

Application-Specific MPSoC Reliability Optimization

Zhenyu (Peter) Gu, Changyun Zhu, Li Shang, and Robert P. Dick

Abstract—This paper presents modeling and estimation techniques permitting the temperature-aware optimization of application-specific multiprocessor system-on-chip (MPSoC) reliability. Technology scaling and increasing power densities make MPSoC lifetime reliability problems more severe. MPSoC reliability strongly depends on system-level MPSoC architecture, redundancy, and thermal profile during operation. We propose an efficient temperature-aware MPSoC reliability analysis and prediction technique that enables MPSoC reliability optimization via redundancy and temperature-aware design planning. Reliability, performance, and area are concurrently optimized. Simulation results indicate that the proposed approach has the potential to substantially improve MPSoC system mean time to failure with small area overhead.

I. INTRODUCTION

Aggressive scaling of CMOS process technology poses serious challenges to application-specific multiprocessor system-on-chip (MPSoC) lifetime reliability. Reduced feature size, increased power density, and increased temperature increase component failure rates. Increasing system integration scale using these vulnerable devices and interconnects results in reduced system reliability. The severities of many reliability problems, such as time-dependent dielectric breakdown in transistors and electromigration in interconnects, increase exponentially with temperature. Optimizing lifetime reliability requires careful planning during MPSoC design and synthesis.

A. Past Work and Contributions

Our work draws from research in the areas of integrated circuit (IC) reliability modeling and reliable system synthesis. Coskun *et al.* [1] and Srinivasan *et al.* [2] provide architectural reliability models and run-time optimization techniques for application-specific MPSoCs and general-purpose microprocessors, respectively. The COFTA hardware-software co-synthesis algorithm produces architectures that achieve the reliability of triple-modular systems at lower cost [3]. Xie *et al.* propose duplicating tasks on idle processors during embedded system synthesis in order to recover from transient faults [4]. Glaß *et al.* propose an evolutionary algorithm that binds tasks to multiple resources with the goal of improving mean time to failure (MTTF) [5]. They consider fault processes with exponential or Weibull distributions; their fault model supports permanent faults. Our system and fault model differs primarily by considering the influence of faults on subsequent fault rates due to the impact of run-time rebinding on temperature profile.

In this paper, we present a reliability model permitting estimation of system MTTF: the expected duration an MPSoC will continue to

Manuscript received September 16, 2006; revised April 29, 2007. This work was supported in part by the National Science Foundation under Award CNS-0347941, by the Semiconductor Research Corporation under Award 2007-HJ-1593, and by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery under Grant 388694-01.

Z. Gu is with Synplicity Inc., Sunnyvale, CA 94086 USA (e-mail: zygu@synplicity.com).

R. P. Dick is with the Electrical Engineering and Computer Science Department, Northwestern University, Evanston, IL 60208-3118 USA (e-mail: dickrp@northwestern.edu).

C. Zhu is with the Electrical and Computer Engineering Department, Queen's University, Kingston, ON K7L 3N6, Canada (e-mail: changyun.zhu@ece.queensu.ca).

L. Shang is with the Electrical and Computer Engineering Department, University of Colorado at Boulder, Boulder, CO 80305 USA (e-mail: li.shang@colorado.edu).

Digital Object Identifier 10.1109/TVLSI.2008.917574

meet its functional and performance requirements. This model takes various design-time and run-time factors into consideration, including multiple failure mechanisms, resource redundancy, and thermal profile. It considers the effect of component wear on fault rate and explicitly models changes to processor allocation, floorplan, task assignment, and schedule. It is efficient enough to permit repeated use during synthesis. We also describe a domain-specific optimization algorithm that improves MPSoC reliability with small area overhead via redundancy and temperature-aware design planning. These ideas allow the proposed MPSoC reliability optimization technique to improve MPSoC system MTTF by an average of 85% with less than 5% area cost and by an average of 436% with less than 25% area cost, compared to area-optimized solutions.

II. PROBLEM DEFINITION

MPSoC reliability optimization requires solutions to the following problems.

- 1) *Modeling*: The lifetime reliability of the MPSoC depends on resource redundancy and failure mechanisms, which in turn depend on numerous design-time and run-time MPSoC characteristics such as floorplan, chip power and thermal profiles, and performance requirements. These effects must be efficiently modeled during MPSoC lifetime reliability analysis.
- 2) *System-Level Optimization*: Numerous system-level design decisions affect MPSoC lifetime reliability. MPSoC lifetime reliability optimization requires the optimization of resource redundancy (via resource allocation) as well as system power and thermal characteristics (via task assignment and scheduling).
- 3) *Physical Design*: MPSoC floorplan directly affects power consumption and temperature profiles. Power distribution should be balanced in order to eliminate local thermal hotspots, thereby improving MPSoC lifetime reliability. Moreover, the use of redundancy at the system and processor core level impacts physical design decisions.

Explaining each component of this synthesis flow in detail is beyond the scope of this paper. We focus on temperature and redundancy dependent reliability modeling and optimization. We will first explain the dominant failure mechanisms and then describe a method of modeling their system-level effects. Our objective is to optimize the areas and MTTFs of a set of MPSoC architectures while honoring functionality and timing constraints.

A. Integrated Circuit Failure Mechanisms

In this section, we characterize IC failure mechanisms. The lifetime reliability of ICs is primarily affected by the following failure mechanisms: electromigration, thermal cycling, time-dependent dielectric breakdown, and stress migration [2].

Electromigration is the gradual displacement of the atoms in metal wires caused by electrical current. It leads to voids and hillocks that cause open and short circuit failures. The MTTF due to electromigration is given by the following equation [6]:

$$\text{MTTF}_{\text{EM}} = \frac{A_{\text{EM}}}{J^n} e^{\frac{E_{\text{aEM}}}{\kappa T}} \quad (1)$$

where A_{EM} is a constant determined by the physical characteristics of the metal interconnect, J is the current density, E_{aEM} is the activation energy of electromigration, n is an empirically-determined constant, κ is Boltzmann's constant, and T is the temperature.

Thermal cycling refers to IC fatigue failures caused by thermal mismatch deformation. In IC chip and package, adjacent material layers

such as copper/low- k dielectric have different coefficients of thermal expansion. As a result, run-time thermal variation causes fatigue deformation, leading to failures. The MTTF due to thermal cycling is given by the following equation [6]:

$$\text{MTTF}_{\text{TC}} = \frac{A_{\text{TC}}}{(T_{\text{average}} - T_{\text{ambient}})^q} \quad (2)$$

where A_{TC} is a constant coefficient, T_{average} is the chip average run-time temperature, T_{ambient} is the ambient temperature, and q is the Coffin–Manson exponent constant.

Time-dependent dielectric breakdown is the deterioration of the gate dielectric layer. This effect depends strongly on temperature and is becoming increasingly prominent with the reduction of gate-oxide dielectric thickness and non-ideal supply voltage reduction. The MTTF due to time-dependent dielectric breakdown is given by the following equation [2], [6]:

$$\text{MTTF}_{\text{TDD}} = A_{\text{TDD}} \left(\frac{1}{V} \right)^{(a-bT)} e^{\frac{A+B/T+CT}{\kappa T}} \quad (3)$$

where A_{TDD} is a constant, V is the supply voltage, and a , b , A , B , and C are fitting parameters.

Stress migration is the mass transportation of metal atoms in metal wires due to mechanical stress caused by thermal mismatch among metal and dielectric materials. The MTTF resulting from stress migration is given by the following equation [6]:

$$\text{MTTF}_{\text{SM}} = A_{\text{SM}} |T_0 - T|^{-n} e^{\frac{E_{a\text{SM}}}{\kappa T}} \quad (4)$$

where A_{SM} is a constant, T_0 is the metal deposition temperature during fabrication, T is the run-time temperature of the metal layer, n is an empirically-determined constant, and $E_{a\text{SM}}$ is the activation energy for stress migration.

Equations (1)–(4) indicate that the lifetime reliability of ICs is strongly influenced by temperature. Therefore, thermal analysis and optimization techniques play important roles in reliability optimization.

B. MPSoC Reliability Modeling and Optimization

We define *system MTTF* as the average amount of time an MPSoC will operate, possibly in the presence of component faults, before its performance drops below some designer-specified constraint or it is no longer able to execute the specified workload. Using system MTTF to characterize reliability has the advantage of taking into account performance; this is important for consumer electronics and most other MPSoC applications.

The system MTTF of an MPSoC is a function of the lifetime reliabilities of all its processing elements (PEs). In this paper, we propose a system-level lifetime reliability model for MPSoCs. Our first step is to derive an efficient modeling method that can accurately predict the lifetime reliability of each MPSoC PE.

1) *Reliability Modeling of On-Chip PEs*: The lifetime reliability of an on-chip PE is influenced by numerous design-time and run-time factors, such as architecture-level and circuit-level redundancy, accumulation of wear, and run-time power temperature. Accurate lifetime characterization of each PE is challenging.

We propose a PE reliability model that is capable of incorporating the effects of multiple fault mechanisms, component-level resource redundancy, and temperature. The dependence of lifetime failure processes on other parameters, such as current density, is not directly considered. Constant values of these parameters resulting in PE MTTFs of 30 years

at 50 °C and 1.8 V are used [2]. For the sake of explanation, our description of PE reliability modeling starts from the simplest case, i.e., a single failure mechanism, single point of failure (no resource redundancy), and constant temperature. These assumptions are later relaxed and the reliability model generalized.

Lognormal Distribution Reliability Model for Single PE, Single Point of Failure: Statistical modeling is commonly used in IC reliability characterization. Researchers have proposed using various statistical models, e.g., exponential, Weibull, and lognormal, to characterize IC lifetime failures. Compared to other commonly-considered statistical models, the lognormal distribution more accurately models the time-dependent degradation processes of ICs, e.g., diffusion, corrosion, migration, and crack propagation [2] caused by the failure mechanisms described in Section II-A. However, using the lognormal distribution complicates the derivation of analytical solutions. Numerical methods, such as Monte Carlo simulation or statistical fitting techniques, are required. These methods are computationally intensive.

Starting from the simplest assumption, for a failure mechanism i , the run-time fault probability density function (PDF) $f_i(t)$ and the corresponding fault cumulative distribution function (CDF) $F_i(t)$ have two parameters: σ_{PE}^i (a shape parameter) and μ_{PE}^i (a scale parameter). The MTTF of an on-chip PE due to a particular failure mechanism i , MTTF_{PE}^i , is then estimated

$$\text{MTTF}_{PE}^i = \int_0^{\infty} t f_i(t) dt = \int_0^1 t dF_i(t) = e^{\mu_{PE}^i + \sigma_{PE}^i{}^2/2}. \quad (5)$$

The overall lifetime reliability of each on-chip PE MTTF_{PE} is modeled by a joint lognormal distribution that depends on the major failure mechanisms described in Section II-A. We assume that the relationships among different failure mechanisms are serial, i.e., each individual failure mechanism can result in the failure of a non-redundant PE. Therefore, for each non-redundant PE, the CDF of its overall lifetime failure probability follows:

$$F_{PE}(t) = 1 - \prod_i (1 - F_i(t)) \quad (6)$$

where i is the index of different failure mechanisms.

Researchers have often used exponential distributions for statistical modeling due to their convenience. Given $F_i(t)$'s with exponential distributions, (6) would yield an easily-computed analytical solution. However, as a consequence of using the more accurate lognormal distribution for each $F_i(t)$, (6) does not allow straightforward estimation of PE MTTF, MTTF_{PE} . In this paper, we use statistical fitting to approximate MTTF_{PE} using a single lognormal distribution, governed by μ_{PE} and σ_{PE} . The parameters for this approximation follow:

$$\mu_{PE} = \frac{1}{2} \log \left(\frac{\left(\int_0^{\infty} t dF_{PE}(t) \right)^4}{\int_0^{\infty} t^2 dF_{PE}(t)} \right) \quad (7)$$

$$\sigma_{PE} = \sqrt{\log \left(\frac{\int_0^{\infty} t^2 dF_{PE}(t)}{\left(\int_0^{\infty} t dF_{PE}(t) \right)^2} \right)}. \quad (8)$$

The derivation of these parameters is omitted due to space constraints.

Reliability Models for Inactive Spare and Active Spare Redundant PEs: PEs may have component redundancy to improve reliability or performance. Such PEs can be designed to continue functioning even after some of their components, e.g., an arithmetic logic unit (ALU) or a cache bank, fail. We define inactive spares to be redundant resources that are not activated until a fault occurs in an active resource. The




Fig. 1. Temperature impact on MTTF.

impact of faults in inactive spares upon the lifetime reliabilities of PEs can be characterized as follows.

Assume a PE contains M types of resources. Each type of resource S_i , $i \in \{1, \dots, M\}$, is comprised of N_i identical elements. Assume the cumulative failure probability of resource element $E_{i,j}$, $i \in \{1, \dots, M\}$, $j \in \{1, \dots, N_i\}$ is $F_{i,j}(t)$. Then, the cumulative failure probability of resource S_i , $F_{S_i}(t) = \prod_j F_{i,j}(t)$. The MIN-MAX approximation [2] may be used to bound the MTTF of a PE with M types of resources as follows:

$$\text{MTTF}_{\text{PE}} = \min_{i=1}^M \left(\int_0^1 t dF_{S_i}(t) \right). \quad (9)$$

Active spares are redundant resources that are actively used even before any faults have occurred. Faults in active spares reduce the performance of the affected PE. Determining the reliability impact of faults that result in changes to observable PE behavior involves system-level design decisions, and will be described in detail in Section II-B2.

Temperature-Dependent Reliability Model for Potentially Redundant PEs: The lifetime reliability of a PE strongly depends on its temperature. After each MPSoC solution is derived, performance and power analysis are conducted. The estimated power profile, MPSoC floorplan, and cooling configuration are provided to a thermal analysis algorithm [7] to determine the thermal profile. Note that (9) is derived under an assumption of constant PE temperature. Next, we discuss temperature-dependent PE MTTF estimation.

The temperature profile of an MPSoC varies as the tasks assigned to it change. Task assignments change whenever migration is used to compensate for a partial or complete PE failure. The impact of temperature variation on MTTF calculation is illustrated in Fig. 1. In this example, T_1 and T_2 are temperatures. The PE is initially hot (T_1) and, at time t_1 , becomes cooler (T_2). Functions $f_1(t)$ and $f_2(t)$ are the fault PDFs given temperatures T_1 and T_2 , respectively. The overall fault distribution of the PE should satisfy the following equation, i.e., the overall cumulative fault distribution equals one

$$\int_0^{t_1} f_1(t) dt + \int_{t_2}^{\infty} f_2(t) dt = 1. \quad (10)$$

When we switch from the fault PDF associated with one temperature, e.g., T_1 , to that associated with another temperature, e.g., T_2 , it is necessary to adjust our start time to the value, in the new time scale, associated with the appropriate amount of wear that had been experienced in the previous time scale, i.e., we must start integrating from the effective age of the PE. For this example, the concept can be summarized as follows: $F_1(t_1) = F_2(t_2)$.

Given that $\{T_0, T_1, \dots, T_{N-1}\}$ denote the PE thermal profile, the overall fault distribution should satisfy the following equation:

$$\int_{t_{s0}=0}^{t_{e0}} f_0(t) dt + \int_{t_{s1}}^{t_{e1}} f_1(t) dt + \dots + \int_{t_{sN-1}}^{\infty} f_{N-1}(t) dt = 1 \quad (11)$$

where $f_i(t)$ denotes the fault PDF of the PE at temperature T_i , $t_{e_i}(t)$ denotes the transition time at which the temperature changes from T_{i-1} to T_i , and $t_{s_i}(t)$ denotes the equivalent age of the PE, starting from $t_{e_{i-1}}$, when the temperature switches to T_i . The value of t_{s_i} can be determined using (11), allowing the MTTF of a PE to be determined using the following equation:

$$\text{MTTF} = \sum_{i=0}^{N-1} \int_{t_{s_i}}^{t_{e_i}} t f_i(t) dt. \quad (12)$$

This has the effect of breaking time into regions ($\sum_{i=0}^{N-1}$) during which the temperature of the PE is uniform and, during each region, weighting each time instant by the probability of failure at that instant ($t \cdot f_i(t)$). Values for t_{s_i} and t_{e_i} are computed based on (11).

Reliability analysis may be conducted numerous times during reliability optimization. Therefore, modeling efficiency is critical. An MPSoC consists of numerous PEs. If the cumulative fault probability distributions $F_i(t)$ are lognormal, then solving (9) requires computationally intensive numerical analysis. To improve computational efficiency, we produce a PE reliability library before reliability optimization by precharacterizing the reliability distributions of PEs as functions of temperature and supply voltage. During MPSoC reliability optimization, when solving (12), the value of $F_i(t)$ is efficiently obtained using table lookups.

2) *Reliability Modeling and Optimization of MPSoCs:* In this section, we discuss MPSoC lifetime reliability estimation and optimization. Many MPSoCs have built-in resource redundancy. In the recent past, techniques to provide both component-level (intra-PE) and PE-level redundancy have been proposed to improve system reliability and performance [2]. For MPSoCs with resource redundancy, faults may or may not cause system failures.

Algorithm 1 System MTTF analysis of an MPSoC solution

- 1: Given an MPSoC solution $MPSoC_{Sol}$, $MTTF_{MPSoC} = 0$, time $t = 0$
 - 2: **while** system schedule is valid **do**
 - 3: $MPSoC_{Func}$ are the functioning PEs in $MPSoC_{Sol}$
 - 4: Fault interval $e_i = \min_{pe \in MPSoC_{Func}} (MTTF_{pe})$
 - 5: $MTTF_{MPSoC} + = e_i$
 - 6: Task migration, scheduling
 - 7: **if** system scheduling is valid **then**
 - 8: Power analysis, thermal analysis
 - 9: Determine the temperature of each PE
 - 10: **else**
 - 11: Return $MTTF_{MPSoC}$
 - 12: **end if**
 - 13: **end while**
-

Algorithm 1 estimates system MTTF based on statistical models of MPSoC run-time fault processes. Starting from time $t = 0$, it determines the minimal MTTF among all the PEs using (12) (line

4). Each fault may result in partial or complete PE failure. In either case, task migration is used to balance system workload and optimize performance. The task migration routine moves tasks from the faulty or partially faulty, and therefore lower performance, PE to other PEs (line 6). After task migration and rescheduling, if the MPSoC still meets its performance requirements, the algorithm moves on to the remaining fault implying the minimal PE MTTF. Task migration results in run-time changes in chip power consumption and temperature profiles, thereby changing the lifetime reliability of each PE. To accurately predict subsequent PE MTTFs, power and thermal analysis are conducted (lines 8 and 9). Steady-state analysis is used but is repeated whenever the assignment of tasks to PEs changes. This process continues until the MPSoC fails to meet its performance or functionality requirements. The system MTTF of the MPSoC solution is then reported (line 11).

Run-time workload migration is used to maintain system functionality and meet performance requirements in the presence of partial and complete PE failures. When an MPSoC fails to meet its performance requirements due to run-time faults, tasks originally assigned to the faulty PE migrate to other PEs using the following policy. Tasks on faulty PEs are first sorted in order of increasing time slack, the difference between the task's latest finish time and earliest finish time. They are then migrated from the PE, to other PEs, in this order until the system performance requirements are met and no tasks are assigned to any totally failed PE. When moving a task from one PE to another, the new PE is selected by Pareto-ranking all PEs in order of increasing utilization ratio (the proportion of time during which the PE is actively executing tasks) and increasing execution time for the task and PE under consideration. If the PE has only partially failed, only a subset of its tasks migrate.

Algorithm 2 Reliability-aware optimization algorithm

```

1: while  $MTTF_{MPSoC} < MTTF_{target}$  do
2:    $\forall_{pe \in MPSoC}$  compute  $MTTF_{pe}$ 
3:   Find vulnerable point:  $PE_{vul}$  is the PE with minimal MTTF
4:   Optimization moves (PE reinforcement, PE swapping, PE
      addition)
5:   Apply the best move based on (13)
6:   System-level synthesis
7:     Task assignment
8:     Scheduling
9:   Physical-level synthesis
10:  Floorplanning
11:  On-chip network synthesis
12:  Performance, power, thermal, reliability analysis
13:  if system MTTF improves and system schedule is valid then
14:    Continue
15:  else
16:    Revert this change
17:  end if
18: end while

```

Starting from area-optimized MPSoC designs, lifetime reliability is optimized using architectural changes that improve redundancy and thermal profile, while maintaining low area overhead. Algorithm 2 shows the actions taken to improve an MPSoC architecture that does not have a sufficient system MTTF, i.e., $MTTF_{MPSoC} < MTTF_{target}$. First, the MTTF of each individual PE is estimated (line 2). The PE with the minimal MTTF is identified as the MPSoC's most vulnerable point, PE_{vul} (line 3). One of the proposed reliability optimization design changes is then used: PE

reinforcement, PE swapping, or PE addition (line 4). *PE reinforcement* introduces component redundancy into the most vulnerable PE. *PE swapping* replaces the most vulnerable PE with a different, more reliable, PE. *PE addition* introduces a new PE into the MPSoC, enabling tasks to migrate from the vulnerable PE to other PEs. These moves consider multiple candidate PEs. Relative reliability gain, defined in (13), is used to determine the best candidate move. This metric takes both power density reduction, resource redundancy improvement, and area overhead associated with the move into consideration

$$G_{rel} = \frac{e^{-P_d} \times MTTF_{ref}}{A}. \quad (13)$$

Note that this value is used only to guide changes. The detailed effect of each tentative change is computed using thermal profile and reliability analysis. MPSoC power profile influences MPSoC temperature profile, which in turn influences reliability. The MTTFs associated with some major fault mechanisms are exponential functions of temperature. Therefore, in (13), an exponential term, e^{-P_d} , is used to characterize the impact of power density reduction on reliability improvement. P_d is the power density reduction resulting from applying a candidate architecture change. In (13), the impact of redundancy is characterized by the second term, $MTTF_{ref}$, the system MTTF improvement resulting from the candidate move. $MTTF_{ref}$ is calculated under the assumption that other design characteristics, e.g., temperature profile and supply voltage, remain the same. The relative reliability gain introduced by each candidate move is the product of these two terms divided by the area overhead. The change with the highest gain is applied (line 5). After each optimization move, system-level and physical-level synthesis algorithms are invoked to update the MPSoC solution. Cost analysis is then conducted to determine the improvement in system reliability, determine the impact on MPSoC area, and validate the system schedule. This optimization process continues until the target system MTTF is achieved.

III. SYNTHESIS RESULTS

This section describes the results of applying the proposed temperature-aware reliability modeling and optimization techniques to a number of MPSoC synthesis benchmarks. Our goal is to determine whether the proposed reliability modeling and optimization techniques have sufficient performance for use within MPSoC synthesis. We will also attempt to draw conclusions on the area costs of improving MPSoC system reliability.

The proposed MPSoC reliability modeling and optimization techniques were evaluated using a number of benchmarks based on the E3S benchmarks suite. E3S contains 17 PEs, e.g., the AMD ElanSC520 and Motorola MPC555. These PEs are characterized based on the measured execution times of 47 tasks commonly encountered in embedded applications, power numbers derived from datasheets, and additional information, e.g., PE areas and prices. The E3S task sets follow the organization of the EEMBC benchmarks [8], with one benchmark for each of the five application suites. The original office automation problem contains only five tasks. Our modified version contains four copies of the original task set. In addition, TGFF [9] was used to generate five random benchmarks, each with 30–50 tasks and task types from E3S.

In E3S, PEs do not have component redundancy. We introduce a redundant version of each PE in E3S by duplicating arithmetic units and register files. Instruction scheduling units and instruction decode units do not have redundancy [2]. Caches have redundancy,




Fig. 2. Solutions produced by reliability optimization algorithm.

i.e., a single fault will reduce performance but the PE will remain operational. We based the impact of redundancy on previous work [2]: area, performance, and power consumption increase by 24%, 25%, and 25%, respectively.

The PEs in E3S have fairly homogeneous energy-delay products. MPSoCs commonly contain heterogeneous PEs. Therefore, for each PE in E3S, we introduced one corresponding PE operating at a higher voltage and another operating at a lower voltage. Note that a maximum of three voltages need to be provided by off-chip regulators. The alpha power law was used to calculate the impact of voltage scaling on performance. A nominal supply voltage of 1.8 V and alpha of 1.3 were used, based on recent short-channel MOSFET characteristics. Therefore, to model high-performance PEs, the supply voltage was scaled to 2.5 V, the performance increased by 25%, and the power consumption increased to $2.4\times$. To model low-power PEs, the supply voltage was scaled to 1.28 V, performance was increased by 25%, and power consumption was decreased to $0.38\times$.

In order to determine the effectiveness and efficiency of the proposed reliability optimization techniques, it was necessary to start from an existing architecture. We used a parallel recombinative simulated annealing system synthesis algorithm, i.e., an evolutionary algorithm, to produce area-optimized solutions that adhere to functionality and timing constraints. The optimization infrastructure was described and validated in a number of publications, the most complete of which is a Ph.D. dissertation [10].

Fig. 2 illustrates the solutions produced by the proposed reliability optimization techniques for all ten benchmarks. In Fig. 2, for each benchmark, the initial area-optimized solution appears at the left-most point of the line associated with the benchmark. We continued to apply the optimization moves described in Section II-B2 until seven subsequent moves did not significantly improve system MTTF. In effect, area is being spent for improvements in thermal profile, PE-level redundancy and MPSoC-level redundancy, thereby identifying solutions near the area-reliability Pareto-optimal curve. Each run finished in less than 1500 s on a 2.2-GHz Athlon XP processor. Table I shows the average system MTTF improvement over initial area-optimized solutions under different area overhead constraints for all ten benchmarks. These results illustrate two key points about the reliable application-specific MPSoC synthesis problem.

As indicated by the super-linear dependence of area on MTTF, reliability comes at some cost in area but this cost is initially small, per year improvement in MTTF. Note that, in Fig. 2, area is plotted on a logarithmic scale. As shown in Table I, improving the average system

TABLE I
SYSTEM MTTF IMPROVEMENT UNDER AREA BOUND

Area bound (%)	MTTF gain (%)	Area bound (%)	MTTF gain (%)	Area bound (%)	MTTF gain (%)
0.0	40.0	15.0	180.0	30.0	457.0
5.0	85.0	20.0	240.0	35.0	468.0
10.0	180.0	25.0	436.0	40.0	470.0

The MTTF improvement for each area bound is computed by selecting the highest-MTTF solution for each benchmark that honors the area bound and computing the average of these MTTF improvements.

MTTF over all benchmarks by 40%, 85%, and 180% results in maximum area overheads of 0.0%, 5.0%, and 10.0%. The proposed techniques are sometimes able to improve MTTF without area overhead because they indirectly result in new floorplans that are more compact than the previous floorplan. However, this is rarely the case and can be viewed as noise. The initial solutions are optimized for area; they tend to have high power densities and temperatures. As a result, the temperature-dependent fault rate is high. Area-optimized solutions also have low resource redundancy, i.e., a single hardware fault will often cause system failure. In addition, vulnerable points, i.e., hot, non-redundant PEs, normally exist in these systems. Therefore, for area-optimized initial solutions, the system reliability can be improved at low area cost. The proposed reliability optimization algorithm introduces PEs with lower power densities and/or replaces non-redundant PEs with redundant ones, thereby optimizing thermal properties and allowing the system to continue operating despite some runtime hardware faults.

As system MTTF increases, the area penalty associated with further improving system reliability increases, i.e., the areas of these application-specific MPSoCs are superlinearly dependent on MTTF. As shown in Table I, the proposed reliability optimization algorithm achieves a significant 436% average system MTTF improvement with a maximum area overhead of 25%. Further improvements to system MTTF become prohibitively costly. This can be explained in the following way. PE failure cumulative distribution functions are non-decreasing. For some large duration, there is a low probability that any PE will operate without a fault. As a result, at very large MTTFs, adding PEs or reinforcing a subset of existing PEs with redundant components has little impact on MTTF.

IV. CONCLUSION

This paper has described techniques for estimating and optimizing MPSoC system MTTF in the presence of temperature-dependent permanent fault processes. The proposed model allows efficient calculation of the expected duration of functionally correct and adequate performance MPSoC operation. It considers the effects of multiple temperature-dependent fault processes, wear, MPSoC architecture, and resource redundancy. Domain-specific reliability optimization algorithms that exploit redundancy and thermal profile optimization were proposed. Experimental results indicate that the resulting system is capable of substantially improving MPSoC system MTTF at a small cost in area, compared to area-optimized solutions.

REFERENCES

- [1] A. K. Coskun, T. S. Rosing, K. Mihic, G. De Micheli, and Y. Leblebici, "Analysis and optimization of MPSoC reliability," *J. Low Power Electron.*, pp. 56–69, Apr. 2006.
- [2] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "Exploiting structural duplication for lifetime reliability enhancement," in *Proc. Int. Symp. Comput. Arch.*, Jun. 2005, pp. 520–531.

- [3] B. Dave, G. Lakshminarayana, and N. K. Jha, "COSYN: Hardware-software co-synthesis of embedded systems," in *Proc. Des. Autom. Conf.*, Jun. 1997, pp. 703–708.
- [4] Y. Xie, L. Li, M. Kandemir, N. Vijaykrishnan, and M. J. Irwin, "Reliability-aware co-synthesis for embedded systems," in *Proc. Int. Conf. Appl.-Specific Syst., Arch., Process.*, Sep. 2004, pp. 41–50.
- [5] M. Glaß, M. Lukasiewicz, T. Streichert, C. Haubelt, and J. Teich, "Reliability-aware system synthesis," in *Proc. Des., Autom. Test Eur. Conf.*, Apr. 2007, pp. 1–6.
- [6] Joint Electron Device Engineering Council, Arlington, VA, "Failure mechanisms and models for semiconductor devices," JEDEC Pub. JEP 122-B, Aug. 2003.
- [7] Y. Yang, Z. Gu, C. Zhu, R. P. Dick, and L. Shang, "ISAC: Integrated space and time adaptive chip-package thermal analysis," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 26, no. 1, pp. 86–99, Jan. 2007.
- [8] EEMBC, Sunnyvale, CA, "Embedded microprocessor benchmark consortium," 1997 [Online]. Available: <http://www.eembc.org>
- [9] R. P. Dick, D. L. Rhodes, and W. Wolf, "TGFF: Task graphs for free," in *Proc. Int. Workshop Hardw./Softw. Co-Des.*, Mar. 1998, pp. 97–101.
- [10] R. P. Dick, "Multiobjective synthesis of low-power real-time distributed embedded systems," Ph.D. dissertation, Dept. Elect. Eng., Princeton Univ., Princeton, NJ, Jul. 2002.