EEE6207 - Dynamic Voltage Scaling

Hamish Sams

October 2019

Abstract

This report explains what dynamic voltage scaling (DVS) is and why it is a useful addition to a circuit. We also discuss how a DVS system can be constructed and how it works along with a case study example of a processor using DVS. Finally we discuss how and when these devices are feasible along with specific applications where DVS would be strongly considered.

1 Introduction

The majority of ICs (Integrated Circuits) and digital circuitry nowadays is constructed using CMOS technology. One appealing factor of CMOS logic is the low static power consumption due to the intrinsic insulated gate.[1] CMOS power dissipation can be split into 2 sections, static and dynamic power. Static power is from leakage currents through the device and dynamic power arises from device capacitors charging and discharging on clock cycles defined by $P_{Static} = I_{leak}V$ and $P_{dyn} = \alpha CV^2 f$ respectively. [2] As devices get faster the frequency increases dynamic power, as devices get smaller the leakage current gets larger increasing static power. [3]Reducing power consumption is a large modern concern and is one of the main features holding back technological advance. Reducing power is important as many devices depend on batteries for power. Smaller systems are strict with their battery sizes and must rely on other techniques to reduce power consumption. Device longevity is also affected by reducing power as a 15 degree rise in temperature doubles the chance of failure. [4]

A DVS (Dynamic Voltage Scaling) system acts to vary the voltage supplied to the processor depending on load allowing for a system to work at the peak processing power whilst reducing power usage at non-peak usage. More advanced DVS systems also implement a variable frequency solution, for now this system is out of scope of a typical DVS.

2 Technical explanation

A common DVS architecture is shown in Figure: 1 [5]. It can be split into 6 distinct sections: A power supply, Ring oscillator, Counter, Binary comparison,
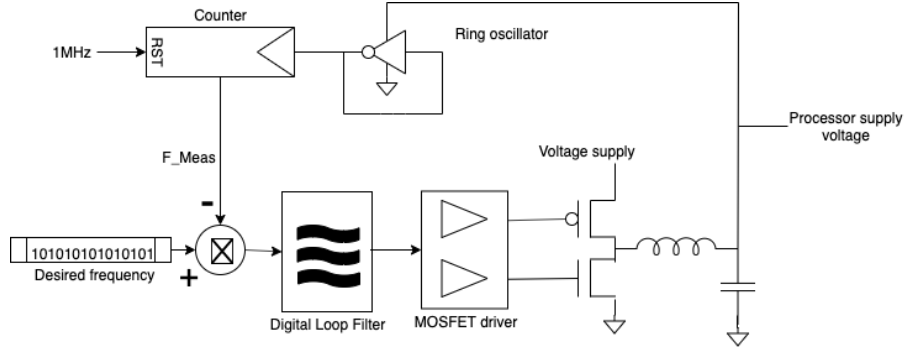
Figure 1: Block diagram of a typical DVS system [5]

Digital Loop Filter and Mosfet driver. Each of these work in a digital and analog loop to generate the processor input voltage based on a value in memory (Desired frequency).

Starting with the power supply at the bottom right of Figure: 1, a buck converter is used to generate the output voltage given the input from a constant supply. If a linear system was used then all power would be dissipated that isn't used meaning no power is saved. A buck converter is used instead of a boost converter as we will always be stepping down the voltage from the supplied line. These inductor and capacitor components must be large enough to reduce ripple on the output as to not effect the processors operation. However the larger the capacitor the larger the transition time between a current voltage and a desired voltage and therefore the energy dissipated in this transition region as seen in Equation: 1 [5]. The larger this capacitance the slower the voltage can propagate from the device, we can work out a rough estimate for this time using $Q = CV$ and $I = \frac{Q}{T}$ giving $\tau = \frac{CV}{I}$ substituting in our circuit values we get Equation 2. The 2 in this equation comes from the triangular current operation of the buck convertor leading to an average current $I = 0.5 \cdot I_{max}$. For a $0.1V$ change using a capacitance of $1000\mu F$ and max current of $1A$ we get a response time of $200\mu S$ meaning no real time changes could be made. This means the best design would minimise capacitance whilst ensuring normal processor operation

$$E_{tran} = (1 - \eta) \cdot C \cdot |V_{Cur}^2 - V_{Des}^2| \tag{1}$$

$$\tau_{tran} = \frac{2 \cdot C}{\cdot} I_m ax |V_{Cur} - V_{Des}| \tag{2}$$

This bucked voltage is fed into a ring oscillator as a power supply. The ring oscillator is simply a CMOS not gate fed into itself leading to the device oscillating as fast as possible defined by the time the CMOS gates charge and discharge. As the voltage supplied to the ring oscillator changes the CMOS capacitances charge and discharge faster leading to a higher frequency for a higher voltage and vice versa.

This variable frequency is then fed into a binary counter, this binary counter adds up each rising edge from the ring oscillator. With a 1MHz signal being fed into the reset pin the counter will add how many oscillations appear in $1\mu s$. This is therefore counting the ring oscillator frequency in MHz. The binary representation of this the then fed out as F_{Meas}.

A defined desired frequency is stored in a memory address available to the DVS system. This value is calculated, likely with a LUT, to generate the lowest possible frequency for the current system load. A logic block then calculates the difference between the desired frequency and F_{Meas} giving a frequency error.

This frequency error is fed into a digital loop filter. This filter is designed to be used in an closed loop feedback system constantly correcting the output generated based on error. This means the system is purely reactive to error present in the system. This loop filter generates a PWM signal output, from its current state modified by the error, to the MOSFET driver.

To close the analysis loop, the MOSTFET driver implements the required dead time and signal inversion required to drive the totem pole without shorting the supply. This totem pole is fed through an inductor and capacitor to create our entire buck converter system of charging and discharging LC circuit.

2.1 Case study

Advanced Micro Devices, Inc. (AMD) is a large chip manufacturing company, many of their products are used in portable and low energy devices such as laptops, in these devices low power consumption is one of the highest priorities. PowerNow! Is a DVFS system, meaning we can dynamically change the voltage and frequency of the processor[6]. PowerNow! Can modify the voltage between $0.925V$ and $2V$ and modify the frequency between $200MHz$ and $500MHz$ leading to an overall minimum power consumption of 8.56% of the maximum processor power. AMD's chipsets that use PowerNow! Are designed for use in laptops and systems with passive cooling needs. Whilst allowing for passive cooling when the chipsets are much more powerful, running a laptop on minimal processing needs such as low level programming or plain text editing the device could run for 10 times longer on the same battery whilst allowing for a more demanding application on the same processor. Realistically an AMD processor may only be used down to around $1.4V$ reducing power to 19.6% of its maximum.

The PowerNow! settings may be set using a 16 byte enhanced power management register (EPMR) model specific register (MSR). This data is set using drivers available for generic operating systems such as windows as well as real time operating systems (RTOS). PowerNow! Also allows for an automatic mode where the voltage and frequency are set dynamically without need for any verbose instruction from the drivers, this is calculated from how busy the processor currently is. The other two options allow for a high performance or power-saver mode depending on what the user would like.

3 Feasibility study

In terms of DVS feasibility we have positives allowing for reduced power consumption in turn leading to increase battery life of portable devices. Whilst we are changing the processor power, we are doing so dynamically having little effect on maximum processing speeds. This does however have some disadvantages with it, although our DVS system allows for a dynamic range of processing ability, the processing ability is reduced in different ways, especially if a job switches between high and low load. A system using DVS also requires a processor with a programmable clock which in turn has its own disadvantages such as accuracy, space and required overheads to name a few. As DVS is another system required to be added into the circuit, the complexity and cost increases as such and as more processors are becoming multi-core and higher frequency the return is diminishing. A DVS system is also only useful if a system runs with different processing needs in its lifetime as if no switching of processor ability is needed for things such as dedicated multimedia or embedded systems the processor can be picked accordingly meaning DVS brings no advantage.

4 Conclusion

Overall a DVS system is very useful and should be a consideration for systems where power management if of a high priority such as portable technology. DVS systems should also be considered when a system has a largely varying load, some examples may be servers with time dependent loads to reduce cost of power given the large possible unused compute. Other devices that require a small form factor or have limited access to cooling such as server rack devices or space systems that cannot dissipate heat would find DVS systems a large advantage. Certain systems such as full time multimedia devices or embedded systems with a constant load may not benefit as much from DVS power savings compared to a good choice of processor.

References

[1] Kahng Dawon. Electric field controlled semiconductor device, August 27 1963. US Patent 3,102,230.

[2] Sparsh Mittal. A survey of techniques for improving energy efficiency in embedded computing systems. *arXiv preprint arXiv:1401.0765*, 2014.

[3] Wolfgang M Arden. The international technology roadmap for semiconductors—perspectives and challenges for the next 15 years. *Current Opinion in Solid State and Materials Science*, 6(5):371–377, 2002.

[4] Dave Anderson, Jim Dykes, and Erik Riedel. More than an interface-scsi vs. ata. In *FAST*, volume 2, page 3, 2003.

[5] Thomas D Burd and Robert W Brodersen. Design issues for dynamic voltage scaling. In *Proceedings of the 2000 international symposium on Low power electronics and design*, pages 9–14. ACM, 2000.

[6] Amd powernow! technology white paper, 2000.