A Heterogeneous Built-In Self-Repair Approach Using System-Level Synthesis Flexibility

Inki Hong, Miodrag Potkonjak, *Member, IEEE*, and Ramesh Karri, *Member, IEEE*

Abstract—Summary and Conclusions —A novel methodology is proposed for designing fault-tolerant real-time multi-processor systems-on-a-chip to achieve optimal productivity. The methodology employs the heterogeneous built-in-self-repair (BISR) based on graceful degradation and yield enhancement techniques as an embedded optimization engine. The technique exploits the flexibility provided in task-level scheduling and algorithm selection steps. A hardware fault model is developed for modern super-scalar processors and multi-processors which enables an efficient treatment of the synthesis and compilation goals. For the first time, heterogeneous BISR is used at the task level. The key idea is to adapt scheduling and algorithm selection to the available nonfaulty resources. If there is a fault in memory, the algorithms that use less memory are selected and the scheduler exploits the other abundant resource, viz, the processors, more vigorously to compensate for the loss of part of memory. Similarly, a fault in a processor is backed up by memory. The synthesis approach minimizes the degradation in performance for single or multiple faults using simulated annealing-based algorithm selection, scheduling, and assignment algorithms. On the large set of examples this adaptive algorithm selection and scheduling technique has achieved important improvement of throughput compared to conventional nonadaptive schemes. The experimental results also indicate that important improvement in productivity can be achieved by using the extra throughput gained from the technique.

Index Terms—Algorithm selection, built-in self-repair, fault-tolerant system, scheduling, system-level synthesis.

ACRONYMS¹

BISR	built-in self-repair
ASIC	application-specific integrated circuit
HSDF	homogeneous synchronous data flow
PIF	productivity improvement factor
s-	statistical(ly)

NOTATION

$Y(n_p, n_m)$	yield when the number of processors is n_p and the number of memory units is n_m
P_p	$\Pr\{\text{a good processor}\}$
P_m	$\Pr\{\text{a good memory unit}\}$

Manuscript received May 30, 2002; revised June 30, 2002. A preliminary version of this paper was presented at the Asia & South-Pacific Design Automation Conference, Yokohama, Japan, 1998 Feb. 10–13, 1998. Responsible Editor: J. D. Healy.

I. Hong is with the Synopsys, Inc. Hillsboro, Oregon 97124 USA (e-mail: Inki@synopsys.com).

M. Potkonjak is with the Computer Science Department, University of California, Los Angeles, CA 90095 USA (e-mail: Miodrag@cs.ucla.edu).

R. Karri is with the Department of Electrical & Computer Engineering, Polytechnic University, Brooklyn, NY 11201 USA (e-mail: Ramesh@india.poly.edu).

Digital Object Identifier 10.1109/TR.2003.819047

¹The singular and plural of an acronym are always spelled the same.

I. INTRODUCTION

A. Motivation

SINCE the size of average application-specific systems and integrated systems has been doubling every year, the focus of behavioral and system-level synthesis has been shifting from single-task applications to multiple-task applications. The increasing levels of integration and aggressive deep-submicron technologies imply a greater need for fault-tolerance. While classical fault-tolerance techniques such as duplication and triplication provide fault tolerance in a straightforward manner, the overhead is usually too high for many cost-sensitive application domains such as consumer electronics and personal communication devices.

Recently, [5] proposed low overhead heterogeneous BISR for single application ASIC. This paper shows that the application domain of heterogeneous BISR can be appreciably enhanced, when intertask relationships are explored. In some sense the technique is also a step forward in the level of heterogeneity with BISR. While the technique due to [5] was able to back up only different types of execution units using a given type of execution unit, the new approach uses synthesis flexibility to back up memory with execution units or vice versa.

To the best of our knowledge, this is the first heterogeneous system-level synthesis approach which explores task-level scheduling and algorithm selection flexibility for design of fault-tolerant systems. The experimental results clearly show the effectiveness of the new technique.

B. Motivational Example

This sub-section provides a motivational example for each of 2 topics:

- adaptive fault-tolerant algorithm selection and scheduling problem,
- design of fault-tolerant system to optimize productivity [5] which is defined to be the ratio of the relative change in yield [20] over the relative change in area.

First, the adaptive fault-tolerant algorithm selection and scheduling problem using the motivational example in Table I are introduced. There are 3 homogeneous processors and 330 units of shared memory that are allocated. Five periodic tasks are given.

Assumption (To Simplify This Example): All the tasks have the same period.

Each task has two different algorithmic choices. Each algorithm requires a certain amount of shared memory to execute so that contention due to shared memory can cause some tasks to wait if there are processors available for their execution.

TABLE I
AN INSTANCE OF THE ADAPTIVE FAULT-TOLERANT ALGORITHM SELECTION AND SCHEDULING PROBLEM;
NUMBER OF ALLOCATED PROCESSORS = 3; AMOUNT OF ALLOCATED SHARED MEMORY = 330

Task, Algorithm	Processing Time	Memory Requirement	Switching Time Between Tasks									
			A,1	A,2	B,1	B,2	C,1	C,2	D,1	D,2	E,1	E,2
A,1	25	90	-	-	11	20	20	6	20	20	21	26
A,2	50	30	-	-	15	25	15	5	18	22	8	20
B,1	50	60	25	20	-	-	20	13	10	29	18	19
B,2	70	40	28	20	-	-	15	8	15	27	22	20
C,1	20	60	8	23	20	30	-	-	25	15	9	26
C,2	15	90	10	21	17	27	-	-	23	14	9	29
D,1	100	80	15	23	23	30	6	30	-	-	19	23
D,2	70	100	10	23	18	27	9	25	-	-	18	22
E,1	25	150	22	17	10	16	25	18	22	16	-	-
E,2	35	80	3	15	12	18	27	16	20	15	-	-

Fig. 1. Execution-time and memory-usage of the algorithm selection and schedule S^* .

To switch from one task to the other, context switching-time overhead occurs. The context switching-time includes the overhead of loading task's code/data from permanent storage to shared memory, state update in the processor, and other operating system over-heads.

Assumption: As soon as the context switch for the new task begins, the peak memory requirement for the task should be reserved. If available memory is less than the memory requirement, then the context switch has to wait until the required memory is available.

Therefore, a task can wait to save the memory usage if there is idle time available on a processor. There is an idle period between the task E and A on a Processor 1 in Fig. 1 due to this policy.

Assumption: A task to be scheduled next is assigned to the first available processor. If there are many processors available, then the processor with the minimum amount of context switching overhead is selected.

Typical real-time operating systems have context switching-time of tens of microseconds, which is usually large compared to typical communication or execution times of one iteration of a periodic hard real-time task.

An optimal algorithm selection and schedule, $S^* = (A1, C2, D2, B1, E1)$ takes 95 units of time per period. Fig. 1

illustrates how the execution time for S^* is computed. When 1 of the 3 processors is down due to a permanent physical fault, then S^* takes 151 units of time, while the other algorithm selection and schedule $S_p = (B1, D2, C2, A1, E1)$ takes 118 units of time. When 110 units of memory becomes unavailable, S^* takes 153 time units of time while the other algorithm selection and schedule, $S_m = (D2, B2, E2, A1, C1)$ takes 114 units of time. In-stead of using the algorithm selection and schedule, S^* , consider using the following adaptive algorithm selection and schedule:

- S^* for fault-free situation,
- S_p when a processor is faulty,
- S_m when a part of the memory is faulty.

Using this adaptive algorithm selection and schedule, the system throughput can be improved by 21.9% when a processor is faulty, and by 25.5% when there is a fault in the memory subsystem. If the required schedule length is in the range [118, 151] for a faulty processor situation and [114, 153] for a faulty memory situation, the graceful degradation of throughput by adapting the algorithm-selection and task-schedule to the currently available resources renders the system operational. On the other hand, in order to tolerate a fault in processor or memory in a similar situation, the classical schemes require duplication of memory and processors.

Next, consider a motivational example for the design of a fault-tolerant system to optimize productivity.

Assumption: A nonredundant system with 8 processors and 8 memory units is the base-system to compare with redundant systems.

Assumption: (For simplicity) The probability of a good processor and the probability of a good memory unit are both 0.9; and a processor and a memory unit take the same sea in a chip. As an initial step, all nonredundant configurations that satisfy throughput constraint using the proposed BISR-approach as an optimization engine are found; they are provided in Table II. # Memory in column 4 in Table II is the number of memory units, where a single memory unit becomes unavailable by a single fault. Consider 2 redundant systems which employ 9 processors and 8 memory units, and 8 processors and 9 memory units, respectively. The yield for the first redundant system becomes 0.334, and the yield for the second one becomes 0.482, while both systems incur the same amount of area penalty. To show how the yield is calculated, use the first redundant system

TABLE II
A MOTIVATIONAL EXAMPLE FOR OPTIMIZATION OF PRODUCTIVITY

# Tasks	Dead-line	nonredundant configuration for throughput constraint			
		# Processor	# Memory	Yield	Schedule Length
10	106	5	15	0.122	104
		6	11	0.167	104
		7	9	0.185	105
		8	8	0.185	106
		12	7	0.135	104

as an example. The yield for this first redundant system is the probability that the system works after manufacturing. There are two such cases:

- i) 9 processors and 8 memory units work, or
- ii) 8 processors and 8 memory units work.

The probability of case i is $0.9^9 \cdot 0.9^8$;

the probability of case ii is $9 \cdot 0.1 \cdot 0.9^8 \cdot 0.9^8$.

Similarly, the yield for the second redundant system is the probability that:

- i) 7 or 8 processors, and 9 memory units work, or
- ii) 8 processors and 8 memory units work.

The productivity of the second redundant system is greater by a factor of 1.44 than that of the first. The details of this methodology to find an optimal design for productivity are in Section V.

C. Paper Organization

Section II reviews the related works on scheduling, algorithm selection, and behavioral and architectural level fault-tolerant techniques. Section III summarizes the selected computation, hardware and fault models, and describes a global synthesis approach for designing a fault-tolerant system to optimize productivity. Sections IV and V are the technical core of the paper; Section IV formally defines the problem of the heterogeneous BISR using the adaptive algorithm selection and scheduling, presents a synthesis method for the problem, and provides comprehensive experimental results. Section V defines the problem of designing fault-tolerant system with optimal productivity, provides a synthesis method for the problem, and presents experimental results.

II. PREVIOUS WORK

Scheduling has been widely studied in many areas, such as behavioral synthesis [13], parallel processing [4], and hard real-time systems [12]. Recently, algorithm selection has been recognized as an important system-level synthesis topic, and several approaches have been proposed [15].

Behavioral synthesis provides the mechanism for design-space exploration so that a variety of design goals can be optimized [2], [13]. Much of the behavioral synthesis research has targeted the optimization of area, speed (throughput), and (more recently) power and testability.

Relatively little work has been reported on behavioral-level synthesis techniques for fault-tolerant design. Reference [18] concentrates on designs with self-recovery from transient faults using micro roll-back and checkpoint insertion. References [7],

Fig. 2. The hardware and fault model; x indicates a fault.

[8], [14] develop a spectrum of behavioral techniques to minimize the fault-tolerance overhead in application specific designs. Recently, [6] developed a high-level synthesis approach to optimize manufacturability using redundant interconnects. Reference [5] develops a heterogeneous BISR behavioral synthesis system, which explores flexibility in scheduling a single task during high-level synthesis, so that the resulting design is operational when there are no more than k faulty units.

For a review of known fault-tolerance techniques, [19] provides comprehensive lists of relevant references.

III. PRELIMINARIES

This section outlines the relevant preliminaries. In particular, it describes the selected hardware, fault, and computation models, and outlines a global design flow of the fault-tolerant system.

A. Hardware and Fault Model

The hardware model being considered is shown in Fig. 2. The processors, interconnect and memory shown in the figure are placed in a single-chip multi-processor platform. It is assumed that the area of a chip is the sum of the areas of processors and memory units, which accounts for the important portion of the overall area of the chip [1].

Assumption: The ratio of the area between processor and memory is provided.

Assumption: The increase in the area of a chip by the addition of a processor or memory is just the area of the added component, regardless of how many processors and memory units are in the current configuration.

It is virtually impossible to replace the faulty processors, memory, and interconnects in the single-chip system. Faults can occur in either a processor, a memory system, or an interconnect. Only the permanent physical faults are considered here. A fault in interconnect can be regarded as a fault in a processor. A faulty interconnect prevents its corresponding processor from receiving data from the memory system, as shown in P_2 . A fault in a memory system causes a part or all of the memory unusable. A fault in a processor causes a single processor down, as shown in P_1 . The identification of the faulty parts can be done by either an offline testing before packaging or an online testing. The controller of the system is reconfigured upon detection of a fault.

B. Computational Model

The syntax of a targeted computation is defined as a hierarchical data-control flow graph (CDFG) [17]. The CDFG represents the computation as a flow graph, with nodes, data edges, and control edges. The semantics underlying the syntax of the CDFG format is homogeneous synchronous data ROW (HSDF) [11], which assumes a semi-infinite or a very long input stream of data arriving and being processed at periodic intervals, imposed by the nature of the specified applications. The HSDF model is well suited for specification of single task computations in numerous application domains such as digital signal processing, video and image processing, broadband and wireless communications, control, information and coding theory, and multimedia.

Assumption: There are no data, control, or timing dependences among tasks. If there are dependences, the tasks are merged in a new composite single task.

Assumption: All tasks are periodic with the same period. This assumption implies that all the tasks should be completed to start the next period. Again note that, with no loss of generality using the least common multiple theorem [10] a set of tasks with arbitrary periods, can be transformed into the equivalent set of tasks with the same period in polynomial time. Since context switching overhead is usually high, we assume no task preemption. To execute a task, there must exist a unassigned processor and a required amount of unassigned operating memory for loading the task's code and data from the permanent secondary storage. Partly due to the loading of code and data, the context switch time is incurred before a new task starts its execution. The context switch times between tasks are described by a matrix between tasks.

IV. HETEROGENEOUS BISR USING ADAPTIVE ALGORITHM SELECTION AND SCHEDULING

This section formulates the adaptive algorithm selection and scheduling problem and provides the computational complexity of the problem. Next, it proposes synthesis method and presents experimental results.

A. Problem Formulation and Complexity

The goal is to achieve the graceful degradation of throughput in the presence of faults in the system using the adaptive algorithm selection and scheduling technique. The adaptive algorithm selection and scheduling problem can be formally stated as follows:

Problem: Throughput Optimization Using Adaptive Algorithm Selection and Scheduling

Instance: Given N tasks with A algorithmic choices that have processor and memory requirements, P homogeneous processors, M memory units, context switch time matrix CST and a period D .

Question: Are there a selection of algorithms and a schedule of tasks such that in the presence of any subset of size at most $K1$ faulty processors and any subset of size at most $K2$ faulty memory units, the resulting schedule length is at most D ?

```

Generate a starting solution  $S$ .
Set the best solution  $S^* = S$ .
Determine a starting temperature  $T$ .
while not yet frozen,
    while not yet at equilibrium for the current temperature,
        Choose a random neighbor  $S'$  of the current solution.
        if  $cost(S') < cost(S)$ 
            Set  $S = S'$ 
        if  $cost(S) < cost(S^*)$ 
             $S^* = S'$ 
    else
        Generate a random number  $r$  uniformly from  $[0, 1]$ .
         $diff = cost(S') - cost(S)$ ;
        if  $r < \exp(-diff/T)$ 
             $S = S'$ 
    Update the temperature  $T$ .
Return the best solution  $S^*$ .

```

Fig. 3. The pseudo-code for simulated annealing algorithm.

Theorem 1: The Throughput Optimization Using Adaptive Algorithm Selection and Scheduling problem is NP-complete.

Proof: The problem is proved to be NP-complete by using Karp's polynomial transformation technique from the traveling salesman problem which is well known to be NP-complete [3]. Consider a special case of the original problem formulation. Suppose there is a single processor with unlimited memory units. There are N tasks with 0 execution times. It takes time only when a processor switches from one task to the other, which is described as the context switch time matrix CST. Assume that the CST matrix is symmetric. Let each task denote a city and the CST matrix denote a distance matrix between the cities. The reduced problem is the traveling salesman problem. Because a special case of the problem is NP-complete, the original problem is NP-complete. Q.E.D.

B. Synthesis Approach

Because the computational complexity of the algorithm selection and scheduling problem forbids an exact or optimal solution, a general combinatorial optimization technique known as simulated annealing (SA) [9] has been used to obtain nearly optimal throughput, given the faults in processors and memory units. The technique can be applied to all the scenarios considered, with some straightforward changes. The SA algorithm is in Fig. 3.

The actual implementation details are presented for each of the following areas; the cost function, the neighbor solution generation, the temperature update function, the equilibrium criterion, and the termination criterion.

- 1) The actual schedule length for the given schedule has been used as the cost function.
- 2) The neighbor solution is generated by the random choice of 2 perturbations: i) the interchange of 2 tasks in the schedule, and ii) the change of algorithms of a task in the schedule.
- 3) The temperature is updated by the function $T_{\text{new}} = \alpha(T_{\text{old}}) * T_{\text{old}}$. α is a function of the tem-

TABLE III
EXPERIMENTAL RESULTS OF ADAPTIVE ALGORITHM SELECTION AND SCHEDULING FOR RANDOM EXAMPLES

Number of Tasks	Synthesis Techniques	(number of processor faults, number of memory faults)									CPU time (min)
		(0,0)	(0,1)	(0,2)	(1,0)	(1,1)	(1,2)	(2,0)	(2,1)	(2,2)	
10	initial	115	145	197	131	150	197	161	185	197	0.66
	adaptive	115	132	148	123	123	156	142	152	169	
	improvement (%)	0.0	9.0	24.9	6.1	18.0	20.8	11.8	17.8	14.2	
20	initial	217	279	340	258	293	340	333	340	363	2.8
	adaptive	217	234	266	235	260	294	303	304	327	
	improvement (%)	0.0	16.2	21.8	9.0	11.3	13.5	9.0	10.6	10.0	
30	initial	276	297	383	344	344	389	434	434	438	7.5
	adaptive	276	278	319	321	328	341	392	392	412	
	improvement (%)	0.0	6.4	17.7	6.7	4.7	12.4	9.7	9.7	6.0	
40	initial	395	475	673	507	516	673	642	642	715	13.8
	adaptive	395	427	473	435	459	520	551	551	592	
	improvement (%)	0.0	10.1	29.7	14.2	11.0	22.7	11.0	11.0	17.2	
50	initial	473	564	742	560	584	744	740	740	810	18.6
	adaptive	473	478	547	541	550	593	676	692	702	
	improvement (%)	0.0	15.3	26.3	3.4	5.8	20.3	8.7	6.5	13.3	

perature. For $T > 200.0$, α is chosen to be 0.1 so that in the high temperature region where every new state has very high chance of acceptance, the temperature reduction occurs very rapidly. For $1.0 < T \leq 200.0$, α is set to 0.95 so that the optimization process explores this promising region more slowly. For $T \leq 1.0$, α is set to 0.8 so that T is relatively quickly reduced to firmly converge to a local minimum. The initial temperature is set to 4 000 000.

- 4) The equilibrium criterion is specified by the number of iterations of the inner loop. The number of iterations of the inner loop is set to 20 times the number of tasks.
- 5) The termination criterion is given by the temperature. If the temperature falls below 0.1, the simulated annealing algorithm stops.

C. Experimental Results

A set of random examples was generated by varying the number of tasks to show the effects of the problem size on the performance. 10, 20, 30, 40, 50 tasks were tried. The number of different algorithmic choices for each task is randomly chosen between 2 and 5. The number of processors allocated is 5, and the number of memories allocated is 1000. The processing times for tasks are chosen randomly between 10 and 100, while the memory requirements are randomly distributed between 10 and 350. The context switching times take values between 5 and 30. Each fault in processors is assumed to cause a processor to be unavailable. When a fault in memory occurs, 200 units of memory are assumed to become unavailable. Table III illustrates the effectiveness of the adaptive scheduling approach.

- The first number in each item represents the schedule length for the nonadaptive schedule which is obtained under the no-fault assumption.
- The second number represents the schedule length for the adaptive schedule, which is obtained under the appropriate fault situation.
- The third number provides the percentage improvement for the adaptive schedule from the nonadaptive one.

The CPU times reported are the running times on SUN Sparc 4. The adaptive scheduling approach achieves an average improvement of 13.1% compared to the nonadaptive approach.

Four different sets of 12 tasks were constructed from the 16 real-life tasks in [16]. The tasks are:

- 2 GE controllers,
- 2 Honda controllers,
- a wavelet filter,
- 4 audio filters,
- an 8×8 discrete cosine transform,
- an NEC digital-to-analog converter,
- 2 components of modems,
- an LMS audio formatter,
- an echo canceler,
- a linear controller for automotive motion control.

For each task, the number of operations and the area of implementation were used as the running time and the memory requirement, respectively. The context switching times are approximated such that they are proportional to memory requirement, the switching between similar tasks costs little overhead and the switching between different tasks requires an important amount of overhead. The context switching costs are distributed between 10 and 173; 4 processors and 160 memory units are allocated. Each fault in memory causes 40 units of memory unusable. Algorithm selection is not considered because it is assumed that only 1 algorithmic choice is available for each task. Table IV presents the experimental results and clearly shows the effectiveness of this technique. The adaptive scheduling approach achieves an average improvement of 10.6% compared to the nonadaptive approach

V. DESIGN OF FAULT-TOLERANT SYSTEM FOR OPTIMIZING PRODUCTIVITY

A. Global Design Flow

The goal is to design a fault-tolerant system with optimal productivity using the heterogeneous BISR-approach (see Section IV). Stage #1 of the design flow designs a nonredundant

TABLE IV
EXPERIMENTAL RESULTS OF ADAPTIVE SCHEDULING FOR REAL-LIFE EXAMPLES WITHOUT ALGORITHM SELECTION

Tasks	Synthesis Techniques	(Number of processor faults, Number of memory faults)								
		(0,0)	(0,1)	(0,2)	(1,0)	(1,1)	(1,2)	(2,0)	(2,1)	(2,2)
Set 1	initial	1156	1769	2299	1276	1769	2299	1744	2009	2352
	adaptive	1156	1452	2096	1187	1472	2115	1494	1587	2136
	improvement (%)	0.0	18.0	8.2	7.0	16.8	8.0	14.3	21.0	9.2
Set 2	initial	1222	1873	2575	1392	1873	2575	1731	1885	2604
	adaptive	1222	1570	2242	1305	1567	2249	1631	1697	2269
	improvement (%)	0.0	16.2	13.0	6.2	17.3	12.7	5.8	10.0	12.9
Set 3	initial	693	693	858	830	830	868	973	973	973
	adaptive	693	693	776	754	754	808	955	956	956
	improvement (%)	0.0	0.0	9.6	9.2	9.2	6.9	1.8	1.7	1.7
Set 4	initial	1059	1739	2105	1107	1739	2115	1524	1789	2170
	adaptive	1059	1403	1885	1084	1403	1885	1321	1503	1900
	improvement (%)	0.0	19.3	10.5	2.1	19.3	10.9	13.3	16.0	12.4

system that satisfies the throughput requirement which is assumed to be given. Stage #2 finds all nonredundant system configurations that can satisfy the throughput requirement using the heterogeneous BISR-approach which uses the task-level scheduling and algorithm selection flexibility. Stage #3 finds an optimal redundant configuration that maximizes productivity by considering both yield-enhancement and area-penalty, which are described in Section V-B.

B. Synthesis Method

The yield and area for a given system configuration must be computed to get the productivity measure. Let the given system use n_p processors and n_m memory units. The probability of a good processor is P_p and the probability of a good memory unit is P_m . All n nonredundant system configurations are $(p_1, m_1), (p_2, m_2), \dots, (p_n, m_n)$; p_i and m_i are the number of processors and memory units, respectively; $p_i \leq p_{i+1}$. Observe that:

- if $p_i \leq p_{i+1}$, then $m_i \geq m_{i+1}$;
- if (p_i, m_i) is a nonredundant system configuration, then (p_j, m_j) is a redundant system configuration when $p_j \geq p_i$ and $m_j \geq m_i$.

The $Y(n_p, n_m)$ is computed using (1) which basically calculates the sum of the probabilities for all the possible functional configurations after manufacturing for a given number of processors and memory units. Each product term represents the probability that there are at least p_k functional processors and m_k memory units for $k = 1, \dots, n$. If $n_p < p_k$, then the product term is 0. The area for a configuration is computed based on an area model in Section III.

Let :

$$P_p \equiv \Pr\{\text{a good processor}\}$$

$$P_m \equiv \Pr\{\text{a good memory unit}\}$$

the n nonredundant system configurations are

$$(p_1, m_1), (p_2, m_2), \dots, (p_n, m_n),$$

$$p_i \equiv \text{number of processors}$$

$$m_i \equiv \text{number of memory units}$$

$$p_i < p_{i+1}$$

TABLE V
ALL NON-REDUNDANT CONFIGURATIONS THAT SATISFY THE THROUGHPUT REQUIREMENT USING THE PROPOSED HETEROGENEOUS BISR-APPROACH

# Tasks	Deadline	# Processor	# Memory	Schedule Length
10	106	5	15	104
		6	11	104
		7	9	105
		8	8	106
		12	7	104
20	183	6	9	182
		7	9	182
		8	8	183
		9	6	173
		14	5	181
30	256	6	10	253
		7	9	247
		8	8	256
		9	7	242
		10	6	255
40	325	6	15	321
		7	9	316
		8	8	325
		9	7	315
50	379	7	13	378
		8	8	379
		10	7	378

The yield equation for the system of n_p processors and n_m memory units is

$$\begin{aligned}
 Y(n_p, n_m) = & \sum_{i=0}^{n_p-p_n} \binom{n_p}{i} \cdot P_p^{n_p-i} \cdot (1-P_p)^i \\
 & \cdot \left(\sum_{j=0}^{n_m-m_n} \binom{n_m}{j} \cdot P_m^{n_m-j} \cdot (1-P_m)^j \right) \\
 & + \sum_{i=\max[n_p-p_n+1, 0]}^{n_p-p_{n-1}} \binom{n_p}{i} \cdot P_p^{n_p-i} \cdot (1-P_p)^i \\
 & \times \left(\sum_{j=0}^{n_m-m_{n-1}} \binom{n_m}{j} \cdot P_m^{n_m-j} \cdot (1-P_m)^j \right) + \dots \\
 & + \sum_{i=\max[n_p-p_2+1, 0]}^{n_p-p_1} \binom{n_p}{i} \cdot P_p^{n_p-i} \cdot (1-P_p)^i \\
 & \cdot \left(\sum_{j=0}^{n_m-m_1} \binom{n_m}{j} \cdot P_m^{n_m-j} \cdot (1-P_m)^j \right). \quad (1)
 \end{aligned}$$

TABLE VI
REDUNDANT SYSTEMS WITH OPTIMAL PRODUCTIVITY WHEN THE PROBABILITIES OF A GOOD PROCESSOR AND MEMORY UNIT ARE LOW

# Tasks	P_m	P_p														
		0.75					0.85					0.95				
		IY	FY	PIF	n_p	n_m	IY	FY	PIF	n_p	n_m	IY	FY	PIF	n_p	n_m
10	0.75	0.010	0.628	164.5	10	12	0.027	0.690	77.8	9	12	0.066	0.603	43.1	8	11
	0.85	0.027	0.708	79.9	10	11	0.074	0.633	40.1	9	10	0.181	0.462	24.9	8	9
	0.95	0.066	0.468	48.4	9	9	0.181	0.496	27.9	8	9	0.440	0.792	12.8	8	9
20	0.75	0.010	0.739	193.8	11	11	0.027	0.397	108.4	10	8	0.066	0.458	94.2	9	8
	0.85	0.027	0.397	108.4	8	10	0.074	0.562	52.6	8	10	0.181	0.645	41.1	9	8
	0.95	0.066	0.458	94.2	8	9	0.181	0.645	41.1	8	9	0.440	0.825	14.0	8	9
30	0.75	0.010	0.697	182.8	11	11	0.027	0.632	88.7	11	9	0.066	0.529	55.8	10	8
	0.85	0.027	0.632	88.7	9	11	0.074	0.684	43.8	10	9	0.181	0.496	27.9	9	8
	0.95	0.066	0.529	55.8	8	10	0.181	0.496	27.9	8	9	0.440	0.792	12.8	8	9
40	0.75	0.010	0.665	174.4	11	11	0.027	0.599	83.8	10	10	0.066	0.473	49.0	9	9
	0.85	0.027	0.599	83.8	10	10	0.074	0.480	43.7	9	9	0.181	0.496	27.9	9	8
	0.95	0.066	0.473	49.0	9	9	0.181	0.496	27.9	8	9	0.440	0.792	12.8	8	9
50	0.75	0.010	0.668	150.1	12	11	0.027	0.748	70.5	11	11	0.066	0.669	36.3	10	10
	0.85	0.027	0.742	69.8	12	10	0.074	0.543	33.7	10	9	0.181	0.398	19.2	8	9
	0.95	0.066	0.528	37.1	11	8	0.181	0.398	19.2	9	8	0.440	0.863	7.7	9	9

a) The area ratio between processor and memory is 1

# Tasks	P_m	P_p														
		0.75					0.85					0.95				
		IY	FY	PIF	n_p	n_m	IY	FY	PIF	n_p	n_m	IY	FY	PIF	n_p	n_m
10	0.75	0.010	0.677	199.8	9	14	0.027	0.517	107.7	8	12	0.066	0.603	64.6	8	11
	0.85	0.027	0.478	99.2	8	12	0.074	0.433	58.0	8	10	0.181	0.462	37.4	8	9
	0.95	0.066	0.261	70.4	8	9	0.181	0.496	41.8	8	9	0.440	0.792	19.2	8	9
20	0.75	0.010	0.460	269.1	8	12	0.027	0.633	133.3	8	12	0.066	0.458	70.7	9	8
	0.85	0.027	0.397	162.6	8	10	0.074	0.562	78.9	8	10	0.181	0.474	39.0	8	9
	0.95	0.066	0.458	141.3	8	9	0.181	0.645	61.7	8	9	0.440	0.825	21.0	8	9
30	0.75	0.010	0.486	228.0	8	13	0.027	0.572	119.8	8	12	0.066	0.611	65.5	8	11
	0.85	0.027	0.454	125.2	8	11	0.074	0.480	65.5	8	10	0.181	0.462	37.4	8	9
	0.95	0.066	0.529	83.7	8	10	0.181	0.496	41.8	8	9	0.440	0.792	19.2	8	9
40	0.75	0.010	0.630	185.7	10	12	0.027	0.369	100.4	8	11	0.066	0.600	64.3	8	11
	0.85	0.027	0.227	88.0	8	10	0.074	0.433	58.0	8	10	0.181	0.462	37.4	8	9
	0.95	0.066	0.261	70.4	8	9	0.181	0.496	41.8	8	9	0.440	0.792	19.2	8	9
50	0.75	0.010	0.750	147.6	12	12	0.027	0.711	75.2	10	12	0.066	0.349	51.0	8	10
	0.85	0.027	0.611	64.2	11	10	0.074	0.492	33.7	9	10	0.181	0.398	28.8	8	9
	0.95	0.066	0.675	31.4	11	9	0.181	0.557	16.6	9	9	0.440	0.616	9.6	8	9

b) The area ratio between processor and memory is 2

IY $\equiv$ initial yield; FY $\equiv$ final yield; PIF $\equiv$ improvement factor in productivity; n_p $\equiv$ # of processors in the final design; n_m $\equiv$ # memory units in the final design; Initially, 8 processors and 8 memory units are used.

The algorithm to maximize productivity is based on enumeration, where, from every configuration, new configurations are generated by adding a processor or memory unit, and the further enumeration from the new configuration stops when the productivity of the new configuration is smaller than that of the previous configuration. The algorithm finds an optimal solution because the relative improvement in yield by adding a processor or memory decreases as the number of processors or memory units increases, and the relative increase in area by addition of a processor or memory stays constant regardless of how many processors and memory units are used in the current configuration.

C. Experimental Results

The same random examples are used here as used in the previous section, with a change that 8 processors and 800 memory units are used, where 100 memory units are affected by a single fault in memory. The number of memory units reported in the paper is in terms of 100 memory units. Using the adaptive algorithm selection and scheduling, all the nonredundant system

configurations have been identified and provided in Table V. The design of fault-tolerant systems with optimal productivity has been performed under various design parameters such as the probability of a good processor, the probability of a good memory unit, the area ratio of processor and memory, as suggested in [20].

The results of redundant system with optimal productivity are provided in Tables VI and VII under ranges of values for the probabilities of a good processor and memory unit, where the first, second, third, fourth, and fifth numbers in each entry represent the initial yield, the final yield, PIF, the number of processors and memory units of the optimal design, respectively. PIF is defined as the relative yield increase divided by the relative area increase.

The relative yield increase is defined as: ‘final yield minus the original yield’ divided by the ‘original yield’. The relative area increase is defined as the ‘new area minus the old area’ divided by the ‘old area’. The running times for the entire design process ranged between several minutes and about an hour on Sun Sparc 4, depending on the number of tasks. The experimental results

TABLE VII
REDUNDANT SYSTEMS WITH OPTIMAL PRODUCTIVITY WHEN THE PROBABILITIES OF A GOOD PROCESSOR AND MEMORY UNIT ARE HIGH

# Tasks	P_m	P_p									
		0.98					0.99				
		IY	FY	PIF	n_p	n_m	IY	FY	PIF	n_p	n_m
10	0.98	0.724	0.955	5.1	8	9	0.785	0.973	3.8	8	9
	0.99	0.785	0.975	3.9	8	9	0.851	0.988	2.6	8	9
	0.995	0.817	0.983	3.2	8	9	0.886	0.993	1.9	8	9
20	0.98	0.724	0.964	5.3	8	9	0.785	0.984	4.1	9	8
	0.99	0.785	0.984	4.1	8	9	0.851	0.990	2.6	8	9
	0.995	0.817	0.992	3.4	8	9	0.886	0.996	2.0	8	9
30	0.98	0.724	0.955	5.1	8	9	0.785	0.975	3.9	9	8
	0.99	0.785	0.975	3.9	8	9	0.851	0.988	2.6	8	9
	0.995	0.817	0.983	3.2	8	9	0.886	0.993	1.9	8	9
40	0.98	0.724	0.955	5.1	8	9	0.785	0.975	3.9	9	8
	0.99	0.785	0.975	3.9	8	9	0.851	0.988	2.6	8	9
	0.995	0.817	0.983	3.2	8	9	0.886	0.993	1.9	8	9
50	0.98	0.724	0.974	2.8	9	9	0.785	0.911	2.6	8	9
	0.99	0.785	0.911	2.6	9	8	0.851	0.993	1.3	9	9
	0.995	0.817	0.948	2.6	9	8	0.886	0.957	1.3	9	8

a) The area ratio between processor and memory is 1

# Tasks	P_m	P_p									
		0.98					0.99				
		IY	FY	PIF	n_p	n_m	IY	FY	PIF	n_p	n_m
10	0.98	0.724	0.955	7.7	8	9	0.785	0.973	5.7	8	9
	0.99	0.785	0.975	5.8	8	9	0.851	0.988	3.8	8	9
	0.995	0.817	0.983	4.9	8	9	0.886	0.993	2.9	8	9
20	0.98	0.724	0.964	5.3	8	9	0.785	0.984	4.1	9	8
	0.99	0.785	0.984	4.1	8	9	0.851	0.990	2.6	8	9
	0.995	0.817	0.992	3.4	8	9	0.886	0.996	2.0	8	9
30	0.98	0.724	0.955	7.7	8	9	0.785	0.973	5.7	8	9
	0.99	0.785	0.975	5.8	8	9	0.851	0.988	3.8	8	9
	0.995	0.817	0.983	4.9	8	9	0.886	0.993	2.9	8	9
40	0.98	0.724	0.955	7.7	8	9	0.785	0.973	5.7	8	9
	0.99	0.785	0.975	5.8	8	9	0.851	0.988	3.8	8	9
	0.995	0.817	0.983	4.9	8	9	0.886	0.993	2.9	8	9
50	0.98	0.724	0.840	3.8	8	9	0.785	0.911	3.8	8	9
	0.99	0.785	0.983	2.0	9	9	0.851	0.920	1.9	8	9
	0.995	0.817	0.948	1.9	9	8	0.886	0.996	1.0	9	9

b) The area ratio between processor and memory is 2

IY $\equiv$ initial yield; FY $\equiv$ final yield; PIF $\equiv$ improvement factor in productivity; n_p $\equiv$ # of processors in the final design; n_m $\equiv$ # memory units in the final design; Initially, 8 processors and 8 memory units are used.

indicate important improvement in productivity and yield by the method.

REFERENCES

- [1] T. D. Burd and R. W. Brodersen, "Processor design for portable systems," *J. VLSI Signal Processing*, vol. 13, no. 2-3, pp. 203-221, 1996.
- [2] G. De Micheli, *Synthesis and Optimization of Digital Circuits*: McGraw Hill, 1994.
- [3] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*: W.H. Freeman and Company, 1979.
- [4] M. J. Gonzales, "Deterministic processor scheduling," *ACM Computing Surveys*, vol. 9, no. 3, pp. 173-204, 1997.
- [5] L. Guerra, M. Potkonjak, and J. Rabaey, "High level synthesis for efficient built-in self repair," in *Int. Workshop on Defect and Fault Tolerance in VLSI Systems*, 1993, pp. 41-48.
- [6] B. Iyer, R. Karri, and I. Koren, "Phantom redundancy: a high-level synthesis approach for manufacturability," in *Int. Conf. Computer-Aided Design*, 1995, pp. 658-661.
- [7] R. Karri and A. Orailoglu, "Transformation-based high-level synthesis of fault-tolerant ASICs," in *Design Automation Conf.*, 1992, pp. 662-665.
- [8] —, "High-level synthesis of fault-secure microarchitectures," in *Design Automation Conf.*, 1993, pp. 429-433.
- [9] S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671-680, 1983.
- [10] E. L. Lawler and C. U. Martel, "Scheduling periodically occurring tasks on multiple processors," *Inf. Process. Lett.*, vol. 12, no. 1, pp. 9-12, 1982.
- [11] E. A. Lee and D. G. Messerschmitt, "Synchronous dataflow," *Proc. IEEE*, vol. 75, no. 9, pp. 1235-1245, 1987.
- [12] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard real time environment," *J. ACM*, vol. 20, no. 1, pp. 46-61, 1973.
- [13] M. C. McFarland, A. C. Parker, and R. Camposano, "The high-level synthesis of digital systems," *Proc. IEEE*, vol. 78, no. 2, pp. 301-317, 1990.
- [14] A. Orailoglu and R. Karri, "Automatic synthesis of self-recovering VLSI systems," *IEEE Trans. Comput.*, vol. 45, no. 2, pp. 131-142, 1996.
- [15] M. Potkonjak and J. Rabaey, "Algorithm selection: a quantitative computation-intensive optimization approach," in *Int. Conf. Computer-Aided Design*, 1994, pp. 90-95.
- [16] M. Potkonjak and W. Wolf, "Cost optimization in ASIC implementation of periodic hard-real time systems using behavioral synthesis techniques," in *Int. Conf. Computer-Aided Design*, 1995, pp. 446-451.
- [17] J. Rabaey et al., "Fast prototyping of data path intensive architecture," *IEEE Design and Test*, vol. 8, no. 2, pp. 40-51, 1991.
- [18] V. Raghavendra and C. Lursinsap, "Automated micro-roll-back self recovery synthesis," in *Design Automation Conf.*, 1991, pp. 385-390.
- [19] D. P. Siewiorek and R. S. Swartz, *Reliable Computer System: Design and Evaluation*, 2nd ed: Digital Press, 1992.
- [20] C. H. Stapper, J. A. Patrick, and R. J. Rosner, "Yield model for ASIC and processor chips," in *Int. Workshop on Defect and Fault Tolerance in VLSI Systems*, 1993, pp. 136-143.

Inki Hong is a Staff R&D Engineer in Physical Synthesis group at Synopsys. He received his Ph.D. in 2001 from the Computer Science Department at the University of California, Los Angeles, and his Master of Science in 1994 in Computer Science from Stanford University. His research interests include system-level synthesis, physical synthesis, intellectual property protection and verification. He has published more than 20 papers in leading journals and conferences on these subjects.

Miodrag Potkonjak is a Professor in the Computer Science Department at UCLA. He received his Ph.D. in 1991 in Electrical Engineering and Computer Science from the University of California, Berkeley. In 1991, he joined C&C Research Laboratories, NEC USA, Princeton, NJ. Since 1995, he has been with UCLA. He received the NSF CAREER award, OKAWA foundation award, UCLA TRW SEAS Excellence in Teaching Award, and several best-paper awards. His research interests include communication systems design, embedded systems, computational security, practical optimization techniques, and intellectual property protection.

Ramesh Karri received his Ph.D. in 1993 in Computer Science from the University of California at San Diego. During 1993–1998 he was an Assistant Professor in the Department of Electrical & Computer Engineering at the University of Massachusetts at Amherst. During 1997–1998 he was a member of the technical staff at Bell Labs, Lucent Technologies. Since 1998 August, he has been an Associate Professor of Electrical Engineering at Polytechnic University, Brooklyn. His research interests include energy-conscious security protocols, high-speed encryption architectures, hardware acceleration of networking protocols, computer aided design, and fault tolerant VLSI systems.

He received an NSF CAREER award and the Alexander Von Humboldt Fellowship. He served as guest editor of IEEE TRANS. CAD, IEEE TRANS. RELIABILITY and *IEEE Design and Test Magazine*. He is serving on several conference program committees and organized the 2002 IEEE Integrated Management of Power Aware Computing, Communications, and Networking (IMPACCT) workshop.