

Nonlinear and Asymmetric Thermal-aware DVFS Control

Sylvain Durand*, Suzanne Lesecq†

* UMI LAFMIA, CNRS - CINVSTAV, Mexico City, Mexico.

† CEA LETI MINATEC Campus, 38000 Grenoble, France.

E-mail: sylvain@durandchamontin.fr

Abstract—Embedded devices using highly integrated circuits must cope with conflicting constraints. They have become more sensitive to variabilities with technology scaling while requiring computational efficiency under even more limited energy storage. Power management techniques, mainly based on Dynamic Voltage and Frequency Scaling (DVFS) algorithms, are hence of great interest. Also, temperature increase/decrease is directly related to the power consumption and platform characteristics highly depend on the temperature. As a result, temperature must be controlled, at least limited. In the present paper, the nonlinearity between power and temperature is analyzed and a thermal-aware DVFS control technique is discussed. The proposed strategy implements a chopped scheme on top of a robust DVFS approach in order to limit the temperature increase. The nonlinearity notably leads to an asymmetry in the temperature behavior, which makes the final attained temperature is higher than in the linear case.

INTRODUCTION

The upcoming generations of embedded integrated systems have reached limits in terms of power consumption, computational efficiency and fabrication yield. One of the problems in Systems on Chip (SoCs) induced by technology scaling is the so-called process variability: the chip performance cannot be ensured from one die to the other one, nor over a single chip. Moreover, clock distribution has become highly difficult, especially for multicore processors for which high speed analog clock signals should be routed around the system, the clock tree suffering from signal integrity issues and layout difficulties, leading to an increase in the design cost and a decrease in the yield. Globally Asynchronous Locally Synchronous (GALS) architectures alleviate this clock distribution difficulty, the chip being split into several frequency islands. Moreover, they allow each synchronous island to be supplied with a different voltage, the island becoming a Voltage and Frequency Island (VFI). As a consequence, GALS architectures are suitable for fine grain power management as the power consumption of the whole platform depends on the supply voltage and the clock frequency applied to each VFI [1], [2]. They also mitigate the impact of (process) variations [3].

Dynamic Voltage and Frequency Scaling (DVFS) methods are designed in order to provide just enough power to a given VFI for executing the tasks that run on it in such a way they meet their deadline. This allows to guarantee the overall performance while minimizing the power consumption. Closed-loop control strategies can hence be efficiently implemented to achieve such an energy/performance tradeoff. However,

for upcoming mobile platforms designed in advanced technologies, the temperature must be as well controlled, at least limited. Actually, the leakage power, which is a significant contributor to the total power consumption, highly depends on the temperature [4], [5]. Consequently, thermal effects have to be taken into account in the DVFS control law.

The present paper describes a thermal-aware DVFS strategy built on top of a robust DVFS control [6]. In section I, an overview of the context is provided: VFI power consumption and its link with the thermal aspects are first described and the system model is given. The main contribution is presented in section II, where a nonlinear thermal-aware DVFS strategy is detailed. We notably show that such a scheme leads to an asymmetry in the temperature dynamics behavior. Simulation results are finally provided in section III in order to highlight the capabilities of the proposed approach. Some discussions conclude the paper.

I. CONTEXT DESCRIPTION

Without lack of generality, and from now on, the DVFS technique is supposed to be implemented with two power modes, whose voltage levels are afterwards denoted V_{high} and V_{low} . ω^{max} and ω_{max} denote the maximal computational speeds – to avoid timing faults [7] – when the system is respectively running at high and low voltage with its maximal associated clock frequency.

A. VFI power consumption and its related thermal aspects

Power consumption in CMOS technology is usually sorted into a dynamic part due to gate switching, and a static one induced by short-circuit and leakage currents [8]

$$\begin{aligned} P &= P_{switching} + P_{short\ circuit} + P_{leakage} \\ &= K_{dyn} f_{clk} V_{dd}^2 + K_{sc} f_{clk} V_{dd} + K_{leak} V_{dd} \end{aligned} \quad (1)$$

where K_{dyn} , K_{sc} and K_{leak} are given parameters. The power consumption can be reduced when the supply voltage V_{dd} and the clock frequency f_{clk} are decreased, to ensure no timing faults (note that scaling down the voltage increases signal delays along the paths through the electronic gates, thus the clock frequency needs to be decreased as well, and inversely) [7]. Therefore, DVFS techniques can be used to efficiently manage the energy consumption of a device.

Naive DVFS techniques apply a constant power level in a given VFI for each task to treat. As a result, if the computational load of one task is higher than the processor capabilities at low voltage and speed ω_{max} , the VFI executes

the whole task with the penalizing high voltage and speed ω^{max} in order not to miss its deadline. To overcome such an intuitive approach, an energy-efficient control has been proposed in [6], where a task is split into two parts in order to really reduce the energy consumption. Firstly, the VFI begins to run at high voltage (if needed) with the maximum available speed ω^{max} in order to go faster than required. Then, the task is finished at low voltage with a speed lower than the maximal speed at low voltage. The task is hence executed with a certain ratio between high and low voltage levels, with a minimum penalizing high voltage running time to fit with the deadline and only one voltage transition (transitions consume more, this is discussed in the sequel) which, consequently, highly reduces the energy consumption. A key point in this strategy is that the (not a priori known) switching time to go from V_{high} to V_{low} has to be suitably calculated in order to ensure good computational performance. This is done thanks to a *fast predictive control* law, previously presented in [6] and not detailed here.

Whereas both naive and energy-efficient strategies reduce the energy consumption, they do not consider the thermal effect. The thermal dynamics of a circuit can be approximated by a first-order model [9], [10], such that

$$\frac{dT(t)}{dt} = aP(t) - b[T(t) - T_{amb}] \quad (2)$$

where T_{amb} is the ambient temperature. It is derived from a RC equivalent electrical model, where $a := 1/C$ and $b := 1/RC$ are some constants related to the thermal system. Actually, when a power mode is applied to the VFI for a long period of time, the temperature within the chip reaches a steady-state value (denoted T_{low} and T_{high}). Also, applying the high power mode for a long time can lead to the occurrence of a temperature hot spot within the chip. Whereas T_{low} cannot be reduced (even by physically decreasing the low voltage level), the temperature achieved at the penalizing high voltage mode can be decreased. The idea previously suggested for instance in [9], [4], [5], [11], [12] consists in switching between both voltage levels in order not to reach the maximum temperature T_{high} . Note that these papers are restrained to periodic tasks whereas the strategy proposed here can be applied to periodic but also non-periodic tasks. The execution of a given task is thus divided into different periods. The side effect of this chopped strategy is a slight increase of the energy consumption, because new transitions are introduced and the system is more consuming during transitions (as discussed in the sequel). The control problem hence naturally becomes a tradeoff between energy, performance and temperature. Furthermore, the power also varies with the temperature due to the leakage term [13], [4]. In fact, the leakage power in (1) can be more precisely defined as follows [4]

$$P_{leakage} = K_{leak1} V_{dd} + K_{leak2} V_{dd} T \quad (3)$$

where K_{leak1} and K_{leak2} are some constants (which grow with the technology scaling). Therefore, the relation between power and temperature becomes a nonlinear expression in

the sense that the power varies with the temperature, which varies in its own with the power. This means that when the temperature increases, the leakage power increases too and so is the temperature. It is hence clearly mandatory to implement a thermal-aware control of the DFVS methods.

B. System architecture

The control system architecture is mainly composed of a VFI and a controller. The voltage and frequency levels of the island is dynamically scaled while managing the tradeoff between energy, performance and temperature, thanks to the controller which monitors the activity of the processor (its so-called computational speed ω , e.g. in number of instructions per second) and adapts the control variables with respect to a computational load to treat (in terms of number of instructions and deadline for each task to execute). One can refer to [6], [14] for further details.

The supply voltage and clock frequency are provided by some actuators to the VFI, such as DC-DC converter or Vdd-hopping [15] in the first case and ring oscillator [16] or digital frequency locked-loop (DFLL) [17] in the second case. Then, the system goes from one level to the other one with a given transition time and dynamics that depend upon an internal control law. Considering that this inner-loop is extremely fast w.r.t. the loop considered in this paper, one can neglect the dynamics of both actuators. However, one knows that the voltage actuator consumes more energy during the transitions than in steady state – as showed in [18] in the case of the Vdd-hopping – and this point cannot be neglected.

C. Computing of the chopping parameters

Even if the power modes are applied in a chopped way, the ratio between high and low voltages has to remain equal (for a given task) to the one applied with the energy-efficient scheme in order to keep the performance. For this reason, this ratio – denoted hereafter *duty ratio* κ – is computed thanks to the same principle. Actually, it was indirectly calculated in the predictive control in [6]. Indeed, the so-called *predicted speed* δ is defined as the speed required to fit the task with its deadline regarding what has already been executed. The energy-efficient speed setpoint is then built in such a way the system runs either at high voltage if δ is higher than the maximum possible speed at low voltage, or at low voltage otherwise. Therefore, the duty ratio κ can be represented as the distance between the predicted speed and the maximal speeds at high and low voltage levels (i.e. ω^{max} and ω_{max} respectively). This gives

$$\kappa = \frac{\delta - \omega_{max}}{\omega^{max} - \omega_{max}} \quad (4)$$

where $0 \leq \kappa \leq 1$ by construction. Moreover, we propose to dynamically update κ at the beginning of each chopping period – denoted hereafter *period of oscillation* τ – in order to be robust to uncertainties in the task information.

The impact of the length of τ on the variation of the temperature for a given duty ratio (this will be analytically shown in the sequel) has to be carefully considered. In

practice, it could be constant or fixed by a given criterion to satisfy. Such a method is treated in [14] in particular, where the product $\kappa\tau$ is fixed, and so is the energy-temperature tradeoff. The suggested idea is based on the fact that *i*) a small period leads to small variations and *ii*) fast oscillations increase the energy consumption (because the transitions consume some power that is not strictly speaking used for the computations). For these reasons, it is interesting to have slow oscillations when the temperature is low (i.e. κ is small) and, inversely, to allow fast oscillations when the temperature becomes high (i.e. when the duty ratio is high). This yields

$$\tau = \frac{\varphi}{\kappa} \quad (5)$$

where φ denotes the *rising temperature time*. Note that this parameter has also to be chosen by the designer but the energy-temperature tradeoff is now self-adjusted. Another solution could be to update the length of τ from the monitoring of the chip internal temperature. This latter solution – studied in [19] for the linear case – is not applied here.

Finally, one could note that establishing the approach on top of the DVFS control proposed in [6] – and more especially on the computation of κ in (4) – allows a strong robustness to technological variability because this seminal strategy decides the power modes without any knowledge on the system parameters (see the reference for more details). However, the present contribution can be extended to any DVFS algorithm.

II. NONLINEAR CHOPPING CONTROL

The RC equivalent electrical model [10] of the thermal equation (2) is given in Fig. 1(a). The components are assumed to be ideal and, therefore, do not dissipate power. The different power values are associated with the different voltage/frequency modes. In the present study case, P_1 and P_2 correspond to the system respectively running at $V_1 := V_{high}$ and $V_2 := V_{low}$ with its associated frequency, assuming $P_1 > P_2$. The temperature increases when P switches from P_2 to P_1 and decreases otherwise, resulting in a rectangular waveform for the power whose average value depends on the duty cycle of the switching control, and so is the average temperature.

The system is supplied by the high power source P_1 (respectively the low power source P_2) during $\kappa\tau$ (resp. $(1 - \kappa)\tau$) when the switch in Fig. 1(a) is in position 1 (resp. 2). The process then cyclically repeats. However, in the present case, the power is no more constant but varies with the temperature. Substituting (3) in (1) yields

$$P_i(t) = A_i + B_i T(t) \quad (6)$$

with

$$A_i := K_{dyn} f_i V_i^2 + K_{leak1} V_i + K_{sc} f_i V_i$$

$$B_i := K_{leak2} V_i$$

where i is used to denote the voltage level $V_i = \{V_1, V_2\}$ and its associated frequency f_i , as well as the corresponding power level $P_i = \{P_1, P_2\}$. Note that the expression in (6) clearly uses the assumption on the neglected dynamics of the

actuators (see subsection I-B). The thermal expression in (2) finally becomes

$$\frac{dT(t)}{dt} = \begin{cases} \alpha_1 - \beta_1 T(t) & \text{if } 0 \leq t \leq \kappa\tau \\ \alpha_2 - \beta_2 T(t) & \text{if } \kappa\tau \leq t \leq \tau \end{cases} \quad (7)$$

with

$$\alpha_i := aA_i + bT_{amb}$$

$$\beta_i := b - aB_i$$

The temperature can then be analytically calculated by solving these thermal differential equations, which gives

$$T(t) = \begin{cases} k_1 e^{-\beta_1 t} + T_1 & \text{if } 0 \leq t \leq \kappa\tau \\ k_2 e^{-\beta_2 t + \beta_1 \kappa\tau} + T_2 & \text{if } \kappa\tau \leq t \leq \tau \end{cases} \quad (8)$$

with $T_i := \frac{\alpha_i}{\beta_i}$

where T_i is the achieved temperature for a given power mode i , with $T_1 > T_2$ by construction. Finally, the result in (8) shows that the nonlinearity leads to an asymmetry in the temperature dynamics. Indeed, the parameter β_i in the exponential, with $\beta_1 \neq \beta_2$ by definition, indicates that the temperature variation will not be the same during the rising and falling behaviors. One could note that the temperature dynamics in its own can be asymmetric. In this case, the parameters a and b in (2) become a_i and b_i , and so are impacted α_i and β_i in (7). This is not developed in this paper but the extension is trivial.

At the end and in order to then sketch the behavior of such a chopped scheme, the parameters k_1 and k_2 in (8) are analytically calculated. This is detailed in the following subsections.

A. Temperature steady-state mean value and ripple analysis

In the steady-state case, the parameters k_1 and k_2 can be analytically obtained from:

- i*) the continuous behavior at $t = \kappa\tau$ for both expressions;
- ii*) the periodicity of the temperature in the steady state, i.e. $T(0) = T(\tau)$.

This gives

$$k_1 = -\Delta_T \frac{1 - e^{-\beta_2(1-\kappa)\tau}}{e^{(\beta_1-\beta_2)\kappa\tau} - e^{-\beta_2\tau}} e^{(\beta_1-\beta_2)\kappa\tau} \quad (9)$$

$$k_2 = \Delta_T \frac{1 - e^{-\beta_1\kappa\tau}}{e^{(\beta_1-\beta_2)\kappa\tau} - e^{-\beta_2\tau}}$$

with $\Delta_T := T_1 - T_2$

with $\Delta_T > 0$ by construction. These expressions can be linearised (this linearization, and the following ones, are only used to simplify the analysis) – applying a first order Taylor expansion – since the oscillating period is very small compared with the time constant of the system, i.e. $\tau \ll RC$

$$k_1 \simeq -\frac{\Delta_T}{\lambda_1} \left[\left(-\beta_1 + \frac{\beta_2}{2} \right) \kappa^2 + \left(\beta_1 + \frac{\lambda_2}{\lambda_1} \right) \kappa - \left(\frac{\beta_2}{2} + \frac{\lambda_2}{\lambda_1} \right) \right] \tau - \frac{\Delta_T}{\lambda_1} (1 - \kappa) \quad (10)$$

$$k_2 \simeq \frac{\Delta_T}{\lambda_1} \frac{\beta_1}{\beta_2} \left[-\left(\frac{\beta_1}{2} \kappa + \frac{\lambda_2}{\lambda_1} \right) \tau + 1 \right]$$

with

$$\lambda_1 := 1 + \left(\frac{\beta_1}{\beta_2} - 1 \right) \kappa$$

$$\lambda_2 := \left(\frac{\beta_1^2}{2\beta_2} - \beta_1 + \frac{\beta_2}{2} \right) \kappa^2 - \frac{\beta_2}{2}$$




Fig. 1. Representation of the chopped control scheme: the RC equivalent model, the representation of the steady-state and transient phases.

A linearized expression of the temperature in (8) in the steady state becomes

$$T(t) \simeq \mu t + \nu \quad (11)$$

$$\text{with } \mu := \begin{cases} -\beta_1 k_1 & \text{if } 0 \leq t \leq \kappa\tau \\ -\beta_2 k_2 & \text{if } \kappa\tau \leq t \leq \tau \end{cases}$$

$$\text{and } \nu := \begin{cases} k_1 + T_1 & \text{if } 0 \leq t \leq \kappa\tau \\ (1 + \beta_1 \kappa\tau)k_2 + T_2 & \text{if } \kappa\tau \leq t \leq \tau \end{cases}$$

A waveform can be sketched such as in Fig. 1(b). The temperature begins at some initial value $T(0)$. It increases during the first subinterval, when the switch is in position 1, with a positive slope given by μ in (11). At time $t = \kappa\tau$, the switch changes from position 1 to position 2, and the temperature decreases until the end of the period, since the slope becomes negative in (11). At time $t = \tau$, the switch changes back to position 1 and the process repeats.

The average value of the temperature as well as its ripple, afterwards denoted T_{avg} and Υ , can then be easily computed from the minimum and maximum temperature peak values which occur at $t = 0$ (or $t = \tau$) and $t = \kappa\tau$ respectively. They are defined as follows

$$T_{avg} = \frac{\Delta T}{2} \frac{1 - e^{-\beta_1 \kappa\tau} - e^{-\beta_2 \tau} + e^{(\beta_1 \kappa - \beta_2)\tau}}{e^{(\beta_1 - \beta_2)\kappa\tau} - e^{-\beta_2 \tau}} + T_2 \quad (12)$$

$$\Upsilon = \Delta T \frac{1 - e^{-\beta_1 \kappa\tau} + e^{-\beta_2 \tau} - e^{(\beta_1 \kappa - \beta_2)\tau}}{e^{(\beta_1 - \beta_2)\kappa\tau} - e^{-\beta_2 \tau}}$$

The linearized versions can also be obtained. However, the results are quite complex to analyze in the present nonlinear and asymmetric case. In order to simplify the expressions, we propose to use in the sequel

$$\beta_1 = \varepsilon \beta_2$$

with $0 < \varepsilon < 1$ by construction, see (7)-(6). This yields

$$T_{avg} \simeq \frac{\Delta T}{2} \varepsilon \kappa \frac{(1 - \varepsilon)\kappa\tau + 2}{(\varepsilon - 1)\kappa + 1} + T_2 \quad (13)$$

$$\Upsilon \simeq \Delta T \frac{\varepsilon \kappa (1 - \varepsilon \kappa)}{(\varepsilon - 1)\kappa + 1} \beta_2 \tau$$

Both expressions depend on the duty ratio and the period. In the symmetric case (i.e. $\varepsilon \approx 1$), the simple relations observed in [14] are confirmed here but the non-linearity highly complexifies the relations. Thus, whereas the average temperature only increases w.r.t. κ in the linear case, it now increases w.r.t. κ and τ , whose increase is respectively lightly/highly impacted with the asymmetry (when ε decreases). On the other hand, whereas the temperature ripple

is maximum when $\kappa = 0.5$ in the symmetric case, it now monotonously increases w.r.t. κ when the difference between β_1 and β_2 grows. In particular, the maximum achieved value dangerously increases when ε becomes close to 0, but this cannot occur by definition of the β_i parameters in (7)-(6) (since β_i is only function of the voltage level V_i , and yet, it is inevitably different for all power modes). Also, the proportional relation between the ripple and τ decreases with the asymmetry.

B. Transient phase analysis

The behavior of the temperature in the transient phase is now analyzed. Suppose that the system runs at V_{low} for a long time, the temperature being stabilized in its lower value T_2 . Then, a new task that needs the processor to (partially) run under V_{high} is applied, and so is the chopped scheme presented above. At the very beginning of this process, the temperature is not in steady-state, i.e. $T(0) \neq T(\tau)$. During the transient phase, the thermal expression in (8) becomes

$$T(t) = \begin{cases} k_1(n)e^{-\beta_1(t-n\tau)} + T_1 & \text{if } n\tau \leq t \leq n\tau + \kappa\tau \\ k_2(n)e^{-\beta_2(t-n\tau) + \beta_1 \kappa\tau} + T_2 & \text{if } n\tau + \kappa\tau \leq t \leq (n+1)\tau \end{cases}$$

where $n \in \mathbb{N}$ is a number incremented at each new period, whose parameters k_1 and k_2 are now depending. In fact, they can be expressed with their steady-state value – previously defined in (9) – plus a term decreasing w.r.t. n , that gives

$$k_i(n) = k_i + \sigma_i \Delta T \frac{1 - e^{\sigma_i \beta_i \kappa\tau}}{e^{(\beta_1 - \beta_2)\kappa\tau} - e^{-\beta_2 \tau}} \xi_i(n+1) \quad (14)$$

$$\text{with } \begin{aligned} \xi_1(n) &:= \zeta(n) e^{(\beta_1 - \beta_2)\kappa\tau} \\ \xi_2(n) &:= \zeta(n) \\ \zeta(n) &:= e^{-n((\beta_1 - \beta_2)\kappa + \beta_2)\tau} \end{aligned}$$

$$\text{and } \begin{aligned} \sigma_1 &:= +1 \\ \sigma_2 &:= -1 \end{aligned}$$

Note that the term $\xi_i(\cdot)$ exhibits an exponential decrease. Moreover, the initial temperature slope $\mu(n=0)$ can be expressed from (11) and (14). Then, at each new oscillating period, the slope is equal to the previous one plus an extra term η defined as follows

$$\mu(n) = \mu(n-1) + \eta(n) \quad (15)$$

$$\text{with } \eta(n) := \sigma_i \Delta T \beta_i (1 - e^{\sigma_i \beta_i \kappa\tau}) e^{(\beta_2 - \beta_1)\kappa\tau} \xi_i(n)$$

The resulting “turn-on” transient phase is depicted in Fig. 1(c), where the initial temperature is equal to T_2 . An

input power P_1 is then applied during the first oscillating period, with the switch in position 1. Hence, the temperature increases with an initial slope equal to $\mu(0) = \beta_1 \Delta_T$. During the second part of the oscillating period (i.e. after the switch has changed from position 1 to position 2), the temperature decreases with an initial slope equal to $\mu(0) = -\beta_2 \Delta_T (1 - e^{-\beta_1 \kappa \tau}) e^{(\beta_2 - \beta_1) \kappa \tau}$. At time $t = \tau$, a new oscillating period occurs and $n = 1$. The temperature during the first part of this second oscillating period increases, but with a smaller slope when compared with the first oscillating period since the added term $\eta(\cdot)$ is negative, see (15). Then, the temperature decreases during the second part of the second oscillating period, but with a higher slope since the added term $\eta(\cdot)$ is positive during the second subinterval. The process is then repeated until the steady-state condition is attained.

III. PERFORMANCE EVALUATION

It is quite difficult to find experimental values for the parameters of the thermal model (2) in the literature. For the simulation presented now, the ones deduced from [13], [20] will be used. The different values are $R = 2^\circ\text{C}/\text{W}$ and $C = 34 \text{ mJ}/^\circ\text{C}$. The ambient temperature is $T_{amb} = 25^\circ\text{C}$ and the supply power levels are 5 and 35 W. In the linear case, this leads to a thermal time constant of 68 ms and the chip temperature varies from 35 to 95°C. In the nonlinear one, the resulting asymmetric behavior runs with some different rising and falling time constants that are of 124 and 74 ms respectively, and the chip temperature varies from 36 to 154°C.

Different preliminary tests are realized, comparing the linear and nonlinear chopped scheme with the classical strategy. For the classical scheme, the task deadline is supposed equal to the time of the whole simulation, i.e. $\bar{\tau} = 0.4 \text{ s}$, and the system runs with the switch in position 1 and 2 during the time interval $\kappa \bar{\tau}$ and $(1 - \kappa) \bar{\tau}$ respectively. As expected, the temperature highly increases. On the other hand, hot spots can be reduced thanks to the chopped scheme. We already showed in [14] that the steady-state average temperature and its ripple depend on both parameters and the maximal attained temperature highly depends on the switching control. Even if the relation with these parameters is more complex in the nonlinear case, such an observation can be done (simulation results are not provided due to space limitation). Moreover, the reached temperature is more important in the nonlinear case than in the linear one.

The control strategy when fixing the product $\kappa \tau$ in order to bound the time when the temperature increases (and consequently the tradeoff between temperature and energy) – as suggested in subsection I-C – is depicted in Fig. 2 where $\varphi = 10 \text{ ms}$. This gives slow and fast oscillations when the temperature is low and high respectively (i.e. when κ is small and large respectively).

The whole thermal-aware control of the energy-performance tradeoff is also tested in simulation. The idea is to reduce the energy consumption while ensuring computational performance, as suggested in [6], taking also into

account the temperature within the chip here. A scenario with three tasks to be executed is run, the number of instructions and deadline for each task being known. The simulation results are depicted in Fig. 3 for the nonlinear case. The plots on the top part show the number of instructions, the deadline, and the laxity. The speeds (average speed setpoint, predicted, measured) are provided on the plot in the middle of the figure, as well as the voltage value. The bottom plot shows the temperature of the chip. The temperature reached with the classical strategy (without any thermal management) is also represented. When the task to be run has a high computational load, the system runs during $\kappa \tau$ and $(1 - \kappa) \tau$ at high and low voltage respectively. Then, this is repeated until the end of the task. As expected, the energy saving is less important than in the classical strategy but the maximal temperature highly decreases (about 8% increase and 47% decrease respectively). Note that the energy increase is more important than in the linear case, but the proposed scheme also results in a higher temperature decrease (about 1% of energy increase and 40% of temperature decrease in the linear case). Therefore, this approach will basically decrease the hot spot appearance.

CONCLUSIONS AND FUTURE WORKS

In the present paper, we presented a nonlinear thermal-aware DVFS technique. The proposed control scheme was implemented on top of a robust DVFS control [6] and a chopped scheme was added in order to limit the temperature growing. The main advantage of this “chopped-DVFS” is that it limits the temperature increase/decrease, since the average temperature values depend on the duty cycle of the switching control and the ripple depends on the duty cycle and the period of oscillation. The duty ratio is calculated from the fast predictive control in [6]. This ensures performance and robustness to uncertainties. Also, the nonlinearity between temperature and power was taken into account. The nonlinearity notably leads to asymmetry in the temperature behavior, with an increased achieved temperature at high voltage (compared with the linear case). Some simulation results have also been provided to show the effectiveness of the proposal and its capability to efficiently reduce the temperature of the chip. Next work will be to implement these schemes in practice.

ACKNOWLEDGEMENTS

This work has been partially performed while S. Durand was post-doc fellow in the joint NeCS Team, INRIA/GIPSA-lab. It has been conducted in cooperation with CEA LETI Minatec Campus, under the CRI PILSI, in Grenoble, France. It has also received partial funding from the Artemis Arrow-Head and Catrene Benefic projects.

REFERENCES

- [1] K. Flautner, D. Flynn, D. Roberts, and D. I. Patel. An energy efficient SoC with dynamic voltage scaling. In *Proceedings of the Design, Automation and Test in Europe Conference and Exhibition*, 2004.

Thermal-aware control -- Variation w.r.t. kappa when the product $\kappa\tau$ is constant




Fig. 2. Simulation results for the nonlinear and asymmetric thermal-aware control scheme when the product $\kappa\tau$ is constant, for different values of κ .

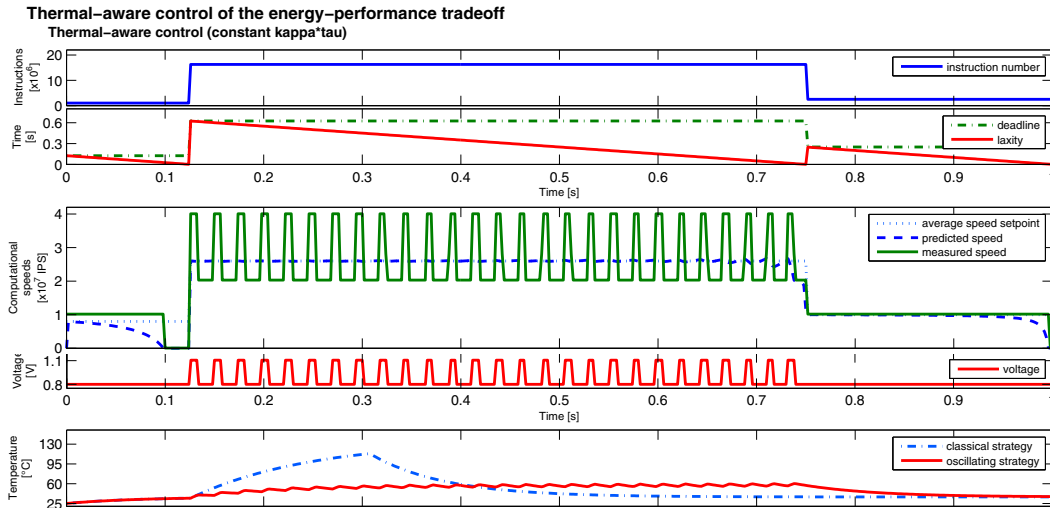


Fig. 3. Simulation results of the global system: nonlinear and asymmetric thermal-aware control of the energy-performance tradeoff.

[2] P. Choudhary and D. Marculescu. Hardware based frequency/voltage control of voltage frequency island systems. In *Proceedings of the 4th International Conference on Hardware/Software Codesign and System Synthesis*, pages 34–39, oct. 2006.

[3] D. Marculescu and E. Talpes. Energy awareness and uncertainty in microarchitecture-level design. *IEEE Micro*, 25:64–76, 2005.

[4] V. Chaturvedi, Huang H., and G. Quan. Leakage aware scheduling on maximum temperature minimization for periodic hard real-time systems. In *Proceedings of the 10th IEEE International Conference on Computer and Information Technology*, 2010.

[5] H. Huang and G. Quan. Leakage aware energy minimization for real-time systems under the maximum temperature constraint. In *Proceedings of the Conference & Exhibition on Design, Automation and Test in Europe*, 2011.

[6] S. Durand and N. Marchand. Fully discrete control scheme of the energy-performance tradeoff in embedded electronic devices. In *Proceedings of the 18th World Congress of IFAC*, 2011.

[7] B. Zhai, D. Blaauw, D. Sylvester, and K. Flautner. The limit of dynamic voltage scaling and insomniac dynamic voltage scaling. *IEEE Transactions on Very Large Scale Integrated Systems*, 13:1239–1252, 2005.

[8] A. P. Chandrakasan and R. W. Brodersen. Minimizing power consumption in digital CMOS circuits. *Proceedings of the IEEE*, 83(4):498–523, 1995.

[9] C. Y. Yang, J. J. Chen, T. Lothar, and T. W. Kuo. Energy-efficient real-time task scheduling with temperature-dependent leakage. In *Proceedings of the Conference & Exhibition on Design, Automation and Test in Europe*, 2010.

[10] S. Zhang and K. S. Chatha. Approximation algorithm for the temperature-aware scheduling problem. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, 2007.

[11] V. Chaturvedi and G. Quan. Leakage conscious DVFS scheduling for peak temperature minimization. In *Proceedings of the 16th Asia and South Pacific Design Automation Conference*, 2011.

[12] H. Huang, G. Quan, J. Fan, and M. Qiu. Throughput maximization for periodic real-time systems under the maximal temperature constraint. In *Proceedings of the 48th ACM/EDAC/IEEE Design Automation Conference*, 2011.

[13] G. Quan and Y. Zhang. Leakage aware feasibility analysis for temperature-constrained hard real-time periodic tasks. In *Proceedings of the 21st Euromicro Conference on Real-Time Systems*, 2009.

[14] S. Durand, S. Lesecq, E. Beigné, C. Fabre, L. Vincent, and D. Puschini. Low-cost dynamic voltage and frequency management based upon robust control techniques under thermal constraints. *Lecture Notes in Computer Science (LNCS), post-proceedings of FMCO 2011*, 7542:334–353, 2012.

[15] C. Albea Sánchez, C. Canudas de Wit, and F. Gordillo. Control and stability analysis for the vdd-hopping mechanism. In *Proceedings of the IEEE Conference on Control and Applications*, 2009.

[16] E. Yahya, O. Elissati, H. Zakaria, L. Fesquet, and M. Renaudin. Programmable/stoppable oscillator based on self-timed rings. In *Proceedings of the 15th IEEE International Symposium on Asynchronous Circuits and Systems*, 2009.

[17] S. Lesecq, D. Puschini, E. Beigné, P. Vivet, and Y. Akgul. Low-cost and robust control of a DFLL for multi-processor system-on-chip. In *Proceedings of the 18th World Congress of IFAC*, 2011.

[18] S. Miermont, P. Vivet, and M. Renaudin. A power supply selector for energy- and area-efficient local dynamic voltage scaling. In *Proceedings of the 17th International Workshop on Power and Timing Modeling, Optimization and Simulation*, 2007.

[19] S. Durand and S. Lesecq. Asynchronous thermal-aware DVFS control. In *Proceedings of the American Control Conference*, 2013.

[20] A. Sridhar, A. Vincenzi, M. Ruggiero, T. Brunschweiler, and D. Atienza. Compact transient thermal model for 3D ICs with liquid cooling via enhanced heat transfer cavity geometries. In *Proceedings of the 16th International Workshop on Thermal Investigations of ICs and Systems*, 2010.