Bounded-Lifetime Integrated Circuits

Puneet Gupta[2] and Andrew B. Kahng[1]
puneet@ee.ucla.edu, abk@cs.ucsd.edu
[1]ECE and CSE Departments, University of California, San Diego
[2]EE Department, University of California, Los Angeles

Abstract. *Integrated circuits with bounded lifetimes can have many business advantages. We give some simple examples of methods to enforce tunable expiration dates for chips using nanometer reliability mechanisms.*

Categories and Subject Descriptors: B.7.2 [Hardware]: INTEGRATED CIRCUITS – Design Aids; J.6 [Computer Applications]: COMPUTER-AIDED ENGINEERING

General Terms: Design, Reliability, Security, Standardization

Keywords: Bounded lifetime, physical IP, integrated circuits

Introduction. We propose physical IP-based enforcement of tunable lifetime bounds on the function of integrated circuits. Our approach exploits circuit physical (reliability) failure mechanisms that are prominent in ≤ 65nm process nodes. Benefits of having a well-defined "expiration date" in semiconductor products include:

- increased IC production volumes potentially resulting from new business models associated with metered or time-based access to IC components;
- reduced support and integration overheads, such as for embedded software, with respect to older product versions (i.e., cost of backward-compatibility); and
- reduced silicon area and power resources when lifetime bounds allow decreased reliability guardbands.

We recognize that an "expiration date" may not be of interest in a life-critical application domain (e.g., pacemaker), or where a central server can disable functionality, or at certain levels of system complexity. On the other hand, with trends to open frameworks, systems implemented on chip, and high-volume platform SOCs – and for existing applications such as personal computing or mobile telephony – bounded chip (and hence product) lifetime may be an attractive proposition. Basic objectives of physical IP for bounded chip lifetime include:

- use of multiple mechanisms to preclude any given attack;
- implementation with drop-in circuit IP in standard process flavors, so as to avoid any process change or non-trivial design methodology changes; and
- use of mechanisms beyond simple 'counters' (timers) that can be subjected to memory-corruption attacks.[1]

Several kinds of lifetime bounding can be contemplated:

1. limiting total time of use (= metering), where the chip can be power-cycled multiple times;
2. limiting the lifetime of the chip starting from the date of manufacture (or first activation) – i.e., a calendar-time limit including time when the chip is shut down; and
3. limiting continuous time of use each time the chip is powered on.

(3) is trivially achieved by use of counters which are reset at power-on. We are not aware of any mechanism for (2) that would continue the timer function while the chip is disconnected from a power supply. Hence, in this paper we focus on (1).

Related Work. Previously-proposed limiters to hardware functionality have included:

- rate-limiting governor circuits have been deployed, e.g., to prevent microprocessor overclocking [1]; and
- expiration dates associated with passwords or other tokens that a given electronic system must present, e.g., to a central server.

Implementation of a 'physical timer' that does not involve a clock and a counter is not obvious. For example, charging an on-chip capacitor at manufacture and bounding the chip lifetime based on its discharge (90% voltage decrease per 2.3RC) would require too large a capacitor, or some method of partial recovery. (A discharging capacitor to measure time in power-off state has been proposed in beyond-chip contexts; see, e.g., [3]). Several methods involving electrolytic solutions or electroplating processes have been proposed for, e.g., automotive contexts. [2] proposed the use of a pair of electrolyte cells wherein application of a current could cause migration of silver ions from one electrode to another in a given cell; migration of all the ions would lead to a high-resistance condition that triggers an alarm.

Proposed Approach. Bounding the lifetime of a circuit requires (1) a circuit disabling mechanism; and (2) an aging detection and lifetime trigger circuit. For disabling the chip, straightforward alternatives exist such as power-gating, clock-gating or excess body biasing. We focus on lifetime triggering methods, and propose exploiting physical (reliability) failure mechanisms such as electromigration or wearout (NBTI, TDDB).[2] In the following, we discuss example approaches for electromigration and NBTI. Criteria for a viable lifetime triggering mechanism include:

- Is the mechanism robust with respect to sensitivities of the underlying reliability mechanism?
- Is the mechanism robust with respect to manufacturing and operating variability?
- Are resource (chip area, power) overheads reasonable?
- Is the mechanism tunable to different lifetime bounds?

Electromigration. With EM, applied current can move metal atoms so as to eventually cause an open fault. The semi-empirical Black's Law for electromigration time to failure states that $t_f = (A / J^n) \exp (E_\alpha / kT)$, where t_f is time to failure, J is current density, T is temperature, and the current density exponent n and activation energy E_α are empirical parameters. A basic life-

DAC 2008, June 8–13, 2008, Anaheim, California, USA

[1] One can count clock ticks whenever the chip is in power-on state, and then terminate the chip function when the total number of ticks exceeds a given limit. This requires a non-volatile counter – and even if embedded NVM is available, many semi-invasive 'reset'attacks are known; see, e.g., [4].

[2] One well-studied reliability mechanism, soft-error (SEU), does not appear useful for our purposes. We do not explicitly discuss hot-carrier degradation or thermal runaway, but analogies to what we describe below can be envisioned.

time (in power-on state) bounding approach is to instantiate a population of wire segments along with a lifetime trigger that is a function of the number of failures that have occurred within the population.

(1) With EM, exponential sensitivity requires that the lifetime-bounding IP be placed in a region of the die (e.g., in a corner) that is not subject to large activity-dependent temperature fluctuations. Further, ambient temperature must be (i) accurately predicted, or (ii) actively compensated on-die (e.g., by switching activity and joule self-heating of wires such that the lifetime is well-bounded across the ambient temperature specification.

(2) Robustness to manufacturing variability can be on one hand achieved by spatially separating elements (wires) of the lifetime-bounding IP, for example, by (x,y) location and/or by layer assignment, as well as by topological separation (connection to different portions of supply grid, etc.). Additionally, the population of wires will affect robustness. We propose that time be measured according to failures of short-lifetime wires that are 'cascaded' (in the spirit of [2]) such that one wire begins to fail after its predecessor has completely failed.[3] The lifetime trigger is the sum of hopefully-independent random variables x_i corresponding to the respective lifetimes of individual wires. If each x_i has normal distribution with parameters (μ,σ), then a population of wires will have sum of lifetimes $L = \Sigma_i x_i$. For N wires, we have that $\mu(L) = N\mu$, and that $\sigma^2(L) = N\sigma^2$. Thus, there is a square-root reduction in the μ/σ ratio (and, correspondingly, improvement in tightness of the lifetime bound) as the population N is increased.

(3) Individual wires must be longer than the Blech length, and incremental area resource depends on the desired population (tightness of lifetime bound). Incremental power is determined by the EM degradation of one short-lifetime wire at a time (ref. (2) above).

(4) To retarget a given tapeout to different lifetime bounds, some reconfigurability (bypassing some wires, changing supply voltage) would be needed.

The figure at right shows a trivial implementation of using EM-based wire resistance degradation as a lifetime trigger. The voltage across the resistor R is used to drive a buffer. The buffer output can be made to switch when R reaches a certain value. This implementation does not address issues mentioned above but gives the flavor of circuits that may be used.

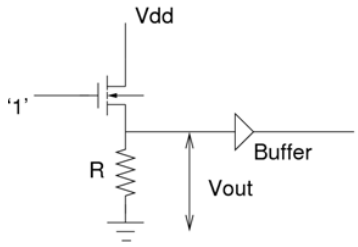

Negative-Bias Temperature Instability (NBTI). With NBTI, PMOS performance can degrade over time. Under static NBTI conditions, 2-5% Ion degradation per year can be assumed [6] and it is likely to worsen with scaling oxide thickness or adoption of different gate oxide materials. Moreover, NBTI is strongly dependent on Vdd, and higher Vdd may be used to worsen NBTI. The temporal dependence of Vt under static NBTI conditions can be expressed as [6]

$$\Delta Vt = (K^2 t^{0.5} + c)^{0.5}$$

[3] An alternative idea is to create a population of identical, relatively long-lifetime wires each with lifetime distribution (μ,σ),then estimate when a certain $(\mu + k\sigma)$ chip lifetime bound has been reached. (E.g., we can estimate a chip lifetime bound of $\mu + 3\sigma$ to have been reached when 99% of the population has failed.) However, implementation seems difficult while maintaining independence (decorrelation) of the wires, and non-use of counters.

The figure below illustrates a simple lifetime trigger circuit to leverage PMOS NBTI effect. The measured current is converted to a voltage (e.g., with a simple resistor or an op-amp based circuit) and then buffered. When the current drops sufficiently over time due to NBTI, the buffer output will switch from 1 to 0. The buffer may be designed to have a low noise margin so that it is sensitive to ~20% input voltage changes. The current-to-voltage conversion can also be used in a variety of ways to tune in the desired lifetime (e.g., keep the initial (time=0) output voltage as close as desired to switching threshold or use a differential amplifier). Another option is to feed the voltage to a forward body bias generator which speeds up hold-critical paths in the design leading to (intermittent) failures. To decrease the power overhead of the circuit while maximizing NBTI degradation, the drain of the PMOS can be tied to VDD under normal operation. Using a power-on-reset and a counter, drain current can be periodically sampled wherein the drain is tied to GND only in the sampling period. To avoid loss of predictability from process variation, instead of using one device, several parallel connected PMOS devices may be used which are strategically placed in different parts of the chip. To avoid Vdd-based fluctuation, the trigger circuit(s) can be placed near the power source (i.e., close to power ring or C4 bumps).

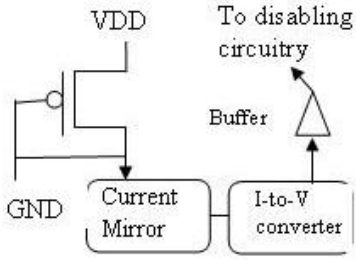

Conclusions.

Finally, we note that any lifetime bounding mechanism will likely be accompanied by an overall savings of area, power and design time due to reduced guardbands in the design process. For example, in 90nm foundry processes the 10-year model for NBTI (PMOS) wearout typically increases NLDM delay table entries by 10-15%. [5] showed an average of 8.7% increase in circuit area to achieve 10-year reliability. Designing to, e.g., a 3-year model that increases delays by 4-5% would reduce this reliability overhead. Our ongoing work, besides pursuing the ideas mentioned in this paper, also investigates possibilities of making counter-based timers usable as secure lifetime triggers. We are also looking into methods to age the circuit even when it is switched off.

References.

[1] D. Poisner, "System for detecting over-clocking uses a reference signal thereafter preventing over-clocking by reducing clock rate", U.S. Patent 6535988, March 18, 2003.

[2] R. Schorsch, "Electronic interval timing device", U.S. Patent 4236145, Nov. 25, 1980.

[3] E. Sakaki et al., "Device for measuring time lapse after turn off of power source and method thereof", U.S. Patent 5500834, March 19, 1996.

[4] R. Anderson and M. Kuhn, "Low cost attacks on tamper resistant devices", M. Lomas et al. (ed.), *Security Protocols*, 5th International Workshop, Paris, April 1997 (Springer LNCS vol. 1361), pp. 125-136.

[5] B.C. Paul et al., "Temporal performance degradation under NBTI: estimation and design for improved reliability of nanoscale circuits", *Proc. DATE* 2006, pp. 780-785.

[6] R. Vattikonda, W. Wang and Y. Cao, "Modeling and Minimization of PMOS NBTI Effect for Robust Nanometer Design", *Proc. DAC 2006,* pp. 1047-1052.