NeuroCGRA: A CGRAs with support for neural networks

***Abstract*—Coarse Grained Reconfigurable Architectures (CGRAs) are emerging as enabling platforms to meet the high performance demanded by modern embedded applications. In many application domains (e.g. robotics and cognitive embedded systems), the CGRAs are required to simultaneously host processing (e.g. Audio/video acquisition) and estimation (e.g. audio/video/image recognition) tasks. Recent works have revealed that the efficiency and scalability of the estimation algorithms can be significantly improved by using neural networks. However, existing CGRAs commonly employ homogeneous processing resources for both the tasks. To realize the best of both the worlds (conventional processing and neural networks), we present NeuroCGRA. NeuroCGRA allows the processing elements and the network to dynamically morph into either conventional CGRA or a neural network, depending on the hosted application. We have chosen the DRRA as a vehicle to study the feasibility and overheads of our approach. Simulation using edge detection reveal that the neural networks can successfully process real-time video for up to 1M pixels. Synthesis results reveal that the proposed enhancements incur negligible overheads (4.4% area and 9.1% power) compared to the original DRRA cell.**

I. Introduction and motivation

Recently, the increasing speed and performance requirements of embedded applications, coupled with the demands for flexibility and low non-recurring engineering costs, have made reconfigurable hardware a very popular implementation platform. The reconfigurable architectures can be classified on the basis of granularity i.e. number of bits that can be explicitly manipulated. Coarse Grained Reconfigurable Architectures (CGRAs), provide operator level configurable functional blocks, word level datapaths, and very area-efficient routing switches. Therefore, compared to the fine-grained architectures (like FPGAs), CGRAs require lesser configuration memory and configuration time (two or more orders of magnitude [1]). As a result, CGRAs achieve a significant reduction in area (from 66 % to 99.06 % [2]) and energy consumed per computation (from 88 % to 98 % [2]), at the cost of a loss in flexibility compared to bit-level operations. Therefore, CGRAs have been a subject of intensive research since the last decade [2], [3].

Today, CGRAs host multiple applications simultaneously on a single platform. Each application can potentially have different requirements (e.g. MPEG4 decoder requires exact calculation while edge detection can tolerate approximations). For the estimation problems, *neural network* promise higher efficiency and scalability compared to conventional processing algorithms [4]. However, the existing CGRAs lack the support to simultaneously provide conventional and the neural networks based processing [4]. As a solution to this problem, we present NeuroCGRA. The proposed architecture allows to efficiently interleave the conventional processing with estimation (using neural networks). For the applications that require exact calculations, the device behaves like a normal CGRA, with MACs/ALUs connected via circuit switched interconnect. When an application, that can tolerate approximate results, enters the platform the device dynamically morphs into a neural network.

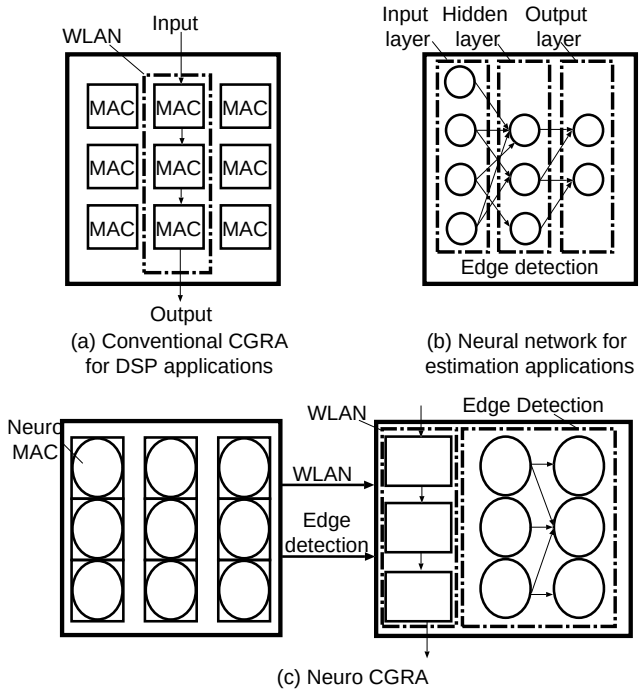

Fig. 1. Motivation for NeuroCGRA

To visualize our technique, consider Fig. 1. Fig. 1 (a) shows simple WLAN transmitter mapped to a conventional CGRA [5], [6]. The figure shows a typical scenario, where multipliers/accumulators are connected in a pipelined fashion to produce outputs. Fig. 1 (b) depicts edge detection realization on a neural network. A neural network typically consists of three elements: (i) neurons (the processing elements), (ii) synapse (the interconnect network), and (iii) weights (data for processing). Usually a neural network consists of three layers of neurons (input layer, hidden layer, and output layer). Each neuron layer performs calculation on the weights based on the neuron model. The results calculated in one layer

are communicated to the neurons in the next layer (to the right). Based on the joint calculation of the neurons, the edges in an image are detected. Fig. 1 (c) shows functionality of the proposed CGRA. Each processing element can act as either as conventional MAC or a Neuron. When WLAN and edge detection request platform resources the platform dynamically creates a different partition for each application. For WLAN and edge detection the platform resources morph into conventional and neural network, respectively.

This paper is organized as follows. In Section II, a brief survey of existing platforms used to realize neural networks is presented. In Section III, an overview of the CGRA platform, used in this paper is described. In Section IV, details of the proposed method are presented. In Section V, we explain how the edge detection is mapped to DRRA. In Section VI, we evaluate the benefits and redundancies imposed by our method on an actual CGRA. Finally, in Section VII, we summarize our contributions and suggest directions for future research.

II. Related work and contributions

Hardware implementations of neural networks has been a subject of intensive research since the last two decades [7]. In this section we will review only the most prominent work dealing with neural network implementation on reconfigurable architectures.

Much of the work implementing neural networks on reconfigurable architectures deals with FPGAs. In particular it focuses on how various algorithms can be realized using neural networks. Krips et al. [8] presented an FPGA based implementation neural networks to track of video images. Yang and Paindavoine [9] proposed a an implementation of neural networks for face tracking. Maeda and Tada [10] developed a neural network model to support online learning. Himavathi et al. [11] used layer multiplexing technique to implement multi-layer feed-forward networks into FPGA.

Recent works have revealed that for estimation or approximate problems, the neural networks in particular offer higher energy efficiency and speedup compared to conventional algorithms. They have shown that neural networks can also be used as accelerators. Chen et al. [12] provide a detailed discussion of hardware and software based accelerators (using neural networks) with special emphasis on memory efficiency. Esmaeilzadeh et al. [13] presented an ASIC based accelerator to efficiently speed up approximate programs. Chakradhar et al. [14] proposed an FPGA based configurable coprocessor that dynamically configures the hardware to provide the best achievable throughput (on the available resources). The main idea of this paper, i.e to interleave conventional processing with neural networks on a unique platform, is inspired from this work. The major difference is that our work focuses on CGRAs, that have significantly different architectures and lesser flexibility, compared to FPGAs. Zhengrong [15] showed that the spiking neural networks (when compared to Sobel and Canny filter) are more efficient in detecting edges (in terms of accuracy), can be easily customized to detect specific parts of an image and can tolerate high levels of noise. Zhengrong's work has motivated access the feasibility to implement edge detection using spiking neural networks.

The related work reveals that most of the work on neural networks deals with their realization on FPGAs. The works that do use neural networks to enhance efficiency, typically employ FPGAs or ASICs. None of the existing works accesses the feasibility of implementing neural networks on CGRAs.
Compared to the related work, this paper has three major contributions:

1) We present NeuroCGRA that allows to interleave the conventional calculations with neural networks;
2) We propose a neural network translator that provides a framework to map neural network on CGRAs; and
3) We evaluate benefits and overheads of implementing the proposed technique on an actual CGRA.

III. System Overview

We have chosen the Dynamically Reconfigurable Resource Array (DRRA) [16] as a vehicle to evaluate feasibility of implementing neural networks on an actual CGRA. Nevertheless, it seems that the results should applicable to most grid based CGRAs as well. The motivation for choosing DRRA is that it is well documented, its bandwidth has been tested on demanding industrial applications (with Huawei) [17], and we have available all its architectural details from RTL codes to physical design.

A. DRRA configuration flow

As shown in Figure 2, DRRA is programmed in two phases (off-line and on-line) [18]. The configware (binary) for commonly used DSP functions (FFT, FIR filter etc.) is written in VESYLA (HLS tool for DRRA) and stored in an off-line library. For each function, multiple versions, with different degree of parallelism, are stored. The library, thus created, is profiled with frequencies and worst case time of each version. To map an application, its (simulink type) representation is fed to the compiler. The compiler, based on the available functions (present in library) constructs the binary for the complete application (e.g. WLAN). Since the actual execution times are unknown at compile-time, the compiler sends all the versions (of each function), meeting deadlines, to the run-time configware memory. To reduce memory requirements for storing multiple versions, the compiler generates a compact representation of these versions. Details of compression algorithm and how it is unraveled are given in [18]. The compact representation is unraveled (parallelized/serialized) dynamically by the run-time resource manager (running on LEON3 processor).

B. DRRA Computation Layer

DRRA computational layer is shown in Fig. 3. It is composed of four elements: (i) Register Files (reg-files), (ii) morphable Data Path Units (DPUs), (iii) circuit-Switched Boxes (SBs), and (iv) sequencers. The reg-files store data for DPUs. The DPUs are functional units responsible for performing computations. SBs provide interconnectivity between different

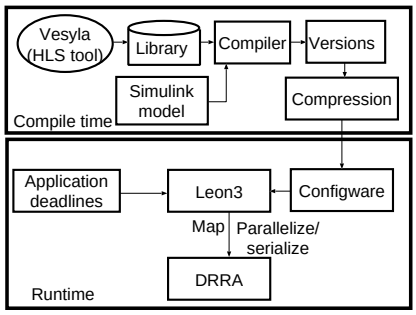

Fig. 2. Configuration Model

components of DRRA. The sequencers hold the configware which corresponds to the configuration of the reg-files, DPUs, and SBs. Each sequencer can store up to 64 36-bit instructions and can reconfigure the elements only in its own cell. As shown in Fig. 3, a *cell* consists of a Reg-file, a DPU, SBs, and a sequencer, all having the same row and column number as a given cell. The configware loaded in the sequencers contains a sequence of instructions (reg-file, DPU, and SB instructions) that implements the DRRA program.

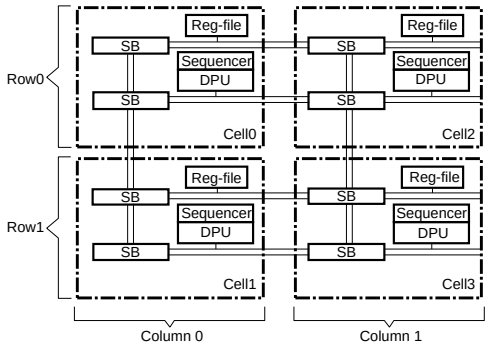

Fig. 3. DRRA computation layer

C. DRRA Storage Layer (DiMArch)

DiMArch is a distributed scratch pad (data/configware) memory that complements DRRA with a scalable memory architecture. Its distributed nature allows a high speed data and configware access to the DRRA computational layer (compared to the global configuration memory). Further discussion of DiMArch is beyond the scope of this paper and for details we refer to [19].

IV. Neural networks/DRRA integration

In this section, we will present the neural network model chosen for this paper followed by its implementation on DRRA.

A. Neural network model

For this work we have chosen Spiking Neural Networks (SNN). The motivation for choosing SNNs is that they are the latest generation of neural networks models that closely emulate the biological neurons. For further details about the SNNs an interested reader can refer to [20]. For understanding, Fig. 4 depicts a simple neuron model. The neurons are connected in one-to-many fashion. In the figure, the neuron 3 receives input spikes from neurons 1, 2, and 3. Each connection is characterized by a weight (Wt1, Wt2, and Wt3). When a spike is generated, depending on the neuron model and weight of the connection, the neuron performs calculations.

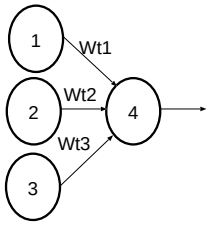

Fig. 4. A simplified neuron model

We have chosen a simple and the widely used model called leaky Integrate and Fire (I & F), to model a neuron. The model is given by the Equation 1:

$$dv/dt = (Wt * I) + a - (b * V). \quad (1)$$

Where the potential V integrates input spikes I and leaks over time with $-bV$ component. Wt is the weight of the connection Coefficient a determines equilibrium point and coefficient b the speed of leakage. When the threshold potential is reached, a neuron outputs a spike and its potential is reset.

B. Neural network realization on DRRA

To realize neural networks on DRRA, we have embedded a dedicated hardware, called neuroDPU, with each DPU of DRRA. As a result of our enhancements, the DPU can be configured to either normal or neuron mode. The details of the normal mode can be found in [2], here we only concentrate on the neuron mode. Fig. 5 shows how the DPU functions in neuron mode. The figure in particular illustrates how simple neural network shown in Fig. 4 is implemented on DRRA. To mimic parallelism in neural networks, we have employed time multiplexing. Where a time slot is reserved for each transmitting neuron (connected to the receiving neuron). In Fig. 5 (a) the transmitting DRRA cell implements three neurons. A register form the reg-file is reserved for each neuron. In every cycle, the neuroDPU receives an input from one of the three registers (representing the neurons) in round robin fashion. The neuroDPU applies the Equation 1 to the input and stores the result in the dedicated register for accumulation. If the result is greater than pre-defined threshold, 1 is stored in the spike register (otherwise 0 is stored). Since the transmitting cell implements three neurons, the value of the spike register is sent to the receiving cell after every three cycle. The receiving cell stores the value of the spike register a reserved location (at first address), of the reg-file. Fig. 5 (b) shows processing done by the receiving neuron. In each cycle, the neuroDPU of the receiving cell extracts a bit from spike register (stored in its reg-file) and the weight of the corresponding connection. It applies the $I\&F$ model to the inputs and generates the outputs.

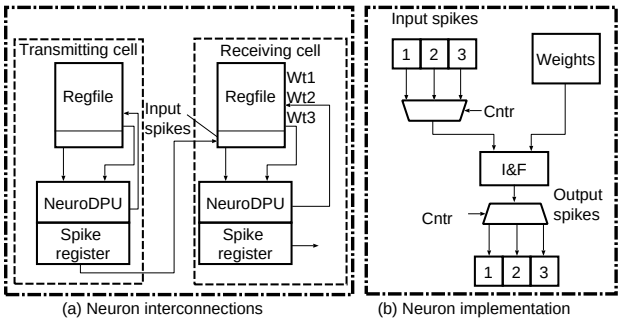

Fig. 5. Neuron realization on DRRA

C. Clustering for scalability

In Section IV-B we showed how a simple neural network (containing three neurons) can be realized on DRRA. In this section, we will explain how a large scale neural network can be implemented in a scalable manner considering the architectural details of DRRA. Neural networks may require one to many communication between large number of neurons. To implement these connections on DRRA, we had to consider two architectural properties: (i) every DRRA component has only two read/write ports and (ii) a DRRA component can be directly connected to a component at most three hops away. Since these architectural characteristics were designed after a careful evaluation of area/power trade-off, we decided not modify them. The one to many connectivity was realized by using time division multiplexing. In the proposed approach a specific time slot is assigned to each pair of neurons. To allow a scalable solution, we have chosen a hierarchical clustered approach shown in Fig. 6.

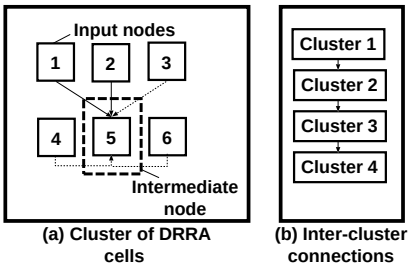

Fig. 6. Inter neural communication realization

Fig. 6 (a) shows a cluster of 6 DRRA cells. In the cluster one of the cells is chosen as an intermediate node. The intermediate node receives data from the 6 cells in the cluster (including itself). To allow connections with 6 cells on a 2-port component, we exploit time division multiplexing combined with partial and dynamic reconfiguration. In the overall process the intermediate cell receives data from 2 cells at a time and then shifts to the other cells. The process continues till the data from all the cells is received. Once every cluster has received inputs, the intermediate nodes communicate with each other serially to ensure one to many connectivity. Fig. 7 shows the instructions in DRRA sequencers (needed to implement the time division multiplexing). The figure shows that 5 cycles are needed to collect information from all the DRRA cells in the cluster. It should be noted that 2 additional cycles are required to reconfigure the circuit switched network. The inter cluster communication takes two cycles for reconfiguration and an additional cycle to transmit data. Each cell in the cluster can represent an arbitrary neurons, depending on the application requirements.

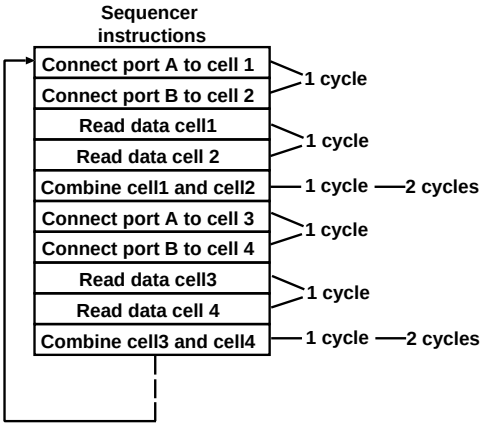

Fig. 7. Sequencer instructions to implement time division multiplexing

D. Architectural integration

To realize NeuroCGRA we have modified the configuration flow shown in the Fig. 2. We have added another block to generate the configware for the estimation algorithms (shown by the dotted box). The key component of the new block is a translator. The translator take three inputs: (i) application, (ii) weights, and (iii) application to generate the reg-file, SB, and DPU instructions for DRRA (see Section III-B).

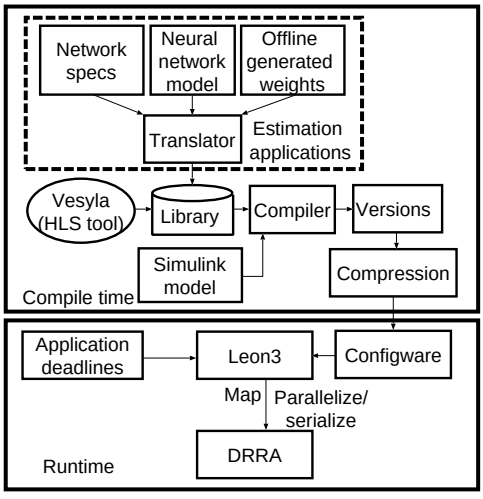

Fig. 8. Modified programming flow

Fig. 9 depicts how each input (to the translator) is represented, considering a neural network composed of 8 neurons. The the neural network application is represented by a *directed*

acyclic graph, similar to [21]. Where each node represents a neuron and the edges represent the interconnection between the two nodes. The motivation for choosing the directed acyclic graph is that it easily represents most of the neural network applications and can be easily converted to linked list for automated application generation. The weights are stores as a look-up table. For each neuron, the weights of all connected neurons are stored. If there is no connection between a pair of neurons (e.g. neuron 4 and 5 in the figure), 0 is stored. Finally, the network specifications in form the cluster size and the in each cluster is provided to the translator. In the simplest case i.e. $clustersize = 1$, only one DRRA cell for each neuron layer (input, hidden and output layer) is sufficient. If $clustersize > 1$, the required the cells, $DRRA_c$, are given by $DRRA_c = clusters_p - clusters_i$. Where $clusters_p$ and $clusters_i$ denote respectively the processing and the intermediate clusters.

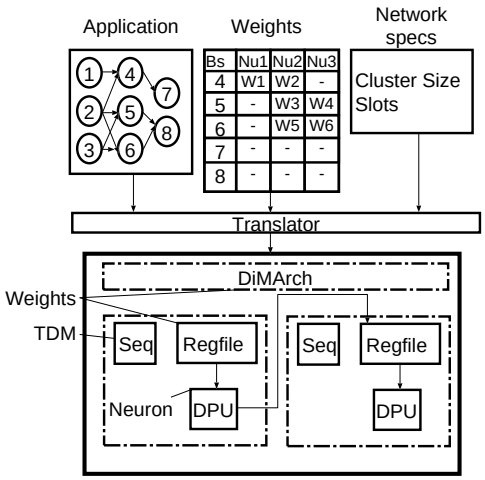

Fig. 9. Translator functionality overview

The translator functionality is shown in Algorithm 1. The translator works in three steps: (i) DPU instruction generation, (ii) Reg-file instruction generation, and (iii) SB instruction generation. The algorithm shows the translator functionality when $clustersize = 1$. If an application requires more clusters, the same algorithm can be applied to generate instructions for each cluster. In the first step, a separate DPU instructions is generated. In the second step, the reg-file for each cell is initialized. One of the register in each reg-file is reserved to store spikes. Since each register of the reg-file is 16-bit wide, the spike register can contain up to 16 spikes. The weights for each neuron connection is loaded to the register files. If a neuron is connected to greater than 63 neurons, the weights are stored in DiMArch memory. The details of the data transfer from DiMArch to reg-files is beyond the scope of this paper and for details an interested reader can refer to [19]. The total number of neurons per reg-file is dependent on the cluster size. In the third step, the switch-box instructions are generated. If the $clustersize = 1$ only three interconnect instructions (one for each layer) are needed. Otherwise, additional interconnect instructions to interconnect the clusters are also generated (similar to Fig. 7).

Input: application, weights, network specifications ;
Output: DPU, reg-file, and SB instructions ;
`/* Generate DPU instructions */`
if *Neural application exists* **then**
  DPU=neural mode ;
else
  DPU=normal mode ;
end
`/* Generate Reg-file instructions */`
for $j \leftarrow 0$ **to** N **do**
  `/* N is the number neurons in application */`
  Spike register delay= con_j ;
  `/* con_j is the connections per DRRA cell */`
  for $k \leftarrow 0$ **to** con_j **do**
    i=62 ;
    $Reg - file_i$= Wt_k ;
    i–;
    if *i¡0* **then**
      store Wt in DiMArch
    end
    ;
    Delay=0 ;
    Reg-file = Delay ;
    Delay ++ ;
  end
end
`/* Generate SB instructions */`
if *Cluster size* < 2 **then**
  Generate SB instructions to connect three layers ;
else
  Generate SB instructions for clusters;
  Generate SB instruction for inter-cluster communications;
end

Algorithm 1: Translator functionality

To mimic the neuron functionality, we have enhanced the functionality of the Data Path Unit (DPU). With each DPU we have added a special unit to perform I & F model. The weights are stored in reg-files and the DiMArch. The reg-file are also used to store spikes from other interconnected neurons. Since the spiked neural networks require timing information, we have added a global counter counter. To realize the complex neural network the sequencers are programmed to implement TDM.

V. Edge detection on DRRA

Edge detection is an image processing technique to identify edges in an image. The edge detection algorithms commonly identify the points in image at which the intensity level

changes dramatically. When the change in intensity is beyond threshold, an edge is identified. Edge detection is mainly used to reduce the image size for many applications e.g. face detection. Conventional edge detection, techniques such as Roberts filter, Sobel filter, or Canny filter employ a mask to the image. In this paper we will evaluate analyze the edge detection using the Sobel filter network and neural networks on DRRA.

A. Edge detection using Spiking Neural Networks

To evaluate the efficacy of implementing neural networks, we have chosen the neural network based edge detection algorithm presented in [22]. The motivation for choosing spiking neural network based edge detection is that it is more flexible (e.g. can be easily customized to detect white lines in an image) resilient to noise compared to Sobel and Canny filters [15]. The overall technique is shown in Fig. 10. The techniques relies on three layers of neurons. The first layer consists of input neurons. The input neurons receive the image and apply the I & F processing to the received pixel values. The result is sent to the middle or intermediate layer of neurons. The intermediate layer has four parallel arrays of neurons. It should be noted that for clarity, the figure only shows one neuron in each array. Each of these layers perform the processing for up, down, left and right edges respectively and is connected to the input layer by differing weights. The processed information is transmitted to the neuron in the output layer. This neuron integrates the outputs from the four neurons of the middle layer to detect an edge. For more details about the theory behind neural network based edge detection, an interested reader can refer to [22], [15].

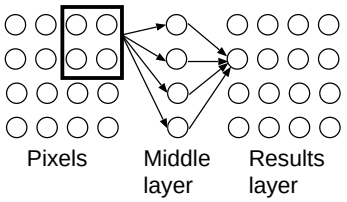

Fig. 10. Edge detection using neural networks

B. Implementation of Edge detection on DRRA

Fig. 11 illustrates how edge detection is mapped to DRRA. For clarity, the figure only shows the mapping for an image with only 2*2 pixels. Four neurons in receptor layer are statically stored in registers 1, 2, 3, and 4 of the cell 1. The pixel values are sequentially transferred to the DPU of cell 1. The DPU applies I & F model to the received pixel value and generates two outputs: (i) spike and (ii) accumulation value. The technique to generate these spikes was already shown in Section IV. The 4 neurons of the middle layer are mapped to the cell 2 of DRRA. In each cycle, the weight of the connection and the spike from the spike register is sent to the DPU. The DPU behaves the same way, as the DPU of cell 1, and sends the spike register value to the output layer. The output layer consists of a single neuron connected by the four neurons of the middle layer. Therefore, four registers are reserved for the four connections. The DPU of the output layer performs the I & F model and generates the final output.

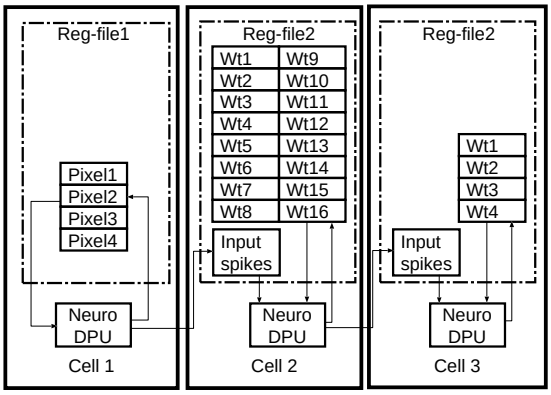

Fig. 11. Edge detection on DRRA

VI. Results

A. Cycle requirement

The main goal of implementing edge detection using spiking neural networks was that it can be easily customized (e.g. to detect white objects in an image) and offers higher tolerance to noise [15]. However, its real-time behavior is also important for its use in robotics and cognitive embedded systems. To analyze the real-time properties of the proposed approach, we simulated the edge detection using neural networks and the Sobel filter using images of different sizes (cf. Table I and Fig. 12). The figure and the table compares the time needed to detect edges using Sobel filter and neural networks. NN1 (column 1 of Table I) indicates a neural network implementation with no offline learning similar to [12]. The table and the figure indicate that with offline learning the neural network requires lesser cycles than the Sobel filter and hence can also be used as an accelerator. Another key thing to notice is that the neural network based edge detection algorithm can determine the edges of an image containing up to in $1M$ pixels in 16000000 cycles. Since DRRA can operate at 500 MHz the edges can be calculated 0.04 secs per image (which is sufficient for high quality TV requiring approximately 24 frames/sec). To evaluate the effect of online learning, we simulated edge detection with different allowed iterations. As shown in Table I and Fig. 12, the obtained results suggest that for up to 4 iterations, neural networks provide better timing. When 7 iterations are allowed, the Sobel filter outperforms the neural network based implementation (in terms of timing). Since complete online edge detection requires thousands of cycle, we conclude that only the edge detection with offline learning can be used as an accelerator.

To study the effect of complexity on the Sobel filter and neural network based edge detection (with online learning), we simulated edge detection with images of different sizes (cf. Fig. 13). The figure shows the additional cycles needed by the

TABLE I
CYCLE REQUIREMENT SOBEL FILTER VS NEURAL NETWORKS

Images	Sobel cycles	NN1 cycles	NN4 cycles	NN7 cycles
5*5	747	400	1600	2800
10*10	5312	1600	6400	11200
15*15	14027	3600	14400	25200
20*20	26892	6400	25600	44800
50*50	191232	40000	160000	280000
100*100	797132	160000	640000	1120000
150*150	1818032	360000	1440000	2520000
200*200	3253932	640000	2560000	4480000
400*400	13147532	2560000	10240000	17920000
1000*1000	82668332	16000000	64000000	112000000

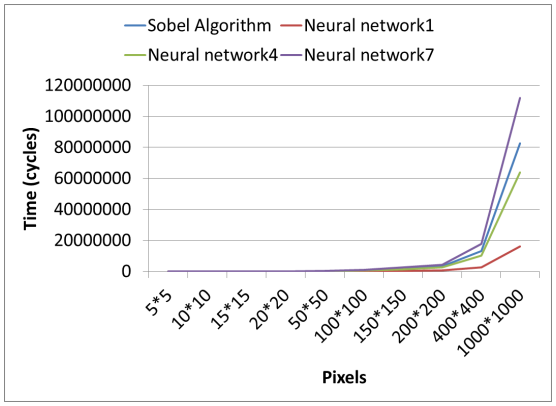

Fig. 12. Cycles required Sobel vs spiking neural network

neural networks compared to the Sobel filter. It can be seen that as the size of the image increases, the average overhead of the neural network based implementation decreases. This trend reveals that very large images, the neural networks with online learning can also be used as accelerators. However, further investigation is needed to quantify the use neural networks with online learning.

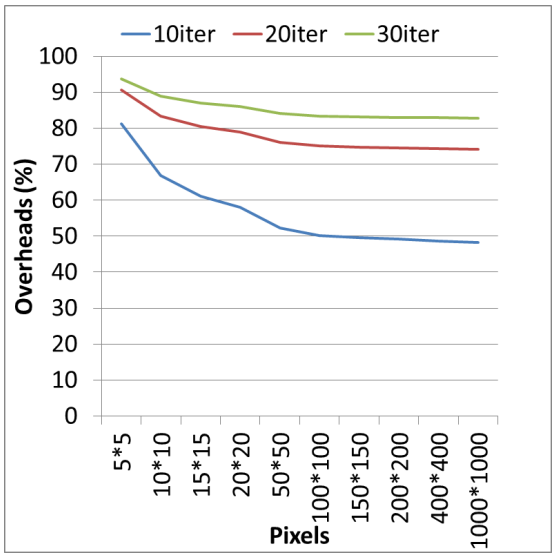

Fig. 13. Overheads with increasing complexity

B. Overhead analysis

To check the overhead imposed by the additional I & F unit, we synthesized it using 65 nm technology with frequency of 500 MHz. Table II shows the obtained results. It can be seen that the proposed enhancement incurs 9.1% area and 4.2% power overheads. Further evaluation revealed that the most of overhead (95% area) results from the fixed-point multipliers used to implement Equation 1. This overhead can be significantly reduced either by using a look up table or reusing the existing multiplier in the DPU (which will be considered in future).

TABLE II
AREA AND POWER CONSUMPTION OF I & F UNIT

	I & F	DRRA cell	Overhead (%)
Power mW	6.44	70.40	9.1
Area μm^2	50920	1199506	4.2

VII. CONCLUSION

In this paper, we have presented NeuroCGRA, to efficiently host estimation algorithms (using neural networks) along side normal calculations algorithms on a CGRA. The overall architecture relies on dedicated hardware blocks (to mimic the neuron) and a modified control flow (to translate the neural network application model to DRRA bitstream). To implement the neural networks, we embedded additional hardware to mimic the neuron functionality. The synthesis/simulation results using edge detection reveal that proposed architecture incurs acceptable overheads (4.4% area and 9.1% power) and process a video stream of up to 1M pixels in real-time. Future research on NeuroCGRA will involve development of other algorithms. Additionally, we also plan to test the feasibility of using neural networks, with online learning, as accelerators (for large holographic images).

REFERENCES

[1] D. Alnajjar, H. Konoura, Y. Ko, Y. Mitsuyama, M. Hashimoto, and T. Onoye, "Implementing flexible reliability in a coarse-grained reconfigurable architecture," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems,*, vol. PP, no. 99, pp. 1–1, 2012.

[2] M. A. Shami, "Dynamically reconfigurable resource array," Ph.D. dissertation, Royal Institute of Technology (KTH), Stockholm, Sweden, 2012. [Online]. Available: web.it.kth.se/~hemani/Athesis15.pdf

[3] Z. ul Abdin and B. Svensson, "Evolution in architectures and programming methodologies of coarse-grained reconfigurable computing," *Microprocess. Microsyst.*, vol. 33, no. 3, pp. 161–178, May 2009. [Online]. Available: http://dx.doi.org/10.1016/j.micpro.2008.10.003

[4] A. Sampson, W. Dietl, E. Fortuna, D. Gnanapragasam, L. Ceze, and D. Grossman, "Enerj: Approximate data types for safe and general low-power computation," in *Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '11. New York, NY, USA: ACM, 2011, pp. 164–174. [Online]. Available: http://doi.acm.org/10.1145/1993498.1993518

[5] S. M. A. H. Jafri, O. Ozbak, A. Hemani, N. Farahini, K. Paul, J. Plosila, and H. Tenhunen, "Energy-aware CGRAs using dynamically reconfigurable isolation cells." in *Proc. International symposium for quality and design (ISQED)*, 2013, pp. 104–111.

[6] S. Jafri, M. A. Tajammul, A. Hemani, K. Paul, J. Plosila, and H. Tenhunen, "Energy-aware-task-parallelism for efficient dynamic voltage, and frequency scaling, in cgras," in *Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIII), 2013 International Conference on*, 2013, pp. 104–112.

[7] J. Misra and I. Saha, "Artificial neural networks in hardware: A survey of two decades of progress," *Neurocomput.*, vol. 74, no. 1-3, pp. 239–255, Dec. 2010. [Online]. Available: http://dx.doi.org/10.1016/j.neucom.2010.03.021

[8] M. Krips, T. Lammert, and A. Kummert, "Fpga implementation of a neural network for a real-time hand tracking system," in *Electronic Design, Test and Applications, 2002. Proceedings. The First IEEE International Workshop on*, 2002, pp. 313–317.

[9] F. Yang and M. Paindavoine, "Implementation of an rbf neural network on embedded systems: Real-time face tracking and identity verification," *Trans. Neur. Netw.*, vol. 14, no. 5, pp. 1162–1175, Sep. 2003. [Online]. Available: http://dx.doi.org/10.1109/TNN.2003.816035

[10] Y. Maeda and T. Tada, "Fpga implementation of a pulse density neural network with learning ability using simultaneous perturbation," *Neural Networks, IEEE Transactions on*, vol. 14, no. 3, pp. 688–695, May 2003.

[11] S. Himavathi, D. Anitha, and A. Muthuramalingam, "Feedforward neural network implementation in fpga using layer multiplexing for effective resource utilization," *Neural Networks, IEEE Transactions on*, vol. 18, no. 3, pp. 880–888, May 2007.

[12] T. Chen, Z. Du, N. Sun, J. Wang, C. Wu, Y. Chen, and O. Temam, "Diannao: A small-footprint high-throughput accelerator for ubiquitous machine-learning," in *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '14. New York, NY, USA: ACM, 2014, pp. 269–284. [Online]. Available: http://doi.acm.org/10.1145/2541940.2541967

[13] H. Esmaeilzadeh, A. Sampson, L. Ceze, and D. Burger, "Neural acceleration for general-purpose approximate programs," in *Microarchitecture (MICRO), 2012 45th Annual IEEE/ACM International Symposium on*, Dec 2012, pp. 449–460.

[14] S. Chakradhar, M. Sankaradas, V. Jakkula, and S. Cadambi, "A dynamically configurable coprocessor for convolutional neural networks," in *Proceedings of the 37th Annual International Symposium on Computer Architecture*, ser. ISCA '10. New York, NY, USA: ACM, 2010, pp. 247–257. [Online]. Available: http://doi.acm.org/10.1145/1815961.1815993

[15] Z. Li, "Aerial image analysis using spiking neural networks with application to power line corridor monitoring," Ph.D. dissertation, Queensland University of Technology, 2011.

[16] M. A. Shami and A. Hemani, "Classification of massively parallel computer architectures," in *Proc. IEEE Int. Parallel and Distributed Processing Symposium Workshops PhD Forum (IPDPSW)*, May 2012, pp. 344–351.

[17] N. Farahini, S. Li, M. A.l Tajammul, M. A. Shami, G. Chen, A. Hemani, W. Ye, "39.9 GOPs/Watt multi-mode CGRA accelerator for a multi-standard base station," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, 2013.

[18] S. Jafri, A. Hemani, K. Paul, J. Plosila, and H. Tenhunen, "Compact generic intermediate representation (CGIR) to enable late binding in coarse grained reconfigurable architectures," in *Proc. International Conference on Field-Programmable Technology (FPT),*, Dec. 2011, pp. 1 –6.

[19] M. A. Tajammul, S. M. A. H. Jafri, A. Hemani, J. Plosila, and H. Tenhunen, "Private configuration environments for efficient configuration in CGRAs," in *Proc. Application Specific Systems Architectures and Processors (ASAP)*, Washington, D.C., USA, 5–7 June 2013.

[20] E. Izhikevich, "Simple model of spiking neurons," *Neural Networks, IEEE Transactions on*, vol. 14, no. 6, pp. 1569–1572, 2003.

[21] Y.-K. Kwok and I. Ahmad, "Static scheduling algorithms for allocating directed task graphs to multiprocessors," *ACM Comput. Surv.*, vol. 31, no. 4, pp. 406–471, Dec. 1999. [Online]. Available: http://doi.acm.org/10.1145/344588.344618

[22] Q. Wu, T. M. McGinnity, L. P. Maguire, A. Belatreche, and B. P. Glackin, "Edge detection based on spiking neural network model," in *ICIC (2)*, 2007, pp. 26–34.