

Dynamic Power Management Optimizes Performance vs. Power in Embedded Applications of Blackfin™ DSPs

by David Katz (david.katz@analog.com)
Rick Gentile (richard.gentile@analog.com)

The rapidly increasing consumer demand for products employing multimedia embedded processors calls for both high performance and low power consumption. But the increased computational complexity and faster clock rates necessary for high performance processing are hard to achieve using tactical power-saving design schemes. What is needed is a strategic way to manage power consumption to optimize performance versus power for the specific embedded application. Such an approach is achievable with the inherent *dynamic power management* capabilities of the Blackfin DSP family.

Blackfin DSPs are fixed-point, dual-16-bit-MAC/dual-40-bit-ALU digital signal processors. They are ideal for power-sensitive multimedia applications because they support a multi-tiered approach to power management that adjusts performance based on system needs. Let's take a look at some of the key power considerations in embedded systems and see how the Blackfin family uses dynamic power management to address them.

What are some typical strategies for saving power?

1. Changing Frequency and Voltage

Modern DSPs are normally designed in a process using CMOS FET switches, which are either fully *on* (and very lightly loaded) or fully *off* (except for leakage currents) during the steady state. The static power dissipation (quiescent power while the processor is idle) is typically much lower than the dynamic power dissipation caused by the charging and discharging of FET load capacitances at very high switching frequencies when the device is actively switching and voltages are slewing.

The charge (Q) stored in the device's equivalent load capacitance equals the capacitance multiplied by the voltage stored across it (which is the DSP's core supply voltage, V_{core}),

$$Q = CV_{core},$$

Since device current to charge this capacitance is defined as the rate of change of charge with respect to time, the dynamic current, I_{dyn} , is given by

$$I_{dyn} = dQ/dt = C(dV_{core}/dt)$$

The rate of capacitor voltage change with respect to time, dV_{core}/dt , is a measure of how fast the capacitor is being charged and discharged. For a given clock frequency, F , the fastest a complete charge or discharge can take place is one clock cycle. Therefore,

$$\begin{aligned} dV_{core}/dt &= V_{core}(F) \\ I_{dyn} &= C(dV_{core}/dt) = CV_{core}F \end{aligned}$$

Finally, dynamic power dissipated is proportional to $V_{core} \times I_{dyn}$, or

$$P_{dyn} \propto CV_{core}^2F$$

Thus, it is apparent that dynamic power dissipation is proportional to the square of operating voltage and linearly proportional to operating frequency. Therefore, lowering F will decrease the dynamic power dissipation *linearly*, while reducing V_{core} will lower it *exponentially* (see Figure 1).

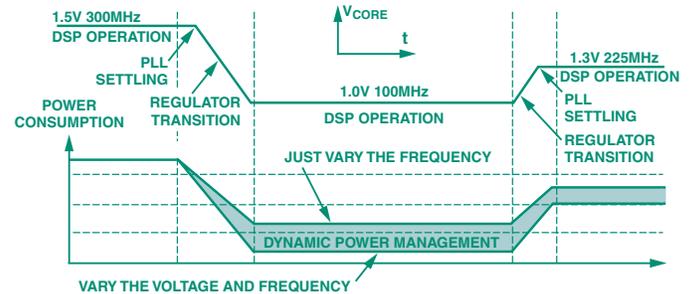


Figure 1. Effects of V and F changes on power consumption.

Consider the three different DSP functions combined in Figure 1, all with very different performance needs:

F0(x)	1.5 V	300 MHz
F1(y)	1.0 V	100 MHz
F2(z)	1.3 V	225 MHz

For instance, F0(x) might be a video-processing algorithm, F1(y) could be a monitoring mode (where the DSP is collecting data and doing minimal processing), and F2(z) might be a process to stream compressed video out of a serial port.

Changing only frequency (and not voltage) in a power-sensitive application is useful when the DSP has an extended period of monitoring activity. That is, if the DSP were waiting for an external trigger, it would not need to run at maximum frequency.

However, in some battery-powered applications, simply changing frequencies may not be enough to save power. For example, if an application is running three sections of code, reducing the operating frequency for any one of these sections means that particular section of code will take longer to execute. But if the DSP is running longer, the same amount of power will be expended when the three sections are complete. If, for example, the frequency is reduced by a factor of two, the code will take twice as long to execute, so no net power savings is realized.

On the other hand, considerable power savings can be achieved by reducing voltage as well as frequency. This savings in power consumption can be modeled by the following equation:

$$P_R/P_N = (F_{CR}/F_{CN})(V_{DDR}/V_{DDN})^2(T_{FR}/T_{FN})$$

where

- P_R/P_N is the ratio of reduced power to nominal power
- F_{CN} is the nominal core clock frequency
- F_{CR} is the reduced core clock frequency
- V_{DDN} is the nominal internal supply voltage
- V_{DDR} is the reduced internal supply voltage
- T_{FR} is the duration running at F_{CR}
- T_{FN} is the duration running at F_{CN}

For example, Figure 2 shows a scenario with the following characteristics:

- $F_{CN} = 300$ MHz
- $F_{CR} = 100$ MHz
- $V_{DDN} = 1.5$ V
- $V_{DDR} = 1.0$ V
- $T_{FR} = 3$
- $T_{FN} = 1$

Thus

$$(P_R/P_N) = (100/300)(1.0/1.5)^2 \times (3/1) = 0.44 \rightarrow 56\% \text{ savings!}$$

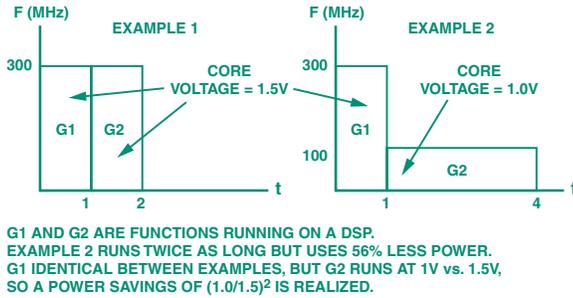


Figure 2. Power dissipation vs. frequency and time.

Since Blackfin DSPs not only have a programmable operating frequency, but also allow core voltage to be changed in concert with frequency changes, less power will be consumed when running a section of code at a lower frequency *and* a lower voltage, even if execution time is longer. The voltage-frequency transition is handled automatically on the ADSP-21532, while for the ADSP-21535, a simple sequence is followed. It is, of course, important to remember that developers must ensure the integrity of peripheral channels connected to external systems during any system clock frequency change.

A videophone application illustrates how the ability to vary both operating frequency *and* operating voltage can be exploited to greatly extend battery life. If, for example, the maximum performance (maximum core clock frequency) is only required during a video connection, the core frequency can be lowered to some preset value when using the phone for a voice-only transaction. For operating time-insensitive value-added features only (e.g., a personal organizer), the frequency can be further reduced. Each of these PLL frequency transitions can be accomplished in less than 40 microseconds on Blackfin DSPs.

IMPLEMENTATION

Blackfin clock generation unit

The *clock-generation unit*, which houses the phase-locked loop (PLL) and associated control circuitry, is an integral element of *dynamic power management* in Blackfin DSPs. The PLL is highly programmable, allowing the user to control the processor’s performance characteristics and power dissipation dynamically.

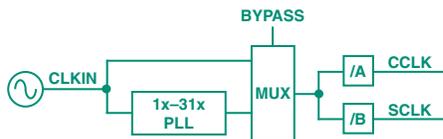


Figure 3. Functional block diagram of ADSP-21532 clock-generation unit.

Figure 3 shows a simplified block diagram of the ADSP-21532 clock generation unit. An input crystal or oscillator signal (10-to-33 MHz)

is applied to the CLKIN pin. The PLL then multiplies this signal by a programmable factor of 1x to 31x. Then, separate A and B dividers independently generate core-clock (CCLK) and system/peripheral-clock (SCLK) frequencies. Control logic ensures that the system clock frequency will not exceed the core clock frequency.

The great advantage in this approach is that CCLK and SCLK can be changed “on-the-fly,” with very little cycle overhead. Thus, designers need not think twice about changing clock frequencies in order to meet different performance requirements for different segments of their code. The resulting linear savings in dynamic power dissipation comes at no implementation cost, from the designer’s perspective.

Another feature of the Clock Generation Unit is that it can be bypassed to allow the CLKIN signal to pass straight through to CCLK. This capability permits use of a very low frequency CCLK during inactive operation intervals, to further reduce overall power dissipation.

2. Flexible Power Management Modes

Many applications involve a set of operating modes that differ markedly with respect to processing needs. Consider the system of Figure 4, where a battery-powered sensor contains a DSP that acts as the central processor. One of the DSP peripherals might be used to sample parameters of the surrounding environment. In this “Mode A,” which requires very low processing power, the DSP might be reading in sporadic packets of telemetry data. When it has read enough data to invoke a computational algorithm, the DSP would then enter “Mode B,” a processing-intensive computational mode. It is likely that a “Mode C” also exists, to provide ultra-low power dissipation when no sensor information is expected and no processing is required.

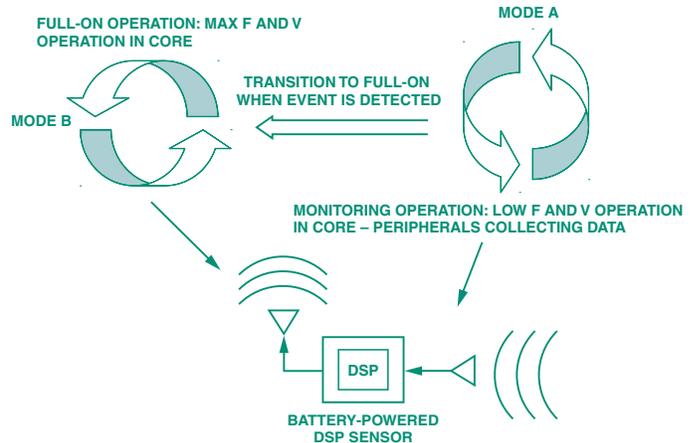


Figure 4. Sample DSP application with different operating modes.

Blackfin DSPs have four distinct operating modes (corresponding to four different power profiles) that provide selectable performance and power dissipation characteristics. Table 1 summarizes the operational characteristics of each mode.

Operating Mode	Core Clock	System Clock	Power Savings
Full-On	Enabled	Enabled	Minimum
Active (PLL bypassed)	Enabled	Enabled	Medium
Sleep	Disabled	Enabled	High
Deep-Sleep	Disabled	Disabled	Maximum

Full-on mode

Full-on is the Blackfin’s maximum performance mode. In this execution state, the processor and all enabled peripherals run at full speed. The PLL is enabled, so CCLK runs at a multiple of CLKIN.

Active mode

In *Active* mode, the PLL is enabled but bypassed, so CCLK comes directly from CLKIN. Because CLKIN is sourced from an external oscillator input no greater than 33 MHz, this mode offers significant power savings. The system-clock (SCLK) frequency is also reduced, because it never exceeds CCLK. With the PLL bypassed in this mode, it is safe to change the PLL multiplier ratio; however, the changes do not take effect until the DSP is back in *Full-on* mode. In Active mode, not only can the PLL be bypassed—it can be disabled, for incremental power savings.

Sleep mode

The *Sleep* mode significantly reduces power dissipation by disabling CCLK, which idles the DSP core. However, SCLK remains enabled so that data transfer can still take place in L2 memory and peripherals. To exit from Sleep mode, the Blackfin provides a DSP core wake-up capability, which operates independently from the core’s *event controller*.

Deep-sleep mode

The *Deep-sleep* mode maximizes power savings by disabling the PLL, CCLK, and SCLK. In this mode, the processor core and all peripherals except the *real-time clock* (RTC) are disabled. In Deep-sleep mode, the DEEP_SLEEP output pin is asserted, to permit external power-mode control. Deep-sleep mode can be exited only by an RTC *interrupt* or *hardware-reset* event. An RTC interrupt causes the processor to transition to *Active* mode; a hardware reset initiates the hardware-reset sequence.

3. Separate Power Domains

Blackfin DSPs support multiple power domains, including a dedicated *phase-locked-loop* (PLL) power domain, a *real-time clock* (RTC) that can be powered by a small, external coin-cell battery, and separate domains for the various peripherals. The core processor also has its own power domain. Using multiple power domains maximizes flexibility while maintaining direct connectivity with a wide variety of commercially available devices, such as SDRAM and SRAM memories. As shown in Figure 5, the separate power domains allow the Blackfin’s core voltage to be varied without disrupting connections to external devices. This is a critical advantage, because—as noted above—the power consumed by a processor is proportional to the square of its operating voltage.

4. Using an Efficient Processor Architecture

Another often-overlooked means of reducing power consumption for a given application is to choose an *efficient processor architecture* for that application. Such features as specialized instructions and fast memory structures can reduce power consumption significantly by lessening overall algorithm execution time. Moreover, power-conscious applications make it imperative to structure algorithms efficiently, taking advantage of native architectural features, such as hardware loop buffers and instruction/data caches. This is

important—complex algorithms often consume more power, since they use more resources. If an algorithm is optimized, it takes fewer instructions to execute. The sooner it completes all its steps, the sooner the core voltage and frequency can be reduced.

Power consumption can be further optimized in architectures that support selective disabling of unused functional blocks (e.g., on-chip memories, peripherals, clocks, etc.).

Blackfin DSPs provide additional power-control capability by allowing dynamic scheduling of clock inputs to each peripheral. This allows finer control of power dissipation. Also, internal clocks are routed only to enabled portions of the device. For example, on the ADSP-21535, the 256KB on-chip L2 memory consists of eight 32KB banks. These banks are only clocked when they are accessed, a feature that can result in significant power savings.

5. Profiling Tools

Providing yet another way to optimize power consumption, the Blackfin VisualDSP++ tool suite can profile applications to determine the exact processing requirements for each section of an algorithm. The tools allow a system designer, in real time, to quantify how much time is spent in any given code segment. Using this technique in battery-powered applications, the core and system frequencies, as well as the core voltage, can be modified to “match” the minimum values required to perform the task.

6. Intelligent Voltage Regulation

Beginning with the ADSP-21532, Blackfin DSPs provide on-chip core-voltage regulation. The first Blackfin DSP, the ADSP-21535, requires an external power management chip to allow dynamic control of the core voltage levels. The ADP3053 is a companion chip that supports power management for the ADSP-21535. The DSP will use up to 3 pins to control the power levels being provided by the ADP3053. The part allows 100-millivolt core-voltage increments, from 0.9 V to 1.5 V. In addition, the ADP3053 provides a low-noise PLL supply.

CONCLUSION

Designers using DSPs do not have to sacrifice power for performance. There are many alternatives to help them balance these often-conflicting demands. By viewing power management strategically, instead of tactically, significant savings can be achieved. The Blackfin DSP family provides an excellent platform for realizing low-power, high-performance embedded applications.

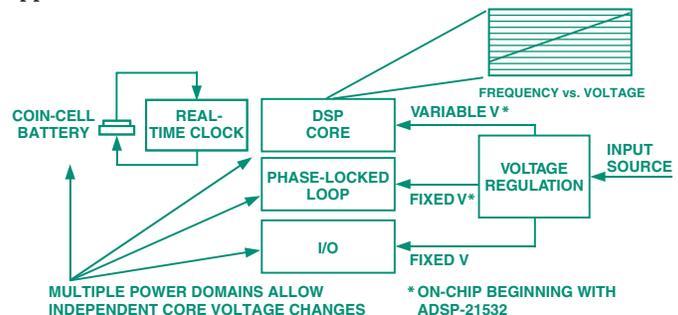


Figure 5. Illustration of multiple power domains within the Blackfin DSP.