

# Energy Efficient SEU-Tolerance in DVS-Enabled Real-Time Systems through Information Redundancy

## Abstract

Concerns about the reliability of real-time embedded systems that employ dynamic voltage scaling has recently been highlighted [1,2,3], focusing on transient-fault-tolerance techniques based on time-redundancy. In this paper we analyze the usage of information redundancy in DVS-enabled systems with the aim of improving both the system tolerance to transient faults as well as the energy consumption. We demonstrate through a case study that it is possible to achieve both higher fault-tolerance and less energy using a combination of information and time redundancy when compared with using time redundancy alone. This even holds despite the impact of the information redundancy hardware overhead and its associated switching activities.

## 1 Introduction

Real-time embedded systems that are employed in defense, space, and consumer applications often have both *energy constraints* and *fault-tolerance requirements*. To address these two issues, dynamic voltage scaling (DVS) and time redundancy are often used. DVS is a popular system-level low-power design technique [4,5], whilst time redundancy is an effective technique to achieve tolerance to transient faults in real-time embedded systems [3,6]. Both techniques require slack time in the system schedule to achieve their goals, i.e. DVS reduces energy by lowering the system operating voltage and frequency, whilst time redundancy improves transient-fault tolerance by performing a number of recovery executions depending on the available slack. If more slack is given to DVS to save more energy, less slack is left for transient-fault-tolerance, and vice versa. This means that there is a resource conflict between DVS and time-redundancy-based fault-tolerance on slack-time which is a limited resource. Furthermore, DVS-enabled systems are more susceptible to transient faults or Single Event Upsets (SEU) (Bit-flips due to the impact of particles on flip-flops), since the rate of SEUs in such systems increases exponentially as supply voltage decreases [3,7]. Such faults have become the major source of concern due to the continuing technology shrinkage [2,8].

The trade-off problem between fault-tolerance and energy consumption in DVS-enabled systems has recently been highlighted [3] and become subject to investigations [1,2]. Non-uniform checkpoint placement policies for the combined purpose of conserving energy and providing fault-tolerance have been proposed in [1]. The technique proposed in [2] uses an adaptive check-pointing scheme to achieve fault-tolerance and energy saving in a unified manner. Although both techniques [1,2] are effective in achieving fault tolerance, the obtained energy savings are limited due to the fact that the time redundancy requires slack time – slack that otherwise could be exploited through DVS to reduce the energy consumption.

As opposed to these approaches, we propose in this paper the usage of information redundancy in fault-tolerant DVS-enabled systems. The aim of using information redundancy is to decouple the fault-tolerance from the slack time and hence to provide more slack to DVS without degrading the fault-tolerance capability of the system. To the best of our knowledge, this paper is the first attempt that addresses energy management through DVS and fault-tolerance through information-redundancy in conjunction.

The rest of the paper is organized as follows. The system, performability and energy models are presented in Section 2. Section 3 compares the performability and the energy consumption of the proposed approach (which uses both time-redundancy and information-redundancy) and the conventional approach (which solely uses time-redundancy), using a set of experiments. Section 4 concludes the paper.

## 2 Preliminaries

### 2.1 System models

In this paper, we will compare and analyze two types of DVS-enabled real-time systems, defined as follows:

**(a) Conventional R system:** This represents a DVS-enabled system which uses pure rollback-recovery, i.e., the conventional time-redundancy based approach (Fig. 1a). In this system, whenever transient faults (i.e. SEUs) occur during the task execution, a recovery execution (re-execution) of the same task is required [3,6]. For instance, as shown in Fig. 1a, during the original task execution three SEUs (2 SEUs in the same clock cycle and 1 single SEU) cause a faulty run, hence necessitating a recovery execution (recovery execution 1). Such executions have to be performed until a non-faulty run happens (e.g. recovery execution 2 in Fig. 1a). In order to achieve a certain degree of fault-tolerance (performability) it is necessary to reserve some system time for recovery executions (slack for recoveries), while the remaining slack until the task deadline  $D$  can be exploited via DVS to reduce the system's energy dissipation.

**(b) Proposed RI system:** These DVS-enabled systems use both rollback-recovery and information redundancy, i.e., the fault-tolerance is achieved through recovery executions as well as through redundant information that can be used to correct faults during execution (i.e. without necessitating a re-execution). Consider Fig. 1b, which demonstrates this approach using the same SEUs as in Fig. 1a. As we can observe, whenever one SEU occurs during a single clock cycle (first and third faults in Fig. 3), the resulting error can be corrected by some additional hardware which is used for information redundancy. Faults that require a recovery execution occur only if two or more SEUs happen during a single clock cycle (for instance, second fault during the original execution in Fig. 1b). Accordingly, the number of necessary recoveries is reduced leaving more exploitable slack to DVS. Suppose a task and its recoveries run at the same frequency  $f$ . Let  $N$  be the number of clock cycles which are needed to execute the task,  $D$  be the deadline (in seconds), and  $\rho_f$  be the probability of having a faulty run. Then, the task execution time is  $N/f$  seconds, and the amount of total slack time is  $D - (N/f)$ . If the original execution fails, the first recovery is executed with the probability of  $\rho_f$ . Similarly, the  $i^{\text{th}}$  ( $i \leq K$ ) recovery will be executed with probability  $\rho_f^i$ . Thus, the expected time required for executing  $K$  recoveries is:

$$T_{K\text{-recoveries}} = \frac{N}{f} \sum_{i=1}^K \rho_f^i \quad (1)$$

Therefore the slack time which is left for DVS is:

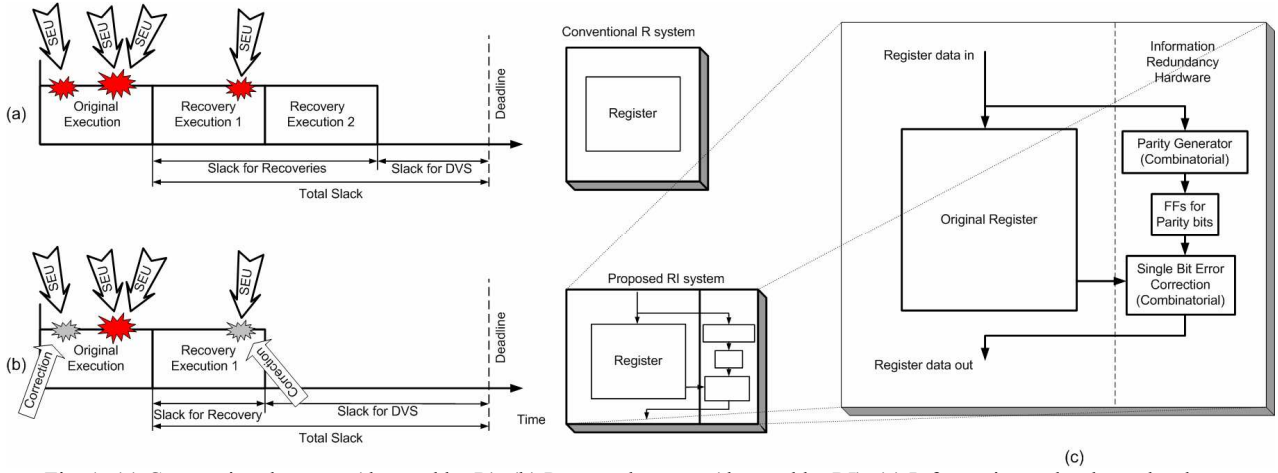


Fig. 1: (a) Conventional system (denoted by R), (b) Proposed system (denoted by RI), (c) Information redundancy hardware

$$T_{DVS} = (D - \frac{N}{f}) - T_{K-recoveries} = D - \frac{N}{f} \sum_{i=0}^K \rho_f^i \quad (2)$$

DVS can use the slack time  $T_{DVS}$  to save energy. It can be seen from Eq. (2) that as  $\rho_f$  (i.e. the probability of having a faulty run) decreases,  $T_{DVS}$  increases. The use of information redundancy decreases  $\rho_f$ , so that  $T_{DVS}$  increases and more slack time becomes available to save energy (compare Fig. 1a and 1b).

Information redundancy in the proposed RI system is obtained by adding some additional hardware to the conventional circuit, as shown in Fig. 1c. This hardware comprises a parity generator (produces parity bits, e.g. Overlapping parity bits [9,10]), flip-flops to store the parity bits, and a single bit error corrector which restores the affected registers to the original content as long as only one bit is corrupted. We will demonstrate in Section 3 that the extra energy associated with the additional hardware can be overcompensated by DVS (because of the  $T_{DVS}$  increase), i.e. the RI systems can yield higher energy savings when compared to the conventional R systems.

## 2.2 Performability Model

In this paper the fault-tolerance capability of DVS-enabled real-time systems is measured by the performability criterion defined as the probability of finishing the task correctly within its deadline in the presence of faults [3,11]. Using this criterion, this section presents an analysis for both the proposed RI and conventional R systems.

In DVS-enabled systems, reducing the supply voltage of a digital circuit requires the reduction of the frequency in order to ensure correct operation. When the conventional R system runs at supply voltage  $V_R$ , the operational frequency can be expressed as [12]:

$$f_R(V_R) = (L_{dR} K_6)^{-1} ((1 + K_1)V_R + K_2 V_{bs} - V_{th1})^\alpha \quad (3)$$

where  $L_{dR}$  is the logic depth of the critical path,  $V_{th1}$ ,  $K_1$ ,  $K_2$ , and  $K_6$  are constants for given process technology,  $V_{bs}$  is the body-to-source voltage, and  $\alpha$  is a measure of velocity saturation whose value has been approximated to be 1 [12].

This paper proposes to use information redundancy which requires some extra hardware logic to process the redundant

information. Suppose that because of the extra hardware logic, the depth of the critical path of the proposed RI system is  $K_c$  times the depth of the critical path of the conventional R system, i.e.  $L_{dRI} = K_c \cdot L_{dR}$ , then the operational frequency of the proposed RI system is:

$$f_{RI}(V_{RI}) = (K_c L_{dR} K_6)^{-1} ((1 + K_1)V_{RI} + K_2 V_{bs} - V_{th1})^\alpha \quad (4)$$

Also, for DVS-enabled systems, the fault rate is determined by the system supply voltage [3]. The arrival process of particle-induced faults (i.e. SEUs) is typically modeled as a Poisson process with an average fault rate  $\lambda$  [2,3,13]. Suppose the supply voltage of the conventional R system can be changed between  $V_{min}$  and  $V_{max}$ . Let  $\lambda_0$  be the fault rate corresponding to  $V_{max}$ , and  $\lambda_0 10^d$  be the fault rate corresponding to  $V_{min}$  (i.e. fault rate at  $V_{min}$  is  $10^d$  higher than fault rate at  $V_{max}$ ). Based on the fault-rate model proposed in [3], when the conventional R system runs at supply voltage  $V_R$ , its fault rate can be expressed as:

$$\lambda_R(V_R) = \lambda_0 10^{\frac{d(V_{max}-V_R)}{V_{max}-V_{min}}} \quad (5)$$

In this paper, it is assumed that  $d=1$  and  $\lambda_0=10^{-6}$  FPS (faults per second), which means that fault rate at the minimum voltage is  $10^{-5}$  FPS and at the maximum voltage is  $10^{-6}$  FPS, which are typical fault rates for particle-induced faults [3].

The use of information redundancy requires some extra flip-flops to store the redundant bits. However, as the number of the flip-flops increases, the rate at which the flip-flops are hit by particles increases linearly [8]. Suppose that because of the redundant bits, the number of the flip-flops of the proposed RI system is  $K_{FF}$  times the number of the flip-flops of the conventional R system, then the fault rate of the proposed RI system is:

$$\lambda_{RI}(V_{RI}) = K_{FF} \lambda_0 10^{\frac{d(V_{max}-V_{RI})}{V_{max}-V_{min}}} \quad (6)$$

Since the particle-induced faults follow a Poisson distribution, in the conventional R system, the probability of having a faulty run (at least one SEU during one of the clock cycles) of the task is:

$$\rho_{f_R} = 1 - e^{-\frac{\lambda_R(V_R) \cdot N}{f_R(V_R)}} \quad (7)$$

In this case, the maximum number of possible recoveries is:

$$k_{f_R} = \left\lfloor \frac{D \cdot f_R(V_R)}{N} \right\rfloor - 1 \quad (8)$$

Based on Eq. (7) and Eq. (8), the performability of the conventional R system is:

$$R_{f_R} = 1 - \rho_{f_R}^{K_{f_R}+1} = 1 - (1 - e^{-\frac{\lambda_R(V_R) \cdot N}{f_R(V_R)}})^{\left\lfloor \frac{D \cdot f_R(V_R)}{N} \right\rfloor} \quad (9)$$

As mentioned in Section 2.1, the proposed RI system has a faulty run if more than one SEU (at least two SEUs) occurs during a clock cycle. Therefore, based on Poisson distribution, the probability of having a faulty run can be expressed as:

$$\rho_{f_{RI}} = 1 - (1 + \frac{\lambda_{RI}(V_{RI})}{f_{RI}(V_{RI})})^N \cdot e^{-\frac{\lambda_{RI}(V_{RI}) \cdot N}{f_{RI}(V_{RI})}} \quad (10)$$

Hence, the performability of the proposed RI system is:

$$R_{f_{RI}} = 1 - \rho_{f_{RI}}^{K_{f_{RI}}+1} = 1 - \left[ 1 - (1 + \frac{\lambda_{RI}(V_{RI})}{f_{RI}(V_{RI})})^N \cdot e^{-\frac{\lambda_{RI}(V_{RI}) \cdot N}{f_{RI}(V_{RI})}} \right]^{\left\lfloor \frac{D \cdot f_{RI}(V_{RI})}{N} \right\rfloor} \quad (11)$$

Eq. (9) and Eq. (11) will be used in Section 3 to compare the performabilities of the conventional R system (based on time-redundancy only) and the proposed RI system (i.e. the proposed approach based on the combination of time and information redundancy).

It is important to note that the performability of both the conventional R system and the proposed RI system increase with increasing supply voltage (and consequently increasing operational frequency), since more recovery executions can be performed within the task deadline. However, the performability of the RI system is in general better than the R system when the same supply voltage is used. This is due to the fact that the additional information redundancy in the RI system, which does not require slack for any recovery execution, covers one SEU per clock cycle, hence leaving more slack for recoveries. This aspect will be clarified in Section 3.

### 2.3 Energy Consumption Model

The energy consumption per cycle of the conventional R system is [12]:

$$E_{cyc_R} = \underbrace{C_{eff} V_R^2}_{Dynamic \ Energy} + \underbrace{\frac{L_{g_R}}{f_R(V_R)} (V_R K_3 e^{K_4 V_R} e^{K_5 V_{bs}} + |V_{bs}| I_j)}_{Static \ Energy} \quad (12)$$

where  $C_{eff}$  is the average switched capacitance per cycle for the whole circuit,  $L_{g_R}$  is the number of the logic gates in the circuit,  $K_3$ ,  $K_4$  and  $K_5$  are constant parameters and  $I_j$  is the current due to junction leakage.

As mentioned in Section 2.1, in the proposed RI system, some extra hardware logic is needed to process the redundant information. Suppose that because of the extra hardware, the number of gates in the proposed RI system is  $K_a$  times the number of gates in the conventional R system, i.e.  $L_{g_{RI}} = K_a \cdot L_{g_R}$ . Let  $C_{eff\_extra}$  be the average switched capacitance per cycle for this extra hardware logic, the energy consumption (per cycle) of the proposed RI system is:

$$E_{cyc_{RI}} = (C_{eff\_extra} + C_{eff}) V_{RI}^2 + \frac{K_a \cdot L_{g_R}}{f_{RI}(V_{RI})} (V_{RI} K_3 e^{K_4 V_{RI}} e^{K_5 V_{bs}} + |V_{bs}| I_j) \quad (13)$$

As mentioned in Section 2.1, both the conventional R and the proposed RI systems use rollback-recovery, i.e. after a faulty run the task has to be re-executed. Such recovery executions consume, just like the original execution, energy. Therefore, to analyze the energy consumption of the proposed RI and conventional R systems, the expected value of energy consumption should be considered. The expected energy consumption is [3]:

$$EE = N \cdot E_{cyc} \sum_{i=0}^{k_f} \rho_f^i = N \cdot E_{cyc} \frac{1 - \rho_f^{k_f+1}}{1 - \rho_f} \quad (14)$$

where  $E_{cyc}$  is given either by Eq. (12) or Eq. (13), depending on which system type is considered.

According to Eqs. (12)-(14), if the conventional R system and the proposed RI system operate at the same supply voltage, the RI system will show higher energy consumption than the R system. However, it is important to note that the RI system has a much better performability than the R system at the same voltage setting, so that it is possible to lower the supply voltage of the RI system via DVS to achieve less energy dissipation than the R system, even though the RI system still provides better performability than the R system.

## 3 Case Study and Experimental Results

In this section we will validate the efficiency and applicability of the proposed approach (i.e. based on the combination of time and information redundancy) as compared to the conventional approach (i.e. based on time-redundancy alone). For this purpose we have performed a Crusoe processor case study as well as some experiments using several ITC'99 benchmarks. Section 3.1 compares the performability and energy dissipation of the conventional R and the proposed RI systems based on the Crusoe processor. Section 3.2 investigates the influence of hardware overhead on the suitability of the proposed approach. Section 3.3 presents synthesis results to clarify the typical hardware overhead required in realistic benchmark circuits.

### 3.1 Case Study: Crusoe Processor

In this section it is demonstrated through a case study that it is possible to achieve both higher performability and less energy consumption using a combination of information and time redundancy techniques (proposed RI system) when compared to using time redundancy alone (conventional R System). We use as a case study a Transmeta Crusoe processor implemented in 0.18 $\mu$ m CMOS technology, for which implementation-relevant parameters are given in [12,14]. These parameters comprises the following constants needed for the evaluation of performability and energy (Eqs. (3)-(14)):  $K_1 = 0.053$ ,  $K_2 =$

0.140,  $K_3 = 3.0 \times 10^{-9}$ ,  $K_4 = 1.63$ ,  $K_5 = 3.65$ ,  $K_6 = 51 \times 10^{-12}$ ,  $V_{bs} = 0$  V,  $V_{th1} = 0.359$  V,  $C_{eff} = 1.11 \times 10^{-9}$  F,  $L_d = 37$ ,  $L_g = 4 \times 10^6$ ,  $I_j = 2.40 \times 10^{-10}$  A.

As an example, a task with a worst-case execution time of  $N=40000$  clock cycles and a deadline at  $D=0.5$ ms is considered here. For this example, the deadline allows 7 recovery executions of the whole task (with  $N=40000$ ) at  $V_{dd}=1.6$ V. Furthermore, for the RI system we assume a hardware overhead as well as increased switching activity of 100% (i.e.  $K_a=2$ ,  $K_{FF}=2$ ,  $C_{eff\_extra}=C_{eff}$ ), and a critical path depth increase of 10% ( $K_c=1.1$ ). This assumption will be generalized in Section 3.2.

First we will investigate the performability and expected energy consumption of the two system types (R and RI) when changing the supply voltage. Suppose one wants to know at which supply voltages, the proposed RI system provides better performability than the conventional R system. To do this, one has to solve the inequality:

$$R_{f_{RI}}(V_{RI}) > R_{f_R}(V_R) \quad (15)$$

Because of the complexity of Eq. (9) and Eq. (11), a numerical method has been used to solve this inequality. Fig. 2, shows the solution. In this figure, the curve (between the shaded and exposed areas) is the geometrical locus of the points at which the performabilities of the two systems are identical and the region above the curve (the shaded area) is the geometrical locus of the points at which the proposed RI system provides better performability than the conventional R system (i.e. the solution of the inequality).




Fig. 2: Solution of the performability inequality

For example, consider the points labeled A, and B in Fig. 2. Point A indicates that when the supply voltage  $V_R=1.5$  V is applied to the conventional R system and the supply voltage  $V_{RI}=1.2$  V is applied to the proposed RI system, the performability of the proposed RI system is better than the performability of the conventional R system (desirable from the fault-tolerance point of view). Point B indicates that when the supply voltage  $V_R=2.5$  V is applied to the conventional R system and the supply voltage  $V_{RI}=1.1$  V is applied to the proposed RI system, the performability of the proposed RI system is less than the performability of the conventional R system (undesirable from the fault-tolerance point of view).

Now, suppose one wants to know at which supply voltages, the proposed RI system provides better expected energy consumption than the conventional R system. To do this, one has to solve the inequality:

$$EE_{RI}(V_{RI}) < EE_R(V_R) \quad (16)$$

Again, because of the complexity of Eq. (7), Eq. (10), Eq. (12), Eq. (13) and Eq. (14), a numerical method has been used to solve this inequality. Fig. 3, shows the solution. In this figure, the curve (between the shaded and exposed areas) is the geometrical locus of the points at which the expected energies of the two systems are identical. The region below the curve (the shaded area) is the geometrical locus of the points at which the proposed RI system provides better expected energy consumption than the conventional R system (i.e. the solution of the inequality).

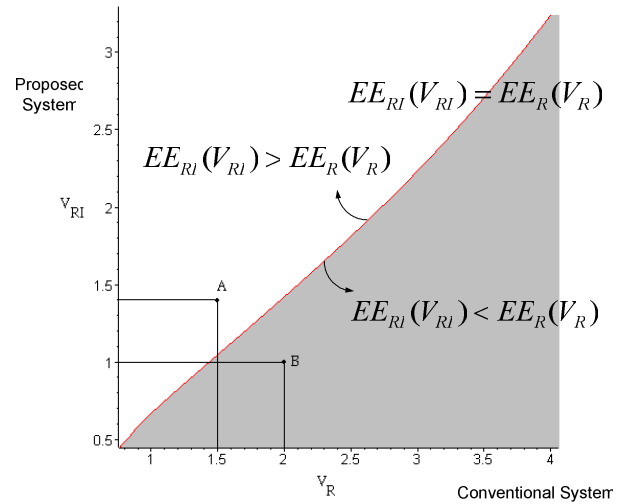


Fig. 3: Solution of the expected energy inequality




Fig. 4: Solution of the expected energy inequality and the performability inequality

For example, consider the points labeled A and B in Fig. 3. Point A indicates that when the supply voltage  $V_R=1.5$  V is applied to the conventional R system and the supply voltage  $V_{RI}=1.4$  V is applied to the proposed RI system, the energy consumption of the proposed RI system is greater than the energy consumption of the conventional R system (undesirable from the energy consumption point of view). Point B indicates that when the supply voltage  $V_R=2.0$  V is applied to the conventional R system and the supply voltage  $V_{RI}=1.0$  V is applied to the proposed RI system, the energy consumption of the proposed RI system is less than the energy consumption of

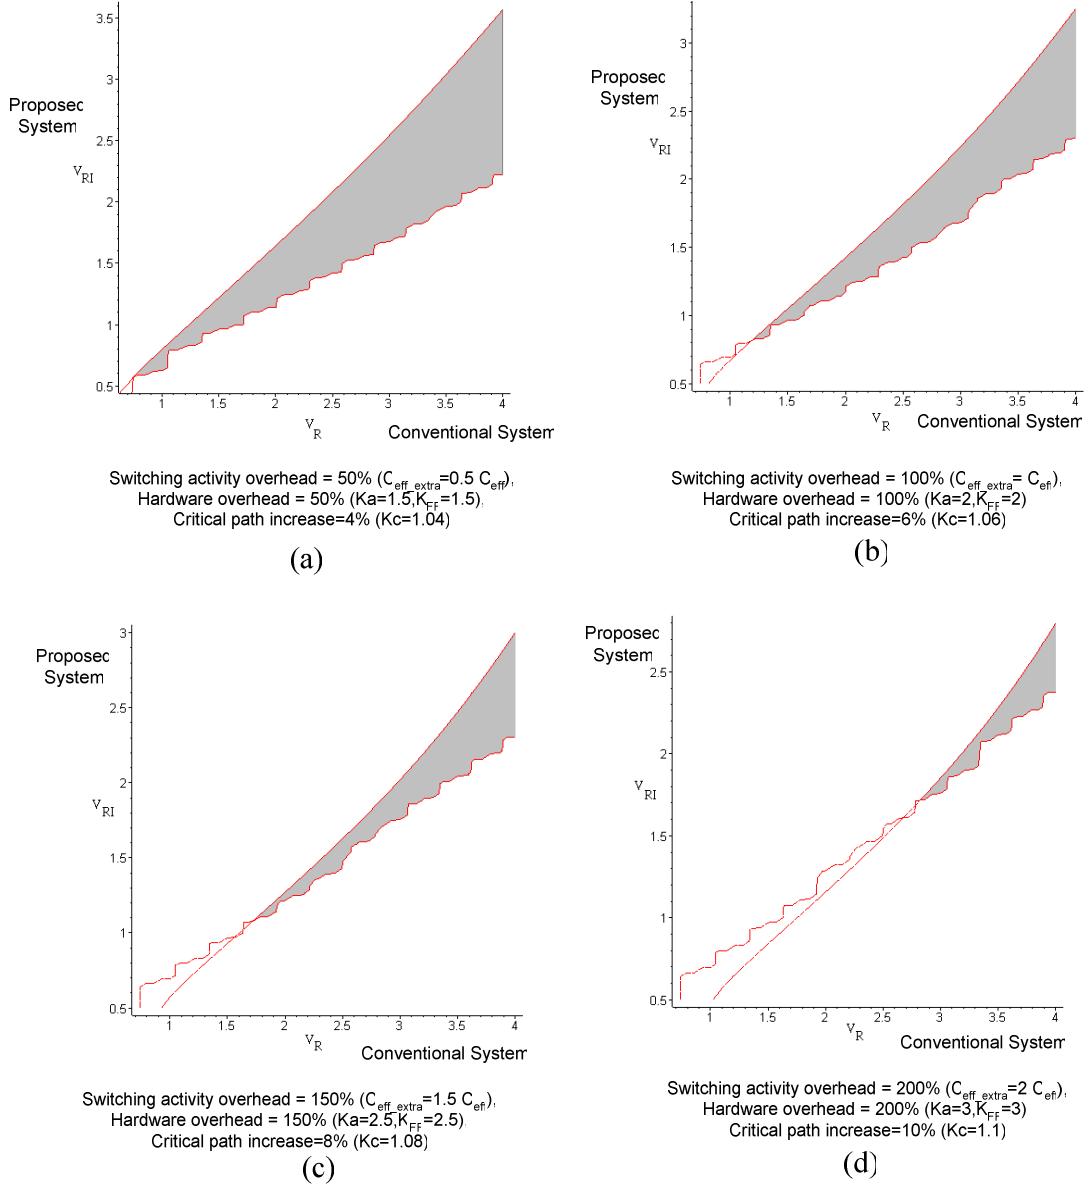


Fig. 5: Impact of information redundancy hardware

the conventional R system (desirable from the energy consumption point of view).

Although Figs. 2 and 3 provide some interesting insight into the performability and energy trade-offs between the R and RI systems, it is not directly apparent in which regions those systems provide better solutions from both performability as well as energy point of view. For this purpose, Fig. 4 shows the solution of the inequalities (15) and (16) together. As shown in this figure, for  $V_R > 1.2V$  and  $V_{RI} > 0.8V$ , the curve of the performability equation is below the curve of the expected energy consumption equation (shaded area). This leads to an interesting conclusion: Given a conventional R system operating at a voltage  $V_R \geq 1.2V$ , it is always possible to find an RI system that offers better performability and, at the same time, lower energy dissipation than the R system. To clarify this assume a conventional R system running at  $V_R = 2.5V$ . An RI system operating at  $V_{RI} = 1.8V$  (Point A in Fig. 4) would require the same energy dissipation, however, it would offer better performability. Similarly, an RI system supplied with  $V_{RI} = 1.5V$  (Point B in Fig. 4) would expose the same

performability than the R system, however it would require less energy. In fact, it can be seen from Eq. (12), Eq. (13) and Eq. (14) that the proposed RI system consumes about 32% less energy than the conventional R system at Point B. For all the points on the vertical line between Points A and B, the proposed RI system offers simultaneously better performability and energy than the conventional R system. In summary, for DVS-enabled systems the RI system is the preferred choice even when considering the energy overheads associated with the additional hardware required for the information redundancy.

### 3.2 Influence of Hardware Overhead

Although the previous analysis has been carried out for the Crusoe processor, most of the parameters (Section 3.1) are independent from Crusoe and are only technology dependent. In fact, the only parameters that depend on the Crusoe processor are, i) number of the gates and flip-flops, ii) average switched capacitance, and iii) depth of critical path. The hardware overhead, which is required to process the redundant

information, influences these three parameters. In order to generalize the result obtained in Section 3.1, and to study the impact of the overhead on the efficiency of the proposed approach, we have regenerated the results in Fig. 5 for different parameters settings, i.e. critical path increase ( $K_c$ ), hardware overhead ( $K_a$  and  $K_{FF}$ ) and switching activity (switched capacitance) overhead ( $C_{eff\_extra}$ ).

As we can observe from Fig. 6a, if the RI system hardware overhead as well as the switching activity are assumed to be 50% higher than in the original R system and the critical path increase to be 4%, then the proposed RI system proves constantly advantageous in terms of performability and energy dissipation. With increasing critical path (up to 10%), hardware and switching overheads (up to 200%), the proposed RI system still provides better performability and energy dissipation for many voltage settings (i.e. shaded area).

### 3.3 Typical Hardware Overheads for RI systems

As we have seen in the previous section, the overheads associated with the additional hardware required for information redundancy have an impact on the suitability of the proposed RI approach. To provide insight into this overhead required for typical circuit designs, we have carried out some synthesis experiments using four circuits from the ITC'99 benchmarks and the Mentor Graphics Leonardo synthesis tool (Version 2003b.35). These experiments were performed for the unmodified circuits (representing the R systems) as well as for the modified circuits that included the extra hardware for the redundant information (representing the RI systems). After synthesis, the total number of signal transitions was used as a criterion to analyze the average switched capacitance and, hence, the dynamic energy consumption. It should be noted that the hardware overhead also accounts for the static energy overhead (see Section 2). The performed experiments indicate a hardware overhead of 42% to 179% and a switching activity overhead of 53% to 168%. Also, it has been found that the critical path length increase is less than 7%. Note that for such overheads the proposed RI system yields better results in terms of energy and performability (Fig. 6).

Overall, the experiments presented in this section have shown that the proposed RI systems offer advantages in terms of energy and performability over conventional R systems. This is particular the case if the hardware overhead for the additional information redundancy can be kept below 200%.

### 4. Conclusions

This paper has presented the first investigation into usage of information redundancy in DVS-enabled systems. Our experimental and analytical studies show that DVS-enabled real-time systems which use a combination of information-redundancy and rollback-recovery can significantly improve the system's reliability as well as energy dissipation, when compared to DVS-enabled systems that rely solely on rollback-recovery, even when considering the imposed hardware overheads.

### References

- [1] R. Melhem, D. Mosse, and E. Elnozahy, "The interplay of power management and fault recovery in real-time systems," *IEEE Transactions on Computers*, vol. 53, no. 2, pp. 217-231, 2004.

- [2] Y. Zhang and K. Chakrabarty, "Dynamic adaptation for fault tolerance and power management in embedded real-time systems," *ACM Transactions on Embedded Computing Systems*, vol. 3, no. 2, pp. 336-360, 2004.
- [3] D. Zhu, R. Melhem, and D. Mosse, "The Effects of Energy Management on Reliability in Real-Time Embedded Systems," in *Proceedings of International Conference on Computer Aided Design (ICCAD 2004)*, pp. 35-40, 2004.
- [4] M. T. Schmitz, B. M. Al-Hashimi, and P. Eles, *System-Level Design Techniques for Energy-Efficient Embedded Systems*, Kluwer Academic Publisher, 2004.
- [5] T. D. Burd, T. A. Pering, A. J. Stratakos, and R. W. Brodersen, "A dynamic voltage scaled microprocessor system," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 11, pp. 1571-1580, 2000.
- [6] F. Liberato, R. Melhem, and D. Mosse, "Tolerance to multiple transient faults for aperiodic tasks in hard real-time systems," *IEEE Transactions on Computers*, vol. 49, no. 9, pp. 906-914, 2000.
- [7] N. Seifert, D. Moyer, N. Leland, and R. Hokinson, "Historical trend in alpha-particle induced soft error rates of the Alpha microprocessor," in *Proceedings of the 39th Annual IEEE International Reliability Physics Symposium*, pp. 259-265, 2001.
- [8] A. Maheshwari, W. Burleson, and R. Tessier, "Trading off transient fault tolerance and power consumption in deep submicron (DSM) VLSI circuits," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 12, no. 3, pp. 299-311, 2004.
- [9] W. V. Subbarao, "Fault tolerant digital system design," in *IEEE Proceedings of Southeastcon '91*, pp. 124-128, 1991.
- [10] D. K. Pradhan, *Fault-tolerant computing: theory and techniques*, Prentice Hall, 1986.
- [11] K. M. Kavi, Hee Yong Youn, B. Shirazi, and A. R. Hurson, "A performability model for soft real-time systems," in *Proc. of the 27th Hawaii International Conference on System Sciences (HICSS)*, pp. 571-579, 1994.
- [12] S. M. Martin, K. Flautner, T. Mudge, and D. Blaauw, "Combined dynamic voltage scaling and adaptive body biasing for lower power microprocessors under dynamic workloads," in *Proceedings of IEEE/ACM International Conference on Computer Aided Design (ICCAD 2002)*, pp. 721-725, 2002.
- [13] J. L. Barth, C. S. Dyer, and E. G. Stassinopoulos, "Space, atmospheric, and terrestrial radiation environments," *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 466-482, 2003.
- [14] "TM5400/TM5600 Data Book," TRANSMETA, 2000.