

IBM Research Report

A Built-In BTI Monitor for Long-Term Data Collection in IBM Microprocessors

Pong-Fei Lu, Keith A. Jenkins
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 208
Yorktown Heights, NY 10598
USA



Research Division

Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

A built-in BTI monitor for long-term data collection in IBM microprocessors

Pong-Fei Lu and Keith A. Jenkins
IBM T.J. Watson Research Center
Yorktown Heights, NY

Abstract

A circuit for long-term measurement of bias temperature instability (BTI) degradation is described. It is an entirely on-chip measurement circuit, which reports measurements periodically with a digital output. Implemented on IBM's z196 enterprise systems, it can be used to monitor long-term degradation under real-use conditions. Over 500 days worth of ring oscillator degradation data from customer systems are presented. The importance of using a reference oscillator to measure performance degradation in the field, where the supply voltage and temperature can vary dynamically, is shown.

A built-in BTI monitor for long-term data collection in IBM microprocessors

Pong-Fei Lu and Keith A. Jenkins
IBM T.J. Watson Research Center
Yorktown Heights, NY

1. Introduction

Negative bias temperature instability (NBTI) has long been known as a problem for pFETs, whereby prolonged bias of the gate electrode with respect to the source and drain results in an increase of the magnitude of the threshold voltage of the device, which leads, in turn, to the reduction of its drain current, and the reduction of frequency of circuits [1]. The effect is accelerated at high temperature and high voltage. Similarly, recent technologies using high-k dielectrics and metal gate structures have shown a similar effect, called positive bias temperature instability (PBTI), on the nFET [2-5]. Lumped together, these bias effects, now called BTI, present a reliability challenge for modern CMOS technology.

The change in drain current produced by BTI degradation is expected to be a few percent over a five to ten year span under normal temperature and voltage conditions. Typically CMOS products take this expectation into account and have some built-in margin to allow for a predicted degradation. However, the prediction is based on degradation measured under accelerated testing, followed by an analytical model which extrapolates the measured degradation to normal conditions. Typically, BTI is measured either by direct measurement of the threshold voltage, or by measuring the frequency of simple ring oscillators. Acceleration is achieved by using higher than normal temperature and voltage, so that measurable threshold voltage changes, or ring oscillator frequency change, can be obtained in a matter of hours, rather than years.

Because of the need to perform accelerated testing for lifetime prediction, there is essentially no data on long-term BTI degradation under nominal conditions. The extrapolation from a few hours of accelerated-condition testing to five to ten years of normal use, may, of course, significantly underestimate or overestimate the degradation. In addition, a product circuit may experience variable use conditions, such as changing temperature and supply voltage, and power-off time, which will change the BTI degradation of the devices in a dynamic way. The latter is hard to be captured in the lifetime projection based on conventional BTI model

This paper describes a simple degradation monitor circuit incorporated in product microprocessors, which can periodically report its measurements to a central data collection database, so that long-term degradation under real-use conditions can be measured.

2. Method

2.1 On-chip monitor circuit




Fig 1. (a) Illustration of the measurement sequence of the on-chip BTI monitors. (b) Block diagram of the ring oscillators, control signals, and counters of the monitors.

The basic concept of the BTI monitor is illustrated in Fig. 1a. Rather than measuring threshold voltage of individual FETs [6], the degradation monitor is based on the change of frequency of ring oscillators which have voltage applied to them continuously whenever the chip is powered on. A pair of matched reference and stressed ring oscillators are designed and placed in close proximity. The stressed ring oscillator is constantly running, while the reference ring is powered off except during short measurement times. The stressed ring frequency degrades with time due to device aging effects, whereas the frequency of the reference oscillator does not change due to degradation, but provides a reference which accounts for any overall chip temperature or voltage variations. The degradation is measured by the difference in the two ring frequencies. The measurement time is kept short to avoid excessive BTI recovery of the stressed oscillator and to minimize any stress of the reference oscillator.




Fig. 2. Schematic of the power-off devices of the reference ring oscillators.

Fig. 1b shows the block diagram of the circuits. Two ring oscillators are built with identical circuit blocks. A pFET header switch is used to turn on or off the power supply by a control signal. The stressed ring oscillator is powered continuously, while the reference ring oscillator is powered on only during measurement. An AC test enable

control provides the capability of DC, continuous or variable duty cycle stress. Each ring oscillator has a dedicated on-chip digital counter; and the two rings are measured simultaneously. The distributed pFET header in each stage is designed to be three times the size of the switching pFET (Fig. 2) to minimize delay impact. Since the header in the stressed ring is always on, it experiences the most stressing. Its size needs to be large enough so that the impact of the delay due to the header degradation will not overwhelm the ring oscillator frequency change. SPICE simulations showed a delay change of less than 0.1% in our design with a threshold voltage shift as much as 40 mV.

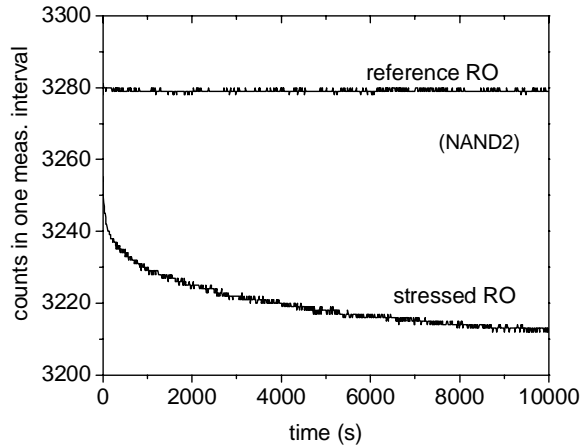


Fig. 3. Demonstration of the effect of accelerated voltage (1.5 V) and temperature (120 °C) stress in a stand-alone prototype of the BTI monitor circuit.

The proposed BTI monitor was first demonstrated in a stand-alone experiment [7] in which stress voltage could be applied to accelerate the degradation. Fig. 3 shows the counter data for two NAND2 ring oscillators in an accelerated voltage and temperature stressing. Under quiet power supply conditions, the reference ring oscillator count is stable for the whole stress period, while the stressed ring oscillator degrades with time.

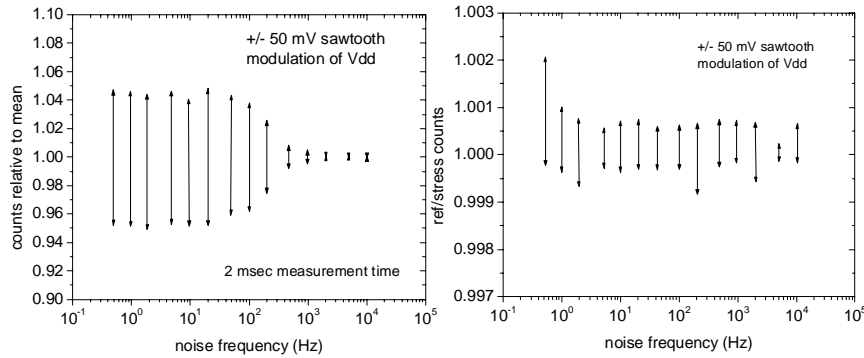


Fig. 4. Demonstration of the effect of using the reference ring oscillator to cancel noise. Noise is applied to the power supply in the stand-alone prototype.

The importance of having a reference ring oscillator is illustrated in Fig. 4. Fig. 4a shows the normalized ring oscillator counts when a saw-tooth noise source was injected into the power supply. A 50 mV modulation at low frequencies can change the ring oscillator frequency by 5%, or about 1% frequency change per 10mV voltage change. (The diminution at high frequency is due to the averaging of the frequency variations during the 2 msec gate time.) For on-chip applications in real systems where the V_{dd} switching noise can be in the order of tens of mV, the uncertainty due to V_{dd} noise can easily overwhelm the BTI effect if the stressed ring oscillator is measured alone. However, by taking the frequency ratio of the two ring oscillators, the impact of the power supply noise is reduced to less than 0.1% (Fig. 4b). Therefore, a reference ring is imperative for noise immunity in on-chip monitors. While the difference between the two rings could be obtained by beat frequency difference, with potentially greater resolution [8-10], the ratio method described here is simple and very compact, and results in sufficient resolution for long-term monitors on product chips. For on-chip monitor in real systems, it is also important to share resource to reduce overhead. The approach described here allows easy integration with other monitors without dedicated hardware, as described below.

2.2 BTI monitor implementation in z196 system




Fig. 5 Block diagram of the implementation of the BTI monitors on an IBM z196 chip.

Built-in BTI monitors are implemented in the IBM zEnterprise 196 (z196) server [10-11] for long-term in-field data collection. The monitors use the same circuits as the stand-alone design described above, with on-chip frequency counters for stressed and reference ring oscillators. As shown in Fig. 5 several BTI monitors consisting of paired reference and stressed ring oscillators are distributed on a chip. They are locally controlled by general register bits which are written via pervasive command bus (PCB) commands in the chip pervasive logic described in [12]. Only one BTI monitor can be enabled and measured at a time by a centralized frequency measurement unit (FMU). Each BTI macro generates a reference and a stressed frequency output. All reference/stressed outputs are ORed together and sent to the corresponding counter in the shared FMU. The BTI monitors are read sequentially by the “select<0:n>” control signal which gates the BTI macro output. The frequency is measured by counting pulses in a fixed timing window. The size of the window is determined by the “start/stop” control in FMU (Fig. 5) based on a preset count of a stable, crystal-based clock. The data acquisition is via firmware which controls the sampling window size and the timing of read and write of the registers.




Fig. 6. Illustration of the placement of the BTI monitors on a central processor (CP) layout [11]. The cores are outlined, and the placement of the five BTI monitors is shown as the small white rectangles.

Each BTI monitor contains ring oscillators with four types of logic: inverter with standard threshold voltage, inverter with high threshold voltage, NAND2, and NOR3, with DC or AC stress mode. The selection of the type of oscillators to be measured is made by additional control signals. The BTI layout needs to be compact for flexible placement on chip. The monitor size is 40x60um in IBM 45nm SOI technology [13]. The z196 96-PU system consists of six central processor (CP) chips and two system controller (SC) chips [10, 11]. There are five BTI monitors in each CP, and five in each SC chip. The placement of the BTI monitors in CP is shown in Fig. 6. There are four in the processor cores which runs at 5.2 GHz (one/per core) , and are the hottest regions of the chip and one near the L2 cache which runs at 2.6 GHz region , and is cooler than the cores [11]. The ring oscillators are unloaded with fan-out (FO) of one. The ring oscillator frequencies are about 1GHz, and are divided by four at the macro output to generate a signal of 250-300 MHz, to match the FMU requirements.

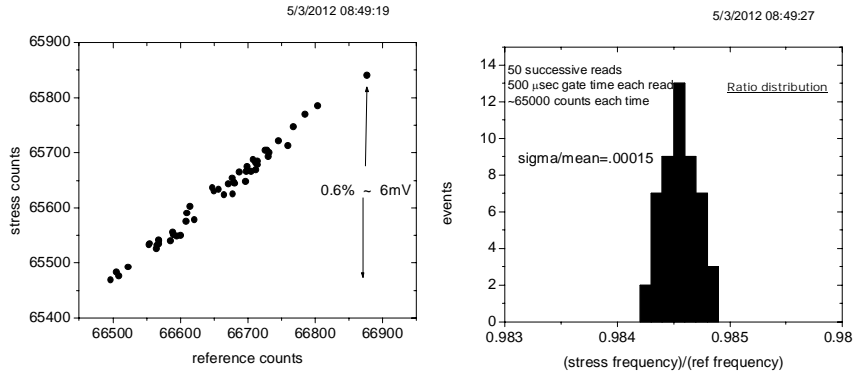


Fig. 7. Demonstration of the measurement resolution and immunity to noise with the BTI monitors on the z196 processor. The highly correlated stressed and reference ring counts in (a) showed a variation of 0.6%; the count ratio in (b) showed a tight distribution due to the noise cancellation.

The current implementation used a setting of 8192 counts of an FMU counter clock frequency at 16 MHz to define the measurement window (i.e. ~ 0.5 ms). The BTI monitors are measured once in a week. The measurement time amounts to 0.26 sec of total on-time of the reference ring oscillator in ten years. To determine the accuracy of the measurement, Fig. 7a shows the counts of 50 consecutive reads, with a few milliseconds apart, for the reference and stressed ring oscillators in a calibration experiment (measurement time ~ 0.25 ms). The variation is about 0.6% in either ring counts, presumably due to small V_{dd} variations (~ 6 mV), whereas the ratio between the two shows a much tighter distribution due to the cancelling of the common noise (Fig. 7b). The sigma-to-mean ratio of 0.015% is about 40X reduction of the single ring count fluctuation of 0.6%.

3. Results

The z196 system reports measurement results when first started, and thereafter on a weekly period. Data are available for many running systems, either in the test facility or at customer sites. All the ring oscillators of all BTI monitors are measured, in sequence, at each weekly reading. The time-stamped raw contents of the counters are available in a central database, and conversion to frequency is trivial. An example of frequency of an inverter ring oscillator in AC mode near the L2 cache region of a customer system taken over a period 150 days is shown in Fig. 8.




Fig. 8. Illustration of raw and ratio measurements of the ring oscillators on a z196 system running customer workloads.

Fig. 8a shows the separate frequencies of the stressed and reference inverter ring oscillators. Variations in both are seen, but when the ratio of stressed frequency to reference frequency is made, shown in Fig. 8b, the common-mode noise of the data is removed, and the stressed frequency shows a gradual decrease, as expected. There are a few spurious points, but otherwise the decrease is monotonic. The ratio clearly shows the importance of using a reference signal to remove common noise in order to detect a small change.




Fig. 9. Long-term measurements of the frequency degradation of ring oscillators on two different z196 systems.

Longer term measurements are shown in Fig. 9. In this graph, frequency *degradation* relative to the first measured point is plotted so that the plots resemble the *increase* of threshold voltage which causes the frequency reduction. Data from ring oscillators on two customer systems are shown. They both qualitatively have the appearance of the familiar power law dependence. Because of a few spurious measurements, the initial point is somewhat uncertain. Nonetheless it is very clear that the system 1 ring oscillator degrades at a faster rate than that of system 2. This difference is entirely consistent with the difference in voltage and temperatures of the two systems. System 1 is operated at a voltage which is about 80mV higher than that of System 2. (In z196 the voltage for each chip is individually tuned to compensate for process variations [14]) The measured temperature, in the center of the chip near the L2 cache where the ring oscillators are placed, is shown in Fig. 10. The average temperature in System 1 is about 13 degrees hotter than System 2 due to the difference in cooling [14]. These differences are consistent with the accelerating effect of temperature and voltage on BTI degradation.




Fig. 10. Measured temperatures of the two systems of Fig. 9.

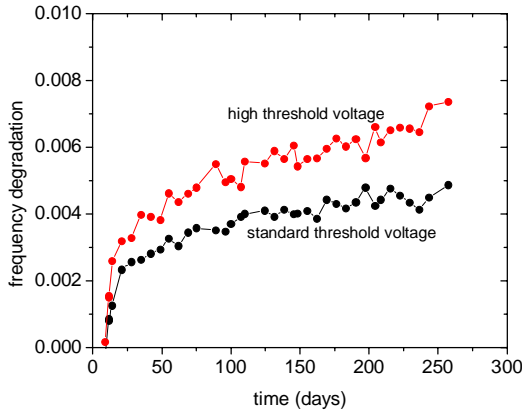


Fig. 11. Illustration of frequency degradation of two types of ring oscillators in the same BTI monitor.

Another example which demonstrates the utility of the measurement circuit, and which corroborates our understanding of BTI effects, is shown in the data in Fig. 11. Plotted in this figure is the reported frequency degradation of two inverter ring oscillators from the same NBTI monitor on the same core. The inverter with higher threshold voltage devices shows more degradation than the one with regular threshold voltages. This is consistent with the understanding that threshold voltages change by about the same amount, but the drive current, and hence the frequency, is more affected by higher threshold devices because of their lower overdrive voltages. The faster aging should be taken into consideration when using higher threshold voltage devices for power reduction.

4. Discussion and analysis

Because of the anticipated variable use conditions in active computer systems, it is not the purpose of these monitors to provide a long-term measurement suitable for modeling BTI effects. Rather, it ensures that the performance margin, which IBM z196 systems have, and which is based on accelerated testing, is neither too small or unnecessarily large. However, it is possible to establish consistency with the models which have been derived from accelerated testing and physical mechanism. In the case of system 1, where the temperature is fairly constant, such a comparison can be made, whereas the system 2 temperature is clearly too variable for a simple analysis.

BTI degradation is generally believed to follow a power-law time-dependence,

$$\Delta f = Ag(T)h(V)t^n \quad (1)$$

where the exponent, n , is predicted to be about 0.16, based on a reaction-diffusion model of hydrogen-silicon bond breaking [1,15]. The functions $g(T)$ and $h(V)$ are temperature and voltage acceleration functions, which affect the magnitude of the degradation, *independent* of the power-law time dependence. Thus, it is expected that the data

obtained by these on-chip BTI monitors will demonstrate the power-law time dependence if the temperature and voltage are constant in the systems.

Simplifying Eq. (1) to remove the acceleration factors results in

$$\Delta f = Bt^n \quad (2)$$

This formulation assumes that the measurement of degradation begins at a known starting time, t_o , before any voltage has been applied to the devices. In general, it is difficult to make a measurement before some degradation has occurred, which leads to erroneous measurement of the exponent [17]. In a complex customer system, such as described here, it is even more difficult, as the chips have been burned in and tested before installed in a system. Thus, this ‘time-zero’ is completely unknown, and the measured degradation will follow the form

$$\Delta f = +B'(t + t_o)^n - Bt_o^n \quad (3)$$

This form can be compared to the measured data shown in Fig. 9. The curve superimposed on the system 1 data in that plot is a fit of the form of Eq. (3), and, ignoring the spurious data, presents very good fit agreement with measured degradation when time-zero is unknown. (The exponent is 0.174 in this fit.) Thus the data collected over 500 days exhibits the same power law dependence on time as accelerated stress data measured over a few hours or less. Such an agreement is very strong corroboration of the power-law dependence from which lifetime predictions are made [1,17]. However, it is noted that the fit parameters obtained from Fig. 9 are not used for such predictions: the high interdependency of the parameters of Eq. (3) can result in many good fits with quite different fit parameters, depending on starting values and constraints in the fitting procedure.

It is also noted that the measurement circuit describe here measures *total* degradation of ring oscillator frequency. Thus hot-carrier degradation, for example, would be indistinguishable from BTI degradation. However, the fit to Eq. (3), and the temperature and voltage dependence of Fig. 9, suggest that BTI is the dominant mechanism in this technology. HCI is known to be insensitive to temperature [9]. Also, since these rings are unloaded with FO=1, the internal switching is fast and the hot carrier injection (HCI) effect, which happens during the time when the device is fully on, is minimal [18,19]. Other ring oscillator topologies have been proposed which can separate hot carrier effects from BTI effects [9]. The measurement and reporting infrastructure described here can easily incorporate any desired ring oscillator type in order to measure these effects over a long time under normal use conditions.

5. Conclusion

The paper has described a built-in BTI monitor in IBM z196 system for long-term field data collection. Over 500 days worth of data were shown. The importance of using a reference ring for noise cancellation for an on-chip monitor in systems under use

conditions has been shown. The in-field data collection capability is important for verifying the lifetime BTI guard-band factored into system design. The monitor design is simple and light-weighted. It can be integrated seamlessly with IBM z196 Pervasive infrastructure, and is versatile through firmware support. Although the field BTI degradation data has a power law dependency of time consistent with that of prevailing BTI model, detailed comparison of the power exponent is difficult due to the “time-zero” problem of the devices that go into the products. Many modern microprocessors are designed with advanced low-power techniques such as power-gating and dynamic voltage and frequency scaling [20,21]. The chip temperature and the supply voltage can vary widely depending on the workload, which also makes it hard to project circuit performance degradation using transistor-based model. However, the on-chip BTI monitor provides a valuable measurement of the performance degradation, and can be used to reduce unnecessary headroom for improved performance, and may be used to provide an alarm when degradation reduces margin to an unacceptable level, or be integrated with on-chip system management to enhance the reliability of computer systems.

Acknowledgements

The authors would like to thank Tobias Webel for support and consultation regarding z196 pervasive infrastructure, Birgit Schubert for FMU design, and Oliver Marquardt and Ralf Schaufler for firmware support for the BTI monitor.

References

- [1] J. H. Stathis, S. Zafar, “The Negative Bias Temperature Instability in MOS devices: A Review.” *Microelectronics and Reliability*, Vol. 46, pp. 270-286, Feb.-Apr. 2006.
- [2] S. Zafar et al., “Threshold voltage instabilities in high- κ gate dielectric stacks”, *IEEE Transactions on Device and Material Reliability*, Vol. 5, No.1, pp. 45-64, Mar. 2005.
- [3] D. P. Ioannou et al., “PBTI response to interfacial layer thickness variation in Hf-based HKMG nFETs”, pp. 1044-1048, *International Reliability Physics Symposium 2010*.
- [4] D.P. Iannou et al., “Burn-in stress induced degradation and post-burn-in high temperature anneal (bake) effects in advanced HKMG and Oxynitride based CMOS ring oscillators”, 5C.2.1-5C.2.5, *International Reliability Physics Symposium 2012*.
- [5] J.-J. Kim et al, “Reliability monitoring ring oscillator structures for isolated/combined NBTI and PBTI measurement in high-k metal gate technologies,” 2B.4.1-2B.4.4, *International Reliability Physics Symposium 2011*
- [6] R. Carlsten, J. Ralston-Good and D. Goodman, “An Approach to Detect NBTI in Ultra-Deep Submicron Technologies,” *International Symposium on Circuits and Systems*, 2007, pp. 1257-1260.

- [7] K. Stawiasz et al., "On-chip circuit for monitoring frequency degradation due to NBTI," pp. 532-535, International Reliability Physics Symposium 2008.
- [8] T.-H. Kim et al., "Silicon odometer: an on-chip reliability monitor for measuring frequency degradation of digital circuits", IEEE Journal of Solid-State Circuits, Vol. 43, No. 4, pp. 874-880, Apr. 2008.
- [9] J. Keane et al., "An all-in-one silicon odometer for separately monitoring HCI, BTI and TDDDB", IEEE Journal of Solid-State Circuits, Vol. 45, No. 4, pp. 817-829, Apr. 2010.
- [10] B. Curran et al., zEnterprise 196 system and microprocessor," IEEE Micro, Vol. 31, No. 2, pp. 26-40, Mar/Apr. 2011.
- [11] F. Busaba et al., "IBM zEnterprise 196 microprocessor and cache subsystem", IBM Journal of Research & Development, Vol. 56, No. 1/2, paper 1, Jan/Mar. 2012.
- [12] T. Webel et al, "Scalable and modular pervasive logic/firmware design", IBM Journal of Research & Development , Vol. 56, No. 1/2, paper 5, Jan/Mar 2012.
- [13] S. Narasimha et al., "High-performance 45-nm SOI technology with enhanced strain, porous low-k BEOL, and immersion lithography," pp. 1-4, International Electron Devices Meeting 2006.
- [14] M. Andes et al., "IBM zEnterprise energy management", IBM Journal of Research & Development, Vol. 56, No1/2, paper 12, Jan/Mar 2012.
- [15] A.-E. Islam et al, "Recent issues in Negative-Bias Temperature Instability: Initial degradation, field dependence of interface trap generation, hole trapping effects, and relaxation," IEEE Transactions on Electron Devices, Vol. 54, No. 9, pp2143-2154, Sept 2007.
- [16] Linder, B.P, Kim, J.-J., Rao, R., Jenkins, K., and Bansal, A. , "Separating NBTI and PBTI effects on the degradation of Ring Oscillator frequency," *Integrated Reliability Workshop Final Report (IRW), 2011 IEEE International* , vol., no., pp.1-6, 16-20 Oct. 2011.
- [17] Velamala, J.B., Sutaria, K.B., Sato, T., and Yu Cao , "Aging statistics based on trapping/detrapping: Silicon evidence, modeling and long-term prediction," *2012 International Reliability Physics Symposium (IRPS)* , vol., no., pp.2F.2.1-2F.2.5, 15-19 April 2012
- [18] K. Hoffman et al., "Highly accurate product-level aging monitoring in 40nm CMOS", Symposium on VLSI Technology Digest of Technical Papers, pp. 27-28, 2010.

[19] C. Schlunder et al., "HCI vs. BTI? – Neither one's out", pp. 2F.4.1-2F.4.6, International Reliability Physics Symposium 2012.

[20] R. Kalla et al, "POWER7: IBM's next-generation server processor," IEEE Micro, Vol. 30, issue 2, pp. 7-15, Mar/Apr., 2010.

[21] M. Ware et al., "Power-performance management on an IBM POWER7 server," pp. 201-206, International Symposium on Low Power Electronics and Design 2010.