Spectral	Energy	Structural
Spectral edge frequency (80%, 90%, 95%)	Hjorth activity, mobility and complexity	Entropy
Dominant spectral peak frequency	RMS amplitude	Fisher info
Variance of the first and second derivatives		Skewness
Zeros crossings of the first and second derivatives	Spectral power	AR model fit
Power in 2 Hz frequency bands (0-2 Hz, 1-3 Hz,... 10-12 Hz)		Kurtosis
Normalised power in 2 Hz bands		Singular value decomposition entropy
Spectral entropy		Spectral entropy
Number of minima and maxima		
Line length		
Zero crossings		
Nonlinear energy		
Wavelet coefficients

Fig. 4: Mean ROC curves and Equal Error Rate (EER) line for adult and neonatal databases.

matrices.

IV. SYSCORE ARCHITECTURE

A. Overview

The 8x4 SYSCORE architecture is shown in Figure 8. There are two main functional elements: Configurable Function Units (CFUs) and RoundAbout Interconnect (RAI). The designer can use as many functional elements as desired depending on the performance targets and area constraints. Two Direct Memory Access (DMA) units inject data into the architecture from the West and North and one DMA unit collects the output data. A column of RAI elements is inserted after every second column of CFUs to facilitate mapping of algorithms with complex dataflow graphs such as the FFT. The array configuration and DMA operations are controlled by the RISC processor.

Figure 9 shows the architecture of a CFU. The CFU has 4 input ports (In0-In3) and 3 output ports (Out0-Out2). It has a Computation Unit (CU) that can perform computational operations. The CU differs from a conventional ALU/MAC in terms of the Set of Operations (SoOs) it can support. All operations can be performed in a single cycle. Two data words can be passed in parallel with the result computed by the CFU via 3 output ports.

Fig. 5: Task computation complexity distribution in REACT.

Fig. 6: Most computationally complex feature extraction tasks

Fig. 7: Computational complexity of other feature extraction tasks

The bitwidth of a datapath can have a significant impact on power consumption [3]. To determine the optimum bitwidth for processing EEG signals, simulations were run using RaCAMS [30]. A variety of biosignal processing algorithms were assessed using EEG data and random 32 bit data. SNR values were calculated with respect to results obtained using 32 bits and are shown in Figure 10 and 11. It can be seen from the graphs that 22-bit is a reasonable trade off point between bitwidth and SNR.

There are 2 General Purpose Registers (GPRs), 2 Coefficient Registers (CERs) and 1 CU register (CU_reg) in a CFU. GPRs are used to store data from the input ports. CERs are used to store coefficients for algorithms such as FIR and FFT. The CU_reg is used to store results from the CU unit. All registers have the same bitwidth.

Each CFU has one 32 bit configuration register (Config_reg). It stores the configuration passed via port In2 when the Config_en signal is high. Table II shows the settings and purpose of the bit fields in the

Fig. 8: A SYSCORE block of 8x4

configuration register. The size of the configuration word is 22 bits, the remaining bits are left for future use.

TABLE II: Configuration register bitfields

Bitfield	Abbreviation	Purpose	Selection list
0 to 2	OP	Operation selection	ADD, SUB, MUL, MUL-ADD, MUL-SUB
3 to 5	ALU0	ALU selection line 0	In0-In3, 2 GPR
6 to 8	ALU1	ALU selection line 1	In0-In3, 2 CER, 2 GPR
9 to 11	ALU2	ALU selection line 2	In0-In3, 2 CER, 2 GPR, ALU_reg
12 to 13	REG0_sel	Input selection for GPR0	In0, In1
14 to 15	REG1_sel	Input selection for GPR1	In2, In3
16 to 17	OP0_sel	Input selection for Out0	ALU_reg, 2 GPR
18 to 19	OP1_Sel	Input selection for Out1	ALU_reg, 2 GPR
20 to 21	OP2_Sel	Input selection for Out2	ALU_reg, 2 GPR

Fig. 9: CFU architecture

Fig. 10: Bitwidth vs SNR for EEG data and

B. Set of Operations (SoO)

As shown in Table II, SYSCORE can perform Addition (ADD), Subtraction (SUB), Multiplication (MUL), Multiply and Addition (MAD) and Multiply and Subtraction (MSU). The later operations, MAD and MSU, are more useful than the traditional MAC operation for mapping algorithms systolically [31]. Because of the feedback from CU_reg to the CU, the CFU can be configured to perform a MAC operation without extra hardware cost. Table III provides a list of the aggregation operations SYSCORE can perform.

SYSCORE can operate in four modes: configuration mode, execution mode, flush mode and power off mode. Table IV provides the list of operational modes of SYSCORE with the relevant control signals where X means that the state of the control signal does not effect the mode of operation. The flush mode is useful for collecting results from the output registers of CFUs after accumulation operations, e.g., matrix multiplication. The shut down mode is useful for saving power when the CFU is not used. The first 3

Fig. 11: Bitwidth vs SNR for random data

TABLE III: List of aggregation functions

Abbreviation	Description
SQACC	Iterative square and accumulate
CRCOR	Iterative multiplication with 2 variable inputs (cross correlation)
MADDC	Iterative multiply and addition with 1 variable input, 1 constant input and variable output location
MULC	Iterative multiplication with 1 constant input and same input-output location
UP1	An repetitive process in SVM which has 4 multiplications and 1 addition, 1 subtraction

modes can be activated at the array level and last mode can be activated at the row level.

TABLE IV: SYSCORE operational modes

Mode	Description	Config_en	Flush_en	Coeff_sel	Global_en
Configuration	SYSCORE configuration	high	low	X	high
Execution	Execution of operations	low	low	X	high
Flush	Allows collection of accumulated result	low	high	X	high
Power off	The CFU is off	X	X	X	low

C. Interconnections

As shown in Figure 8, all CFUs are connected to their nearest neighbours to the East and West. To avoid dense interconnections, cross interconnections are only introduced at odd numbered columns. Cross interconnections are useful to perform non-systolic functions, such as the FFT butterfly. Cross interconnect functionality is provided by RAI elements that allow data to pass from any Westerly CFU to any Easterly CFU. The conceptual structure of a RAI element is shown in Figure 12. Each RAI element has 6 input ports (I0-I5), 6 output ports (O0-O6) and a 16 bit configuration register. As in a CFU, the RAI element can be reconfigured when SYSCORE is in configuration mode. The output ports of the RAI element can be configured to take data from the input ports. Figure 13 shows the available output port options in RAI. There are no global interconnections, except control signals (as described in the previous section), which saves chip area and reduces power consumption and control overheads.

V. IMPLEMENTATION

An 8x8 SYSCORE array was built from two 8x4 array blocks, each as shown in Figure 8. The architecture was modelled and algorithm mappings were performed for the algorithms listed in Table

Fig. 12: Conceptual diagram of RA interconnect

Output	Input from
0	2-5
1	2-5
2	0-3
3	0-3
4	0-5
5	0-5

Fig. 13: Input port selection options for output ports

V using the software modeller and cycle accurate simulator RaCAMS [30]. The algorithms were mapped manually as mentioned in [31]. RaCAMS simulation outputs were verified by comparison with Matlab. The hardware architecture was implemented in Verilog and all algorithms were mapped using SystemVerilog. The results provided below were obtained by averaging over 10,000 iterations. The Verilog implementation was synthesized using Synopsys tools and a 90nm CMOS technology library. The maximum operating frequency of DSP and SYSCORE were 95.14 MHz and 174.82 MHz respectively. For comparison purposes below, it is assumed that both run on the same clock frequency. The area of a CFU, a RAI element, SYSCORE and DSP were 38k, 6k, 2500k and 60k gates, respectively, when synthesized for a 100 MHz clock (NAND2 gate count equivalent). Because of the differences in technology libraries, it was not possible to directly compare SYSCORE's power metrics with those of an existing conventional DSP processor. So, for the purpose of energy comparison, a typical DSP processor was implemented. The processor architecture had 1 single cycle MAC unit, 24 bit registers, a fetch and decode unit, Program RAM and Data RAM. The ISA of this DSP can execute all the instructions that a typical DSP can execute. Cycle accurate performance results were obtained by running simulations using RaCAMS model and energy consumption results were obtained by running gate level simulations and power analysis.

To assess computational complexity, the REACT algorithm was implemented in C for Analog Devices Blackfin BF-537. The computational complexity distribution of the overall REACT algorithm was measured in terms of number of cycles.

In the REACT-SYSCORE system, the operations listed in Table II and III were offloaded to SYSCORE if they occurred repetitively in any part of the REACT algorithm. The feature extraction stage of REACT algorithm was targeted as a whole since it is 74% of total computational complexity distribution. Since the input signals to the various feature extraction functions are the same, these functions can be concatenated in the CGRA, reducing the load on the DMA. The classification stage contains decision-making tasks which are irregular and hence are difficult to map to systolic arrays. The energy consumption savings and speed up for operations listed in Table II and III were calculated individually and then overall energy savings and speed up for REACT processes were calculated based on their proportional contribution to the overall computation complexity.

VI. RESULTS

A. Energy Consumption

Table V shows the energy consumption of SYSCORE and the DSP processor for various DSP algorithms. The results are for 1000 iterations. Using these figures, the energy savings achieved by off-loading portions of REACT to the CGRA were calculated and are shown in Table VI. The overall energy savings and speed up depends on the computation complexity and frequency of use of algorithms in the REACT system. It also shows the overall energy savings for the algorithms depending on their proportional contributions to REACT algorithm. SYSCORE gives overall energy consumption savings of 38%. Reconfiguration energy is included in the SYSCORE figures.

TABLE V: DSP and SYSCORE energy consumption comparison

Algorithm	Type	Energy (pJ)		Energy savings (%)
		DSP	SYSCORE	
FIR filter	5	197	49	75
	25	768	146	78
	63	1852	388	79
	127	3678	776	79
Matrix multiplication	4x4	4037	2824	30
	8x8	8074	5648	30
Wavelet transform	db2	254	112	55
DFT	8 Point	99627	791	99
FFT	8 point,radix-2	5345	2886	46
Aggregate Functions	SQACC	31	18	42
	CRCOR	31	31	0
	MADDC	31	31	0
	MULC	31	18	44
UP1	-	36	33	8

TABLE VI: Energy savings compared to DSP processor in REACT system using REACT-SYSCORE system (* includes reconfiguration penalty)

Feature	Algorithm	Contribution to overall energy(%)	Energy Savings (%)	Overall Energy savings(%)
SVD Entropy and Fisher Information	SSQACC	2	0.6	28
	CRCOR	45	0	
	MADDC	46	38.6	
	MULC	2	0.8	
	UP1	1	0.1	
	Total	96	40.2	
AR Model Fits	MADDC	7	5.4	5
	MULC	92	38.5	
	Total	99	43.9	
FFT-derived	-	2	46	1
Wavelet Coefficient Energy	-	0.5	55	0.5
RMS Amplitude	-		42	
Pre-processing Filtering	-	4	79	3
Total	37.5			

DVFS was not taken into consideration in the analysis but if it was used then, for the CMOS technology used, we estimate that scaling from 1.3V to 0.7V can further reduce energy consumption by up to 71%. This is facilitated by the speed ups achieved using the SYSCORE architecture.

The majority of the energy savings in REACT-SYSCORE system are due to reductions in the number of RAM accesses. Figure 14 shows a comparison of RAM Data Reuse (RDR) between the DSP processor and SYSCORE architecture, where RDR is given by:

$$RDR = \frac{\text{Number of unique RAM addresses accessed}}{\text{Number of RAM accesses}} \quad (1)$$

A RDR value close to 1 indicates that RAM locations are only accessed once. A value close to 0 indicates that the same RAM locations are accessed many times.

It is clear from the results that data reuse in SYSCORE is considerably higher than in the DSP processor, reducing the overall energy consumption. RDR for CRCOR and MADDC on the DSP and REACT-SYSCORE systems are the same. However, during SVD feature extraction in REACT, CRCOR and

Fig. 14: RDR comparison of commonly used biosignal algorithms

MADDC use the same data. Thus double reads are avoided and RDR is increased. Figure 15 shows a comparison of energy consumption between REACT and REACT-SYSCORE systems.

Fig. 15: Energy consumption comparison between REACT and REACT-SYSCORE systems

B. Performance

The energy savings and speed improvement was achieved by off-loading the most computational complex tasks from the feature extraction phase in the REACT system. The least intensive features were not off-loaded because the computational complexity distribution was less than 1% of the total feature extraction process. The off-loaded features were SVD, AR model, FFT and RMS amplitude. Table VII compares the cycle count on SYSCORE and the DSP for various biosignal processing algorithms. The speed up improvement obtained by off-loading the algorithms to the CGRA in the REACT-SYSCORE system is given in Table VIII. The overall speed up gained for the REACT-SYSCORE system was 60%

assuming the DSP and SYSCORE are running at the same clock frequency. When the higher clock frequency of the SYSCORE is taken into account the speed up becomes 111%. Figure 16 illustrates the performance of REACT on the DSP and on SYSCORE. Reconfiguration time is included in the SYSCORE figures. This speed up can be traded for further reductions in energy consumption by means of voltage scaling. The large speed-up opens the door to aggressive techniques such as sub-threshold operation.

TABLE VII: DSP and SYSCORE performance comparison

Algorithm	Type	DSP (cycles)	SYSCORE (cycles)	Performance improvement (%)
FIR filter	5 taps	6000	1008	83
	25 taps	26000	1028	96
	63 taps	64000	1064	98
	127 taps	128000	1144	99
Matrix multiplication	4x4	64000	8000	88
	8x8	512000	16000	97
Wavelet transform	db2	8000	4001	50
DFT using correlation	8 point	80000	8008	90
FFT 8 point	8 point, complex numbers	144000	8016	94
Aggregate Functions	SQACC	4000	1011	74
	CRCOR	3000	1004	67
	MADDC	3000	1019	67
	MULC	3000	1019	67
UPI	-	3000	255	91

TABLE VIII: Performance comparison of REACT system on DSP processor and on SYSCORE (includes cycles for reconfiguration)

Feature	Algorithm	Contribution to overall complexity(%)	Performance improvement (%)	Overall performance improvement(%)
SVD Entropy and Fisher Information	SSQACC	2	1	45
	CRCOR	45	44	
	MADDC	46	45	
	MULC	2	2	
	UPI	1	1	
	Total	96	93	
AR Model Fits	MADDC	8	8	8
	MULC	92	89	
	Total		97	
FFT-derived	-	2	94	2
Wavelet Coefficient Energy	-	0.5	50	0.5
RMS Amplitude	-		98	
Pre-processing Filtering	-	4	98	4
Total				59.5

VII. CONCLUSIONS

This paper investigates the effectiveness of a novel CGRA architecture, SYSCORE, for implementation of low power biosignal processing applications. The architecture allows systolic mapping of DSP algorithms to reduce memory accesses and so reduces power consumption. RAI interconnect elements were introduced to increase the flexibility of the architecture in supporting algorithms which cannot be

Fig. 16: Performance comparison between REACT and REACT-SYSCORE systems

easily mapped systolically. The architecture was evaluated for the REACT system. The REACT-SYSCORE system gives 38% energy savings and 60% speed up compared to a conventional single-processor DSP implementation of the REACT algorithm. This illustrates the effectiveness of CGRA architectures for low energy biosignal processing applications.

ACKNOWLEDGMENT

This research was funded as a part of the Efficient Embedded Digital Signal Processing for Mobile Digital Health (EEDSP) cluster, 07/SRC/I1169, by Science Foundation Ireland (SFI).

Dr Kunjan Patel Biography text here.

PLACE
PHOTO
HERE

Dr Chris Bleakley Biography text here.

PLACE
PHOTO
HERE

REFERENCES

- [1] J. Sander and S. Shorvon, "Incidence and prevalence studies in epilepsy and their methodological problems: a review." *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 50, no. 7, pp. 829–839, 1987.
- [2] E. Waterhouse, "New horizons in ambulatory electroencephalography," *Engineering in Medicine and Biology Magazine, IEEE*, vol. 22, no. 3, pp. 74–80, 2003.
- [3] K. Patel, C. P. Chua, S. Faul, and C. J. Bleakley, "Low power real-time seizure detection for ambulatory EEG," in *Pervasive Computing Technologies for Healthcare, 2009. PervasiveHealth 2009. 3rd International Conference on*, April 2009, pp. 1–7.
- [4] A. Temko, R. McEvoy, D. Dwyer, S. Faul, G. Lightbody, and W. Marnane, "REACT: Real-Time EEG Analysis for Seizure Detection," in *AMA-IEEE Medical Tech. Conf. on Individualized Healthcare*, 2010.
- [5] S. Faul, "Dynamic channel selection to reduce computational burden in seizure detection," in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, 31 2010-sept. 4 2010, pp. 6365–6368.
- [6] R. Hartenstein, "Coarse grain reconfigurable architecture (embedded tutorial)," in *Proceedings of the 2001 conference on Asia South Pacific design automation*. ACM, NY, USA, 2001, pp. 564–570.
- [7] S. Faul, G. Boylan, S. Connolly, and W. M. G. Lightbody, "An evaluation of automated neonatal seizure detection methods," *Clinical Neurophysiology*, vol. 116, no. 7, pp. 1533–1541, July 2005.
- [8] C. P. Young, C.-H. Hsieh, and H.-C. Wang, "A low-cost real-time closed-loop epileptic seizure monitor and controller," in *Instrumentation and Measurement Technology Conference, 2009. I2MTC '09. IEEE*, 5-7 2009, pp. 1768–1772.
- [9] N. Bhavaraju, M. Frei, and I. Osorio, "Analog seizure detection and performance evaluation," vol. 53, no. 2, feb. 2006, pp. 238–245.
- [10] A. Temko, I. Korotchkova, W. Marnane, G. Lightbody, and G. Boylan, "Validation of an automated seizure detection system on healthy babies," in *Proceedings of the INSTICC/IEEE International Joint Conference on Biomedical Engineering Systems and Technologies*, Valencia, Spain, January 2010, pp. 312–317.
- [11] S. Faul, A. Temko, and W. Marnane, "Age-independent seizure detection," in *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 09)*, Minneapolis, September 2009, pp. 6612–6615.
- [12] A. Temko, E. Thomas, G. Boylan, W. Marnane, and G. Lightbody, "An SVM-based system and its performance for detection of seizures in neonates," in *Proceedings of the 31st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 09)*, Minneapolis, September 2009, pp. 2643–2646.
- [13] I. Al Khatib, D. Bertozzi, F. Poletti, L. Benini, A. Jantsch, M. Bechara, H. Khalifeh, M. Hajjar, R. Nabiev, and S. Jonsson, "MPSoC ECG biochip: a multiprocessor system-on-chip for real-time human heart monitoring and analysis," in *CF '06: Proceedings of the 3rd conference on Computing frontiers*. New York, NY, USA: ACM, 2006, pp. 21–28.
- [14] J. Kwong and A. Chandrakasan, "An energy-efficient biomedical signal processing platform," in *ESSCIRC, 2010 Proceedings of the*, 2010, pp. 526–529.
- [15] W. Massagram, N. Hafner, M. Chen, L. Macchiarulo, V. M. Lubecke, and O. Boric-Lubecke, "Digital Heart-Rate Variability Parameter Monitoring and Assessment ASIC," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 4, no. 1, pp. 19–26, Feb. 2010.
- [16] X. Zhang, H. Jiang, L. Zhang, C. Zhang, Z. Wang, and X. Chen, "An energy-efficient ASIC for wireless body sensor networks in medical applications," *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 4, no. 1, pp. 11–18, Feb. 2010.
- [17] H. Singh, M.-H. Lee, G. Lu, F. Kurdahi, N. Bagherzadeh, and E. Chaves Filho, "MorphoSys: an integrated reconfigurable system for data-parallel and computation-intensive applications," *Computers, IEEE Transactions on*, vol. 49, no. 5, pp. 465–481, May 2000.
- [18] S. Goldstein, H. Schmit, M. Budiu, S. Cadambi, M. Moe, and R. Taylor, "Piperench: A reconfigurable architecture and compiler," *Computer*, vol. 33, no. 4, pp. 70–77, 2002.
- [19] C. Ebeling, D. Cronquist, P. Franklin, and C. Fisher, "RaPiD-a configurable computing architecture for compute-intensive applications," *University of Washington Department of Computer Science & Engineering Tech Report TR-96-11-03*, 1996.
- [20] B. Mei, S. Vernalde, D. Verkest, H. De Man, and R. Lauwereins, "ADRES: An architecture with tightly coupled VLIW processor and coarse-grained reconfigurable matrix," *Lecture notes in computer science*, vol. 2778, pp. 61–70, 2003.
- [21] Y. Park, H. Park, and S. Mahlke, "CGRA express: accelerating execution using dynamic operation fusion," in *CASES '09: Proceedings of the 2009 international conference on Compilers, architecture, and synthesis for embedded systems*, 2009, pp. 271–280.
- [22] H. Park, Y. Park, and S. Mahlke, "Polymorphic pipeline array: a flexible multicore accelerator with virtualized execution for mobile multimedia applications," in *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2009, pp. 370–380.
- [23] A. Sudarsanam, R. Barnes, J. Carver, R. Kallam, and A. Dasu, "Dynamically reconfigurable systolic array accelerators: A case study with extended Kalman filter and discrete wavelet transform algorithms," *Computers Digital Techniques, IET*, vol. 4, no. 2, pp. 126–142, 2010.
- [24] C. Liang and X. Huang, "SmartCell: A power-efficient reconfigurable architecture for data streaming applications," in *Signal Processing Systems, 2008. SiPS 2008. IEEE Workshop on*. IEEE, 2008, pp. 257–262.
- [25] K. Sarriorgidis and J. Rabaey, "A Scalable Configurable Architecture for Advanced Wireless Communication Algorithms," *The Journal of VLSI Signal Processing*, vol. 45, pp. 127–151, 2006, 10.1007/s11265-006-9762-9.
- [26] Y. Kim and R. Mahapatra, "Dynamic Context Compression for Low-Power Coarse-Grained Reconfigurable Architecture," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 18, no. 1, pp. 15–28, jan. 2010.
- [27] S. D. Faul, "Dynamic channel selection to reduce power consumption in seizure detection," in *Accepted for 32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2010)*, 2010.
- [28] EPILEPSIAE. (2008, Jan.) EPILEPSIAE - Evolving Platform for Improving Living Expectation of Patients Suffering from IctAl Events, A EU project EPILEPSIAE started in January 2008. [Online]. Available: <http://www.epilepsiae.eu/>
- [29] R. McEvoy, S. Faul, and L. Marnane, "Ambulatory REACT: Real-time seizure detection with a DSP microprocessor," in *32nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC2010)*, 2010.
- [30] K. Patel, S. McGettrick, and C. J. Bleakley, "Rapid functional modelling and simulation of coarse grained reconfigurable array architectures," *Journal of Systems Architecture*, vol. 57, no. 4, pp. 383–391, 2011.

- [31] K. Patel and C. Bleakley, "Systolic Algorithm Mapping for Coarse Grained Reconfigurable Array Architectures," *Reconfigurable Computing: Architectures, Tools and Applications*, pp. 351–357, 2010.