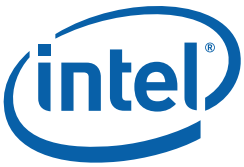

White Paper

Dirk Blevins
Technical Marketing Engineer
Michael Loewe
Technical Marketing Engineer
Jack R. Johnson
Technical Marketing Engineer
Intel Corporation

Designing Systems Without a Suspend Supply

December 2008

321082

Executive Summary

Designing with Intel components allows a lot of power management flexibility. One of the down sides of this flexibility is the need for independent power supply wells driving the chip. A large number of embedded customers do not require this flexibility and therefore are sacrificing cost and board space. If the rails are not needed, adding them does not provide a positive return on investment. For example, if the chip requires a 3.3V_main power rail and a 3.3V_standby power rail. This paper focuses on how to collapse these into a single power rail. For the purposes of this paper, this concept will be referred to as "associated power rails".

Collapsing the main power rail and the standby power rail into a single power rail can save money and provide more board space.

A common question is, "How can I design an Intel-based system by collapsing these power rails?" The Intel design documents at a platform level do not take this type of design into consideration. This document focuses on the process for removing the extra components associated with the suspend wells. It focuses on the I/O Controller Hub (ICH) documents showing power and reset sequencing that includes both standby and main power rails. The abstract principles in this paper are illustrated in a practical example using an Intel® Xeon® Processor 5000 sequence CPU's, Intel® 5100 Memory Controller Hub Chipset, and ICH9 I/O Controller Hub-based solution.

Notes to set the stage:

1. It is imperative to note that in an ICH9 power sequence, the general rule of thumb is that the higher voltage rails must ramp

ahead and stay above the next higher voltage. For powering down, the same principle applies but in the reverse order.

2. The VccSus supplies should never be active while the VccRTC supply is inactive.
3. The associated 3.3 V and 1.05 V supplies are assumed to power up or down 'together'.

Contents

What Are Power States?

Before we can explore the remaining sections and what makes the combining of the power rails possible, we must first comprehend power state operation.

Power states in the context of this document are referring to that of the Advanced Configuration and Power Interface (ACPI) specification Revision 3.0B. In this specification a concept of Global State control is used to allow further power management by the I/O, CPU, and Sleep Power. In addition, allowances are made to support non-ACPI systems, known as Legacy Mode.

Figure 1 shows a state diagram taken from the ACPI specification of the global power states.

Figure 1. Global States and Their Transitions

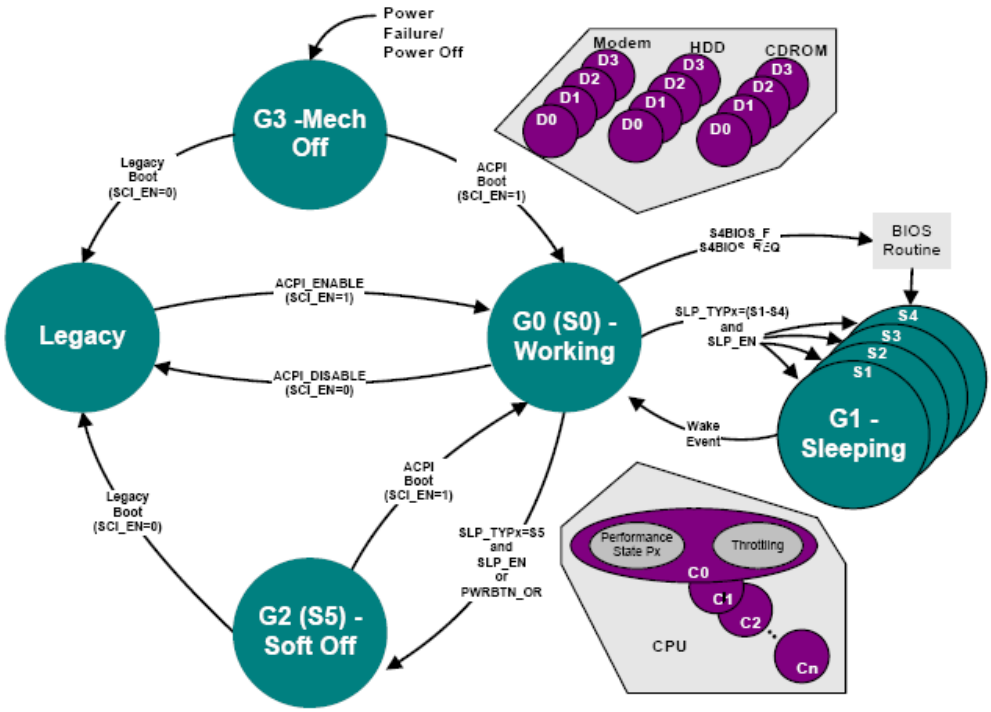

The four global states in the context of the power rail usage are:

- G0 (S0) – Working State. This is the operational state when a computer is functioning. I/O power states (known as the D states) and CPU power states (known as the C states) are being managed dynamically. All power rails are active.

- G1 (S1, S2, S3, S4) – Sleep State. This is a series of sub states which allow various amounts of latency to be obtained before restoring the system back to the Working State. Some power rails are on but not all of them.
- G2 (S5) – Soft Off. This is the state where the machine appears that the system is off. However, resume well power (also known as Standby Power) is still maintained but all other power is removed including power to the DDR, with the exception of the Real Time Clock (RTC) well.
- G3 – Mechanical Off. This is the state were all power is removed from the system. Thinking of a desktop Personal Computer this would be when the plug is removed from the wall. All power rails are removed. The only power active in this configuration is the RTC power rail supplied by the battery.

Legacy Mode, the D states, and the C States are outside the scope of this paper and are not discussed.

The Sleep States subsection of G1 and G2 are various low power states which the system can be placed in under user control. These sleep states are defined as follows:

- S1 – Lowest Latency wake state, CPU is basically halted to save power. All rails are active.
- S2 – Similar to the S1 state but the CPU and system caches are cleared and disabled. Some power rails would be removed but not all.
- S3 – Sleep state where all system context is lost except system memory. OS context is maintained within the memory and is not reloaded at startup. Some wells are on such as DDR wells, suspend wells, etc.
- S4 – This sleep states powers down the memory and all I/O. OS context is stored off to non-volatile memory so that an entire reload of the OS is not required. All wells are off but the resume well and the RTC well.
- S5 – This sleep state looks very much like S4 with the exception of the fact that the OS context is not stored off. The OS must be reloaded. The resume well and the RTC wells are still active in this state.

Various transitions described in the ACPI specification are allowed among these S states.

The key point to the discussion on the S and G states for this paper is that designs desiring to remove power wells shall only experience what appears as a G3 to G0(S0) transition and a G0(S0) to G3 transition. G1 and G2 never are truly entered and therefore the resume wells are not required for this type of operation.

The Goal

This section describes the desired outcome in combining power rails from the I/O Controller Hub's perspective. The designer of a combined power rail system is predominately concerned with three sequences from the I/O Controller Hub's specifications:

- Power-Up Sequence
- G3 to S0 Sequence
- S0 to G3 Sequence

Many other timing sequences exist, but the intent of this design is to make the system look either on (S0) or off (G3). The remainder of this section takes these three timing diagrams and shows how they look in the non-combined (or Normal Operation) operation and the combined operation.

Power-Up Sequence

In a normal system (see Figure 2), the V5REF_SUS, VccSus3_3 and the VccSus1_1 suspend power rails become active prior to their associated main well supplies (from a voltage level perspective) V5REF, VCC3_3 and the VCC1_1 (and other power) to allow the ICH to work properly.

For a solution not requiring standby functionality, these rails are combined. Care must be taken to assure the power rails power up from the higher voltage (5V) to the lower voltage (1.1V). Failure to do so may lead to excessive leakage and/or latch-up causing part failure. Figure 3 shows the desired timing diagram of such a solution.

Figure 2. Power Sequencing and Reset Signal Timings - Normal

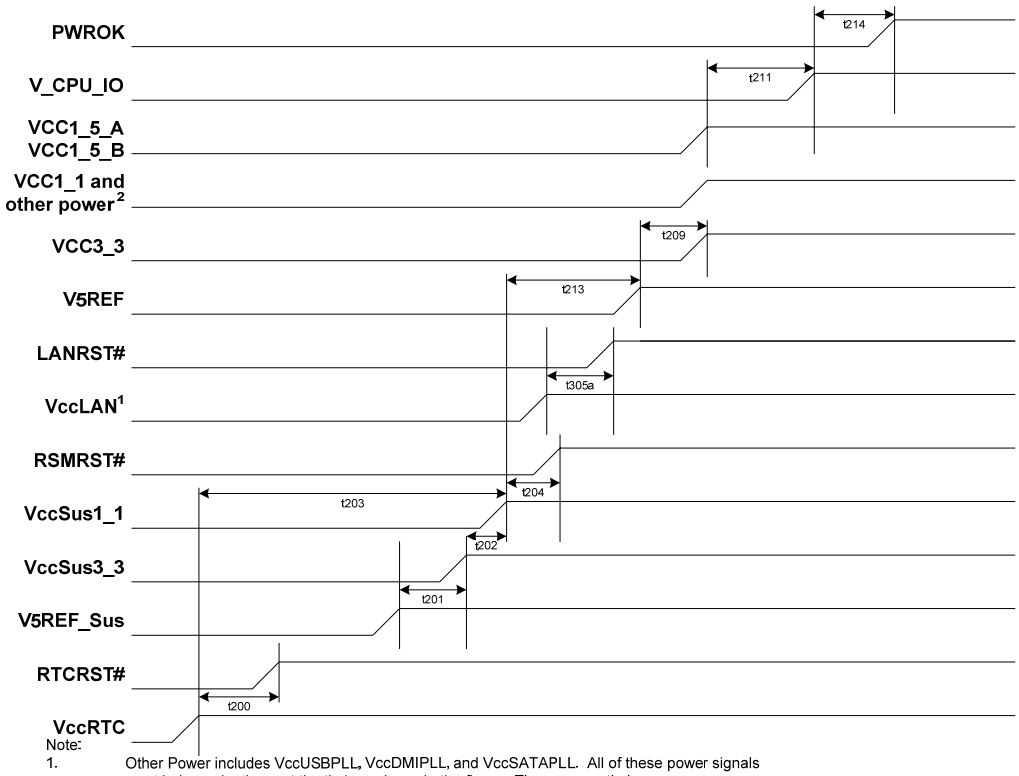

Note:
1. Other Power includes VccUSBPLL, VccDMIPLL, and VccSATAPLL. All of these power signals must independently meet the timings shown in the figure. There are no timing interdependencies between Vcc1_1 and these other power signals. There are also no timing interdependencies for these power signals, including Vcc1_1, Vcc3_3, and Vcc1_5_a/ Vcc1_5_B.
2. PWROK must not glitch, even if RSMRST# is low.

As stated previously, care needs to be taken to assure power is sequenced from high to low. The following is recommended:

- The VccRTC well should come up first and must be to an active power level prior to either the VCC3_3 or the VccSus3_3 rails becoming active.
- 5VREF and 5VREF_Sus are combined and are allowed to come up together.
- VccSus3_3 and VCC3_3 (and 3.3V VccLAN signals) are combined and become active only after all 5V power rails are active.
- All remaining lower voltage supplies become active simultaneously.

The sequence shown in Figure 3 is allowed due to many of the minimum timing values having a requirement of 0ms. In theory, this requirement allows the power rails to power up simultaneously. In fact, the only signals which have non-zero minimums are RTCRST#, RSMRST#, LANRST# and PWROK. Care must be taken with the above statement as the specification includes footnotes. These notes are as follows:

- VCC1_5 must power up before V_CPU_IO or after V_CPU_IO within 0.7V. Also V_CPU_IO must power down before VCC1_5 or after VCC1_5 within 0.7V.
- 5VREF must be powered up before VCC3_3, or after VCC3_3 within 0.7V. Also, 5VREF must power down after VCC3_3, or before VCC3_3 within 0.7V. See the appropriate Platform Design Guide for details.
- The VccSus supplies must never be active while the VccRTC supply is inactive.

The above notes show that in theory, these signals could be powered up simultaneously. In practice, it is recommended that they be slightly sequenced from high to low on the power up and from low to high on the power down to assure these two notes requirements are met.

As stated previously, the VccRTC well must always be active and needs a back-channel to tie this supply to the VccSus supplies in the event of the battery dying. Please reference the Platform Design Guide for the circuit to meet this requirement.

Figure 3. Power Sequencing and Reset Signal Timings - Combined

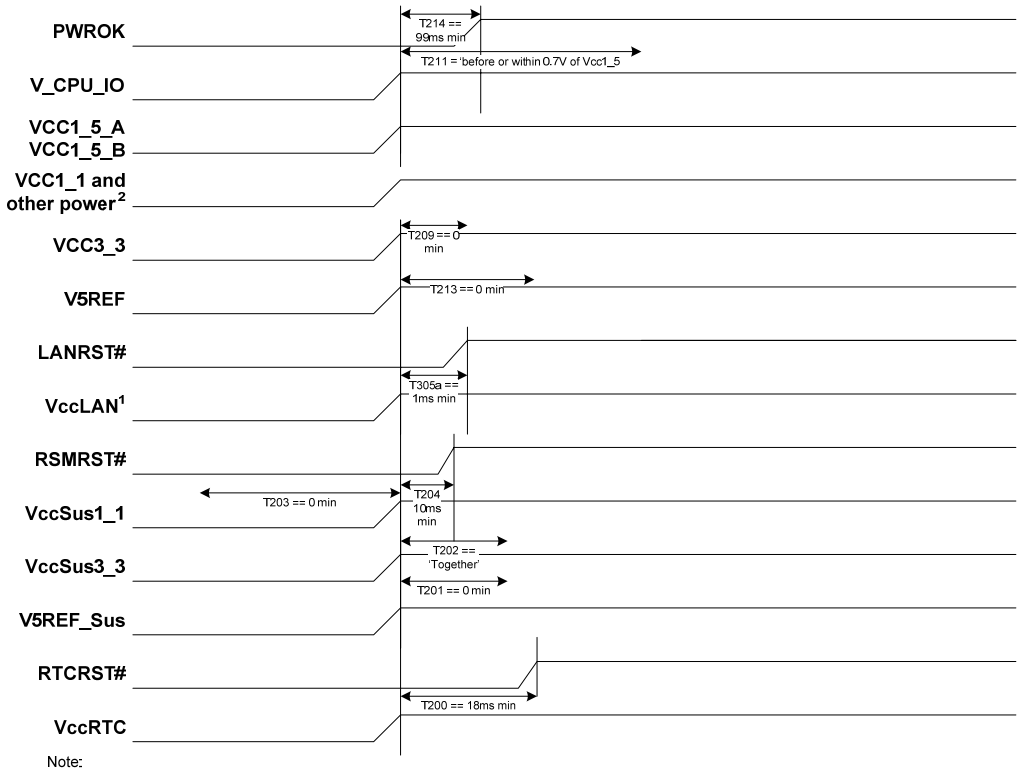

Note:
1. Other Power includes VccUSBPLL, VccDMIPLL, and VccSATAPLL. All of these power signals must independently meet the timings shown in the figure. There are no timing interdependencies between Vcc1_1 and these other power signals. There are also no timing interdependencies for these power signals, including Vcc1_1, Vcc3_3, and Vcc1_5_a/Vcc1_5_B.
2. PWROK must not glitch, even if RSMRST# is low.

G3 to S0 Sequence

Now that a firm grasp of the power rail sequencing is understood, it is important to understand additional relationships and how those relationships change due to the deviations from the Platform Design Guide recommended solution.

Figure 3 demonstrates that per specification, all of the Suspend (Sus) and Main power rails can be combined. Figure 2 illustrates the behavior of a standard system when normally connected. This concept is extended to create a solution combining the associated power rails.

Figure 4. G3 (Mechanical Off) to S0 Timings - Normal

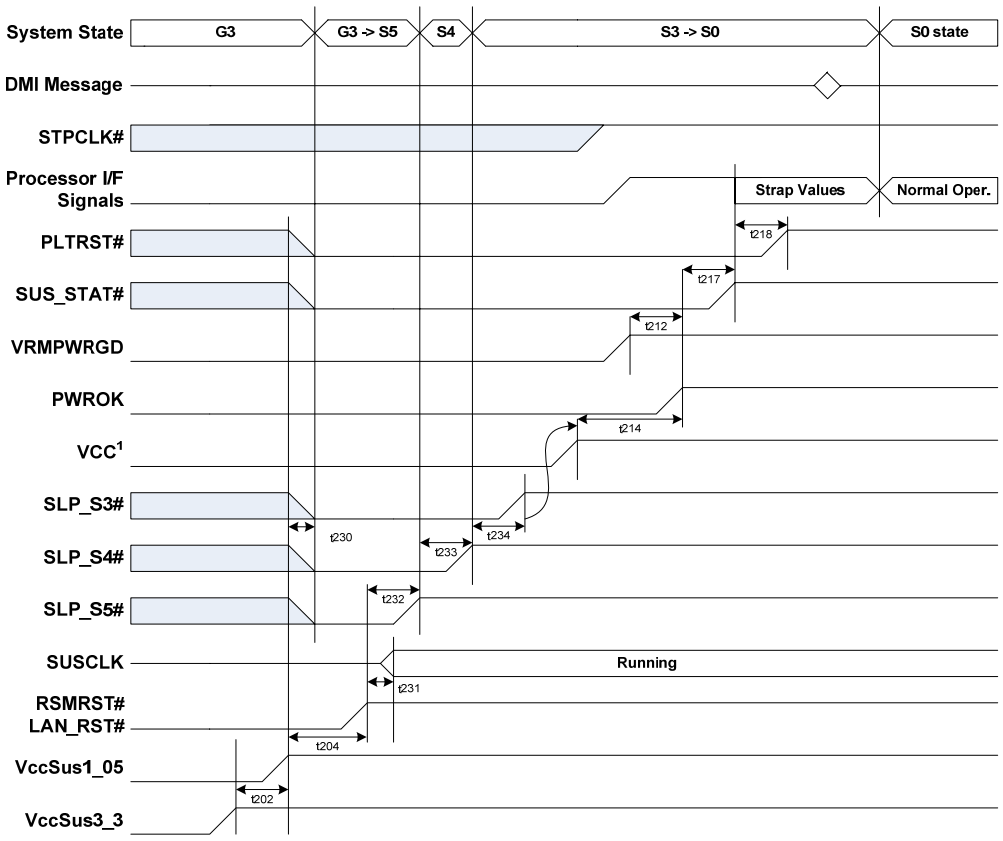

Note:
1. Vcc includes Vcc1_5_A, Vcc1_5_B, Vcc3_3, Vcc1_05, VccUSBPLL, VccDMIPLL, VccSATAPLL, V5REF, and V_CPU_IO

Figure 4 illustrates the modifications to the normal power up sequence. In Figure 5, the VCC rails come up simultaneously with the VccSus3_3/VccSus1_05. The system must use something other than the SLP_S3# signal to enable the VCC voltage rails. The enabling of VRMPWRGD along with PWROK is handled separately from the other power rails and still relies on SLP_S3#.

Enabling these power rails earlier than the Platform Design Guide's recommended solution does not compromise system integrity, as long as the sequencing and other critical relationships remain correct. For example, SUSCLK must be running for S5, S4 and S3 to operate properly. VccSus3_3 and SUSCLK are inter-dependant. Therefore, VccSus3_3 must be up for SUSCLK to start running.

Figure 5 illustrates the theoretical timing diagram modifications for the system to come up using combined the combined rail approach.

Figure 5. G3 (Mechanical Off) to S0 Timings - Combined

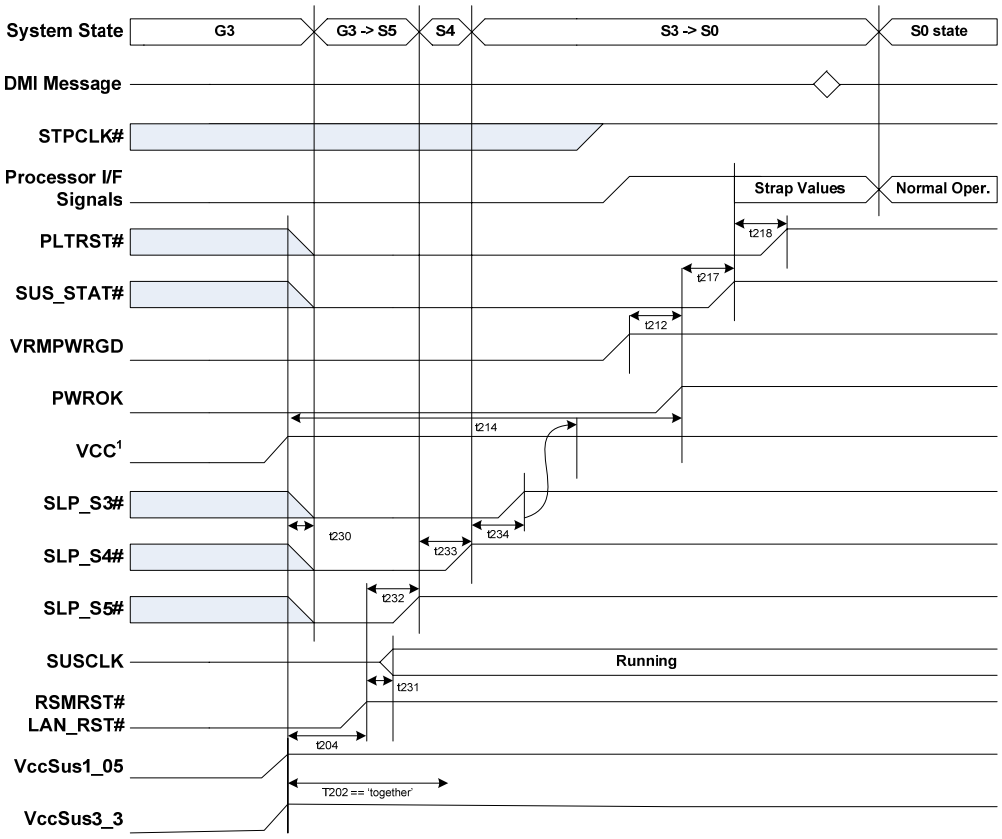

Note:
1. Vcc includes Vcc1_5_A, Vcc1_5_B, Vcc3_3, Vcc1_05, VccUSBPLL, VccDMIPLL, VccSATAPLL, V5REF, and V_CPU_IO

S0 to G3 Sequence

Figure 6 illustrates the key requirement when powering off your system. Notice the denotation (Desktop Only) in the figure's title. The I/O Controller Hub components are designed to be utilized with Server, Desktop, and Mobile applications. This denotation is there on certain diagrams and it is important to know the product being used.

Figure 6 re-iterates the requirements for turning off the voltage rails from Low to High. Figure 6 shows that the PWROK signal as well as the VRMPWRGD signals should both go inactive before the remaining VCC supplies go inactive. The timing for T294 is only 20ns so in reality so they are almost simultaneous. The key is assuring the ordering is handled correctly and since the Low voltage rails come up last it makes sense for them to go down first.

There is no change to this timing relationship between the normal and combined solutions.

Figure 6. S0 to G3 Timings (Desktop Only) – Normal Operation

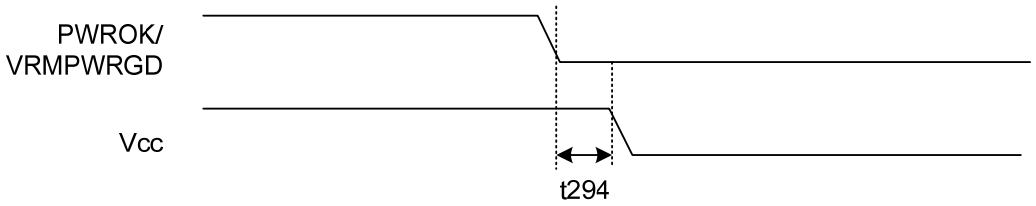

Solution

There are several solutions which meet the system timing requirements and allow the design to function without suspend wells. This section of the paper describes the guidelines and pitfalls for a manufacturer to properly design a solution.

The first assumption is the availability of a higher voltage rail prior to the system being enabled. The most common voltage used is +12 volts. This voltage is sufficient to allow good power efficiencies along with acceptable current capacity. A power sequencer can be used to generate all of the derivative voltages in sequence. The sequencer is also used to check for correct operation of each rail prior to bringing up the next rail. Sequencers are available from a variety of manufactures.

The primary +12V volt supply is used to drive the sequencer and all of the remaining voltages are generated from +12 volt supply. The sequencer is enabled by an external signal typically generated by a system controller. Once the sequencer begins operation, the goal is to transition to the S0 state where the CPU begins normal operation.

The simplification of a design not using standby rails is that several parts of a standard solution are not required. For instance, in a standard design the DIMMs may need to be powered even though the CPU is no longer running. This is known as the S3 standby mode and in many embedded systems do not require this functionality. Several deeper sleep modes can also be

omitted. This design basically assumes the system is either 'on' (S0) or 'off' (G3).

Once all of the voltage rails are valid, the other control signals like RSMRST#, PWRBTN#, SYSRST#, etc. can be released in the appropriate time. Depending on the system design the exact order and sequence may be adjustable. A few important items to keep in mind when bringing up a system are shown below:

- The I/O Controller Hub must receive a valid wake signal in order to bring the system out of sleep mode. There are several valid wake signals including PWRBTN# and WAKE#. Each one has different characteristics so it is necessary to review their function to make sure the system begins operation correctly.
- The SLP_S3# signal may be used to enable the VRM power supply. This supply is specific to the CPU and uses VID signals from the CPU in order to set the proper voltage. The VRM may also be brought up by the sequencer.
- An interesting part of the startup sequence involves the BSEL signals from the CPU. These signals indicate which BCLK frequency needs to be supplied to the CPU for correct FSB operation. The BSEL signals are routed to the clock chip as well as the Intel® 5100 MCH Chipset. These signals are not valid until VTT is stable and then they are latched into the clock chip and finally into the Intel® 5100 MCH Chipset. The clock chip needs to be up and running before the Intel® 5100 MCH Chipset is enabled so it is important work thru the details of the ordering. In addition, the frequency select inputs to the clock chip become clock outputs after a few milliseconds so care must be taken in this area of the design.

Summary

In conclusion, it is possible to design an Intel-based solution merging the suspend wells with the main wells. This solution reduces overall cost and board space of designs not requiring those features. The customer needs to keep in mind that deviation from Intel's validated design guidelines must be weighed against the need to reduce cost and/or save board space. With careful design practices, using the detailed numbers in the chip specifications and the above template as a guide, the desired results can be achieved for any Intel platform.