

Economic Analysis of Testing Homogeneous Manycore Chips

Lin Huang, *Student Member, IEEE*, and Qiang Xu, *Member, IEEE*

Abstract—The employment of a large number of structurally identical cores on a single silicon die is generally regarded as a promising solution for tera-scale computation, known as manycore chips. To ensure the product quality of such complex integrated circuits before shipping them to final users, extensive manufacturing tests are necessary and the associated test cost can account for a large share of the total production cost. By introducing spare cores on-chip, the burn-in test time can be shortened and the defect coverage requirements for core tests can be also relaxed, without sacrificing quality of the shipped products. If the above test cost reduction exceeds the manufacturing cost of the extra cores, the total production cost of manycore chips can be reduced. In this paper, we develop novel analytical models to study the above tradeoff and we verify the effectiveness of the proposed test economics model for hypothetical manycore chips with various configurations.

Index Terms—Analytical model, manycore chip, product quality test economics.

I. INTRODUCTION

ADVANCEMENTS in semiconductor technology enable integration of a large number of cores on a single silicon die. This technique has been employed by many state-of-the-art computing systems [1]–[3], known as multicore or manycore chips. They have the benefits of power-efficiency and short time-to-market and, therefore, have become increasingly popular in the industry [4], [5]. To improve the manufacturing yield of such complex circuits, typically a few *yield-driven redundant cores* are placed on-chip and the system can be reconfigured to bypass those defective cores [6]. For instance, the 192-core Cisco Metro network processor [7] contains four spare cores while the 128-core Nvidia Geforce 8800 GPU [2] can be degraded to a 96-core version.

Meanwhile, customers have high expectations for the quality and reliability of semiconductor products, and typically

Manuscript received September 25, 2009; revised December 30, 2009. Date of current version July 21, 2010. This work was supported in part by the General Research Fund from the Hong Kong SAR Research Grants Council (RGC), under Grants CUHK417807 and CUHK418708, in part by the National Science Foundation of China (NSFC), under Grant 60876029, and in part by the NSFC/RGC Joint Research Scheme, under Grant N_CUHK417/08. A preliminary version of this paper was published in the *Proceedings of the International Test Conference*, 2009, Paper 12.3. This paper was recommended by Associate Editor N. K. Jha.

The authors are with the CURE Laboratory, Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong (e-mail: lhuang@cse.cuhk.edu.hk; qxu@cse.cuhk.edu.hk). Q. Xu is also affiliated with Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAD.2010.2049052

only a few hundred defective parts per million (DPPM) are allowed with a lifespan of several years. Various types of manufacturing test are performed at different stages of the integrated circuit (IC) manufacturing process to achieve this daunting objective. On one hand, sophisticated automatic test pattern generation techniques are used to achieve adequately high-defect coverage. As technology advances, test patterns that target delay faults and many other kinds of subtle errors (e.g., signal integrity faults) are also essential to guarantee test quality, in addition to the traditional stuck-at vectors. The associated large number of test patterns not only require long testing time on the automatic test equipment (ATE), but also indirectly result in more false rejects and thus lower the manufacturing yield of the ICs. On the other hand, accelerated testing methods such as burn-in test are used to screen out those chips with early-life failures to enhance product reliability. For ICs fabricated with latest technology, it is increasingly difficult to setup and control appropriate stress conditions for the circuits during the burn-in process [8], which makes it the bottleneck of the manufacturing process [9]. Due to the above, test cost can account for a great share of the total production cost of complex ICs. In particular, the cost of burn-in test may range from 5% to 40% of the total cost of the product, as pointed out in [10].

Homogeneous manycore chips are inherently defect-tolerant. This property provides us an opportunity for test cost reduction without sacrificing products' quality and reliability. That is, we introduce one or more dedicated spare cores (namely *test cost-driven redundant cores*) [11], [12], in addition to those yield-driven spares. By doing so, the defect coverage requirement for the core tests can be lowered and burn-in test can be also reduced or eliminated. Consider a manycore chip that functions with 16 cores, to guarantee that all 16 cores work well provided that they have passed manufacturing tests, we need core tests with very high-defect coverage to identify chips containing *killer defects* and sufficient burn-in to weed out chips with *latent defects*. If, by contrast, we add two test cost-driven redundant cores on-chip, since off-line manufacturing test is responsible for 16 out of 18 cores (instead of all the 18 cores) to work after the so-called *infant mortality* period, the defect coverage requirement can be relaxed and the burn-in test can be partially or fully eliminated (note, the functioning cores can be identified in-field with online testing techniques, if necessary). If the associated test cost reduction exceeds the manufacturing cost increment for the spare cores, we are able to cut down the total production

cost for the homogeneous manycore chips without sacrificing the required quality and reliability of the shipped products.

In this paper, we propose comprehensive analytical model to study the test economics for homogeneous manycore chips. Because of the many complicated factors involved in our model, the problems are formulated and solved progressively. That is, for the sake of simplicity, we first consider the case of introducing spare cores for partial/no burn-in test and fix the manufacturing defect coverage of core tests, assuming no false rejects. Sequentially, we relax the defect coverage constraint and take the impact of false rejects into consideration to study the complex relationship among test coverage, test escapes and false rejects, partial/no burn-in test, yield-driven redundancy, test cost-driven redundancy, and product quality. Next, product binning based on the number of functioning cores are considered. Experimental results on hypothetical manycore chips with various configurations in terms of defect density and test cost distributions show the effectiveness of the proposed analytical model.

The remainder of this paper is organized as follows. Section II presents preliminaries and motivations of this paper. Section III formulates the problems studied in this paper. The proposed analytical models to solve these problems are then detailed in Sections IV–VI. Next, experimental results are presented in Section VII to show the effectiveness of the proposed analytical model. Finally, Section VIII concludes this paper.

II. PRELIMINARIES AND MOTIVATION

A. Preliminaries

Integrated circuit fabrication is an extremely complex process. It is inevitable that some manufactured chips are defective. Prior work has proposed several methods to model the spatial distribution of defects on the wafer [13], and it was shown that negative-binomial distribution fits quite well with the actual defect distribution [14], as IC defects typically feature “clustering” effects.

In terms of defect type, *killer defects* reveal themselves as long as a proper testing strategy is conducted to activate them. In contrast, *latent defects* occur on weak ICs after a relatively short period of usage time, thus causing early-life reliability failures. As a result, IC products usually undergo a period with high-failure rate in their early lifetimes, commonly referred to as *infant mortality* (as shown in Fig. 1). In particular, the failure rate in this stage decreases with usage time, and the corresponding reliability function follows two-parameter Weibull distribution [15], which is related to the core structural properties and usage-related factors (e.g., operational voltage and frequencies) [16], [17]. By stressing the circuit at elevated temperature and voltage during burn-in test, the number of latent defect-induced failures increases and these weak chips can be identified. Experiments conducted at Intel over a wide range of yield values showed that the two kinds of defects have a linear relationship, and typically, for every 100 killer defects present, one expects, on average, 1–2 latent defects [18].

B. Motivation

A notable feature in IC testing is that most test patterns are applied to achieve the last several percentages of defect coverage. For example, according to the model presented in [19], to

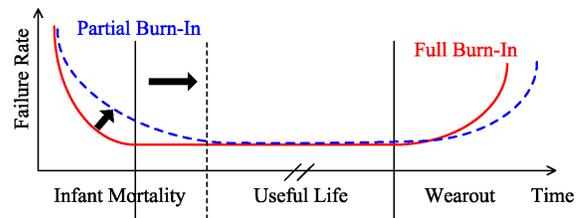


Fig. 1. Bathtub curve.

improve defect coverage from 99% to 99.9%, the number of test patterns increases for about 50%. In addition, testing with extremely high-defect coverage leads to more false rejects and thus lower the ICs manufacturing yield. Consequently, if we are able to relax this coverage requirement, the manufacturing test cost can be dramatically reduced. The semiconductor industry, however, tries to increase defect coverage for their products as high as possible. The reason is simple: if we do not remove these defective chips during the manufacturing test phase, we need to repair them at the board level or even system level and the cost would be much higher.

The above argument holds true for circuits with irregular structures. However, for the emerging homogeneous manycore chips that contain a large number of structurally identical cores, since they are inherently defect-tolerant and it is quite easy to conduct online test and reconfiguration, we are given an opportunity to apply the above test cost reduction strategy. In addition, considering the high cost of burn-in test due to their lengthy testing time, introducing spare cores on-chip can also alleviate the burden for burn-in test to reveal all latent defects, which facilitates us to conduct partial burn-in or even eliminate burn-in test completely. Obviously, the manufacturing cost per fabricated chip increases with more redundancy, however, if the associated test cost reduction exceeds the manufacturing cost increment for the spare cores, we are able to cut down the total production cost for the homogeneous manycore chips. This has motivated the test economics model studied in this paper.

In [11], we proposed the concept of test cost-driven redundancy for homogeneous manycore chips and used a case study to show its potential advantage, but the detailed analysis is missing due to space limitation of the paper. Shamshiri *et al.* [12] presented a cost analysis framework for manycore chips with spares and advocated to introduce redundant cores to eliminate burn-in test. The authors considered the case to ship ICs with high DPPMs to customers and a large service cost is modeled for replacement. The test economics model presented in that work, however, has several limitations.

- 1) It neglected the correlation among parameters that are strongly related. For example, the manufacturing yield for cores are determined by the defect density and the “clustering” defect distribution parameter, but they are set as arbitrary values in [12].
- 2) The impact of defect coverage on testing cost is not considered in that work.
- 3) This paper only analyzed the case for either full burn-in or no burn-in, without considering partial burn-in.
- 4) The difference between yield-driven redundancy and test cost-driven redundancy is not considered in [12].

Different from [12], we consider to introduce test cost-driven redundancy on-chip without sacrificing the required quality and reliability for shipped products and it is hence not necessary to consider the service cost for replacement of defective chips. More importantly, our proposed test economics model captures the complex relationship among test coverage, test escapes and false rejects, partial burn-in test, yield-driven redundancy, test cost-driven redundancy, and product quality. A preliminary version of this paper was published in [20]. We extend it by taking product binning into consideration in the proposed economical analytical framework, i.e., those products that cannot meet original design specification can be sold as degraded ones at reduced prices.

C. Online Testing

The fault-free cores on a manycore chip can be identified by online testing techniques. There is a rich literature in this field, covering various aspects. For example, software-based self-testing techniques for processor cores have been extensively studied [21], [22]. Design-specific solutions have also been explored, see [23] for UltraSPARC T1 microprocessor and [24] for ARM CORTEX-A8. In addition, since network-on-chip becomes a promising solution to tackle the on-chip communication problem, online testing in this context is also a research hot topic. As an example, [25] presents a protocol for testing the embedded cores without interrupting the functionality of other cores and on-chip interconnects. Moreover, it is worth to note that online testing for homogeneous manycore processors can be conducted easily. That is, by running two copies of the same test programs (or real applications) on two processor cores and comparing their results, we are able to tell whether a particular core contains faults.

More importantly, we believe online test and reconfiguration solutions will be included in future manycore chips despite of whether test cost-driven redundancies are available or not, due to the ever-increasing reliability threats with technology scaling [26]. That is, it is essential to include online test circuitries to detect not only defects introduced during manufacturing process but also soft errors and wearout-related effects occurred over the ICs' useful lifetime. Due to the above, the cost for online testing used to identify defective cores in-field is not considered in our test economics model.

III. PROBLEM FORMULATION

We model the test economics for homogeneous manycore chips progressively in this paper.

Firstly, with given defect coverage for core tests, we consider to introduce t redundant cores into a homogeneous manycore chip that functions if no less than m cores are defect-free for partial/no burn-in. In this context, in total we fabricate $u = m + t$ cores on a chip. As some latent defects are not detectable because of insufficient burn-in, only chips with all u cores pass test are sold to ensure product quality and reliability. Eventually we only need to guarantee m cores are defect-free at the end of infant mortality period. Only a very small amount of shipped products are allowed to be defective

or containing unrevealed latent defects and the percentage is set to τ .

This problem involves several manufacturing steps, each has a series of parameters. To clarify, we use negative-binomial distribution to capture the "clustering" effects of IC defects produced during fabrication, which has two parameters: λ_K the average number of killer defects per core and α the clustering parameter. The quantity of latent defects is proportional to that of killer defects and we denote by γ the ratio. That is, the average number of latent defects per core $\lambda_L = \gamma \cdot \lambda_K$. Then, burn-in test is applied to screen out the chips with early-life failures. To speedup the early-life stage (also called infant mortality), a certain time duration of burn-in test can be applied. During this process, latent defects gradually reveal themselves and the reliability due to these defects follows Weibull distribution with shape parameter β . Under burn-in condition, similar to prior work, we assume that all latent defects become detectable after time T_{IM} . Also, the role of manufacturing test is of importance such that a sufficient high-defect coverage B_r should be guaranteed to ensure the product quality, especially when partial/no burn-in is applied.

As to the parameters used in test economic modeling, we use unified ratios among various cost factors to normalize these factors toward a reference case. That is, the ratio between ATE cost and manufacturing cost per fabricated core is ρ , and the ratio between fully burn-in process cost and manufacturing cost per fabricated core is ξ . These parameters are regarded as the inputs into our model. With these notations, the problem can be formulated as follows.

Problem 1 (Partial Burn-In): Given the following.

- 1) *Design and material-related parameters.*
 - a) The required number of cores m for the homogeneous manycore chip to function.
 - b) The wafer dimension d , and the silicon area of a core A_r .
- 2) *Defect distribution parameters.*
 - a) The average number of killer defects per core λ_K and the clustering parameter α , assuming defects to follow the negative-binomial distribution.
 - b) The ratio between the average number of latent defects per core and that of killer defects per core γ .
- 3) *Burn-in related parameters.*
 - a) Infant mortality time with full burn-in that reveals all latent defects T_{IM} .
 - b) Shape parameter of reliability function in infant mortality duration β .
- 4) *Test-related parameters.*
 - a) The product quality requirement for the maximum test escape percentage τ (e.g., $\tau=0.0005$ for 500 DPPM).
 - b) Sufficiently high-defect coverage of manufacturing test B_r to ensure product quality.
- 5) *Cost-related parameters.*
 - a) The ratio between ATE cost for applying manufacturing test patterns per core and its manufacturing cost ρ .

- b) The ratio between the burn-in cost per core and its manufacturing cost ξ .

Determine the number of burn-in driven spares t to achieve the minimum production cost per sold chip under product quality constraint and the associated burn-in time T .

Next, we consider to introduce n redundant cores to not only enable partial/no burn-in test but also relax the defect coverage for core tests in our test economics model. The imperfect manufacturing test process leads to both *test escapes* (i.e., bad chips pass the test) and *false rejects* (i.e., good chips fail the test, also known as *test overkill*), which is related to the effectiveness of the test decision criterion μ [27]. Generally speaking, the more patterns applied to the circuits (for higher defect coverage), the more false rejects occur, and the steepness is described by parameter ν . Therefore, by relaxing defect coverage, we can also achieve cost savings with less test overkills and this effect is considered in our model. Because some defects do not reveal themselves due to insufficient burn-in and some defects are not detected because of low-test defect coverage, n *test cost-driven spares* are used to ensure the product quality and reliability for no less than m defect-free cores on-chip at the end of infant mortality period. In addition, we consider to have s *yield-driven spares* placed on-chip for yield enhancement. Consequently, we have totally $w = m + n + s$ homogeneous cores on-chip and only those chips containing no less than $m + n$ pass-test cores are shipped to the market.

Problem 2 (Partial Burn-In and Relaxed Defect Coverage): Given all parameters as specified in Problem 1, except that the test-related parameter set becomes as follows.

- 1) *Test-related parameters.*
 - a) The effectiveness of the test decision criterion μ .
 - b) The parameter ν that sets the steepness of the fallout curve that shows the probability for defects to be detected when test patterns are gradually applied [27].
 - c) The product quality requirement for the maximum test escape percentage τ (e.g., $\tau = 0.0005$ for 500DPPM).

Determine the number of test cost-driven spares n , the number of yield-driven redundant cores s , the defect coverage for core test B_r , and the burn-in time T such that the production cost per sold chip is minimized under the product quality and reliability constraints.

In practice, those products that cannot meet the design specification might not be discarded. Instead, they can be sold as products with degraded performance at a lower price, known as *product binning*. For example, the 128-core Nvidia GeForce 8800 GPU [2] can be degraded to a 96-core version. We, therefore, consider to separate the chips into b bins and sell the products in different bins at different prices. In particular, the chips with no less than $m + n$ pass-test cores can be viewed as in the first bin and sold at the highest price because we expect m defect-free cores at the end of infant mortality while have lower expectation for the remaining chips. By doing so, the production cost can be further reduced.

It is worth noting that the chips are dropped into different bins according to the number of pass-test cores but eventually

their qualities are evaluated with the number of defect-free cores at the end of infant mortality. That is, those chips whose quantities of pass-test cores are no less than f_ℓ but less than $f_{\ell-1}$ are shipped to the market as the ℓ th-bin products, and the probability of no less than g_ℓ defect-free cores at the end of infant mortality should be no less than $1 - \tau$. Clearly, we have $f_1 = m + n$, $f_0 = m + n + s$, and $g_1 = m$.

Problem 3 (Product Binning): Given all parameters as specified in Problem 2.

- 1) *Binning-related parameters.*
 - a) The number of bins b .
 - b) For each bin ℓ , the required number of cores g_ℓ for the homogeneous manycore chip to function.

Determine the binning criteria f_ℓ for all bins to minimize the production cost per sold chip under the product quality and reliability constraints.

IV. TEST ECONOMICS WITH PARTIAL/NO BURN-IN

In this section, we present our analytical model that captures the impact of introducing burn-in driven redundancy for partial/no burn-in on the total production cost of homogeneous manycore systems. It is worth noting that the defect coverage in this model is assumed to be a sufficient high value, and hence the ATE cost for each fabricated core is fixed.

A. Impact of Partial Burn-In

The reliability function of latent defects follows two-parameter Weibull distribution [15], which has the form

$$R(T) = \exp\left(-\left(\frac{T}{\theta}\right)^\beta\right). \quad (1)$$

The parameter β and θ determines the shape and scale of Weibull function, respectively. Assuming that all latent defects reveal themselves after full burn-in time T_{IM} [15], we are enabled to eliminate the scale parameter from (1). To be specific, by the assumption, the reliability induced by latent defects is given by

$$R(T_{IM}) = \exp(-\lambda_L). \quad (2)$$

By this normalization, we can obtain the scale parameter θ with given shape parameter β as

$$\theta = \frac{T_{IM}}{(\lambda_L)^{\frac{1}{\beta}}}. \quad (3)$$

In addition, as killer defects (whose quantity per core is λ_K) and latent defects are linearly related with ratio γ , we have

$$\lambda_L = \gamma \cdot \lambda_K. \quad (4)$$

Substituting (3) and (4) into (1) yields the reliability function with partial burn-in time T

$$R(T) = \exp\left(-\lambda_L \left(\frac{T}{T_{IM}}\right)^\beta\right) = \exp\left(-\gamma \cdot \lambda_K \left(\frac{T}{T_{IM}}\right)^\beta\right). \quad (5)$$

B. Product Quality and Chip Test Yield

To meet the product quality requirement, the probability that a sold chip actually functions (i.e., contains no less than m good cores) after all early infant mortality failures have been revealed should be higher than a threshold $(1 - \tau)$. This conditional probability (given a chip is sold) is referred as product quality Q hereafter. Recall that a chip will be sold when all $(u = m + t)$ cores pass test. Let us use X_1 to represent the event that no less than m cores on a chip is defect-free at the end of infant mortality, and use X_2 to denote the event that all $(m + t)$ cores on a chip pass manufacturing test after burn-in time duration T , the product quality can be expressed as conditional probability, that is

$$Q = \Pr\{X_1|X_2\} = \frac{\Pr\{X_1 X_2\}}{\Pr\{X_2\}}. \quad (6)$$

Let us start with the calculation of $\Pr\{X_2\}$. We define the events that i -out-of- u cores are defect-free after burn-in time T as $C_{i,u,T}$ ($i = 0, \dots, u$), and thus the events that a chip containing u cores passes manufacturing test given it contains i defect-free cores as $[X_2|C_{i,u,T}]$. By the theorem of total probability, we obtain

$$\Pr\{X_2\} = \sum_{i=0}^u \Pr\{X_2|C_{i,u,T}\} \Pr\{C_{i,u,T}\}. \quad (7)$$

The event $C_{i,u,T}$ can be further divided into two sub-events: $M_{j,u}$ represents that j -out-of- u cores are initially defect-free, and $N_{i,j,T}$ indicates that i cores among them remain defect-free at the end of burn-in time T . Apparently, i must be no greater than j , as shown in Fig. 2. Assuming that the occurrence of killer defects and that of latent defects are independent, we have

$$\Pr\{C_{i,u,T}\} = \sum_{j=0}^u \Pr\{M_{j,u}\} \Pr\{N_{i,j,T}\}. \quad (8)$$

Thus, with $i \leq j$, (7) can be rewritten as

$$\Pr\{X_2\} = \sum_{j=0}^u \sum_{i=0}^j \Pr\{X_2|C_{i,u,T}\} \Pr\{M_{j,u}\} \Pr\{N_{i,j,T}\}. \quad (9)$$

This equation detaches three influential factors from the event X_2 : $M_{j,u}$ which is determined by manufacturing defect distribution only, $N_{i,j,T}$ that is up to latent defect-induced reliability, and $[X_2|C_{i,u,T}]$ which depends on manufacturing test quality. These aspects are discussed in the following separately. According to the negative-binomial defect distribution, $\Pr\{M_{j,u}\}$ can be derived as [28]

$$\Pr\{M_{j,u}\} = \binom{u}{j} \sum_{\ell=0}^{u-j} (-1)^\ell \binom{u-j}{\ell} \left(1 + \frac{(j+\ell)\lambda_K}{\alpha}\right)^{-\alpha}. \quad (10)$$

Since infant mortality can be characterized by Weibull distribution with shape parameter $0 < \beta < 1$, we express $\Pr\{N_{i,j,T}\}$ in terms of reliability function $R(T)$ that is defined by (5) as

$$\Pr\{N_{i,j,T}\} = \binom{j}{i} R^i(T) (1 - R(T))^{j-i}. \quad (11)$$

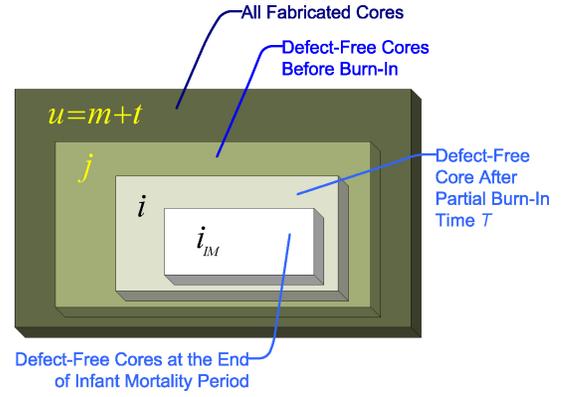


Fig. 2. Defect-free core sets at various time points.

For the sake of simplicity, we assume no false rejects (i.e., all good cores pass manufacturing test) for the time being in the calculation of $\Pr\{X_2|C_{i,u,T}\}$. This assumption will be lifted later. Thus, all i defect-free cores after (insufficient) burn-in pass test. Note that, it is possible that some cores in this set contain unrevealed latent defects. Due to imperfect manufacturing test, q cores with revealed defects (including killer defects and revealed latent defects) out of $(u - i)$ also pass test, while the remaining $(u - i - q)$ are rejected. Since a chip is shipped to customers if all its $(m + t)$ cores pass test, we have $q = u - i$. Therefore, denoting by B_r the defect coverage of manufacturing test, we have

$$\Pr\{X_2|C_{i,u,T}\} = (1 - B_r)^{u-i}. \quad (12)$$

We then move to the computation of $\Pr\{X_1 X_2\}$, which is more complicated. Denoting by i_{IM} the number of defect-free cores at the end of infant mortality, we always have $i_{IM} \leq i$ because a core with no revealed defects after partial burn-in may still contain latent defects. To compute $\Pr\{X_1 X_2\}$, we define the event that i -out-of- u cores do not contain revealed defects after burn-in time T and exactly i_{IM} cores are defect-free at the end of infant mortality as $D_{\{i,T\},\{i_{IM},T_{IM}\},u}$. The probability of both X_1 and X_2 occurrence is, therefore

$$\Pr\{X_1 X_2\} = \sum_{i=m}^u \sum_{i_{IM}=m}^i \Pr\{X_1 X_2|D_{\{i,T\},\{i_{IM},T_{IM}\},u}\} \cdot \Pr\{D_{\{i,T\},\{i_{IM},T_{IM}\},u}\} \quad (13)$$

where $i = m, \dots, u$ because event X_1 should be hold, that is, no less than m cores should be defect-free.

The event $D_{\{i,T\},\{i_{IM},T_{IM}\},u}$ can be divided into two sub-events: $M_{j,u}$, which has been introduced before [see (10)]; and $P_{\{i,T\},\{i_{IM},T_{IM}\},j}$, meaning that i cores contain no revealed defects after burn-in time T (i.e., event $N_{i,j,T}$) and then i_{IM} cores are eventually defect-free at the end of infant mortality (i.e., event $N_{i_{IM},j,T_{IM}}$). Since j must be no less than i and $i \geq m$, we get $j \geq m$. That is

$$\begin{aligned} \Pr\{D_{\{i,T\},\{i_{IM},T_{IM}\},u}\} &= \sum_{j=m}^u \Pr\{M_{j,u}\} \Pr\{P_{\{i,T\},\{i_{IM},T_{IM}\},j}\} \\ &= \sum_{j=m}^u \Pr\{M_{j,u}\} \Pr\{N_{i,j,T} \cap N_{i_{IM},j,T_{IM}}\}. \end{aligned} \quad (14)$$

Note that, the event that i_{IM} cores are defect-free at T_{IM} is not independent of the event that i cores contain no defects at T , where $T \leq T_{IM}$. By the multiplication rule, we are able to obtain

$$\Pr\{N_{i,j,T} \cap N_{i_{IM},j,T_{IM}}\} = \Pr\{N_{i,j,T}\} \Pr\{N_{i_{IM},j,T_{IM}} | N_{i,j,T}\} \quad (15)$$

where $\Pr\{N_{i,j,T}\}$ has been expressed by (11). To compute the conditional probability $\Pr\{N_{i_{IM},j,T_{IM}} | N_{i,j,T}\}$, we start from the conditional reliability of a single core given it contains no revealed defects after partial burn-in. It is given by

$$R_c(T_{IM}|T) = \frac{R(T_{IM}IM)}{R(T)}. \quad (16)$$

Substituting (5) into this expression yields

$$\begin{aligned} R_c(T_{IM}|T) &= \frac{\exp(-\gamma \cdot \lambda_K)}{\exp\left(-\gamma \cdot \lambda_K \left(\frac{T}{T_{IM}}\right)^\beta\right)} \\ &= \exp\left(-\gamma \cdot \lambda_K \cdot \frac{T_{IM}^\beta - T^\beta}{T_{IM}^\beta}\right). \end{aligned} \quad (17)$$

With this notation, the conditional probability can be computed by

$$\Pr\{N_{i_{IM},j,T_{IM}} | N_{i,j,T}\} = \binom{i}{i_{IM}} R_c^{i_{IM}}(T_{IM}|T) (1 - R_c(T_{IM}|T))^{i - i_{IM}}. \quad (18)$$

We then consider the conditional event $\{X_1 X_2 | D_{[i,T],\{i_{IM},T_{IM}\},u}\}$ in (13). Because of $i_{IM} \geq m$, the condition $D_{[i,T],\{i_{IM},T_{IM}\},u}$ guarantees that no less than m cores is defect-free at the end of infant mortality (i.e., X_1). Also, the event $D_{[i,T],\{i_{IM},T_{IM}\},u}$ implies that i cores are without revealed defects after partial burn-in and the remaining $(u - i)$ have defects that are severe enough to result in faults. Based on these two considerations, the conditional event comes down to the event that all $(u - i)$ defective cores pass test after burn-in time T . Its probability can, therefore, be computed by the similar argument of (12) and given by

$$\Pr\{X_1 X_2 | D_{[i,T],\{i_{IM},T_{IM}\},u}\} = (1 - B_r)^{u-i}. \quad (19)$$

By using these equations, we conclude that

$$\begin{aligned} \Pr\{X_1 X_2\} &= \sum_{j=m}^u \sum_{i=m}^j \sum_{i_{IM}=m}^i \Pr\{X_1 X_2 | D_{[i,T],\{i_{IM},T_{IM}\},u}\} \\ &\quad \cdot \Pr\{M_{j,u}\} \Pr\{N_{i,j,T}\} \Pr\{N_{i_{IM},j,T_{IM}} | N_{i,j,T}\}. \end{aligned} \quad (20)$$

Substituting (9) and (20) into (6) results in the expression for product quality. Note that, it should be no less than the predefined threshold $(1 - \tau)$, namely

$$Q = \frac{\Pr\{X_1 X_2\}}{\Pr\{X_2\}} \geq 1 - \tau. \quad (21)$$

Test yield indicates the probability that no less than $(m + t)$ cores on the homogeneous manycore chips pass test.¹ These

¹Test yield reveals the influence of test quality on manufactured chips, while true yield, having been well studied in the literature, reflects the probability that less than m cores on a chip contain defects. True yield in our case can be simply computed as $Y_{true} = \sum_{i=m}^u \Pr\{C_{i,u,T_{IM}}\}$.

chips will be shipped to customers as quality products. We have

$$Y_{test} = \Pr\{X_2\}. \quad (22)$$

With this analytical model, we first consider full burn-in case, that is, setting $T = T_{IM}$. As no burn-in driven redundancy is introduced, t is set to zero. Then, we gradually reduce T , with which some latent defects do not reveal themselves before functioning in-field and then cannot be detected during the manufacturing test process. Without introducing burn-in driven redundancy, the product reliability decreases with the reduction of burn-in time. To meet product quality constraint, some burn-in driven redundant cores need to be introduced on-chip.

C. Proposed Cost Model

Various cost models for IC manufacturing and testing have been presented in [29]–[31], and [32]. Different from the above models that involve a great amount of manufacturing and test parameters, we present a simple yet effective cost model to capture the key impact of introducing burn-in driven redundancy into homogeneous manycore chips. That is, instead of obtaining concrete values for different cost factors, the input to our model is unified ratio parameters among various cost factors.

The production cost per sold chip² can be calculated by the following equation to evaluate different redundancy configurations:

$$C_{prod}^{chip} = \frac{(C_{manu}^{core} + C_{ATE}^{core} + C_{burn-in}^{core}) \cdot (m + t)}{Y_{test}} \quad (23)$$

where C_{manu}^{core} , C_{ATE}^{core} , and $C_{burn-in}^{core}$ indicate the manufacturing cost, ATE cost for applying test patterns, and burn-in cost per fabricated core, respectively. Note that, test cost includes both ATE cost and burn-in cost.

1) *Manufacturing Cost*: We set the manufacturing cost of each core for homogeneous manycore chips without redundancy to be 1 unit and we normalize manufacturing cost for chips with redundancy to this base value accordingly. Superscript *wo* is used to distinguish the “without redundancy” case from the “with redundancy” case.

Without redundancy, since m cores are fabricated on the same die, the gross die per wafer can be modeled as a function of the number of on-chip cores [33], that is

$$N^{wo} = \frac{\pi(d/2)^2}{A_r \cdot m} - \frac{\pi \cdot d}{\sqrt{2} \cdot A_r \cdot m} \quad (24)$$

where A_r is the area of each core, d is the dimension of the wafer.

Manufacturing cost per fabricated chip is then simplified as fabrication cost per wafer F divided by the gross die per wafer. We have

$$C_{manu}^{chip,wo} = \frac{F}{N^{wo}}. \quad (25)$$

²Other cost factors (e.g., research and development cost) are excluded without loss of the model's accuracy as they are independent to the redundancies introduced on-chip.

Thus, the manufacturing cost per core in this “without redundancy” case is given by

$$C_{\text{manu}}^{\text{core,wo}} = \frac{F}{N^{\text{wo}} \cdot m}. \quad (26)$$

With the normalization stated earlier (i.e., $C_{\text{manu}}^{\text{core,wo}} = 1$), we obtain the normalized fabrication cost per wafer F as

$$F = N^{\text{wo}} \cdot m. \quad (27)$$

We then consider the “with redundancy” case, which fabricates ($u = m + t$) cores on a die. By the model presented in (24), the gross die per wafer becomes

$$N = \frac{\pi(d/2)^2}{A_r \cdot u} - \frac{\pi \cdot d}{\sqrt{2} \cdot A_r \cdot u}. \quad (28)$$

Then, similar to (26), we obtain the manufacturing cost per chip with redundancy to be

$$C_{\text{manu}}^{\text{core}} = \frac{F}{N \cdot u}. \quad (29)$$

Substituting (27) into this expression yields the manufacturing cost per fabricated core of the chip with redundancy

$$C_{\text{manu}}^{\text{core}} = \frac{N^{\text{wo}} \cdot m}{N \cdot u}. \quad (30)$$

2) *ATE Cost:* As mentioned before, the defect coverage of manufacturing test has been assumed to be a fixed value B_r . Recall that the manufacturing cost per core in “without redundancy” case is normalized to be unit 1. We simply set the ATE cost per fabricated core as ρ units, where ρ is a unified ratio parameter between the ATE cost and the manufacturing cost. That is

$$C_{\text{ATE}}^{\text{core}} = \rho \cdot C_{\text{manu}}^{\text{core,wo}}. \quad (31)$$

3) *Burn-In Cost:* The burn-in cost highly depends on the occupation time of chips on the test equipments (e.g., burn-in ovens), because the service life of these expensive equipments is limited. We, therefore, assume the burn-in cost to be proportional to the burn-in time T . By normalizing the cost of fully burn-in process as $\xi C_{\text{manu}}^{\text{core}}$, we model the burn-in cost with duration T as

$$C_{\text{burn-in}}^{\text{core}} = \xi \cdot C_{\text{manu}}^{\text{core,wo}} \cdot \frac{T}{T_{\text{IM}}}. \quad (32)$$

D. Case Study for Partial/No Burn-In

Consider a homogeneous manycore chip that functions with no less than 32 defect-free cores (i.e., $m = 32$). To meet the product quality requirement $\tau = 500$ DPPM for the homogeneous manycore chip without any redundancy given full burn-in process, we set defect coverage for core test as a sufficient high value 99.9%. In addition, the killer defect density $\lambda_K = 0.05$, latent-to-killer defect density ratio $\gamma = 0.02$, ATE cost ratio $\rho = 10\%$, and burn-in cost ratio $\xi = 20\%$. Other parameter setups are provided in Section VII-A. The experimental results for this case study are shown in Fig. 3.

When we perform full burn-in process with adequately high-defect coverage, it is not necessary to introduce any burn-in

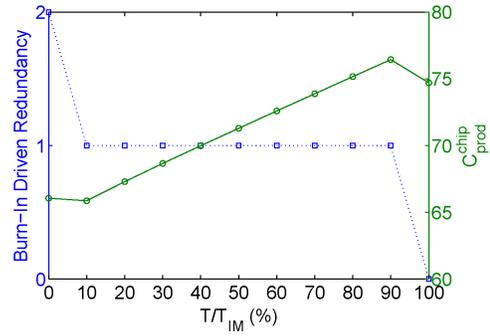


Fig. 3. Production cost with partial/no burn-in test.

driven redundancy (i.e., $t = 0$). With the shortening of burn-in time T , introducing more burn-in driven redundant cores becomes a must for meeting the product quality requirement. To be specific, when T is in the range from $90\%T_{\text{IM}}$ to $10\%T_{\text{IM}}$ one burn-in driven redundant core is enough, while if no burn-in test is provided (i.e., $T = 0$) the system needs one more redundant core to guarantee product quality (see the dotted line).

If we introduce one burn-in driven redundancy but do not reduce burn-in time much (i.e., $T = 90\%$ or $80\%T_{\text{IM}}$), the production cost does not decrease. This is because the manufacturing cost increment caused by burn-in driven redundancy exceeds the burn-in cost reduction. But if the burn-in time is further shortened, we observe significant benefits in terms of total cost. In this experiment, the minimum production cost is achieved at $T = 10\%T_{\text{IM}}$. The cost reduction compared with full burn-in case is close to 10%. When the burn-in time drops from $T = 10\%T_{\text{IM}}$ to 0, more redundant cores are introduced (i.e., t increases). As a result, the increment of manufacturing cost and ATE cost per fabricated chip is 1.1 units, while the burn-in cost reduction is only 0.66, which increases the total production cost.

V. TEST ECONOMICS WITH PARTIAL MANUFACTURING TEST

Similar to partial burn-in process, insufficient manufacturing tests can lead to product quality decrease but the quality loss could be recovered by introducing some redundant cores. As can be observed in Section IV-D, if the test cost reduction exceeds the manufacturing cost increment, the total production cost can be reduced. We, therefore, study the impact of introducing redundancy for relaxed defect coverage requirement. As both partial/no burn-in and partial manufacturing test result in product quality decrease and the corresponding spares are used for recouping the product quality loss, we do not distinguish these two types of redundancy deliberately in the rest of this paper. That is, instead of representing them separately, we use n to denote the total number of test cost-driven redundant cores.

A. Impact of Test Decision Criterion

The primary objective of manufacturing test is to obtain low-test escapes in order to ensure the quality of the shipped

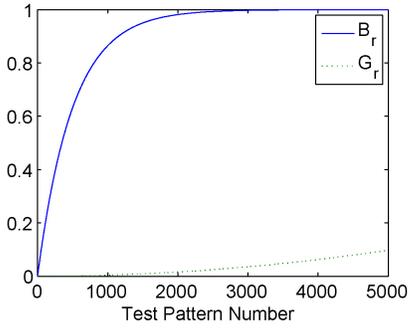


Fig. 4. Test escapes versus false rejects.

products, and a limited number of false rejects are considered as acceptable loss. With the ever advancement in semiconductor technology, however, it has been reported that the number of false rejects has dramatically increased [34], [35]. The associated test yield loss (i.e., the fraction of chips that fail manufacturing tests but would work in application) may have significant impact on manufacturing cost. We, therefore, examine the influence of false rejects and test escape in this section.

To capture the above effects, let us use G_a , G_r , B_a , B_r to denote the conditional probability that a defect-free core passes test, that a defect-free core is rejected, that a core containing defects escapes from the manufacturing test, and that a bad core is rejected, respectively. Apparently, $G_a + G_r = 1$ and $B_a + B_r = 1$. According to [19], the fraction of bad cores to be detected after applying k test patterns can be modeled as

$$B_r = 1 - e^{-vk}. \quad (33)$$

Also, depending on the effectiveness of the decision criterion μ , the correlation between B_r and G_r can be expressed as [27]

$$B_r = 1 - e^{-\mu\sqrt{G_r}}. \quad (34)$$

Note that, this equation is used to describe the right general shape instead of an accurate representation of any decision process, as pointed out in [27]. To model a particular decision process, we can resort to curve fitting techniques to obtain the parameter μ .

Combining (33) and (34), we can express G_r , the probability for false rejects, in terms of the number of applied test patterns k as

$$G_r = \left(\frac{vk}{\mu}\right)^2. \quad (35)$$

Ideally, a perfect manufacturing test is able to reject all bad cores while accept all defect-free ones, i.e., $B_r \equiv 1$ and $G_r \equiv 0$. It can be achieved by taking the limit of the above equations as μ goes to ∞ . In reality, because of various challenges in decision-making (e.g., the overlap between the good and the bad populations [27]), μ is a finite value. Fig. 4 shows a typical G_r and B_r versus the test pattern count, where the decision criterion μ is set to be 32 and v is set to be 0.002 [19]. Generally speaking, the better the decision method, the more square the plot of B_r versus G_r .

B. Product Quality With False Rejects

With test cost-driven redundancy, the problem comes down to determine s and n values such that the production cost for sold chips is minimized under the product quality constraint. In total ($w = m + n + s$) homogeneous cores are fabricated on the chip. Among them, if no less than $(m + n)$ cores pass test after burn-in time T , the chip will be shipped to the market. We need to guarantee that the probability that a sold chip contains no less than m defect-free cores at the end of infant mortality is higher than the given threshold τ . We, therefore, redefine X_2 as the event that no less $(m + n)$ cores among all w cores pass the partial manufacturing test given burn-in time T , denoted by \tilde{X}_2 . Again, by the total probability theorem, we compute $\Pr\{\tilde{X}_2\}$ in a divide-and-conquer manner, that is

$$\Pr\{\tilde{X}_2\} = \sum_{i=0}^w \Pr\{\tilde{X}_2|C_{i,w,T}\} \Pr\{C_{i,w,T}\}. \quad (36)$$

The computation of $\Pr\{C_{i,w,T}\}$ is similar to that of $\Pr\{C_{i,u,T}\}$ in Section IV, that is

$$\Pr\{C_{i,w,T}\} = \sum_{j=0}^w \Pr\{M_{j,w}\} \Pr\{N_{i,j,T}\}. \quad (37)$$

Given $C_{i,w,T}$, taking both false reject and test escape into account, event \tilde{X}_2 is the union of a series of mutually exclusive events $[A_{p,i} \cap B_{q,w-i}]$, meaning that exactly p good cores and q bad cores pass test, where $A_{p,i}$ represents the event that among i good cores on a chip, p pass the test while $(i - p)$ fail the test, and $B_{q,w-i}$ is the event that q -out-of- $(w - i)$ bad cores pass test. To explore all possible combinations for $[\tilde{X}_2|C_{i,w,T}]$ event to be true, we need to determine all possible values for p and q . Apparently, due to false rejects and the fact that the number of good cores that pass test cannot exceed that of good cores i , p can be $0, \dots, i$. As for q , a sold chip should satisfy two conditions: (i) $q = 0, \dots, w - i$; (ii) $p + q \geq m + n$, that is, only the chips contain no less than $(m + n)$ pass-test cores are sold. We, therefore, have $q = \max\{0, m + n - p\}, \dots, w - i$. For the ease of discussion, let $\omega \equiv \max\{0, m + n - p\}$. Based on the above, we have

$$\Pr\{\tilde{X}_2|C_{i,w,T}\} = \sum_{p=0}^i \sum_{q=\omega}^{w-i} \Pr\{A_{p,i} \cap B_{q,w-i}|C_{i,w,T}\}. \quad (38)$$

Since the two events $A_{p,i}$ and $B_{q,j}$ are conditionally independent given event $C_{i,w,T}$, we have [36]

$$\Pr\{A_{p,i} \cap B_{q,w-i}|C_{i,w,T}\} = \Pr\{A_{p,i}|C_{i,w,T}\} \cdot \Pr\{B_{q,w-i}|C_{i,w,T}\}. \quad (39)$$

Substituting it into (38) yields

$$\Pr\{\tilde{X}_2|C_{i,w,T}\} = \sum_{p=0}^i \sum_{q=\omega}^{w-i} \Pr\{A_{p,i}|C_{i,w,T}\} \cdot \Pr\{B_{q,w-i}|C_{i,w,T}\}. \quad (40)$$

And then, substituting (37) and (40) in (36) results in

$$\Pr\{\tilde{X}_2\} = \sum_{j=0}^w \sum_{i=0}^j \sum_{p=0}^i \sum_{q=\omega}^{w-i} \Pr\{A_{p,i}|C_{i,w,T}\} \cdot \Pr\{B_{q,w-i}|C_{i,w,T}\} \Pr\{M_{j,w}\} \Pr\{N_{i,j,T}\}. \quad (41)$$

In this equation, the last two terms have been defined by (10) and (11), respectively. When calculating $\Pr\{A_{p,i}|C_{i,w,T}\}$ and $\Pr\{B_{q,w-i}|C_{i,w,T}\}$, for the sake of simplicity, we assume that the event that good cores rejected by test and that bad cores accepted by test follow Poisson distribution [28]. This assumption is acceptable because: we apply the same test patterns on all the cores that suffer from random manufacturing defects, the test escapes of bad cores can be hence regarded as mutually independent, so are the false rejects of good cores. Therefore, we have

$$\Pr\{A_{p,i}|C_{i,w,T}\} = \binom{i}{p} (1 - G_r)^p \cdot G_r^{i-p} \quad (42)$$

and

$$\Pr\{B_{q,w-i}|C_{i,w,T}\} = \binom{w-i}{q} (1 - B_r)^q \cdot B_r^{w-i-q} \quad (43)$$

where G_r and B_r are the functions of the applied number of test patterns k [see (33) and (35)]. Note that, by the definition, $\Pr\{\tilde{X}_2\}$ can also be viewed as the test yield of fabricated chips (i.e., \tilde{Y}_{test}).

To compute the product quality \tilde{Q} , it is also necessary to redefine (20). By similar argument, we have

$$\begin{aligned} \Pr\{X_1 \tilde{X}_2\} &= \sum_{j=m}^w \sum_{i=m}^j \sum_{i_{\text{IM}}=m}^i \sum_{p=0}^i \sum_{q=0}^{w-i} \Pr\{A_{p,i}|C_{i,w,T_{\text{IM}}}\} \\ &\quad \cdot \Pr\{B_{q,w-i}|C_{i,w,T_{\text{IM}}}\} \Pr\{M_{j,w}\} \\ &\quad \cdot \Pr\{N_{i,j,T}\} \Pr\{N_{i_{\text{IM}},j,T_{\text{IM}}}|N_{i,j,T}\} \end{aligned} \quad (44)$$

and, therefore, the product quality constraint can be written in terms of $\Pr\{X_1 \tilde{X}_2\}$ and $\Pr\{\tilde{X}_2\}$, that is

$$\tilde{Q} = \frac{\Pr\{X_1 \tilde{X}_2\}}{\Pr\{\tilde{X}_2\}} \geq 1 - \tau. \quad (45)$$

C. Cost Model

With the analytical model of product quality and test yield, we move to discuss the impact of partial manufacturing test and test cost-driven redundancy on the total production cost. The total production cost for the homogeneous manycore chip is given by

$$\tilde{C}_{\text{prod}}^{\text{chip}} = \frac{(C_{\text{manu}}^{\text{core}} + \tilde{C}_{\text{ATE}}^{\text{core}} + C_{\text{burn-in}}^{\text{core}}) \cdot (m + n + s)}{\tilde{Y}_{\text{test}}}. \quad (46)$$

In the above equation, $C_{\text{manu}}^{\text{core}}$ and $C_{\text{burn-in}}^{\text{core}}$ remain the same as derived in Section IV-C.

For $\tilde{C}_{\text{ATE}}^{\text{core}}$, the ATE cost per fabricated core is determined by the number of test patterns applied on ATE, which is constrained by the test quality requirement. We calculate this value as follows.

Given the ATE cost ratio parameter ρ , we set the ATE cost per fabricated core with an arbitrary defect coverage η as ρ units. The actual ATE cost, since its corresponding defect coverage may not be η , is normalized to the reference case. To be specific, we first compute the test pattern count for achieving defect coverage η by (33) as

$$k_\eta = \frac{\ln(1 - \eta)}{-v}. \quad (47)$$

Thus, we obtain the normalized average cost for applying a single test pattern, that is

$$\tilde{C}_{\text{ATE}}^{\text{pattern}} = \frac{\rho \cdot C_{\text{manu}}^{\text{core,wo}}}{k_\eta} = \frac{\rho}{k_\eta}. \quad (48)$$

Similar to (47), the test pattern count for achieving defect coverage B_r is given by

$$k = \frac{\ln(1 - B_r)}{-v}. \quad (49)$$

The ATE cost for each core is, therefore

$$\tilde{C}_{\text{ATE}}^{\text{core}} = \tilde{C}_{\text{ATE}}^{\text{pattern}} \cdot k. \quad (50)$$

With this model, clearly, the proposed strategy is preferred when the number of processor cores in the original design (i.e., m) is not very small. Otherwise, the manufacturing cost increment induced by redundancy might be very considerable, and hence exceed the associated test cost reduction.

VI. TEST ECONOMICS WITH PRODUCT BINNING

In this section, we examine the impact of an important economic activity—product binning—on the effectiveness of the proposed strategy.

A. Product Binning and Product Quality

Recall that the problem has been formulated as: the chips whose quantities of pass-test cores are within the range from f_ℓ to $f_{\ell-1}$ are sold as the ℓ th-bin products, and the probability of no less than g_ℓ defect-free cores at the end of infant mortality should be no less than $1 - \tau$. To model these events, we use superscript ℓ to indicate the notations for the ℓ th bin. To be specific, we denote by X_1^ℓ the event that no less than g_ℓ cores on a chip is defect-free at the end of infant mortality, and X_2^ℓ the event that the quantity of pass-test cores after burn-in time duration T is in the range $[f_\ell, f_{\ell-1} - 1]$ (both inclusive). In particular, for the first bin (namely, $\ell = 1$) the range is $[f_1, f_0]$ instead of $[f_1, f_0 - 1]$, where $f_1 = m + n$ and $f_0 = m + n + s$.

With these notations, the percentage of ℓ th-bin products among all the products is given by $\Pr\{X_2^\ell\}$. In addition, assuming the same product quality requirement for all bins, we need to guarantee that

$$Q^\ell = \Pr\{X_1^\ell | X_2^\ell\} \frac{\Pr\{X_1^\ell X_2^\ell\}}{\Pr\{X_2^\ell\}} \geq 1 - \tau, \quad \forall \ell. \quad (51)$$

To calculate Q^ℓ , again, we start with the calculation of $\Pr\{X_2^\ell\}$. Similar to the argument in Section V, the number of good cores that pass manufacturing test cannot exceed that of all good cores on the chip, therefore, $p = 0, \dots, i$. For a chip in bin ℓ , its number of bad cores that pass test q has constraints that $0 \leq q \leq w - i$ and $f_\ell \leq p + q \leq f_{\ell-1} - 1$. We, therefore, have $q = \max\{0, f_\ell - p\}, \dots, \min\{f_{\ell-1} - 1 - p, w - i\}$. Thus, (41) is redefined as

$$\begin{aligned} \Pr\{X_2^\ell\} &= \sum_{j=0}^w \sum_{i=0}^j \sum_{p=0}^i \sum_{q=\max\{0, f_\ell - p\}}^{\min\{f_{\ell-1} - 1 - p, w - i\}} \Pr\{A_{p,i}|C_{i,w,T}\} \\ &\quad \cdot \Pr\{B_{q,w-i}|C_{i,w,T}\} \Pr\{M_{j,w}\} \Pr\{N_{i,j,T}\}. \end{aligned} \quad (52)$$

We then move to redefine (44). By similar argument, since the number of pass-test cores is in the range of $[f_\ell, f_{\ell-1} - 1]$, the possible values for p and q are the same as that in (52). In addition, because of the product quality requirement, the number of defect-free cores at the end of infant mortality i_{IM} should be no less than g_ℓ . Moreover, as depicted in Fig. 2, we have $i_{IM} \leq i \leq j \leq w$. With these constraints, the probability of event $[X_1^\ell X_2^\ell]$ is, therefore, given by

$$\Pr\{X_1^\ell X_2^\ell\} = \sum_{j=g_\ell}^w \sum_{i=g_\ell}^j \sum_{i_{IM}=g_\ell}^i \sum_{p=0}^i \sum_{q=\max\{0, f_\ell-p\}}^{\min\{f_{\ell-1}-1-p, w-i\}} \Pr\{A_{p,i}|C_{i,w,T_{IM}}\} \Pr\{B_{q,w-i}|C_{i,w,T_{IM}}\} \cdot \Pr\{M_{j,w}\} \Pr\{N_{i,j,T}\} \Pr\{N_{i_{IM},j,T_{IM}}|N_{i,j,T}\}. \quad (53)$$

A special case is $\ell = 1$, namely, the first bin. In this case, the number of pass-test cores is in the range of $[f_1, f_0]$, where $f_1 = m + n$ and $f_0 = m + n + s$. Also, $g_1 = m$. Thus, (52) and (53) are simplified as

$$\Pr\{X_2^1\} = \sum_{j=0}^w \sum_{i=0}^j \sum_{p=0}^i \sum_{q=\max\{0, m+n-p\}}^{\min\{w-p, w-i\}} \Pr\{A_{p,i}|C_{i,w,T}\} \cdot \Pr\{B_{q,w-i}|C_{i,w,T}\} \Pr\{M_{j,w}\} \Pr\{N_{i,j,T}\} \quad (54)$$

and

$$\Pr\{X_1^1 X_2^1\} = \sum_{j=m}^w \sum_{i=m}^j \sum_{i_{IM}=m}^i \sum_{p=0}^i \sum_{q=\max\{0, m+n-p\}}^{\min\{w-p, w-i\}} \Pr\{A_{p,i}|C_{i,w,T_{IM}}\} \Pr\{B_{q,w-i}|C_{i,w,T_{IM}}\} \cdot \Pr\{M_{j,w}\} \Pr\{N_{i,j,T}\} \Pr\{N_{i_{IM},j,T_{IM}}|N_{i,j,T}\}. \quad (55)$$

Because of $p \leq i$, these equations are exactly the same as (41) and (44), respectively.

B. Cost Model

There are a few metrics to evaluate the production cost of the homogeneous manycore chip given product binning. One of them is with the assumption that every shipped product has the same production cost, that is

$$\widehat{C}_{\text{prod}}^{\text{chip,avg}} = \frac{(C_{\text{manu}}^{\text{core}} + \widetilde{C}_{\text{ATE}}^{\text{core}} + C_{\text{burn-in}}^{\text{core}}) \cdot (m + n + s)}{\sum_{\ell=1}^b \Pr\{X_2^\ell\}} \quad (56)$$

where the derivation of $C_{\text{manu}}^{\text{core}}$, $\widetilde{C}_{\text{ATE}}^{\text{core}}$, and $C_{\text{burn-in}}^{\text{core}}$ are the same as that in Section V-C. Given there is only one bin, (56) can be simplified as (46).

VII. EXPERIMENTAL RESULTS

A. Experimental Setup

To evaluate the effectiveness of the proposed strategy, we perform extensive experiments for the production cost of a homogeneous manycore chip that functions with no less than 32 defect-free cores (i.e., $m = 32$), varying the number of test cost-driven redundancy n and burn-in time T . In our work, the best n , T and s combination in terms of production cost

is determined by exploring the possible n , T and s solution space. Note, this is not a time-consuming process because the computation time for each single configuration is quite small. Also, a large value of n and/or s will increase the production cost due to the extra silicon area and hence the possible combinations to be explored are quite limited.

We set the system parameters based on prior work as follows: $v = 0.002$ [19], $\mu = 18$ [27], $\alpha = 0.3$ [13], $\beta = 0.3$, $\xi = 0.2$ [15], $d = 300$ mm, $A_r = 10$ mm², $\eta = 95\%$ unless otherwise specified. The product quality requirement is set to 500 DPPM (i.e., $\tau = 5 \times 10^{-4}$) unless specified otherwise.

B. Results and Discussion

Four sets of experiments are conducted to analyze as follows.

- 1) The tradeoff between burn-in cost and ATE cost under a certain product quality constraint.
- 2) The effectiveness of the proposed strategy.
- 3) The impact of defect clustering effects and product quality constraints on the effectiveness of the proposed strategy.
- 4) The impact of product binning on the effectiveness of the proposed strategy.

In the first three sets of experiments, one product bin is assumed and we utilize the economic model proposed in Section V for analysis. For the last experiment, we employ the model presented in Section VI to analyze the effect of product binning.

1) *Tradeoff Between Burn-In and ATE Cost:* First of all, we demonstrate the tradeoff between burn-in cost and ATE cost under the product quality constraint. As both partial burn-in process and partial manufacturing test may scarify some product quality, we introduce a few test cost-driven redundant cores to recap this loss. From another point of view, given certain test cost-driven spares, to meet the same product quality requirement, if we shorten the burn-in time, the test coverage must be increased and hence the ATE cost increases if we save burn-in cost. Fig. 5(a) illustrates this trend for the cases when $n = 1$, $n = 2$, and $n = 3$, with high-defect density ($\lambda = 0.05$, $\gamma = 0.05$) and high-ATE cost ratio ($\rho = 20\%$). No yield-driven spares are introduced in this experiment (i.e., $s = 0$).

As shown in this figure, with the shrinking of burn-in time, the ATE cost increases dramatically when $n = 1$, but if more redundant cores are added (that is, $n = 2$ and $n = 3$), the ATE cost only increases slightly. This is because, the responsibility for manufacturing test can be significantly relaxed by introducing more than one test cost-driven redundant cores, in spite of the large number of fabricated cores on a chip ($m = 32$). Thus, although a great percentage of latent defects cannot be revealed because of partial burn-in test, we still do not need very high manufacturing test coverage. However, if only one test cost-driven redundancy is involved (i.e., $n = 1$), the product quality requirement is not relaxed much. As a result, we need to trade a small percentage of burn-in time reduction with a significant test pattern count increment. Another interesting observation is that, with the decrease of burn-in

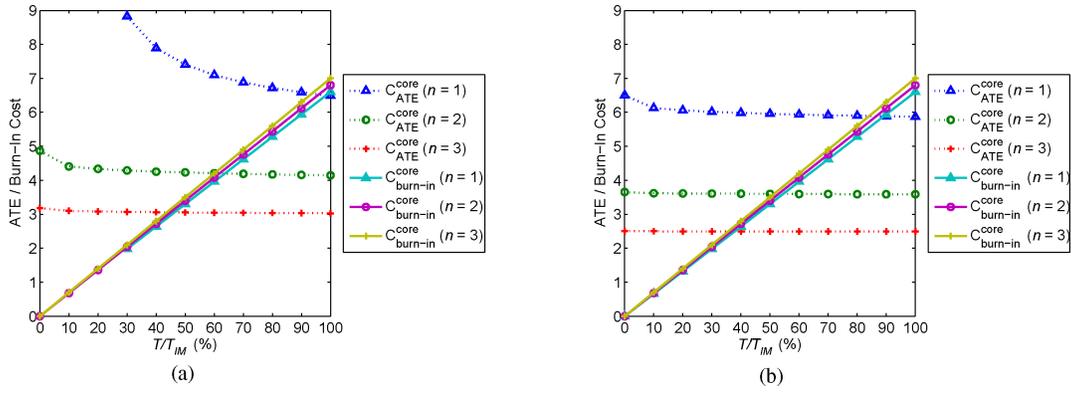


Fig. 5. Tradeoff between ATE cost and burn-in cost. (a) $\lambda = 0.05$ and $\gamma = 0.05$. (b) $\lambda = 0.02$ and $\gamma = 0.02$.

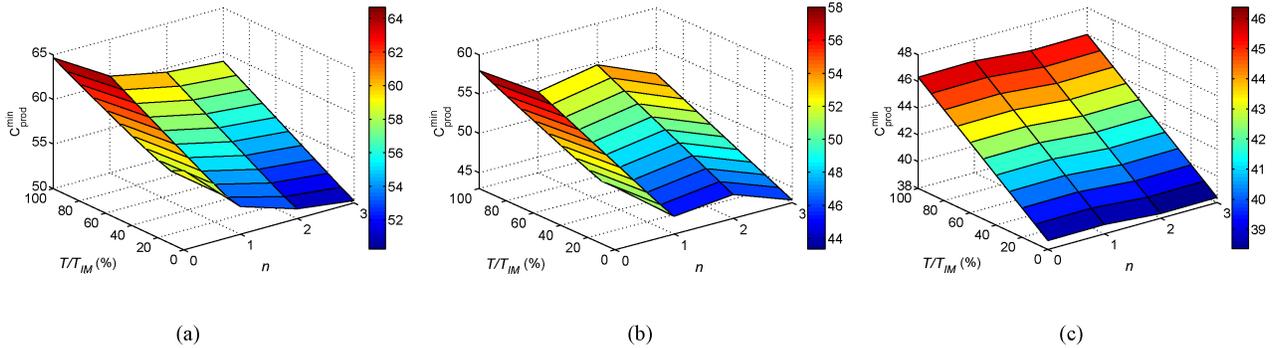


Fig. 6. Minimum production cost. (a) $\lambda = 0.05, \gamma = 0.05$, and $\rho = 20\%$. (b) $\lambda = 0.02, \gamma = 0.02$, and $\rho = 10\%$. (c) $\lambda = 0.01, \gamma = 0.01$, and $\rho = 10\%$.

time, ATE cost grows sharply. We attribute this phenomenon to the slowdown of failure rate decrease with respect to the last of burn-in time. Also, we observe that if the burn-in time is shorter than $30\%T_{IM}$, we cannot satisfy the product quality constraint by increasing manufacturing test coverage. The above observations indicate that employing very limited number of test cost-driven redundant cores cannot reduce test cost much, when the defect density is high. Introducing a few more test cost-driven spares can be much beneficial.

Fig. 5(b) shows the results when the defect density is relatively low ($\lambda = 0.02$ and $\gamma = 0.02$). In this case, the test pattern count and hence the ATE cost increment with respect to burn-in time decline is very small. We, therefore, tend to achieve better results in terms of production cost by reducing burn-in time. If the ATE cost ratio is even lower than $\rho = 20\%$, since the test pattern count does not vary with ρ , more benefits can be obtained. This observation can also be used to explain the phenomenon shown in Fig. 3.

2) *Effectiveness of the Proposed Strategy:* Next, we compare the traditional strategy that does not introduce any test cost-driven redundancy (i.e., $n = 0$) but includes a few yield-driven spares (s can be more than zero), and the proposed approach with both redundancies, for various burn-in time T . For fair comparison, given T , and n , we vary s to find the minimum production cost per sold chip, and record this value and the corresponding s .

Fig. 6(a) shows the minimum production cost given various burn-in time T and test cost-driven redundancy n , for killer defect density $\lambda = 0.05$, latent-to-killer defect density ratio

$\gamma = 0.05$, and ATE cost ratio $\rho = 20\%$. Introducing test cost-driven spares results in significant cost reduction when compared with the traditional approach with full burn-in time without such redundant cores. In particular, the maximum cost reduction is as high as 22.28%, obtained when $n = 3, s = 6$ and $T = 0$. With the shrinking of burn-in time, the production cost gradually decreases. We take a closer observation and consider $n = 1$ case as an example. In this case, in spite of the variation of burn-in time T , the number of yield-driven redundant cores s that results in the minimum production cost remains the same value ($s = 5$). Thus, the manufacturing cost does not increase with the decrease of burn-in time. Since the burn-in cost reduction is more significant than the ATE cost increment, the total production cost keeps declining. A special case is observed when $n = 0$. When the burn-in time drops from $T = 10\%T_{IM}$ to 0, the total production cost increases because of the diminishing test yield and the increment of ATE cost.

We observe even more production cost reduction when $\lambda = 0.02$ and $\gamma = 0.02$, as shown in Fig. 6(b). We achieve up to 25.26% production cost reduction by using three test cost-driven spares (i.e., $n = 3$) and no burn-in test. This is mainly because, due to the low-defect density in this experimental setup, the test coverage requirement and the associated ATE cost do not increase much with the shortening of burn-in time [see Fig. 5(b)]. Therefore, with the decrease of T , the test cost (including both ATE cost and burn-in cost) reduces dramatically. It is interesting to observe that the production cost increases when n moves from 1 to 2 and $T = 0$. This

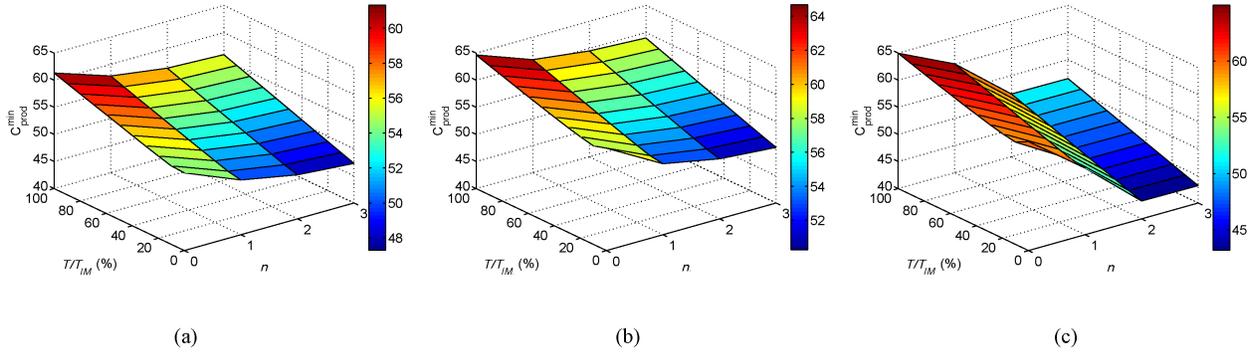


Fig. 7. Minimum production cost with defect clustering effects ($\lambda = 0.05$, $\gamma = 0.05$, and $\rho = 20\%$). (a) $\alpha = 0.1$. (b) $\alpha = 0.3$ [reprinted Fig. 6(a)]. (c) $\alpha = 0.5$.

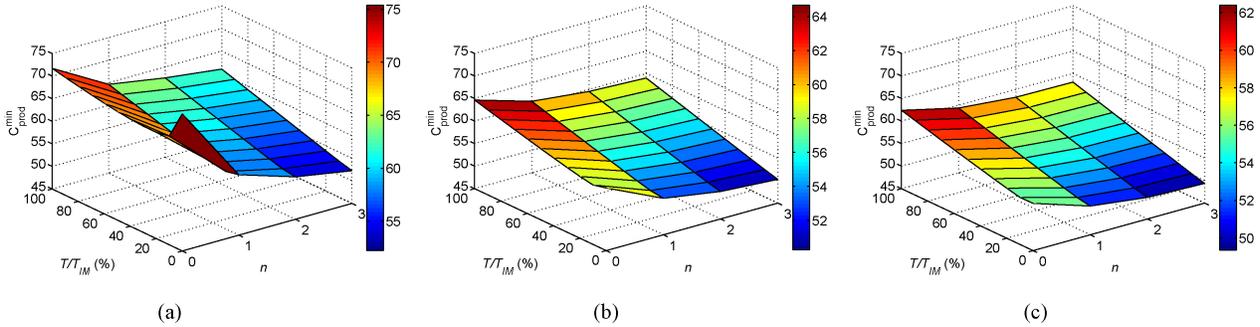


Fig. 8. Minimum production cost with product quality requirements ($\lambda = 0.05$, $\gamma = 0.05$, and $\rho = 20\%$). (a) $\tau = 100$ DPPM. (b) $\tau = 500$ DPPM [reprinted Fig. 6(a)]. (c) $\tau = 1000$ DPPM.

is due to the slight increment of ATE cost and the significant test yield reduction. To be specific, when n increases from 1 to 2, both manufacturing cost and burn-in cost per fabricated chip remain the same, while ATE cost decreases. As a result, the production cost per fabricated chip decreases by 2.70%. The test yield, on the other hand, is reduced by 3.66%. Similar phenomenon does not occur in Fig. 6(a) because the reduction of production cost per fabricated chip and that of test yield are 4.63% and 0.76%, respectively. A closer examination on these two cases shows that the difference can be attributed to the fact that ATE cost takes a greater share in case of heavier defect density (namely, λ and γ are larger values).

Fig. 6(c) presents a case that the minimum production cost occurs when two test cost-driven redundant cores are involved (i.e., $n = 2$). For certain burn-in time T , only modest variation in terms of production cost can be observed. This is because when n increase from 0 to 1, and then to 2, the number of fabricated cores on a chip in these cases remains the same (i.e., $m + n + s = 38$ for all cases). Thus, the manufacturing cost and burn-in cost per chip are fixed. Since the defect density is quite low ($\lambda = 0.01$ and $\gamma = 0.01$), there is no significant ATE cost reduction with respect to the product quality relaxation. If more test cost-driven redundant cores are introduced, as it might result in the increment of number of fabricated cores on chip, we cannot obtain more benefits.

3) *Sensitivity Analysis*: We then discuss the influence of defect clustering effects on the effectiveness of the proposed strategy shown in Fig. 7 with $\lambda = 0.05$, $\gamma = 0.05$, $\rho = 20\%$. The chips with clustering parameter $\alpha = 0.1$ have the least production cost in these three cases given no redundancies

are introduced, although the average quantities of defects per core are the same. This is due to the highest test yield which is obtained with the strongest clustering effects. We compare three cases for $n = 0$, $T = T_{IM}$ as an example, where $\alpha = 0.1$ case has test yield 89.23% while other two cases are no higher than 84.30%.

For weak clustering effect case (i.e., $\alpha = 0.5$), we observe the most significant cost reduction by adding test cost-driven spares, especially when $n \geq 2$. The difference is up to 27.98%, thanks mainly to ATE cost reduction and dramatic test yield increment. To clarify, suppose the IC products are tested conventionally, that is, no test cost-driven redundancy and with full burn-in test, the required test pattern quantity increase but the test yield decrease with respect to the increment of α because the defects are more dispersive. However, once $n > 0$, the defect coverage requirement is relaxed almost to the limit: the ATE cost for $n = 2$ is only 13.00% of that for $n = 0$, and the yield increase from 83.02% to 95.95%. In this sense, the proposed strategy has more remarkable effectiveness for the products with weaker clustering defects.

We are also interested in what if more stringent product quality criterion is required, which is of importance in some special domain (e.g., aerospace engineering, automobile industry). The requirement τ is set to 100 DPPM, 500 DPPM and 1000 DPPM for comparison, as shown in Fig. 8. Other parameters are set to $\lambda = 0.05$, $\gamma = 0.05$, and $\rho = 20\%$. As expected, tighter product quality requirement results in higher production cost. When $n = 0$ and $T = T_{IM}$, we observe high-production cost up to 71.72 for $\tau = 100$ DPPM [see Fig. 8(a)]. By using the proposed partial burn-in and relaxed defect cover-

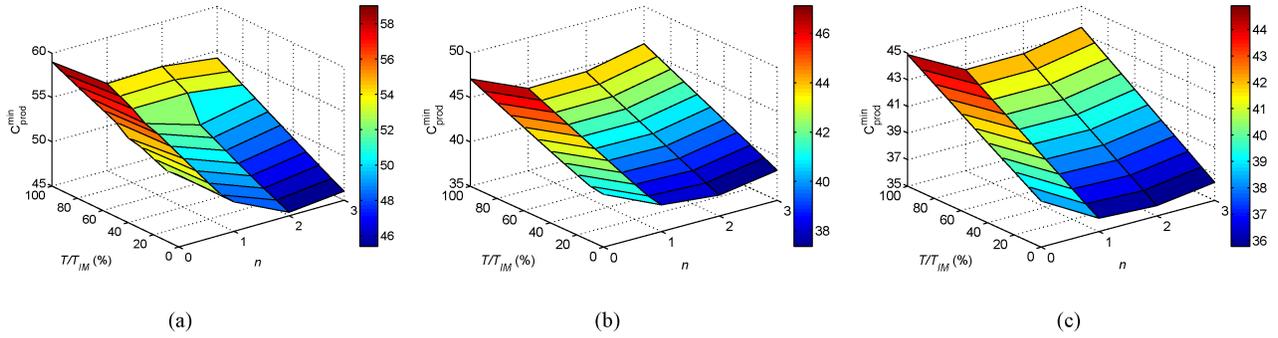


Fig. 9. Minimum production cost with product binning. (a) $\lambda = 0.05$, $\gamma = 0.05$, and $\rho = 20\%$. (b) $\lambda = 0.02$, $\gamma = 0.02$, and $\rho = 10\%$. (c) $\lambda = 0.01$, $\gamma = 0.01$, and $\rho = 10\%$.

age strategy, it is cut down to 52.25 when three test cost-driven redundant cores are introduced on-chip and thus the difference is as high as 27.15%. By contrast, considering the other two cases, the differences are 22.28% and 20.88%. We have also examined the cases with even tighter product quality 50 DPPM and 10 DPPM. The trends are similar. The differences in those cases are 30.29% and 40.62%, respectively. This observation suggests that the proposed method brings greater benefits for a more stringent product quality criterion.

Another interesting phenomenon illustrated in Fig. 8 is the minimum production cost 68.55 occurs when $T = 30\%T_{IM}$, setting $\tau = 100$ DPPM and $n = 0$. If we completely eliminate burn-in test in this case, the production cost increases to 75.37, even higher than the cost with full burn-in, 71.72. We attribute it to the combined factors of ATE cost, burn-in cost, and test yield. To be specific, when the burn-in time decreases from T_{IM} to $30\%T_{IM}$, and then to 0, the ATE cost per fabricated chip increases from 12.55 to 13.34, and then to 14.95. This effect is dominated by the burn-in cost reduction, which decreases from 7.6 to 0. In other words, the growth rate of ATE cost is lower than the falling rate of burn-in cost with respect to the decrease of burn-in time, eliminating burn-in test results in minimum production cost per fabricated chip. At the same time, however, no burn-in test case has a much lower test yield when comparing to the partial or full burn-in test cases. That is, the test yield in case of no burn-in test is only 70.26%. In contrast, we have 78.22% and 81.07% test yield in partial and full burn-in test cases, respectively. These two aspects combine to result in the occurrence of minimum production cost per shipped chip at $T = 30\%T_{IM}$ instead of no burn-in case.

4) *Impact of Product Binning:* Finally, Fig. 9 shows the impact of product binning on the effectiveness of the proposed strategy, wherein two bins are assumed (i.e., $b = 2$) and g_2 is set to 28. The remaining parameters are the same as that in Fig. 6. We demonstrate the minimum production cost per sold chip (namely, $\hat{C}_{prod}^{chip,avg}$) on the figures, given burn-in time T and test cost-driven redundancy n .

It is interesting to notice that the minimum production cost is obtained when $n = 2$ and $T = 0$ in all three cases, indicating that two test cost-driven spares and no burn-in is the best combination. This is different from the case shown in Fig. 6, within which the minimum production cost occurs at $n = 3$. To

understand the cause of this difference, let us examine the effect of n for the two-bin cases. We take Fig. 9(a) as an example and set T to zero. When n increases from 0 to 2, and finally to 3, the total test yield (i.e., $\sum_{\ell=1}^b \Pr\{X_2^\ell\}$) reduces from 93.52% to 88.15%, and to 86.25%, while the production cost per fabricated chip decreases from 49.65 to 40.02, and then 39.70. More intuitively, when n moves from 2 to 3, the decrease of production cost per fabricated chip slows down, while that of the total yield does not. This phenomenon is different from our observation in Fig. 6(a). This is because the two-bin strategy results in more yield increment in the cases with small n when compared with the one-bin case. We then set $n = 2$ to examine the effect of burn-in time. The product binning strategy does not affect the required number of redundant cores and the required test pattern count. Rather, it influences the total test yield only. However, we observe similar yield increment when varying the burn-in time while keeping the number of test cost-driven spares. Thus, similar to the one-bin case, the production cost reduction with respect to the shrinking of burn-in time comes from the significant burn-in cost reduction.

VIII. CONCLUSION

Test cost can account for a large share of the total production cost for IC products, mainly due to critical coverage requirement to ensure product quality. In this paper, we proposed a novel test cost-driven redundancy concept for homogeneous manycore systems to reduce their production cost. By doing so, the test cost was likely to be decreased dramatically and exceeds the manufacturing cost increment for the extra cores, thus cutting down the total production cost of the system. An analytical model was presented to capture the key impact of the proposed approach on two cost factors: manufacturing cost and test cost. In addition, the impact of product binning on the test economics of homogeneous manycore chips was also discussed. With a set of experiments for hypothetical manycore processors with various configurations, we validated the effectiveness of the proposed strategy.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their constructive comments.

REFERENCES

- [1] Y.-H. Lee and C. Chen, "A two-level scheduling method: An effective parallelizing technique for uniform nested loops on a DSP multiprocessor," *J. Syst. Softw.*, vol. 75, nos. 1–2, pp. 155–170, 2005.
- [2] *Geforce 8800 Graphics Processors*. Nvidia. Santa Clara, CA [Online]. Available: <http://www.nvidia.com/page/geforce8800.html>
- [3] *Tile64 Processor Family*. Tiler. San Jose, CA [Online]. Available: <http://www.tiler.com/products/processors.php>
- [4] D. Geer, "Chip makers turn to multicore processors," *IEEE Comput.*, vol. 38, no. 5, pp. 11–13, May 2005.
- [5] K. Olukotun, B. A. Nayfeh, L. Hammond, K. Wilson, and K. Chang, "The case for a single-chip multiprocessor," *SIGOPS Oper. Syst. Rev.*, vol. 30, no. 5, pp. 2–11, Dec. 1996.
- [6] L. Zhang, Y. Han, Q. Xu, and X. Li, "Defect tolerance in homogeneous manycore processors using core-level redundancy with unified topology," in *Proc. DATE*, 2008, pp. 891–896.
- [7] *Cisco and IBM Collaborate to Design and Build World's Most Sophisticated, High-Performance 40 Gb/s Custom Chip*. Cisco. San Jose, CA [Online]. Available: http://newsroom.cisco.com/dlls/partners/news/2004/pr_prod_06-09.html
- [8] W. C. Riordan, R. Miller, and E. R. S. Pierre, "Reliability improvement and burn in optimization through the use of die level predictive modeling," in *Proc. IEEE Int. Reliab. Phys. Symp.*, 2005, pp. 435–445.
- [9] M. F. Zakaria, Z. A. Kassim, M. P.-L. Ooi, and S. Demidenko, "Reducing burn-in time through high-voltage stress test and Weibull statistical analysis," *IEEE Design Test Comput.*, vol. 23, no. 2, pp. 88–98, Mar. 2006.
- [10] A. W. Righter, C. F. Hawkins, J. M. Soden, and P. Maxwell, "CMOS IC reliability indicators and burn-in economics," in *Proc. IEEE ITC*, 1998, pp. 194–203.
- [11] L. Huang and Q. Xu, "Is it cost-effective to achieve very high fault coverage for testing homogeneous SoCs with core-level redundancy," in *Proc. IEEE ITC*, 2008, p. 1.
- [12] S. Shamshiri, P. Lisherness, S.-J. Pan, and K.-T. Cheng, "A cost analysis framework for multi-core systems with spares," in *Proc. IEEE ITC*, 2008, pp. 1–8.
- [13] W. Kuo and T. Kim, "An overview of manufacturing yield and reliability modeling for semiconductor products," *Proc. IEEE*, vol. 87, no. 8, pp. 1329–1344, Aug. 1999.
- [14] I. Koren, Z. Koren, and C. H. Stapper, "A unified negative-binomial distribution for yield analysis of defect-tolerant circuits," *IEEE Trans. Comput.*, vol. 42, no. 6, pp. 724–733, Jun. 1993.
- [15] T. S. Barnett and A. D. Singh, "Relating yield models to burn-in fall-out in time," in *Proc. IEEE ITC*, 2003, pp. 77–84.
- [16] L. Huang, F. Yuan, and Q. Xu, "Lifetime reliability-aware task allocation and scheduling for MPSoC platforms," in *Proc. DATE*, 2009, pp. 51–56.
- [17] L. Huang, F. Yuan, and Q. Xu, "On task allocation and scheduling for lifetime extension of platform-based MPSoC designs," *IEEE Trans. Parallel Distrib. Syst.*, to appear.
- [18] W. C. Riordan, R. Miller, J. M. Sherman, and J. Hicks, "Microprocessor reliability performance as a function of die location for a 0.25 μm , five layer metal CMOS logic process," in *Proc. IEEE Int. Reliab. Phys. Symp.*, 1999, pp. 1–11.
- [19] F.-F. Ferhani, N. R. Saxena, E. J. McCluskey, and P. Nigh, "How many test patterns are useless," in *Proc. IEEE VTS*, 2008, pp. 23–28.
- [20] L. Huang and Q. Xu, "Test economics for homogeneous manycore systems," in *Proc. IEEE ITC*, 2009, pp. 1–10.
- [21] A. Krstic, W.-C. Lai, K.-T. Cheng, L. Chen, and S. Dey, "Embedded software-based self-test for programmable core-based designs," *IEEE Design Test Comput.*, vol. 19, no. 4, pp. 18–27, Jul.–Aug. 2002.
- [22] M. Psarakis, D. Gizopoulos, and M. Hatzimihail, "Systematic software-based self-test for pipelined processors," in *Proc. ACM/IEEE DAC*, 2006, pp. 393–398.
- [23] P. J. Tan, T. Le, K.-H. Ng, P. Mantri, and J. Westfall, "Testing of UltraSPARC T1 microprocessor and its challenges," in *Proc. IEEE ITC*, 2006, pp. 1–10.
- [24] T. L. McLaurin, "The challenge of testing the ARM CORTEX-A8 microprocessor core," in *Proc. IEEE ITC*, 2006, pp. 1–10.
- [25] P. S. Bhojwani and R. N. Mahapatra, "A robust protocol for concurrent on-line test (COLT) of NoC-based systems-on-a-chip," in *Proc. ACM/IEEE DAC*, 2007, pp. 670–675.
- [26] S. Borkar, "Designing reliable systems from unreliable components: The challenges of transistor variability and degradation," *IEEE Micro*, vol. 25, no. 6, pp. 10–16, Nov.–Dec. 2005.
- [27] P. M. O'Neill, "Statistical test: A new paradigm to improve test effectiveness and efficiency," in *Proc. IEEE ITC*, 2007, pp. 1–10.
- [28] I. Koren and C. H. Stapper, "Yield models for defect-tolerant VLSI circuits: A review," in *Proc. Int. Workshop Defect Fault Tolerance VLSI Syst.*, 1988, pp. 1–22.
- [29] S.-K. Lu and C.-Y. Lee, "Modeling economics of DFT and DFY: A profit perspective," *IEE Proc., Comput. Digital Tech.*, vol. 151, no. 2, pp. 119–126, Mar. 2004.
- [30] J.-M. Lu and C.-W. Wu, "Cost and benefit models for logic and memory BIST," in *Proc. DATE*, 2000, pp. 710–715.
- [31] P. K. Nag, A. Gattiker, S. Wei, R. D. Blanton, and W. Maly, "Modeling the economics of testing: A DFT perspective," *IEEE Design Test Comput.*, vol. 19, no. 1, pp. 29–41, Jan.–Feb. 2002.
- [32] K. Sundararaman, S. Upadhyaya, and M. Margala, "Cost model analysis of DFT based fault tolerant SoC designs," in *Proc. ISQED*, 2004, pp. 465–469.
- [33] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*, 4th ed. San Mateo, CA: Morgan Kaufmann, 2006.
- [34] J. Rearick, "Too much delay fault coverage is a bad thing," in *Proc. IEEE ITC*, Nov. 2001, pp. 624–633.
- [35] J. Saxena, K. Butler, V. Jayaram, and S. Kundu, "A case study of IR-drop in structured at-speed testing," in *Proc. IEEE ITC*, 2003, pp. 1098–1104.
- [36] A. P. Dawid, "Conditional independence in statistical theory," *J. R. Statist. Soc.*, vol. 41, no. 1, pp. 1–31, 1979.



Lin Huang (S'08) received the B.S. degree in electronic engineering from Shanghai Jiaotong University, Shanghai, China, in 2007. She is currently pursuing the Ph.D. degree from the Reliable Computer (CURE) Laboratory, Department of Computer Science and Engineering, Chinese University of Hong Kong, Shatin, Hong Kong.

Her current research interests include reliability analysis of multicore systems and fault-tolerant computing.



Qiang Xu (M'06) received the B.E. and M.E. degrees in telecommunication engineering from the Beijing University of Posts and Telecommunications, Beijing, China, in 1997 and 2000, respectively, and the Ph.D. degree in electrical and computer engineering from McMaster University, Hamilton, ON, Canada, in 2005.

Since 2005, he has been an Assistant Professor with the Department of Computer Science and Engineering, Chinese University of Hong Kong (CUHK), Shatin, Hong Kong. He also leads the CUHK Reliable Computer (CURE) Laboratory in the same department. His current research interests range from testing and debugging of system-on-a-chip integrated circuits to fault tolerance and reliable computing. He has published more than 50 technical papers in these areas.

Dr. Xu was a recipient of the Best Paper Award in the 2004 IEEE/ACM Design, Automation and Test in Europe Conference. He is a member of the ACM SIGDA. He has served as a technical program committee member for a number of conferences on very large scale integrated design and testing.