Powering VLSI Based Systems

By Kevin Parmenter

Exar Corporation

System Evolution

Not long ago the typical system required voltages consisting of 5 or 3.3 volts supplied by a switching regulator or an external brick power supply which directly supplied the required voltage to the system with few or even no other point of load power processing required. The geometries of the process technologies used to fabricate the core silicon of our systems were much greater than they are today and thus much more forgiving of accuracy, sequencing and slew rate applied to the power rails. Today modern devices such as microprocessors, microcontrollers, FPGA's, network processors, memory, ASICs and more are requiring more accurate and dynamically controlled power systems of greater complexity. Additionally as systems on a chip or SOC's) have become a reality where memory, Graphics processors, multiple cores and a great deal of I/O are on in one IC the demands and requirements for powering these highly integrated solutions have drastically increased.

Overcoming Constraints with Intelligent System Power

A recent study by MIT and the University of California at San Diego (hotchips conference 2011) concluded that with each successive generation of new VLSI devices the percentage of a chip that can actively switch drops exponentially due to power constraints. The issue is that Power is now more expensive than area in IC design. The movement to "green" electronics and designing in power saving technologies must be approached now at a system level instead of simply pressuring the power supply group to create more efficient power systems – the system itself must be intelligently designed to self modulate the power system. This technique has been used in portable system designs such as cell phones for years. This concept is not widely used outside of the portable space as in these areas battery capacities have limited power and every microwatt must be sparingly used. Thus the system must be capable of sensing the demand requirements and modulating the power system to turn off and on subsystems, ramp voltages to a data retention state and so forth to use power wisely. The intention is to save energy by changing modes when the system is not fully utilized, and optimize frequency and voltage to match the needs of the workload or to stay within the constraints of the operating environment. This is a system level performance aware strategy of energy management which can be used to increase performance or reduce power consumption without compromising the performance level. These techniques can also be used to extract all the processing performance from silicon without damaging it (MIPS per watt optimization).

The latest VLSI devices are being designed with internal multiple points (across the die) temperature sensors, voltage sensors used to sense conditions on the silicon of the VLSI device. In turn the VLSI

devices have built into them an interface (I2C for example) to communicate off chip to the power system. This allows monitoring then communicating system wide voltages to take into account effects of leakage, clock speed, temperature, die voltage (multiple points) and system utilization requirements as previously mentioned. This information is used to precisely control the power rails under various operating requirement.

Optimizing Accuracy of Power Delivery

The convergence of two industry trends are coming together to necessitate this fundamental change in how we power VLSI based systems. First are the mentioned requirements for all systems regardless of where the power comes from – including the main AC wall supply - to be more energy efficient and or to meet global regulatory requirements. The second is physics not legislative. Vcc levels have been dropping over the last few years – however it's not mentioned why very often or if it is, only part of the story is told.

As the geometries for VLSI process technology have dropped from 90nm to 65nm to 45nm heading to 28nm and below the process leakage increases, the need for precise accurate power delivery increases and the ability to slap any old regulator on it and forget about it goes away at about 65nm. This approach simply won't work properly as the process dips south of 65 nm. The reason for this is that the process technology is sensitive to latch up conditions within the device itself. Moreover when the device is connected off chip to other devices such as memory, or other VLSI devices of like geometries, care must be taken to not create sneak paths or latch up conditions of applied bias at the system level. The sequencing of the power up including the applied slew rates must be within guard banded conditions or the device will draw excessive current or even be damaged. Slew rate control and sequencing is often required on the shut down as well. As was previously mentioned the devices are becoming more system aware measuring clock speed, utilization of the system, die temperatures, die voltages and the device must communicate to the power supply to dynamically margin the applied voltages or change sequencing and slew rates as needed for a variety of system conditions This information can also be used to shut down unutilized blocks of the system until needed – then ramping the power supply back up.

Studying a recent datasheet for a multicore microprocessor SOC device it was noted that the on chip GPU has a specification to dynamically vary with clock speed, the GPU core voltage proportionally to the SRAM Vdd in .005 volt increments within a range of 0.8 to 1.150 volts on the core and 1.150 to 1.325 volts on the SRAM depending on the processor conditions these voltages must be varied dynamically in relation to varying system conditions. Add to this the need proper sequencing of the voltages, slew rate control and time delays. This particular processor has 8 voltage rails itself which need to be independently controlled for slew rate, sequencing and dynamic voltage control. As these voltages have gone down the need for accuracy has actually gone up. Today we see voltages as low as 0.7 volts specified with accuracies of 1% over temperature listed as a requirement. It is timely to mention the differences between accuracy, precision and resolution as these terms are often confused. Accuracy is

the exactness of the measurement compared to a known standard of reference such as NIST standards – also known as the variation or uncertainty of the measured value or closeness of the result from a known and accepted standard. Precision is the fineness to which an instrument can be read repeatedly and reliably. Or simply stated precision is equal to repeatability and or reproducibility. Resolution is the least discernable or smallest change detectable in the quantity – example the least significant digit or least significant bit causing an output change.

VLSI Specific Semiconductor Devices: Power Requirements, Achieving Precise Sequencing & Control

Let's take a look at specific product categories and the power needs coming along for each one. FPGA's have been pushing the need for multi-rail complex power for some time. This has created a plethora of "solutions" from every power semiconductor company wanting to promote devices into this market. The problem with most if not all of them is that they assume the system consists of nothing but the FPGA. The designer needs a power solution which is reusable – reconfigurable and is able to handle the complexity created when the FPGA is connected to the rest of the system which creates the need for slew rate control, multiple rails, sequencing and delays to prevent latch up and sneak paths in the FPGA or the rest of the system –which vary depending on what is being connected to the FPGA. Some of the trends in FPGA technology driving them to more of a SOC or system on a chip include, The move to parallel and multicore processing for power efficiency The shift of FPGAs to the leading edge of new semiconductor process technology, Increasing use of FPGAs in embedded systems, The economic realities of Moore's law, The consolidation of CPU architectures and a drive for computing power efficiencies. Most of these trends dictate the need for precision power management under system aware intelligent control. The current state of the art process technologies in FPGA's are around 28 nm – due to the complexity and integration of the devices which can have memory subsections, processors on board and so forth. Power management is critical in these devices and the need to shut down or change power supply levels depending on what the device is doing can save much energy – or even be the difference between the device working properly in the system or not and of course saving as much energy and reducing power dissipation in the VLSI devices. Partial reconfiguration is a technique where the FPGA blocks are reconfigured dynamically when needed or not needed or a different function is needed – under these circumstances it is often required to change the power supply when all this is taking place. The cores and other subsections require low voltages at high accuracies – 1% over temperature is not uncommon. Again, these devices along with the rest of the system require precise sequencing and control in addition to great transient performance for the system to function properly when the FPGA is connected to the rest of the system – I/O, memory, other VLSI devices and such, the power gets multi-rail and complex in a hurry.

Another category of VLSI semiconductor devices are network processors used in communications and networking equipment. A recent network processor has the following statistics: 45nm -ultra-low k BEOL process, die size is 410 sq mm, 1.43B Transistors, 3.2 million latches, 5 PLLs, 10 concurrent frequencies. The datasheet indicates that the system level Adaptive Power Supply Strategy REQUIRES 8 voltages domains with sequencing, slew rate control on start up and shut down and– continuous

communications with the power system to modulate system rails under various system conditions. The datasheet does not mention what is needed when it's connected to the rest of the system – memory and what have you however you can expect it to be complex and required to avoid issues. Memory often leads the pack with density issues and DDR memory is no exception. – DDRx – a quick scan of manufacturer's datasheets indicate 1.0-1.2 volts becoming more common – high precision and accuracy is required. Less than.075V accuracy today on the supply rails which will be will be below a volt with .05V accuracy needed over temperature in upcoming DDR4. A recent memory scheme in a use case of connecting a microprocessor to a memory system requires Frequency power scaling – or more simply stated, changing voltage vs. freq of operation. The memory suppliers datasheet indicates that VDD and VDDQ must be within 300mV of each other at all times, and VREF must not be greater than 0.6 × VDDQ. Again Memory does not do much by itself so when the memory is connected to the rest of the system – the power system complexity increases in order to optimize and protect the system. When a system is created from multiple VLSI devices this will create a system map similar to figure 1 which is a system level power requirement that a customer recently approached us in a panic recently and needed some help to solve. This example is only a few rails and is not very complex – we have been seeing 10-12-20 rails needed by systems recently.

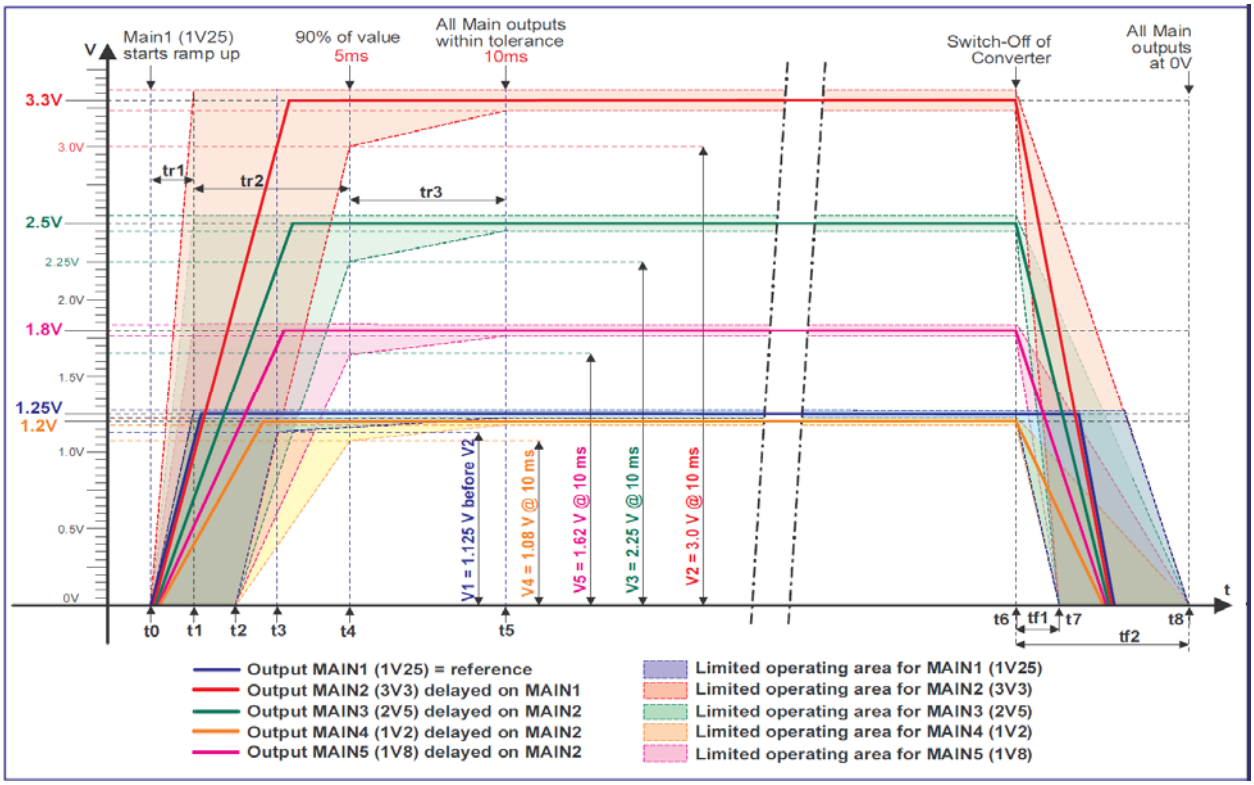

Figure 1: System Power requirements graph voltage vs time Multiple VLSI Devices

Design Options: Build It Yourself, System Monitors, Programmable Solutions

The question becomes what are the options available to a designer when something like this is presented to the power supply team from the digital-software engineers (who typically debug and develop their systems with a bank of lab power supplies run by lab view on a PC. This of course has to be duplicated in something affordable that can be incorporated into the product – usually without much design time. A build it yourself design will usually required a handful of DACS, ADC's, digital pots, microcontrollers, precision resistors, capacitors, PWM controllers and lots of design time. A designer could also choose to use one of the complex "system monitor" devices which incorporate some of this functionality yet still require external PWM controllers to be connected to it in order to begin to provide our system level power requirements. These parts tend to be rather expensive and you may pay for features, functions and complexity you may not use or need. Another approach could be to use a software programmable power solution one of which is available from Exar called the PowerXR family of products This incorporates digital power control which gives a great deal of flexibility to the designer and a graphical user interface to allow the design engineer to easily set up the system. Figure 2-4 shows the ease of use. No software programming knowledge is required.

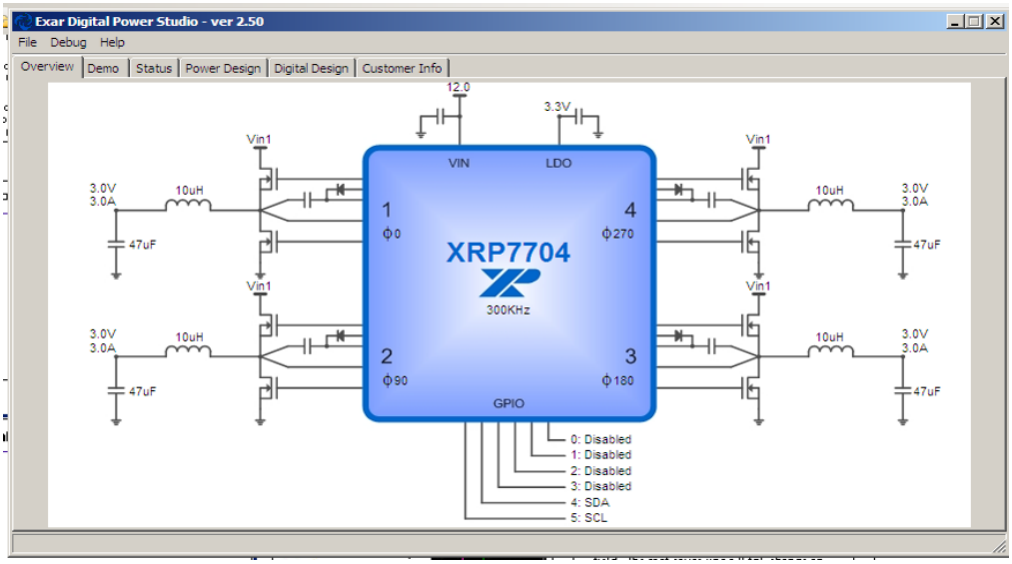

Figure 2: one circuit can be repeated and software defined

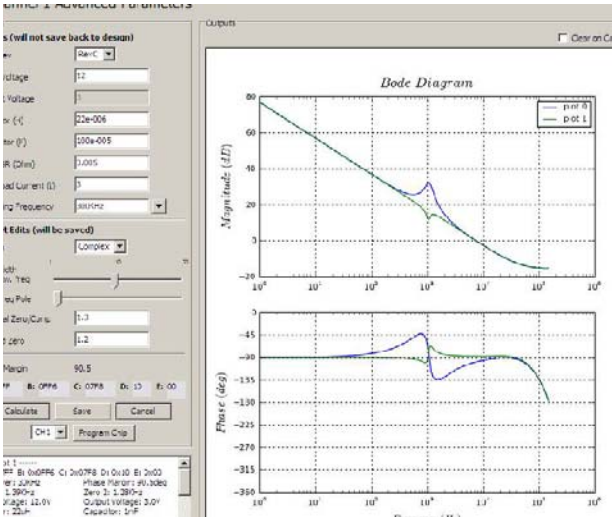

Figure 3: traditional analog information such as bode plots are shown

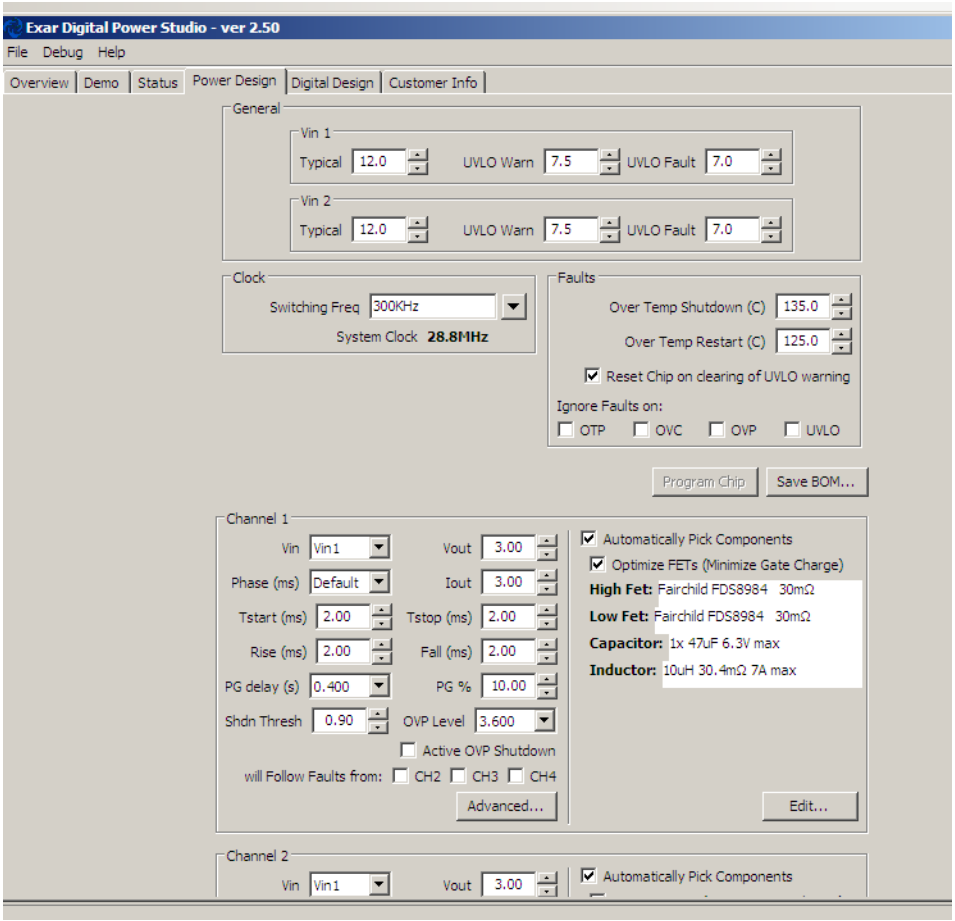

Figure 4: The GUI provides intuitive interface to configure the power converters

This allows systems power designers to simply answer questions and fill in boxes with information to create a successful initial design – the I2C interface on the IC allows continuous communications with the system so dynamic changes with operational system needs can be adjusted on the fly or reconfigured as needed The GUI can be used to develop and direct register control via the I2C interface is used for direct system control from the microprocessor- controller in the system itself. The GUI is available online at no charge. The resulting design is easier to implement, proven, with less design time and a great deal fewer parts. Since no compensation network is needed with the usual resistors and capacitors – initial accuracy and drift over time and temperature high accuracy is reduced greatly. Additionally complex slew rate control delays and sequencing is simple to implement. Importantly as system needs change, the power system can be reconfigured without any hardware changes even remotely so reconfigurability over the entire product life cycle is possible and in fact easy.

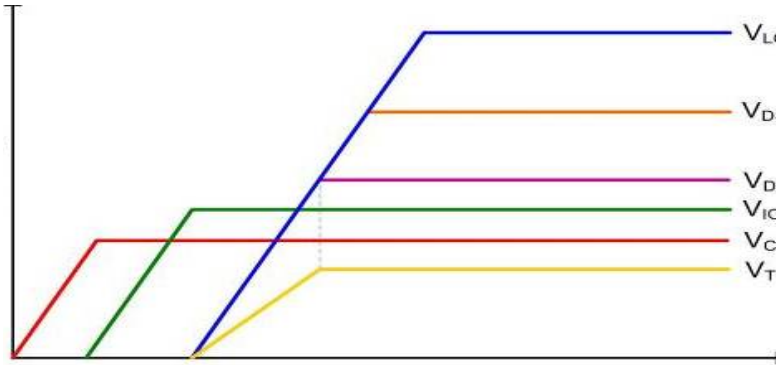

Figure 5: sequencing and ramp slew rate control is simple to configure

One of the benefits of the system level approach to power implemented in this way is that the hardware is completely re-usable. The hardware design stays constant and reusability is easier than ever before – differentiation is achieved by defining the power system in simple software and saving it on the IC. Additionally if 3-4 or 20 or more channels are needed it's a simple matter to daisy chain more channels and extend the system via I2C to control more channels by adding essentially more of the same and repeating the design until the desired numbers of channels are achieved in hardware. Current levels per channel are defined in software and component selection. Differentiation is done via software so the implementation can be reused across multiple products and product lines and economies of scale are possible where reusing the hardware is accomplished and different products are differentiated only in software. This simplifies not only design but the supply chain, logistics, field service all areas of the operation can benefit from this approach. The implementation allows telemetry so once a system operating data is available for debug or even remotely in the field the system can be reconfigured and monitored for system health or field service diagnostics.

Power System Solution Approach

The needs of VLSI devices are undeniably becoming more complex – a power system solution approach can make implementation easier and at the same time create customer benefits previously not possible

at the same time. Not the least of which includes saving energy by allowing the system to modulate and optimize its own power supply to extract the best performance per watt and at the same time making it possible for the VLSI system to work properly while adding features and functions that end users will benefit from. Digital designers have had reprogrammable and software defined systems for years now – it's time that power system designs can take advantage of the same benefits cost effectively. It's easy to get started designing system level power rails with low cost evaluation boards and the free software available. It will be no time at all until you are surprising yourself with what is possible under software control to save design time, money and implement reuse in a simple and cost effective manner to save time and money – yet providing the necessary system power environment for successful VLSI based products to work reliably over time and environmental extremes. For more information to learn more and obtain an evaluation board go to www.digikey.com.