

# A Fuzzy Logic Reconfiguration Engine for Symmetric Chip Multiprocessors

Muhammad Yasir Qadri, Klaus D. McDonald-Maier  
School of Computer Science and Electronic Engineering  
University of Essex, CO4 3SQ  
Colchester, United Kingdom  
yasirqadri@acm.org, kdm@essex.ac.uk

**Abstract**— Recent developments in reconfigurable multiprocessor system on chip (MPSoC) have offered system designers a great amount of flexibility to exploit task concurrency with higher throughput and less energy consumption. This paper presents a novel fuzzy logic reconfiguration engine (FLRE) for coarse grain MPSoC reconfiguration that facilitates to identify an optimum balance between power and performance of the system. The FLRE is composed on two levels of abstraction layers. The system selects an optimal configuration of Level 1 / Level 2 cache size and Associativity, processor operating frequency and voltage, the number of cores based on miss rate, and energy and throughput information of the system both at core and SoC level. An 8-core symmetric chip multiprocessor has been used to evaluate the proposed scheme. The results show an overall decrease of energy consumption with not more than 30% decrease in the throughput.

**Keywords**- *Fuzzy Logic; Reconfigurable Hardware; Symmetric Chip multiprocessors; Energy; Performance*

## I. INTRODUCTION

Fuzzy logic is a popular soft computing technique that is suited for runtime hardware adaptive systems as it is straightforward and lightweight to implement [1]. Fuzzy based systems are particularly well suited for applications where absolute precision is less important in making a general improvement to the configuration. With the advent of modern reconfigurable multiprocessor architectures it has now become possible to apply such techniques to find a better compromise between energy and throughput of a system.

This paper presents a novel fuzzy logic reconfiguration engine (FLRE) for coarse grain MPSoC reconfiguration to find an optimum balance between power and performance of the system. The FLRE is composed on two levels of abstraction layers. The FLRE at primary level or core level reconfigures core frequency/voltage, level 1 (L1) cache size and Associativity based on the miss rate, throughput and energy consumption of the core. At the secondary level or SoC level, Level 2 (L2) cache Associativity and size, along with number of active cores are selected on the basis of L2 miss rate, throughput and energy consumption of the entire system. This reconfigurable cache architecture can be

implemented in a runtime reconfigurable FPGA platform or alternatively by using hardware cache partitioning schemes such as the one proposed by Settle et al. [2].

This paper is organized into five sections. The following section covers the related research in the field of reconfigurable architectures and hardware adaption techniques. The section 3 describes the experimental setup and the FLRE. The subsequent two sections present the results, conclusion and future research directions.

## II. RELATED WORK

The majority of research in reconfigurable architectures has been done on field programmable gate arrays (FPGAs) as the target platform. Some of these FPGAs also provide runtime reconfiguration option which is largely being used in adaptive hardware scenarios to provide optimized energy consumption and performance, and to support larger configurations than the available area on the device. However FPGAs suffer from higher reconfiguration latencies as discussed in [3-5]. This issue has largely been addressed by preempting task implementation schemes for all possible locations at design time and then during the runtime a placement algorithm selects proper bitstream to be downloaded. One such approach was presented by Resano et al. in [5] using a dynamically reconfigurable hardware (DRH) model which facilitates task migration and inter-task communication. These authors have also proposed a hardware reconfiguration manager adopting prefetch scheduling and task replacement scheme designed to hide reconfiguration latency by more than 93% even for highly dynamic applications. Another example of such an approach was proposed by Kalte et al. [6] in which they have implemented each task to a single location in FPGA and manipulated the configuration data stream to relocate the task in the FPGA. As a result of their work a tool called REPLICA2Pro was developed to facilitate the reconfiguration task for the Virtex-II/Pro FPGAs. Danne et al. [7] proposed two periodic real-time tasks scheduling algorithms for full FPGA reconfiguration. The first one is based on the earliest deadline first (EDF) concept, termed EDF-Next Fit (EDF-NF) and the second is based on the concept of servers that reserve area and execution time for other tasks called Merge Server Distribute Load (MSDL).

The system utilization of EDF-NF algorithm was found to be better than MSDL, however MSDL was proven to be more feasible for larger real-time tasks sets. Other examples of scheduling techniques and their applications could be found in [8-11].

MPSoCs have been investigated for runtime energy aware scheduling in several research articles. Thread scheduling is generally classified into 1) balanced and 2) unbalanced scheduling categories. Balanced scheduling has an equal amount of threads distributed among the cores. DeVuyst et al. [12] have analyzed performance of various scheduling schemes considering both energy and performance, and have shown that uneven thread scheduling often outperforms balanced scheduling as greater throughput can be achieved by combining certain threads together on one core rather than by distributing these among several cores. Reconfigurable multicore platforms also pose challenges in handling the communication between dynamically changing tasks and their synchronization. Li [13] has performed a detailed analysis of the performance of various task scheduling algorithms for minimizing schedule length combined with an energy consumption constraint and for minimizing energy consumption combined with a schedule length constraint on Dynamic Voltage and Frequency (DVF) supported multiprocessor systems. Yang et al. [14] have proposed a task scheduling method for concurrent tasks on a multicore platform that combines offline and online scheduling to exploit the energy-performance trade-off at runtime. This work is an extension of a proposed framework based on grey box modeling for improved concurrency and lower energy consumption by Prayati et al. [15]. Ma et al. [16] discuss a design time and runtime scheduling scheme for concurrent task management for real-time applications on a heterogeneous multicore platform. At design time, a set of schedules and assignments for each task was defined using Pareto curves, and at runtime a lightweight scheduler was used to select optimal working points exploiting dynamic and nondeterministic behavior of the system. For an MPEG4 texture decoder application their approach has shown significant improvement in performance while maintaining lower energy consumption. Other reference to related work in this field can be found in [17-19].

Considering memory as the best candidate for optimization in an energy constrained multicore scenario, Ahn et al. [20] presented a simplified approach of grouping DRAM chips into multiple virtual memory devices that receive separate address and control signals on a shared command path. This approach reduces energy consumption by minimizing the number of bits activated per memory access and by replacing the memory register with a demultiplexor register for routing command signals to the appropriate memory module instead of mere transmission on the path. Another candidate for power optimization in an MPSoC is Networks-on-Chip (NoC). Kim et al. [21] have presented a novel low latency router architecture with two stage pipeline employing adaptive routing scheme for congestion aware flow control. The router architecture was termed as path sensitive, as it utilizes look-ahead routing for

selecting the next route based on the four possible quadrants and routes the packet to the corresponding virtual channels assigned to that quadrant. Additionally, based on this partitioned approach a decomposed crossbar switch was proposed that results in a reduction of size for its connections and lower packet conflicts. Their work also includes a complete solution safeguarding against both the traditional link faults and internal router upsets, without incurring any significant latency, area and power overhead. Park et al. [22] have provided a detailed analysis of various logic errors and have proposed data recovery mechanisms. Individual cases were analyzed such as link errors occurring during flit traversal between routers, deadlocks, intra router errors in router pipeline such as errors caused by virtual channel allocators, routing units, switch allocators, and crossbars.

### III. SYSTEM DESCRIPTION AND EXPERIMENTAL SETUP

An 8-core symmetric chip multiprocessor (SCMP) platform was designed to be evaluated on the proposed reconfiguration scheme. The SCMP is based on the Intel x86 architecture, with a customised shared memory architecture. The platform is comprised of L1 and L2 caches with configurable size and Associativity. The number of cores and processor frequency/voltage could be adjusted for energy and throughput regulation. The system configuration and parameters are described in Table 1. The energy consumption and voltage/frequency information was obtained from the Intel 486 GX embedded processor datasheet [23]. Each core of the system is connected to a CMOS switch, which allows to turn off the core, and prevent the leakage energy of the core from contributing to the overall energy consumption of the system. The default size and Associativity for L1 and L2 caches are 8KB, and 4-way set associative; and 128KB and 8-way set associative. The miss rate for L1 cache was assumed to be 10 cycles and that for L2 cache was set as 30 cycles. Each core in the SCMP is linked through a router that provides a seamless communication interface to the other members of this NoC. The router architecture was selected as a virtual channel arrangement with 5-stage pipeline similar to the one proposed by Peh et al. [24].

TABLE I. SYSTEM PARAMETERS

Parameter	Value
Processor Type	Intel x86
Number of Cores	8
Operating Frequencies	[16, 20, 25, 33] MHz
Operating Voltages	[2, 2.2, 2.4, 2.7]V
Energy Consumption per cycle	[13.1,15.4,18.7,22.9]nJ

In order to search for an optimal configuration, a fuzzy logic reconfiguration engine based on Mamdani's [25] inference technique was employed. A detailed description of each membership function for input variables such as L1 and L2 miss rate, normalized energy consumption, and throughput; and output parameters such as L1 and L2 Cache

Associativity and size, operating frequency and number of cores is given in Table 2.

To establish the relationship between the input variables and output parameters of the SoC, fuzzy logic rules were defined as shown in Appendix B. The rules were formed in way that a balanced throughput and energy consumption ratio could be achieved. For primary or core level configuration the FLRE keeps track of the average L1 miss rate, energy consumption and throughput for all the cores and strives to find an optimum cache size, Associativity and operating frequency. The cache size and Associativity does not only affect the miss rate but also have an impact on the throughput and energy consumption of the device. Similarly for the secondary or system level configuration the FLRE strives to find an optimal number of cores and L2 cache size and Associativity while taking into account the L2 miss rate and total throughput and energy consumption of the SoC.

TABLE II. INPUT AND OUTPUT MEMBERSHIP FUNCTIONS OF FLRE

$\mu_A = \begin{cases} 0 & \text{if L1/L2 Miss rate} \leq 0\% \text{ or } \geq 40\% \\ \frac{40-x}{40} & \text{if } 0\% \leq \text{L1/L2 Miss rate} \leq 40\% \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if L1/L2 Miss rate} \leq 0\% \text{ or } \geq 75\% \\ \frac{x-25}{25} & \text{if } 25\% \leq \text{L1/L2 Miss rate} \leq 50\% \\ \frac{75-x}{25} & \text{if } 50\% \leq \text{L1/L2 Miss rate} \leq 75\% \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if L1/L2 Miss rate} \leq 60\% \text{ or } \geq 100\% \\ \frac{x-60}{40} & \text{if } 60\% \leq \text{L1/L2 Miss rate} \leq 100\% \end{cases}$	
<b>(a) L1 and L2 Miss rate</b>	
$\mu_A = \begin{cases} 0 & \text{if Energy cons.} \leq 0 \text{ or } \geq 0.35 \\ \frac{0.35-x}{0.35} & \text{if } 0 \leq \text{Energy cons.} \leq 0.35 \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if Energy cons.} \leq 0 \text{ or } \geq 0.8 \\ \frac{x-0.2}{0.3} & \text{if } 0.2 \leq \text{Energy cons.} \leq 0.5 \\ \frac{0.8-x}{0.3} & \text{if } 0.5 \leq \text{Energy cons.} \leq 0.8 \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if Energy cons.} \leq 0.65 \text{ or } \geq 1.0 \\ \frac{x-0.65}{0.35} & \text{if } 0.65 \leq \text{Energy cons.} \leq 1.0 \end{cases}$	
<b>(b) Normalized Energy Consumption</b>	
$\mu_A = \begin{cases} 0 & \text{if Throughput} \leq 0 \text{ or } \geq 0.35 \\ \frac{0.35-x}{0.35} & \text{if } 0 \leq \text{Throughput} \leq 0.35 \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if Throughput} \leq 0 \text{ or } \geq 0.8 \\ \frac{x-0.2}{0.3} & \text{if } 0.2 \leq \text{Throughput} \leq 0.5 \\ \frac{0.8-x}{0.3} & \text{if } 0.5 \leq \text{Throughput} \leq 0.8 \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if Throughput} \leq 0.65 \text{ or } \geq 1.0 \\ \frac{x-0.65}{0.35} & \text{if } 0.65 \leq \text{Throughput} \leq 1.0 \end{cases}$	
<b>(c) Normalized Throughput</b>	
$\mu_A = \begin{cases} 0 & \text{if L1/L2 Cache Assoc.} \leq 0 \text{ or } \geq 2 \\ 1 & \text{if } 0 \leq \text{L1/L2 Cache Assoc.} \leq 2 \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if L1/L2 Cache Assoc.} \leq 1 \text{ or } \geq 8 \\ 1 & \text{if } 1 \leq \text{L1/L2 Cache Assoc.} \leq 8 \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if L1/L2 Cache Assoc.} \leq 4 \text{ or } \geq 16 \\ 1 & \text{if } 4 \leq \text{L1/L2 Cache Assoc.} \leq 16 \end{cases}$	
<b>(d) L1 and L2 Cache Associativity</b>	
$\mu_A = \begin{cases} 0 & \text{if L1 Cache size} \leq 0KB \text{ or } \geq 3.5KB \\ \frac{3.5-x}{3.5} & \text{if } 0KB \leq \text{L1 Cache size} \leq 3.5KB \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if L1 Cache size} \leq 0KB \text{ or } \geq 7KB \\ \frac{x-2}{2.5} & \text{if } 2KB \leq \text{L1 Cache size} \leq 4.5KB \\ \frac{7-x}{2.5} & \text{if } 4.5KB \leq \text{L1 Cache size} \leq 7KB \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if L1 Cache size} \leq 5.5KB \text{ or } \geq 8KB \\ \frac{x-0.65}{0.35} & \text{if } 5.5KB \leq \text{L1 Cache size} \leq 8KB \end{cases}$	

<b>(e) L1 Cache Size</b>	
$\mu_A = \begin{cases} 0 & \text{if Operating freq.} \leq 16MHz \text{ or } \geq 20MHz \\ 1 & 16MHz \leq \text{Operating freq.} \leq 20MHz \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if Operating freq.} \leq 20MHz \text{ or } \geq 25MHz \\ 1 & \text{if } 20MHz \leq \text{Operating freq.} \leq 25MHz \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if Operating freq.} \leq 25MHz \text{ or } \geq 33MHz \\ 1 & \text{if } 25MHz \leq \text{Operating freq.} \leq 33MHz \end{cases}$	
<b>(f) Operating Frequency</b>	
$\mu_A = \begin{cases} 0 & \text{if L2 Cache size} \leq 0KB \text{ or } \geq 50KB \\ \frac{50-x}{50} & \text{if } 0KB \leq \text{L2 Cache size} \leq 50KB \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if L2 Cache size} \leq 20KB \text{ or } \geq 100KB \\ \frac{x-20}{40} & \text{if } 20KB \leq \text{L2 Cache size} \leq 60KB \\ \frac{100-x}{40} & \text{if } 60KB \leq \text{L2 Cache size} \leq 100KB \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if L2 Cache size} \leq 80KB \text{ or } \geq 128KB \\ \frac{x-80}{48} & \text{if } 80KB \leq \text{L2 Cache size} \leq 128KB \end{cases}$	
<b>(g) L2 Cache size</b>	
$\mu_A = \begin{cases} 0 & \text{if No. of cores} \leq 1 \text{ or } \geq 3 \\ 1 & 1 \leq \text{No. of cores} \leq 3 \end{cases}$	
$\mu_B = \begin{cases} 0 & \text{if No. of cores} \leq 2 \text{ or } \geq 6 \\ 1 & \text{if } 2 \leq \text{No. of cores} \leq 6 \end{cases}$	
$\mu_C = \begin{cases} 0 & \text{if No. of cores} \leq 5 \text{ or } \geq 8 \\ 1 & \text{if } 5 \leq \text{No. of cores} \leq 8 \end{cases}$	
<b>(h) Number of Cores</b>	

The fuzzy logic reconfiguration scheme was tested on the proposed MPSoC using Simics full system simulator. Simics [26] facilitates instruction level simulations and is capable to run unmodified OS such as VxWorks, Solaris, Linux, Tru64, and Windows XP virtually on the target platforms. The simulator is targeted to provide fairly accurate timing profile, but at present does not support energy profiling of a target system. Simics also provides a fairly accurate cache profiling utility making it ideal for memory system research. Fedora v10 Linux operating system was used on the target platform, as it supports Advanced Configuration and Power Interface (ACPI) for hot-plugging (i.e. turning on/off) a CPU core on the go which is a vital feature for reconfigurable MPSoC scenarios such as the one presented here. Cache energy information was obtained from CACTI [27] which is an open source standard tool for highly accurate cache energy and timing analysis. However it is not a trace driven simulator, so energy consumption resulting in number of hits or misses is not accounted for a particular application. The information obtained from CACTI along with cache hit and miss rate profile provided by Simics was used to find the total cache energy consumption for the sample interval. The core interconnect energy consumption was estimated from Orion a power-performance simulator for interconnection networks [28-30]. Three OpenMP [31] based Class B, benchmark applications namely BT (Block Tridiagonal), CG (Conjugate Gradient), and LU (Lower-Upper symmetric Gauss-Seidel algorithm [32]) from the NAS parallel benchmark suite were executed on the target platform to evaluate the performance of the proposed scheme. The thread scheduling was done statically, and the thread profile was collected using Intel Concurrency Checker [33] which provided data such as core utilization, thread distribution, percentage of parallelism and timing of the applications. All applications were sampled for the first five seconds and then reconfiguration was done through decisions made by the fuzzy logic engine. Simics provides a facility of check-pointing through which the

current state of the system can be saved and then machine parameters can be modified in the checkpoint file. When this file is reloaded into the simulator, the simulation resumes from the state at which it was previously saved, however with new parameters. This feature was exploited for each iteration to modify parameters such as cache size, and Associativity, and operating frequency of the target platform. The number of cores in the SoC was adjusted by using the Linux hotplug feature. The applications were re-executed for each iteration, since, as the application proceeds in execution, the sampled average cache miss rate keeps on changing, so a clear impact of cache reconfiguration could not be judged and the same is the case for dynamic thread scheduling.

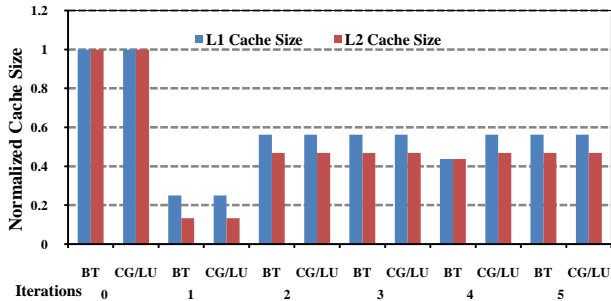


Figure 1. FLRE Results for Cache Sizing

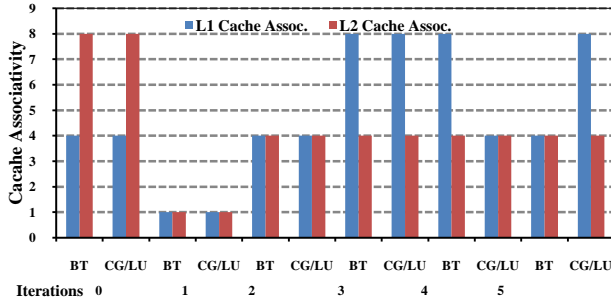


Figure 2. FLRE Results for Cache Associativity

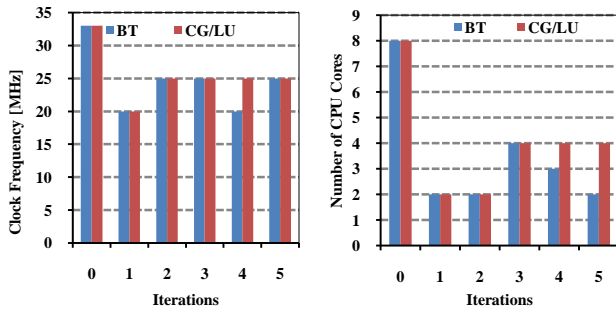


Figure 3. FLRE Results for Clock Frequency and Number of Cores

#### IV. RESULTS

The FLRE started optimizing the configuration parameters by first taking the data from the un-optimized core (Iteration 0). For each iteration, core and system level optimizations were passed to the target system and the benchmark applications were re-executed for a consistent miss rate profiling. The main objective of the system is to search for an optimum solution for performance and energy

of the MPSoC. The inference engine completed the system configuration in five iterations and results were found to be invariant for all the subsequent iterations. The operation of the FLRE is shown in Figures 1, 2, and 3; where for input variables such as core utilization, throughput, and energy consumption; parameters such as L1 and L2 cache size (Figure 1) and Associativity (Figure 2), processor frequency and number of cores (Figure 3) are being optimized. For each of the iterations, the resultant core utilization, throughput, L1 and L2 miss rate along with the total energy consumption are shown in Figures 4, 5, and 6. The L1 cache size was finally configured as 4.5KB, for all applications while the Associativity was selected as 4-way and 8-way set Associative for BT and CG/LU benchmarks respectively, as compared to the original 8KB, 4-way set associative cache. The L2 cache size and Associativity was optimized as 60KB, 4-way set associative cache; while the original configuration was 128KB, 8-way set associative cache. The number of cores has been reduced from 8 to 2 for BT, and for CG/LU to 4, while the frequency of operation has been selected as 25MHz where the default was 33MHz. The optimizations resulted in an overall 100% increase in core utilization for BT, and around 38% increase for CG, and 67% for LU applications. The energy consumption has been reduced by more than 6 times for BT, and by approximately two times for CG and LU applications. The system throughput has been decreased to 74% for the BT, 91% for the CG, and 82% for the LU benchmarks when compared with the one for default configuration. A significant increase in the miss rate in L1 and L2 could be observed but that is a result of finding an optimal balance in the system's performance and energy consumption. An infinitely large cache with highest Associativity is an ideal solution for the least possible miss rate. However the cache energy and throughput greatly varies with its size and Associativity. Therefore the miss rate was compromised to an extent in order to permit greater energy savings for the overall system.

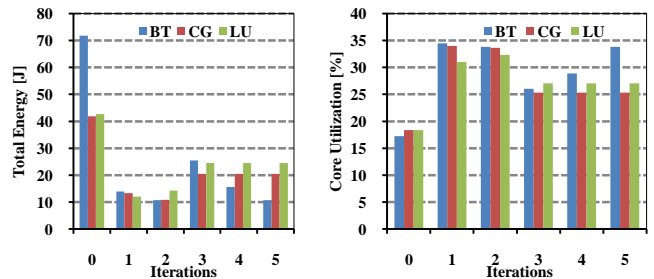


Figure 4. Impact of optimizations on Total Energy and Core Utilization

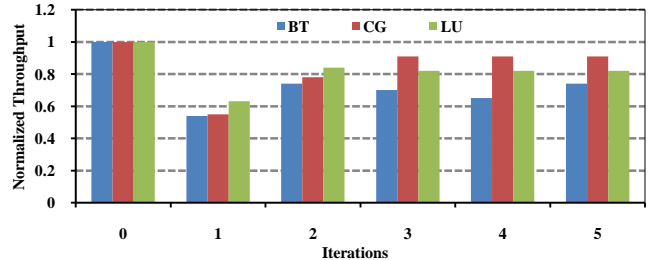


Figure 5. Impact of optimizations on Throughput

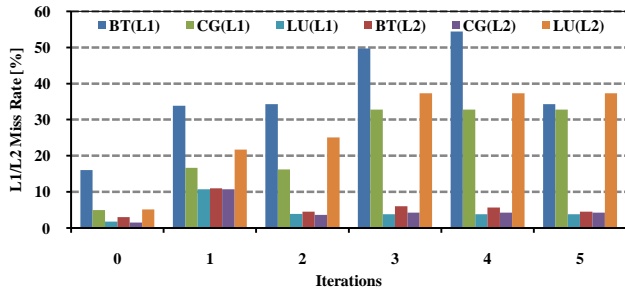


Figure 6. Impact of optimizations on L1 and L2 Miss Rates

## V. CONCLUSION AND FUTURE DIRECTIONS

This paper has presented a novel fuzzy logic based MPSoC reconfiguration scheme. The fuzzy reconfiguration engine was used to find an optimal balance between energy consumption and performance of the system. To evaluate the proposed scheme an Intel x86 based multicore SoC with 8 processor cores and a shared memory architecture, was simulated using Simics full system simulator. The SoC architecture included power gated cores for minimal energy consumption whilst not in use. A detailed analysis of core, cache, and interconnect power consumption was conducted and a significant amount of energy saving with increased core utilization has been observed. However due to these optimizations, the device throughput was reduced with an increase in cache miss rate.

The system in general validated the use of the proposed Fuzzy Logic based technique for MPSoC reconfiguration; therefore this technique can be adapted for a variety of architectures to search a good compromise for throughput and energy under user defined constraints. The proposed MPSoC architecture can be tailored to be used in variety of applications such as NoC research, dynamic thread scheduling, operating system development and high performance computing. Future work will see dynamic thread scheduling applied to the system for it to be able to reconfigure while executing a task.

## REFERENCES

- [1] V. KADIRKAMANATHAN, "FUZZY LOGIC AND CONTROL: SOFTWARE AND HARDWARE APPLICATIONS. MOHAMMAD JAMSHIDI, NADER VADIEE AND TIMOTHY J. ROSS (EDS.)," *ARTIFICIAL INTELLIGENCE REVIEW*, VOL. 13, PP. 337-339 1999.
- [2] A. SETTLE, D. CONNORS, E. GIBERT, AND A. GONZALEZ, "A DYNAMICALLY RECONFIGURABLE CACHE FOR MULTITHREADED PROCESSORS," *JOURNAL OF EMBEDDED COMPUTING*, VOL. 2, PP. 221-233, 2006.
- [3] K. COMPTON AND S. HAUCK, "RECONFIGURABLE COMPUTING: A SURVEY OF SYSTEMS AND SOFTWARE," *ACM COMPUTING SURVEYS (CSUR)*, VOL. 34, PP. 171-210, 2002.
- [4] T. MARESCAUX, A. BARTIC, D. VERKEST, S. VERNALDE, AND R. LAUWEREINS, "INTERCONNECTION NETWORKS ENABLE FINE-GRAIN DYNAMIC MULTI-TASKING ON FPGAs," IN *LECTURE NOTES IN COMPUTER SCIENCE*. VOL. 2438 MONTPELLIER, FRANCE, 2002, PP. 795-805.
- [5] J. RESANO, D. MOZOS, D. VERKEST, AND F. CATTHOOR, "A RECONFIGURABLE MANAGER FOR DYNAMICALLY RECONFIGURABLE HARDWARE," *IEEE DESIGN & TEST OF COMPUTERS*, VOL. 22, PP. 452-460, 2005.
- [6] H. KALTE AND M. PORRMANN, "REPLICA2PRO: TASK RELOCATION BY BITSTREAM MANIPULATION IN VIRTEX-II/PRO FPGAs," IN *PROCEEDINGS OF THE 3RD CONFERENCE ON COMPUTING FRONTIERS ISCHIA, ITALY*: ACM, 2006, PP. 403-412.
- [7] K. DANNE AND M. PLATZNER, "A HEURISTIC APPROACH TO SCHEDULE PERIODIC REAL-TIME TASKS ON RECONFIGURABLE HARDWARE," IN *PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON FIELD PROGRAMMABLE LOGIC AND APPLICATIONS (FPL) 2005* TEMPERE, FINLAND, 2005, PP. 568-573.
- [8] K. DANNE, R. MIHLENBERND, AND M. PLATZNER, "EXECUTING HARDWARE TASKS ON DYNAMICALLY RECONFIGURABLE DEVICES UNDER REAL-TIME CONDITIONS," IN *PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON FIELD PROGRAMMABLE LOGIC AND APPLICATIONS, 2006. FPL '06*. MADRID, SPAIN, 2006, PP. 1-6.
- [9] K. DANNE AND M. PLATZNER, "AN EDF SCHEDULABILITY TEST FOR PERIODIC TASKS ON RECONFIGURABLE HARDWARE DEVICES," IN *PROCEEDINGS OF THE 2006 ACM SIGPLAN/SIGBED CONFERENCE ON LANGUAGE, COMPILERS, AND TOOL SUPPORT FOR EMBEDDED SYSTEMS* OTTAWA, ONTARIO, CANADA: ACM, 2006, PP. 93-102.
- [10] P. SAHA AND T. EL-GHAZAWI, "EXTENDING EMBEDDED COMPUTING SCHEDULING ALGORITHMS FOR RECONFIGURABLE COMPUTING SYSTEMS," IN *PROCEEDINGS OF 3RD SOUTHERN CONFERENCE ON PROGRAMMABLE LOGIC, 2007. SPL '07*. 2007 MAR DEL PLATA, ARGENTINA, 2007, PP. 87-92.
- [11] P. SAHA AND T. EL-GHAZAWI, "SOFTWARE/HARDWARE CO-SCHEDULING FOR RECONFIGURABLE COMPUTING SYSTEMS," IN *PROCEEDINGS OF 15TH ANNUAL IEEE SYMPOSIUM ON FIELD-PROGRAMMABLE CUSTOM COMPUTING MACHINES, 2007. FCCM 2007*. NAPA VALLEY, CALIFORNIA, 2007, PP. 299-300.
- [12] M. DEVUYST, R. KUMAR, AND D. M. TULLSEN, "EXPLOITING UNBALANCED THREAD SCHEDULING FOR ENERGY AND PERFORMANCE ON A CMP OF SMT PROCESSORS," IN *PROCEEDINGS OF THE 20TH IEEE/ACM INTERNATIONAL PARALLEL AND DISTRIBUTED PROCESSING SYMPOSIUM* RHODES ISLAND, GREECE: IEEE/ACM, 2006.
- [13] K. LI, "PERFORMANCE ANALYSIS OF POWER-AWARE TASK SCHEDULING ALGORITHMS ON MULTIPROCESSOR COMPUTERS WITH DYNAMIC VOLTAGE AND SPEED," *IEEE TRANS. PARALLEL DISTRIB. SYST.*, VOL. 19, PP. 1484-1497, 2008.
- [14] P. YANG, C. WONG, P. MARCHAL, F. CATTHOOR, D. DESMET, D. VERKEST, AND R. LAUWEREINS, "ENERGY-AWARE RUNTIME SCHEDULING FOR EMBEDDED-MULTIPROCESSOR SOCS," *IEEE DESIGN & TEST OF COMPUTERS*, VOL. 18, PP. 46-58, 2001.
- [15] A. PRAYATI, W. CHUN, P. MARCHAL, N. COSSEMENT, F. CATTHOOR, R. LAUWEREINS, D. VERKEST, H. DE MAN, AND A. BIRBAS, "TASK CONCURRENCY MANAGEMENT EXPERIMENT FOR POWER-EFFICIENT SPEED-UP OF EMBEDDED MPEG4 IM1 PLAYER," IN *PROCEEDINGS OF INTERNATIONAL WORKSHOPS ON PARALLEL PROCESSING, 2000*. TORONTO, CANADA, 2000, PP. 453-460.
- [16] Z. MA, C. WONG, P. YANG, J. VOUNCKX, F. CATTHOOR, I. M. CENTER, AND B. LEUVEN, "MAPPING THE MPEG-4 VISUAL TEXTURE DECODER: A SYSTEM-LEVEL DESIGN TECHNIQUE BASED ON HETEROGENEOUS PLATFORMS," *IEEE SIGNAL PROCESSING MAGAZINE*, VOL. 22, PP. 65-74, 2005.
- [17] F. A. BOWER, D. J. SORIN, AND L. P. COX, "THE IMPACT OF DYNAMICALLY HETEROGENEOUS MULTICORE PROCESSORS ON THREAD SCHEDULING," *IEEE MICRO*, VOL. 28, PP. 17-25, 2008.
- [18] Y. JIANG, X. SHEN, J. CHEN, AND R. TRIPATHI, "ANALYSIS AND APPROXIMATION OF OPTIMAL CO-SCHEDULING ON CHIP MULTIPROCESSORS," IN *PROCEEDINGS OF THE 17TH INTERNATIONAL CONFERENCE ON PARALLEL ARCHITECTURES AND COMPILATION TECHNIQUES* TORONTO, ONTARIO, CANADA: ACM, 2008, PP. 220-229.
- [19] T. LI, D. BAUMBERGER, D. A. KOUFATY, AND S. HAHN, "EFFICIENT OPERATING SYSTEM SCHEDULING FOR PERFORMANCE-ASYMMETRIC MULTI-CORE ARCHITECTURES," IN *PROCEEDINGS OF THE 2007 ACM/IEEE CONFERENCE ON SUPERCOMPUTING* RENO, NEVADA: ACM, 2007, PP. 1-11.
- [20] J. H. AHN, J. LEVERICH, R. SCHREIBER, AND N. P. JOUPPI, "MULTICORE DIMM: AN ENERGY EFFICIENT MEMORY MODULE WITH INDEPENDENTLY CONTROLLED DRAMS," *COMPUTER ARCHITECTURE LETTERS*, VOL. 8, PP. 5-8, 2009.

- [21] J. KIM, D. PARK, T. THEOCHARIDES, N. VIJAYKRISHNAN, AND C. R. DAS, "A LOW LATENCY ROUTER SUPPORTING ADAPTIVITY FOR ON-CHIP INTERCONNECTS," IN *PROCEEDINGS OF THE 42ND ANNUAL DESIGN AUTOMATION CONFERENCE ANAHEIM, CALIFORNIA, USA: ACM*, 2005, pp. 559-564.
- [22] D. PARK, C. NICOPoulos, J. KIM, N. VIJAYKRISHNAN, AND C. R. DAS, "EXPLORING FAULT-TOLERANT NETWORK-ON-CHIP ARCHITECTURES," IN *DEPENDABLE SYSTEMS AND NETWORKS, INTERNATIONAL CONFERENCE ON*. VOL. 0 LOS ALAMITOS, CA, USA: IEEE COMPUTER SOCIETY, 2006, pp. 93-104.
- [23] INTEL, "EMBEDDED ULTRA-LOW POWER INTEL486™ GX PROCESSOR," IN *DATASHEET: INTEL CORPORATION*, 1997, p. 48.
- [24] L.-S. PEH, N. AGARWAL, N. JHA, AND T. KRISHNA, "GARNET: A DETAILED ON-CHIP NETWORK MODEL INSIDE A FULL-SYSTEM SIMULATOR," IN *INTERNATIONAL SYMPOSIUM ON PERFORMANCE ANALYSIS OF SYSTEMS AND SOFTWARE (ISPASS)*, 2009.
- [25] E. H. MAMDANI AND S. ASSILIAN, "AN EXPERIMENT IN LINGUISTIC SYNTHESIS WITH A FUZZY LOGIC CONTROLLER," *INTERNATIONAL JOURNAL OF MAN-MACHINE STUDIES*, VOL. 7, pp. 1-13, 1975.
- [26] P. S. MAGNUSSON, M. CHRISTENSSON, J. ESKILSON, D. FORSGREN, G. HALLBERG, J. HOGBERG, F. LARSSON, A. MOESTEDT, AND B. WERNER, "SIMICS: A FULL SYSTEM SIMULATION PLATFORM," *IEEE COMPUTER*, VOL. 35, pp. 50-58, 2002.
- [27] S. J. E. WILTON AND N. P. JOUPEI, "AN ENHANCED ACCESS AND CYCLE TIME MODEL FOR ON-CHIP CACHES," HP LABS %L WRL RESEARCH REPORT 93/5, 1994.
- [28] X. CHEN, L.-S. PEH, AND S. MALIK, "LEAKAGE POWER MODELING AND OPTIMIZATION IN INTERCONNECTION NETWORKS," IN *PROCEEDINGS OF THE INTERNATIONAL SYMPOSIUM ON LOW POWER AND ELECTRONICS DESIGN (ISLPED)* SEOUL, KOREA, 2003.
- [29] H. WANG, L.-S. PEH, AND S. MALIK, "ORION: A POWER-PERFORMANCE SIMULATOR FOR INTERCONNECTION NETWORKS," IN *PROCEEDINGS OF MICRO 35* ISTANBUL, TURKEY, 2002.
- [30] H. WANG, L.-S. PEH, AND S. MALIK, "A POWER MODEL FOR ROUTERS: MODELING ALPHA 21364 AND INFINIBAND ROUTERS," *IEEE MICRO*, VOL. 23, pp. 26-35, 2003.
- [31] H. JIN, M. FRUMKIN, AND J. YAN, "THE OPENMP IMPLEMENTATION OF NAS PARALLEL BENCHMARKS AND ITS PERFORMANCE," NASA AMES RESEARCH CENTER, 1999.
- [32] R. V. D. WUNGAART, "NAS PARALLEL BENCHMARKS VERSION 2.4," NASA ADVANCED SUPERCOMPUTING (NAS) DIVISION, NASA AMES RESEARCH CENTER, MOFFETT FIELD, CA 2002.
- [33] INTEL, "INTEL CONCURRENCY CHECKER v2.1," INTEL CORPORATION, 2008.
- [34] J. S. R. JANG AND N. GULLEY, "FUZZY LOGIC TOOLBOX FOR USE WITH MATLAB," *THE MATH WORKS INC*, 1995.
- [35] L. A. ZADEH, *FUZZY SETS, FUZZY LOGIC, AND FUZZY SYSTEMS: SELECTED PAPERS BY LOTFI A. ZADEH* VOL. 6: WORLD SCIENTIFIC, 1996.

## APPENDIX A: FUZZY LOGIC INTRODUCTION

Fuzzy logic uses a collection of membership functions defining input and output variables and specifies their corresponding relationship by IF-THEN based conditional statements called rules.

### a) Membership Functions

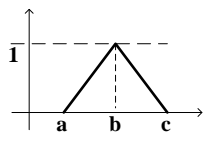
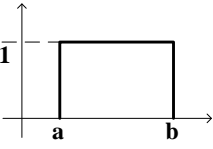
In contrast to a crisp-set, in which an element can belong to a set or not (i.e. having membership value of 1 or 0), a fuzzy logic membership function is a curve that defines the mapping of input values to a membership value between 0 and 1 [34]. This, in turn makes it convenient to represent linguistic labels such as slow, fast, medium, heavy etc. Although there are many types of fuzzy membership functions such as pi, bell, trapezoidal etc, we only describe in Table A1, rectangular or discrete, and triangular functions as they are used in this article.

### b) Logical Operations

In fuzzy logic, logical operations such as AND, OR, NOT have corresponding equivalents such as min, max, and complement and are defined as following [35]

$$\begin{aligned}\mu_{A \cup B}(x) &= \max[\mu_A(x), \mu_B(x)] \\ \mu_{A \cap B}(x) &= \min[\mu_A(x), \mu_B(x)] \\ \mu_{\bar{A}}(x) &= 1 - \mu_A(x)\end{aligned}$$

TABLE A1. TRIANGULAR AND RECTANGULAR MEMBERSHIP FUNCTIONS

Membership Function	Definition
	$\mu(x; a, b, c) = \begin{cases} 0 & \text{for } x < a \\ \frac{x-a}{b-a} & \text{for } a \leq x < b \\ \frac{c-x}{c-b} & \text{for } b \leq x \leq c \\ 0 & \text{for } x > c \end{cases}$
	$\mu(x; a, b) = \begin{cases} 0 & \text{for } x < a \text{ or } x > b \\ 1 & \text{for } a \leq x \leq b \end{cases}$

### c) If-Then Rules

The fuzzy logic rules comprise of if-then statements operating on the fuzzy sets using fuzzy operators e.g.

**If** temperature is **high** **then** put the fan at **high** speed

**If** temperature is **very low** and humidity is **high** **then** put the heater at **high** temperature.

Generally, a single rule cannot specify the relationship among the inputs and outputs so two or more rules are required. The output of each rule is a fuzzy set, which are aggregated to find a single output fuzzy set. The resulting fuzzy set is then defuzzified to get a crisp number output which could be applied to the physical world.

### d) Fuzzy Inference Systems

The process of formulating the mapping from input to output using fuzzy logic is the function of fuzzy inference systems. This article has used Mamdani's fuzzy inference system [25] which was among the earliest implementations of fuzzy logic in a control system. For defuzzification centeroid of the curve method was used to find the crisp output.

## APPENDIX B. RULES FOR FUZZY LOGIC INFERENCE ENGINE

TABLE B1. CORE LEVEL RULES FOR FLRE

L1 Miss Rate	Energy Cons.	Thput.	L1 Cache Assoc.	L1 Size	Clock Freq.
L	L	L	-	-	H
L	L	M	-	-	M
L	L	H	-	-	-
L	M	L	M	M	H
L	M	M	M	M	M
L	M	H	M	M	M
L	H	L	L	L	M
L	H	M	L	L	L
L	H	H	L	L	L
M	L	L	H	M	H
M	L	M	H	M	M
M	L	H	H	M	-
M	M	L	M	M	H
M	M	M	M	M	M
M	M	H	M	M	M
M	H	L	H	L	M
M	H	M	H	L	L
M	H	H	H	L	L
H	L	L	H	H	H
H	L	M	H	H	M
H	L	H	H	H	-
H	M	L	H	M	H

H	M	M	H	M	M
H	M	H	H	M	M
H	H	L	M	M	M
H	H	M	M	M	L
H	H	H	M	M	L

**Legend:** L = Low, M = Medium, and H = High

TABLE B2. SOC LEVEL RULES FOR FLRE

L2 Miss Rate	Energy Cons.	Thput.	L2 Cache Assoc.	Cache Size	No. of Cores
L	L	L	-	-	L
L	L	M	-	-	L
L	L	H	-	-	M
L	M	L	M	M	L
L	M	M	M	M	L
L	M	H	M	M	M
L	H	L	L	L	L
L	H	M	L	L	M
L	H	H	L	L	L
M	L	L	H	M	L
M	L	M	H	M	L
M	L	H	H	M	M
M	M	L	M	M	L
M	M	M	M	M	L
M	M	H	M	M	M
M	H	L	L	H	L
M	H	M	L	H	M
M	H	H	L	H	L
H	L	L	M	H	L
H	L	M	M	H	L
H	L	H	M	H	M
H	M	L	H	H	L
H	M	M	H	H	L
H	M	H	H	H	M
H	H	L	M	M	L
H	H	M	M	M	M
H	H	H	M	M	L

**Legend:** L = Low, M = Medium, and H = High