



## DESIGN CONSIDERATIONS FOR SIZE, WEIGHT, AND POWER (SWAP) CONSTRAINED RADIOS

Presented at the  
2006 Software Defined Radio Technical Conference  
and Product Exposition  
November 14, 2006

### ABSTRACT

For battery powered radios, average power consumption of the radio dictates the operational life of the radio in field use. Many of these radio applications are also severely size and weight limited. Therefore, the design of both hardware and software must take into account Size, Weight, and Power (SWAP) as primary design considerations. This paper discusses best practices for hardware and software design in SWAP constrained radios. Topics covered are dynamic clock scaling, power supply scaling, modes of operation, duty cycle impact, static versus dynamic power, hardware power reduction considerations, and programmable logic versus ASIC designs. The focus is on practical design methods and trade-offs, especially for networking waveforms. Hardware only power and space saving techniques are discussed first. Waveform impact on SWAP and some software techniques for power savings are discussed. Software interaction for dynamically controlled power savings are addressed along with the requisite hardware interfaces. Software topics include impact of SDR waveforms on hardware power and space.

### 1. INTRODUCTION

Today's communications solutions demand the capabilities and flexibility provided by Software Defined Radios (SDR) with their associated demand for processing power. Adapting these power-hungry digital designs for use in size and power constrained radios is one of the most challenging aspects of modernizing tactical communications. Fortunately, there are techniques that can be applied that will reduce or minimize power lost as thermal energy. Efficient power regulation design and use of frequency and voltage scaling [1] are examples of low power design practices. But to really achieve high levels of power efficiency requires a proactive effort to design the waveform software to use the hardware in the most power efficient way possible. This paper describes a range of methodologies that can be used to provide a flexible SDR digital architecture, while minimizing its size and power consumption.

### 2. SPACE AND POWER ON DIGITAL HARDWARE DESIGNS

The biggest technical challenge of SDRs such as those being developed for the US Joint Tactical Radio System (JTRS) Program is achieving the functionality desired within the size, weight and power specified. According to a Government Accountability Office (GAO) report [2], "To realize the full capabilities of the Wideband Networking Waveform [WNW], including transmission range, the Cluster 1 [renamed to Ground Mobile Radio (GMR)] radio requires significant amounts of memory and processing power, which add to the size, weight, and power consumption of the radio." Given that the GMR sets are targeted for ground vehicles, which generate their own power, the GAO report is even more relevant to dismantled, battery powered applications such as Handheld and Manpack radios. It is the responsibility of both the radio's hardware and software design teams to use integrated design techniques so the platform consumes as little power as possible.

As customers demand the increased functionality and flexibility that can only be provided by SDRs it is precisely in the digital and not the RF circuits that power consumption is increasing dramatically. The increased bandwidth and networking performance are the factors driving the digital baseband requirements for more memory and processing power. Both waveform and networking processing are being driven to higher levels of demand by waveforms like WNW and SRW (Soldier Radio Waveform) and the requirements of the JTRS Program's Software Communications Architecture (SCA). Balancing against the processing demands of these high throughput networking waveforms, is the requirement that many military radios must still operate in a power efficient manner to provide long operational life from battery packs.

There are many proven techniques to improve the power and size efficiency of modern electronic circuits. Lower core voltages, more efficient interconnections, and the general trend for more efficient design of devices all contribute to the ability to provide more processing performance for less power. These methods can be applied to any design and do not require anything special of the system software. The remainder of this paper discusses

hardware design concepts and software assisted hardware design approaches for low power baseband radio designs.

### 3. REGULATION AND POWER TREES

Traditional power distribution schemes call for input power to be supplied at a voltage higher than the voltages required by the radio's circuitry, minimizing the requirement for inefficient up-converters. Typical circuits utilized to provide this power distribution include:

1. Switching Regulators
2. Pass Transistor Regulators

Each of these two basic regulator types has its strengths and weaknesses, and in SWAP-constrained designs, both have a place in size and power consumption reductions. There are many design variations of both types but, in general, Pass Transistor Regulators tend to be more space efficient while Switching Regulators are more power efficient.

The efficiency of a Pass Regulator is dependent upon the down-conversion step size of the regulator and the current draw through the regulator. The regulator's efficiency can be described by the relationship of:  $\text{Efficiency} = (P_{in} - P_{out})/P_{in} = (V_{in} - V_{out}) * I^2$ . As shown by the formula, the larger the down-conversion step size, the more power is dissipated within the regulator as heat and the lower the regulator's efficiency. Also, since the regulator's efficiency is exponentially related to the supply current, efficiency is reduced as the supply current increases.

For example, in an implementation with an output of 2.5 V (Volts) and an input of 5 V, there is a down-conversion step size of 2.5 V across the regulator. Utilizing the relationship of step size to power efficiency:  $(5V - 2.5V) * I^2 = 2.5V * I^2$ . Whatever the current draw is across the regulator,  $(2.5V * I^2)$  Watts will be dissipated as heat from the regulator. This example wastes as much power in heat as the amount of power needed by the circuit.

As shown above, the closer the input voltage is to the regulated output voltage the less power is wasted across the regulator; therefore Low Dropout Regulators (LDO) dominate the Pass Transistor class of regulators. LDO regulators allow very small  $V_{in}/V_{out}$  ratios that minimize power loss. Good design practice calls for the use of space-saving Pass Transistor Regulators where the voltage down-conversion and current requirements are relatively small.

Switching Regulators are generally much more efficient than Pass Transistor Regulators, especially when the  $V_{in}/V_{out}$  ratio is much greater than 1 and at higher current draws. Switchers work on some form of pulse width

modulation and in conjunction with other passive components with efficiency ranges from the sixties well into the upper ninety percent range. Switching regulators require relatively large inductors and capacitors to filter the pulse width modulation. It is this extra support circuitry that causes a switching regulator to take much more space than a pass regulator.

The ideal system design approach is to use a hybrid of both Pass Regulators and Switching Regulators with the use of Pass Regulators for space efficiency and Switching Regulators for power efficiency. A classic method for portable radios is to use switching regulation to step down the battery voltage to a lower voltage distributed throughout the digital portions of the radio supplemented by the use of pass transistor regulators to further step the voltage down where lower voltages are needed.

For example, many dismounted systems operate using the 12 or 24 V outputs of the standard BA-5590 batteries. Using the power efficient Switching Regulators, this nominal 12 or 24V input power can be regulated down to 3.3V. Pass Regulators can then be used to further step the 3.3V down to supply digital circuits and processors with a range of voltages, such as 2.5V, 1.8V, and 1.2V. Note that this example assumes that these lower voltages have low current requirements. If not, the designer should consider the use of a Switching Regulator for efficiency optimization.

Although there are many other designs specific trade offs when it comes to power regulation and distribution in a size and power constrained design; the general rule is to optimize space with the use of Pass Transistor Regulators where the  $V_{in}/V_{out}$  ratio and current draw allow, and to optimize efficiency where the  $V_{in}/V_{out}$  ratio or current draw is too high to efficiently implement a Pass Transistor Regulator.

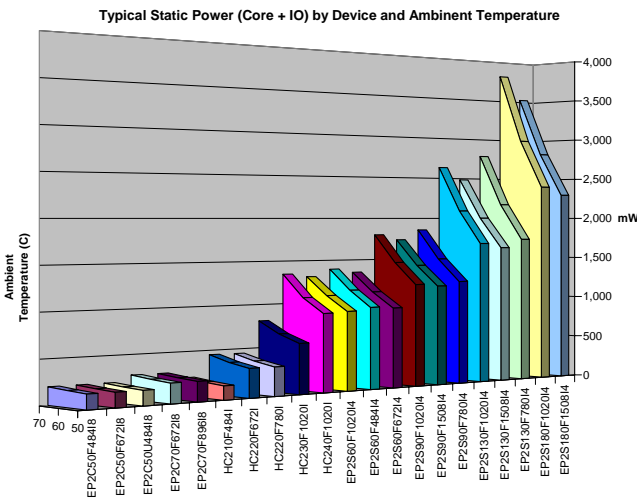
### 4. FPGA STATIC AND DYNAMIC POWER

Modern networking waveforms are pushing the limits of the performance of even the fastest Digital Signal Processors. FPGAs provide an attractive alternative since they offer reprogrammability and the simple advantage of high levels of parallelism that cannot be achieved by the essentially sequential DSP (even with pipelining and parallel execution units). The transition of signal processing functions that require up to multi-megabit per second data rates for networking radios are moving from the traditional DSP to parallel logic implementations.

This is an excellent solution for non-battery powered systems. Though, for battery powered systems FPGA

devices present a significant power penalty. That penalty is most pronounced in the static power consumption of modern FPGA devices. Static power consumption is attributed to transistor leakage current due primarily to parasitic diodes in gate junctions [6]. This static current is power dissipated as heat simply when the device is powered up and no logic is being clocked. This is especially undesirable for battery powered radios.

As FPGA devices move toward smaller transistor geometry to achieve higher chip density and faster dynamic speed, the leakage current in each transistor goes up substantially. Figure 1 shows a sampling of typical 90 nanometers geometry FPGA device power loss due to leakage current. The devices on the right of the graph are resource rich with plenty of gates, memory, and specialized DSP function blocks. These are the devices most appropriate in resources for networking waveforms; unfortunately, they are also the devices with the highest quiescent current drain.



**Figure 1: FPGA Quiescent Power Consumption**

Although FPGA devices allow complete in-circuit reprogramming, there is a penalty for this advantage. Reprogrammable FPGA devices need additional signal routing resources and much faster switching logic to compensate. This has driven the FPGA industry toward implementing the smallest transistor geometry possible. Current generation FPGA devices are at 90 nanometers and quickly moving to 65 nanometers in next generation devices.

This fine geometry combined with the extra support logic has substantially driven up the quiescent current of FPGA devices. As geometry drops below 70 nanometers the leakage current will begin to dominate and account for more than 50% of the power consumption [5]. The clocked

(dynamic) power is simply added to the base static current when logic is clocked. Both quiescent and dynamic current increase substantially as device junction temperature increases. Most battery powered handheld, manpack, and small form factor radios are environmentally sealed and have no airflow at all. This compounds the power problem for small, military implementations since quiescent and dynamic power are very sensitive to thermal conditions such as ambient temperature, heat sinking, and air flow.

Although the industry has been a bit cavalier about the tendency of Moore's law to solve these types of problems, it is precisely the drive for faster and denser FPGA devices that will continue to provide devices that also have substantial higher leakage current.

As an alternative to reprogrammable FPGAs, typical CMOS ASIC devices in the 130-250 micron geometry hold more total logic (gates+memory) with quiescent power in the milliwatt to microwatt range. Dynamic power is also substantially less in a CMOS ASIC versus a FPGA device. FPGA devices require additional routing resources and other support circuitry for reprogramming that ASIC devices do not need. These add substantially to both dynamic and static power consumption [7]. This translates into very large increases in power consumption of logic implemented in an FPGA as opposed to an ASIC. There have been some recent investigations of power reduction methods in FPGAs and have some benefit in reducing power consumption [6] but not nearly the reduction by moving logic from a FPGA to an ASIC.

Any design would benefit from a full evaluation of all waveform signal processing that a platform must host and parsing these functions into two categories, those that are common across the waveforms, and those that are specific to a particular waveform. Those functions that have some degree of commonality across waveforms could be hosted within an ASIC allowing a low power implementation. Only those functions that call for true reprogrammability on a waveform-by-waveform basis should be hosted within an FPGA. In addition, processing functionality should be evaluated for stable routines that could be hosted within an ASIC and selected when required by the waveform. These design techniques ensure that only those processing functions that must be hosted within a FPGA are allocated to the power-hungry devices, reducing the size and current drain of the FPGA device required. A combination of ASIC devices with configurable functions and smaller FPGA devices provide the flexibility required for complex waveforms in a power-efficient implementation. This type of approach will ultimately be necessary to achieve SWAP goals for battery powered portable radios.

## 5. DYNAMIC POWER SAVING METHODS

### 5.1. Frequency Scaling

Controlling the frequency of a digital circuit has a direct linear relationship to power consumption. As clock rate increases, power consumption increases in a linear relationship. Most waveforms require high processing speeds only during a portion of their activity. A clever design technique takes advantage of this variation by adjusting the clock rate over time according to the processing needs.

Certain microprocessors cannot provide dynamic clock scaling due to the complex, often phase locked relationship of internal clock distribution. For example, no current version of the PowerPC supports dynamic clock scaling. But many modern RISC processors do not have this issue. As an example, the ARM family of RISC processors has supported some form of clock scaling for many years.

The graph in Figure 2 shows the clock versus frequency relationship for a typical ARM microprocessor core power. This device has an internal clock divider that is settable under software control to the discrete values shown in the graph. There is also a companion chip for controlling core voltage. Figure 2 shows both V-F scaling and frequency scaling only (non-voltage scaled).




Figure 2: V-F Scaled Power Consumption

### 5.2. Simple Sleep Mode

Sleep mode is the simplest method of clock scaling. A sleep mode is simply a mode with extremely low clock rate (0 to a few kilohertz range) to allow very low power consumption. Sleep is generally a hardware supported state in which the processor code execution is effectively halted until an event occurs. The sleep mode typically has power consumptions in the low milliwatts to even microwatts of power consumption. Sleep state is exited on some form of interrupt such as from a periodic interval timer or external communication port.

The use of sleep mode can have substantial impact on reducing power consumption. For example, a RISC processor running at 400 MHz may consume 1 watt while active, but have a sleep state power consumption of only 500 microwatts. The average power savings is directly proportional to the amount time spent in sleep state. If one spent half the time in sleep state the average consumption would be about ½ watt.

### 5.3. Voltage Scaling

Voltage scaling is nothing new given the on-going reduction in core and bus signaling voltages in modern digital systems. These both result in substantial power reduction. Power consumption is proportional to  $V^2$  due to the simple equation  $P=V^2/R$ . Effectively, as voltage decreases, current decreases proportionally and therefore power decreases exponentially. More specifically, for CMOS circuits [3] the relationship is  $P = A \cdot C_L \cdot V^2 \cdot F$  where A is the activity constant,  $C_L$  is the load capacitance, V is the supply voltage, and F is the switching frequency. As can be seen from this equation, power is proportional to frequency and exponential to voltage.

There is a direct relationship between minimum supply voltage levels and the switching frequency of the digital circuit. As clock frequencies are reduced, the required  $V_{\text{threshold}}$  levels decrease. This is a sometimes complex issue but related to the probability of error and somewhat to edge rate of digital switching circuits. It is sufficient for voltage scaling purposes to understand that as the clock frequency decreases it may be possible to decrease the core supply voltage. Thus, we may incorporate some form of voltage-frequency (V-F) scaling to microprocessors and other digital CMOS circuits. The upper line in Figure 2 shows the power for just frequency scaling with a fixed core voltage. The lower line (V-F Scaled) shows the power improvement of scaling both voltage as well as frequency in a microprocessor.

Saewong and Rajkumar [4] point out that although this relationship is strictly true, it may not have the intended effect. The instantaneous power is lower with lower V-F, but it may not be true that any given code segment will take less power at a lower V-F setting since, at lower clock speeds, the code will take longer to run. If a code segment takes 500 milliseconds to run at frequency A and it takes half the power over a period of 1 second at frequency A/2, then the power consumed is the same.

The other significant issue with microprocessor power scaling is that the software environment for a SDR radio is decidedly real time. In other words, there are likely to be critical code segments that must execute at high clock

speed. This means that either frequency or V-F scaling algorithms must have extremely fast response times for critical code sections. The monitoring and reaction for dynamic power control is computationally expensive [5].

## 6. SHOULD SOFTWARE CARE ABOUT POWER?

The previous section discussed dynamic methods for reducing power consumption in hardware. Although these techniques can occasionally have closed loop control by hardware circuitry, the dominant method is to use software to control these primitive hardware functions to manage and minimize average power consumption.

### 6.1. Static Software Methods

Software defined radios by their very nature should have software control of circuitry for both voltage and frequency scaling. But in order to take advantage of this control, the waveform code should have the concept of and support for power management.

Simple waveform awareness of its own processing load and performance requirements will allow the radio to operate in the most power efficient manner allowable. Battery powered radio users expect the software to use the hardware resources in a way that prolongs battery life.

### 6.2. Simple State Monitoring

As already noted, dynamic continuous V-F scaling can be computationally expensive. A simpler but potentially more effective method is to have simpler predetermined state transitions for V-F scaling. Section 5.2 discussed simple sleep mode clock scaling. This could be extended further to V-F scaling with little overhead other than V-F rescaling on idle state routine entry and exit. All other external routines can then run at maximum V-F.

Sleep states are generally hardware supported states in many microcontrollers. This method could be extended to software only states such as idle process or suspend states. These could be incorporated into the operating system such that the default condition is to maintain low V-F states. Of course this makes the speed exiting the slow state to be potentially problematic in that the end result could be longer interrupt latency. But this can be simply analyzed and resolved as compared to a dynamic V-F scaling algorithm.

### 6.3. Radio State Transitions

The prior section discussed generic system states available in any software system. The modern day military radio may have multiple distinct radio states as well. For example, a half duplex radio would have at least the states of transmit, receive and sleep. The radio spends time in

each state and the waveform implementation has this knowledge implicitly. Using these states to do V-F scaling allows the radio to minimize power in each state. This combined with the other idle and sleep states can substantially reduce power consumption.

These state control methods can be extended from the processor control to other hardware in the system. A software system and radio waveform that is fully aware of states and transitions can use this to effectively reduce power consumption and greatly prolong battery life.

## 7. CONCLUSIONS

The increasing demands of modern networking waveforms on digital radio electronics challenge small form factor and battery powered digital hardware designers. Although Moore's law helps reduce the space and power issues, SWAP sensitive radio designs for high bandwidth networking radios will not succeed without the interaction of power-aware software. Proactive hardware design combined with integrated waveform support for managing power can greatly prolong battery life and thus the operational value of software defined radios for networking.

## 8. REFERENCES

- [1] G. Magklis et al., "Dynamic Frequency and Voltage Scaling For A Multiple Clock Domain Microprocessor," *IEEE Micro*, IEEE Computer Society Press, 2003, pp. 62-68.
- [2] GAO-05-669, "Resolving Development Risks in the Army's Networked Communications Capabilities Is Key to Fielding Future Force" *United States Government Accountability Office*, Defense Acquisitions, June 2005.
- [3] B.M.G.R. Gonzalez and M. A. Horowitz. "Supply and Threshold Voltage Scaling for Low Power CMOS", *IEEE Journal of Solid Stat Circuits*, 32(8), August 1997.
- [4] S. Saewong and R. Raikumar, "Practical Voltage Scaling for Fixed-Priority RT-Systems", Ninth IEEE Real-Time and Embedded Technology and Application Symposium, May 2003, pp. 106-115.
- [5] A. Andrei et. al., "Quasi-Static Voltage Scaling for Energy Minimization with Time Constraints", *Design Automation and Test in Europe*, Vol 1, Mar 2005, pp.514-519.
- [6] Andres Garcia et. al., "Power Consumption Management On FPGAs", *Proceedings of the 15<sup>th</sup> International Conference on Electronics, Communications, and Computers*, 2005
- [7] A. Gayasen et. al., "Tools and Architectures for Power Minimization: Reducing Leakage Energy in FPGAs Using Region-Constrained Placement", *Proceedings of the 2004 ACM/SIGNA 12<sup>th</sup> Int. Symp. On FPGAs*, Feb 2004