White paper energy efficiency AC DC. Lee Harrison November 2009

The Drive Towards Energy Efficient Computing

The effects of global warming are recognized by us all as rising temperatures, melting ice, extreme weather events and greenhouse gasses from the use of carbon based energy. Most of us have our own opinions on the subject, but the fact remains we consume electricity at a faster rate now, than at any other time in the history of our planet. It is now believed that computer technology in the form of PC's and Servers are responsible for more than 2% of all C02 emissions, 2% is the same amount as the worlds total air traffic.

The reasons for increasing the efficiency of computers and servers are based not only on global warming, the increasing costs of oil and energy, along with rising technology demands are also key indicators to the energy efficiency drive. Real estate costs are now outweighed by electricity costs as utility bills account as the primary expense for data center managers, with power consumption per data center ranging from 2MW to 22MW. With the recent fluctuations and instability in the cost of energy, the data centers are in a volatile world as they strive to budget for energy costs, and balance power requirements versus power availability

In 2007 estimates were carried out on the total energy consumption worldwide for the Internet. Data from Stanford University show these figures represent 9.4% of total US electricity consumption, and 5.3% of global electricity consumption. Networking equipment such as modems, routers, hubs, switches etc account for about 25% of the electricity demand in an average office. Therefore if 200 kWh is needed for computers and servers then 50 kWh is needed for the networking components of that infrastructure.

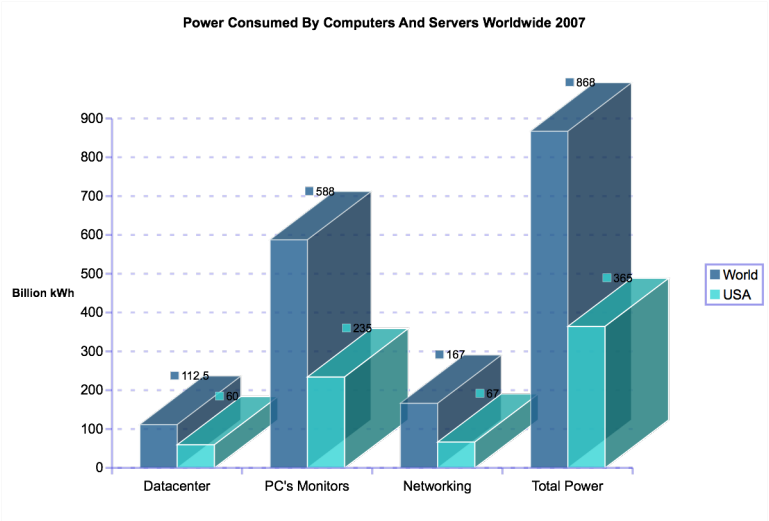

The current path of energy consumption in the data center is unsustainable as computers and servers become more powerful. Peak power on the grid in the US from servers and data centers is estimated above 7GW, equating a total output of 15 base load power plants. Current estimates, based on the estimated rise in power would see 7GW rise to over 12GW before 2011, requiring a further 10 new power plants to be built. The situation deteriorates when 45% of the power consumed is for air conditioning and cooling. Within the modern data center, performance per watt has become more critical than performance per processor.

Infrastructure changes and Renewable Energy

Many are still evaluating the effectiveness of renewable energy; with estimates of 120,000 square feet of solar panels to produce 1MW of electricity cost becomes a factor. Combined heat and power (CHP, where waste heat is converted to electricity), micro turbines, fuel cells and wind turbines have all fallen under the watchful eyes of the industries determined to cut energy usage and produce clean efficient power.

This technology is further advanced than some may think.

- In 2007 Google used wind turbines to provide a portion of the power requirement in the Netherlands.
- Intel made use of a 10kW photovoltaic system in New Mexico,
- Fujitsu, Chevron Texaco and the Hamilton Sundstrand data center in Connecticut have implemented fuel cell technology.

Real estate that can support large quantities of renewable energy are being acquired by large corporate companies intent on improving profit margins by controlling the energy they consume, they all have the same aim, to find long term cheaper alternatives to grid power.

Renewable energy will play a major part in data center power delivery, and this is likely to become a DC based world. Consider the power requirements in a modern data center and DC power is the lifeblood. Servers, LED lighting and cooling systems all require DC, the traditional AC to DC conversion is an efficiency drain and within 2 years we will see a consolidated drive towards DC power. There are skeptics who do not believe this will happen that fast, with some companies waiting for the ball to start rolling at the infrastructure level before they consider engagement in supporting the DC technology.

Nonetheless, industry is gaining momentum,

- ETSI (The European Telecommunications Institute), a non-for profit organization seeking to produce telecommunication standards throughout Europe and beyond. ETSI had a scheduled meeting in Italy in October 2009, discussing 400VDC and ETSI EN 300 132-3 V1.2.1
- Korean Telecom DC Internet Computing Center at http://www.kt-icc.com
- IBM are planning one of the worlds most energy efficient data centers in New York, making the use of "green" technology to reduce energy consumption by more than 50%. This data center will be a $12.4 million, 6000 square foot data center, featuring it's own electrical tri-generation system. This project specifically includes a DC power distribution system generated on site, thereby eliminating transmission and conversion losses. The full article can be seen at http://www-304.ibm.com/jct03004c/press/us/en/pressrelease/27612.wss
- NTT have presented at many conferences, including the Green Grid TGG Forum in St. Louis in September 2009.
- Many Telecom companies are expected to start issuing 400VDC RFI's within the next 6 months.
- Nextek Power Systems at http://www.nextekpower.com is a founding member of the Emerge Alliance, launched to create new standards for power distribution in commercial buildings. They are pioneers in DC networking solutions, and as such understand this market, and the industry drive better than most.

Why is power demand increasing?

The increase in power density is caused entirely by higher computing demands, leading to higher power consumption. The increasing cost of energy from the cost of carbon based fuels and utility bills increases the total cost of ownership directly to the datacenters and end users.

All users of computers and servers demand higher performance year on year. The technology industry, along with server and PC manufacturers respond with larger or faster disk drives, faster memory, multi core processors and more I/O devices. Whilst this satisfies customer demand, the new designs require additional cooling; consuming more power and the trend towards higher power consumption is sustained. Average power consumption per server has increased from 150W - 250W in year 2000, to 450W - 800W in 2009. Power increases per rack are expected to grow by at least eightfold from 2000 to 2010, while many datacenters are already concerned that power demand is outstripping supply. Average kW per rack of servers in 2000 was close to 1kW, rising to 6-8 kW in 2006, and expected to top 20kW in 2010.

Reasons for increased efficiency

It is not just the obvious reasons of climate change that the energy drive took hold. Any electronic equipment that runs cooler will have improved MTBF, increasing reliability. Data centers not only save money as the servers require less power, but the infrastructure cooling and air conditioning costs are also impacted. Large server users can benefit from relatively small reductions in electricity cost. Google, for example, may be able to save close to $1M per annum in utility bills by reducing power consumption between 2% and 3%. Some believe the financial cost benefits are seen as more important than the climate issues, in a time of reducing costs across all businesses this may be one of the reasons the energy efficiency drive is becoming successful.

The Efficiency Drive

The efficiency drive has many directions in the computing space. AC DC, DC DC, VRM and DDR power in addition to power management through software. This article focuses on the AC DC power supply as this item was identified as an area where big improvements could be made initially.

The current targets as defined by the Climate Savers Initiative and the EPA for single output power supplies are shown in Table 1, multiple output supplies are shown in Table 2.

Table 1

Loading	CSCI Silver And Energy Star v1			CSCI Gold And Energy Star v1			CSCI Platinum		
	Efficiency	PFC (≤1000W)	PFC (>1000W)	Efficiency	PFC (≤1000W)	PFC (>1000W)	Efficiency	PF (≤1000W)	PF (>1000W)
10%	75%	0.65	0.80	80%	0.65	0.80	82%	0.65	0.80
20%	85%	0.80	0.0	88%	0.80	0.90	90%	0.80	0.90
50%	89%	0.90	0.90	92%	0.90	0.90	94%	0.90	0.90
100%	85%	0.95	0.95	88%	0.95	0.95	91%	0.95	0.95

Table 2

Loading	CSCI Bronze Energy Star for Servers v1 Energy Star for computers v5		CSCI Silver		CSCI Gold	
	Efficiency	PFC	Efficiency	PFC	Efficiency	PFC
20%	82%	0.80	85%	0.80	87%	0.80
50%	85%	0.90	88%	0.90	90%	0.90
100%	82%	0.95	85%	0.95	87%	0.95

What the efficiency drive means to power supply manufacturers.

The AC DC power supplies used in general computers were relatively inefficient before the efficiency initiatives began. Server power supply efficiency was a little higher, but there remained much room for improvement. Initially the gains were relatively easy for power supply manufacturers. Increasing a single output power supply from 70% to mid 80's efficiency was not too complicated. General topology tweaks, attentive selection of silicon and the hardest part of the design was taken care off.

The difficulties began when efficiency targets introduced 10% load points. The Silver category defined in table 1 was challenging for some power vendors, but Gold and the next generation Platinum levels are in a different league. Some power supply vendors have achieved the Gold standard, and some servers are currently available with Gold power supplies, but it is likely to be the end of 2010 before the Gold standard becomes more established.

Increasing efficiency to Gold and Platinum has led to many different engineering directions. The sections that follow describe some of the issues that are being, or have been dealt with in the drive to produce more efficient power, and particularly the direction that needs to be taken to meet Platinum power.

Interleaved Power Factor Control

The entire power industry understands interleaving techniques in multiphase buck converter designs. This design is commonly used to meet fast load transient demands for DSP's and processors. Some designs interleave multiple synchronous power stages to increase power delivery to the load and decrease input and output capacitance, maximizing the benefits of smaller output inductors in each phase of the design due to the ripple current cancellation effects at the output capacitors, transient demands in excess of 350A/us can be realized with this topology.

Interleaving PFC boost stages benefits from the same principles. PFC boost can be defined as one of three modes of operation. CCM (continuous conduction mode), DCM (discontinuous conduction mode), or CRM (critical conduction mode). High efficiency designs are likely to use CCM mode of operation in the server space, as this mode of operation suits power supplies >350W output power. Although CCM makes the design of the EMI filter easier, the boost inductor is generally larger than alternative design techniques. Interleaved CCM is not without its own risks; the design inherently suffers from higher switching losses and reverse recovery losses in the rectifier, although these can be partially overcome with Silicon Carbide (SiC) diodes. The frequency is not generally adjustable either in this type of design, CCM operates as an average current mode, PWM control, fixed frequency design, forcing average input current to be proportional to the rectified AC line cycle.

Efficiency targets primarily drive interleaved PFC, but as enclosures and footprints for power become smaller, the better EMI performance helps reduce the size of EMI filters. Power density is excellent for this topology, but the design is difficult. High power designs and interleaving PFC has to be approached carefully due to the many different modes of operation, safe operation has to be guaranteed at start up and shutdown, auto ranging wide input sources, and fault conditions. Digital techniques are assisting in this field.

Interleaved or bridgeless PFC does assist with high efficiency of this stage of a power supply. There are still issues for the unwary and experienced alike. In order to achieve high efficiency at loads less than 15%, particularly around the 5% - 10% range, it is likely that some form of phase shedding, pulse skipping or burst modes would be implemented. These can be fully autonomous or selectable via software and DSP dependant on load. Either way round it is not the actual phase shedding, skipping or burst control that causes the problems; it's the fact that there can be other functions occurring at the same time. Some designs

implement pulse skipping or burst modes, whilst these attain the desires of increased efficiency at light loads, they wreak havoc when trying to accurately measure PFC in these conditions, cause EMI issues, and can cause strange harmonic effects in racks and across servers connected to the same AC source supply.

Members of the Climate Savers Initiative, Ecos consulting and EPRI have studied measuring PFC at light loads. The solution, although not fully understood, is to measure PFC using a LISN and capacitor between the AC source and power supply. Further investigation will be ongoing to fully understand the cause and effects of measuring PFC under these conditions. Questions are also being raised about strange harmonic effects in high efficiency power conversion and more work needs to be done in this area to fully understand the issues. So high efficiency comes at a price, we may be able to improve efficiency to unsurpassed performance but we are introducing new and previously unseen side effects, we will need to work with universities and true experts in AC power to fully understand these issues.

Further problems caused by PFC in burst or skipping modes are the digital sensory Vin, Pin and Iin outputs. It is likely that only fully digital controlled power supplies will be able to cope with these instances, with the addition of more complex algorithms and look up tables. It is virtually impossible to obtain valid data from a control signal that is constantly changing. Many designs aim to turn off power reporting at these lighter load conditions, but with Energy Star determined to tighten up on the already difficult to meet specifications, solutions must be found to report data at light loads.

Digital techniques in PFC have recently produced stunning performance capabilities at light loads, but at the expense of THD. This is likely to become a discussion point in the near future.

Bridgeless PFC

Bridgeless PFC offers loss savings in the 0.4W to 1.5W region at high load conditions, depending on output power. The benefits in a low line mode of operation are much higher, and suit countries that still run servers in low line conditions. The standard bridge rectifier suffers due to the series connected semiconductors, but remains a low cost and reliable solution to AC rectification. In some designs today, standard bridge rectifies can run dangerously close to their thermal ratings. Airflow is typically not ideal in the location of a 1U dense server power supply, and requests have been made to power semiconductor manufacturers for increased thermal capabilities from 150C to 175C.

Bridgeless designs have not found there way into the mainstream due to complexity of design and cost. Current efficiency improvements have been realized with the use of interleaved designs and careful component selection. Main roadblocks against the use of bridgeless designs remain the same as they have done for years, these being primarily cost, complexity in gate drive circuits and due to the circuit configuration the difference between input and output ground references has caused problems. Prototypes of these designs have shown high EMI common mode switching elements at the negative bulk capacitor connection, and with virtually no low frequency path back to the AC source the EMI problem is complicated further.

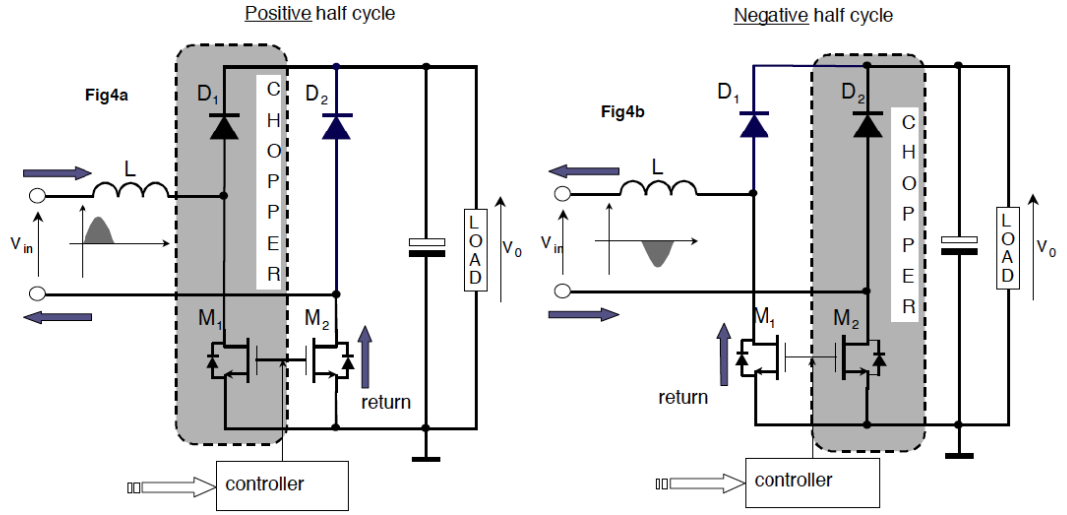

Various semiconductor companies now offer controllers for bridgeless rectification, although some will only operate in a high line condition. Most are fixed frequency around 100kHz to 150kHz, but costs still surround the implementation. Estimates of $20 - $60 are common to implement bridgeless front ends, and a $20 increase per power supply is too high a price to pay for approximately 1W power saving in a server at high line. The fact remains, the bridge rectifier accounts for one of the biggest losses in the current high efficiency designs. There will come a time, when a semiconductor manufacturer will release a mass produced FET controlled bridge rectifier at a reasonable cost, within the same standard 4 pin connection we have become accustomed to with the standard rectification package.

For now bridgeless designs have been sidelined for interleaved CCM designs, careful use and choice of silicon and light load power saving techniques. Unless there is a major break through in the cost, complexity and confidence of bridgeless technology, it may be destined to remain a distant idea for some to come. With the advent of DC based infrastructures, the AC DC Bridgeless technology may become redundant before it ever got off the ground, and will go down in history as one of those good ideas that never really went anywhere.

Resonant Topologies

Wide use of resonant converters may provide the edge to realizing Platinum designs due to advantages in zero voltage switching losses and other benefits this technology offers. Switching frequencies generally run higher for these topologies, reducing footprint sizes in semiconductors and magnetic components, while parasitics, which are problematic for alternative topologies become advantageous in resonant topologies as Table 3 shows. PWM designs lose power when switching frequency is too high, resonant mode topology begins where PWM topologies end. There are many options available in waveform shaping to eliminate switching losses, Current waveform (ZCS), Voltage waveform (ZVS) and Quasi Resonant (uses only a part of the sinusoid waveform).

LLC resonant converters look set to become the choice of the power designers. These offer good regulation and frequency control, with efficiencies for the LLC stage recording up to 97%, further assisting the Platinum compliance. Some power design houses are reporting concerns with LLC in N+1 redundant designs. Investigations into these continue. Other features of LLC are

- Primary transformer inductance is reduced in value such that it has an impact on the resonant network.
- Excellent frequency control over wide load range.
- Excellent line and load regulation down to zero loads.
- LLC overcomes disadvantages of series resonant converter.
- LLC offers ZCS control of secondary rectifiers.

Table 3

Converter	Advantages	Disadvantages
Series Resonant	Behaves as a current source and best suited to high voltage high power designs. Overload and fault conditions are handled well. Requires small or no output filter.	Very poor regulation at light and no load conditions. Continuous current mode (below resonance) causes high component stress and premature failure if not carefully designed.
Parallel Resonant	Behaves as a voltage source and suitable for low voltage outputs. No load and light load regulation is good. Discontinuous mode of operation is very similar to standard PWM buck conversion.	Requires snubbers consuming efficiency. Will not self protect in overload or short circuit conditions requiring additional protection. Constant resonant circulating current at all load conditions can cause problems.
Zero Current Switching, Quasi Resonant Fixed On Time.	Turn off losses within silicon is very low. Any remaining leakage current is re-cycled.	Power components contain high peak currents.
Zero Voltage Switching, Quasi Resonant Fixed Off Time.	Switching losses are virtually eliminated. Any remaining leakage current re-cycled.	Operates above resonant range. Power devices require at least 3X supply voltage. Multi resonant techniques have been proven to offset this.
Series Parallel Resonant.	Excellent performance at light load, no load and short circuit or fault conditions.	Parasitic elements are produces that can be troublesome requiring additional passive components.
Multi Resonant, Zero Voltage Switching Class E.	Virtually eliminates switching losses, virtually no frequency shift compared to other methods.	Can cause stability issues and generates high peak voltage and current across semiconductors.
Clamped PWM.	Simple control loop with low voltage stress. Current mode capable and fixed frequency of operation.	Turn on losses can be high, operates best at medium to full load conditions.
PWM, Zero Voltage Switching.	Fixed frequency design with switching losses maintained at low levels across all power devices. High efficiency at high switching frequency.	Creates problematic switching noise at light loads and does not achieve zero voltage switching in the light load mode of operation.
LLC Resonant	High Efficiency with excellent line/load regulation Out performs Series Resonant Mode.	Performance is superior to other modes of operation, but design is critical. Possible N+1 redundancy concerns.

Silicon improvements.

Improvements in silicon packaging over time have seen power devices shrink from large T03P or T0220 devices all the way down to ball grid array and other surface mount components. Whilst this change has helped with the miniaturization of power supplies, the increase in efficiency is key. High efficiency equates to lower power dissipation, smaller or eliminated heat sinks and therefore smaller components. In the 1970's industry typically used the bipolar FET, with HEXFETS online in the 1980's these offered better performance. Then came Trench FETS, super junctions and other high performing silicon.

The biggest improvements are in silicon material advances. These advances promise higher efficiency and lower running costs, so they sound ideal. For a while, Gallium Arsenide (GaAs) components promised power savings and improved designs with faster production than similar silicon devices. GaAs has an energy gap 4 orders of magnitude greater than Silicon and can be produced semi insulating, leading to reduced parasitics. It is a direct band gap material, producing the ability to emit light, an

interesting characteristic, but no benefit to the power supply. They do however make excellent solar panels and electricity light converters.

Overtime however, GaAs has not out competed Silicon for a few very important reasons. Arsenic is very toxic, handling and disposal of these devices is causing problems in current manufacturing plants. GaAs is more rare than Silicon, expensive to locate, and where Silicon uses an oxide layer as a mask to protect the Silicon, GaAs is unable to do this, further increasing the cost of GaAs devices. The thermal conductivity of silicon is 2.75 times greater than GaAs and packaging densities are therefore less. Arsenide related semiconductors would remain expensive, as the conductor is gold as opposed to aluminum. There is a place in technology for GaAs, and that is in space electronics. In orbit there will be no cross contamination between arsenide and silicon, and where gold is toxic to silicon, arsenide and silicon are dopants for each other.
For the foreseeable future, Silicon will remain the dominating material.

The next most likely choice from the industry is likely to be the Gallium Nitrite (GaN) material. This offers hope for higher efficiency and lower running costs that GaAs originally identified. GaN has been proven to be 10 times better than silicon and will likely be the next step in long path of FET improvements. The path is set to be a lengthy one, even with GaN coming on line within the next 12 months, and promising reductions in RDSon tenfold for 50V devices before 2013, it is unrealistic to expect >10% of applications to be widely implementing it before 2015.

Magnetic components.

Magnetic components have had their part to play in this subject. Anything that can be done to minimize core and copper losses has to be done; another valuable 0.5W to 2W can be saved here. Planar transformer technology and integrated magnetics with power components as part of the transformer assembly are ways in which this is being realized. Planar magnetics are still favorable due to lower switching and copper losses, with their low profile, high power density and high frequencies of operation they can be built into the main PCB, or assembled as sub assemblies for daughter cards.

Integrated magnetics can reduce footprint size by up to 50% when compared to separate magnetic and rectification stages. High frequency power supply noise spikes and EMI are lower and more controlled due to lower commutation loops, and with reduced copper losses, efficiency increases are guaranteed. Generally, integrated magnetic designs offer repeatability and cost reductions in manufacturing, tolerances of magnetic characteristics and output stages show closer results to each other when compared to separate discrete stages. This technology will find its way into future high efficiency designs; the option of increased power density alone is a big enough reason to adopt the technology. Efficiency gains may even become secondary to this requirement.

The magnetic industry has not seen the same improvements in technology as semiconductors. Passive components account for >75% of power supply real estate; these have to be the target for size reduction and increases in power density. Magnetic components, including EMI filters account for the biggest real estate consumer, followed closely by bulk capacitor electrolytics.

Digital Control.

It is important to note that a power supply implementing an I2C digital bus does not qualify under the guise of a digital power supply. I2C busses in power supplies are purely digital management, and not digital control. Digital power control is defined as direct control of switching devices via a digital processor. Voltage and currents are digitized and outputs controlled on past and present digital input commands. Many supplies today use both digital and analog designs; pure digital power supplies are on their way and will assist with overall system control. By 2012 it is expected that digital power will become the dominating power design, with a widespread use of primary and secondary DSP's, further improving efficiency and reducing cost.

Digital control within the power supply has helped with the efficiency drive in a number of ways. There are two sides to the digital conversations, the first being digital control and optimization, the second being sensor accuracy for voltage and current reporting from the power supply.

The recent Energy Star requirement at a system level requires ±10% accuracy with a cutoff at ± 10 watts (i.e. accuracy never needs to be better than ± 10 watts). End users however are asking today for accuracies at a power supply level better than 2% accuracy at 20% load, with 5% accuracy at 10% load. The voltage reporting accuracy is relatively easy, the current accuracy reporting is difficult. The sensing component has to be extremely low in value, otherwise it is detrimental to the efficiency, consequentially, this means that at low loads the current waveform used for determining the current reporting is too distorted to be meaningful. Input current is typically sensed using the PFC current sense. Signal to noise ratios can be poor and it is not unusual to need signal amplifiers to improve the signal. To achieve meaningful readings at lighter load can require calibration controlled by the same DSP. The current Energy Star requirement is hard to meet, the requirements expected in the Tier 2 specification in October 2010 make this requirement one of the hardest to meet as it is set to require ±5% accuracy with a cutoff at ± 5 watts (i.e. accuracy never needs to be better than ± 5 watts).

Input voltage is typically sensed through a rectifier separate from the main power rail. A primary side DSP using correction factors will sense peak or average voltages and report the status. Accuracy varies across the industry from 1% to 5%.

Input power can be reported as a function of Iin and Vin. Again calibration is usually required to prevent errors, and ageing of components has to be built into this calibration. Without digital filtering and carefully controlled algorithms, these requirements would be almost impossible to meet. Even with today's technology it is a difficult task.

Digital control not only assists with the latest energy sensor reporting requirements, but it assists in the control of the power conversion and optimization of efficiency. It allows direct control over phase control, burst or skipping modes, switching frequency control, transient response, voltage and current limits, dead times and sequencing of internal waveforms and voltages. Optimization of switching frequencies is a complex task, this almost becomes a four way-juggling act between power supply performance vs. efficiency vs. magnetic component footprint sizes vs. actual power supply enclosure limitations.

Conclusions

Improving power supply efficiency from 90% to 94% is not a task to be taken lightly. It is difficult, and the path is littered with for and against arguments for just about every stage within the power supply. Losses have to be carefully understood; making small improvements in multiple parts of the power supply is the only way to gradually creep up the efficiency performance. There will not be one single area where the final 2% to 4% can be gained. Power supply manufacturers have to build in between 0.3% and 0.7% margin for unit variability and CPK requirements for the industry standards. Worse case, they may need to add up to 0.7% efficiency to the requirement to guarantee compliance, taking the 94% Platinum requirement to nearly 95%. Improving PSU efficiency from Gold to Platinum in a 1000W supply means reducing losses by around 30%. System builders must weigh up the performance advantages versus the cost increases for their final products. Even if the OEM's consider this a worthwhile task, the data center managers will have to be able to justify the cost increases of the servers vs. the energy cost difference between Gold and Platinum designs.

The Platinum specification is relatively new and the impact is still being evaluated. Power supply manufacturers are quoting 15% to 20% cost increases for this level of performance, with some companies posing the question, "are we coming close to the point whereby further increases in efficiency are detrimental to the cost of the product, or has that that point already been reached with the Platinum requirement?" I don't think that point has been reached and will explain why in a moment. Some companies believed Gold was a step too far, complaining about the constant need to stay ahead of an ever-increasing efficiency trend, at the same time others embraced the ideas, and in seeing the bigger picture released products to market with a technology advantage and marketing edge.

Whilst all the Power Vendors are currently working to make Platinum Power a reality, Power One recently announced their Platinum Plus Power Supply, capable of efficiency levels that meet and exceed Platinum. The cost of that power supply is in line with the expected 15-20% cost increases over Gold, so I decided to work out the cost recovery times for this technology. The results are interesting, assuming a rack of 40 X 1U servers using this supply, with loads varying as they would in a real world scenario;

- The additional cost increase of Platinum over Gold is recoverable within 4 months running time, just outside most vendors 90 Day billing period.
- The Platinum supply actually pays for itself in 23 months in utility cost savings.

Those same companies that were pessimistic about the increasing efficiency requirements relative to power would still argue that the Platinum power supply is an unnecessary upgrade, missing the point altogether that the end user with 40 Platinum servers in a rack would have saved nearly $14,000 in utility bills over a 5 year life of those servers. These companies fail to see the bigger picture. Take a company like Google for example who are energy conscious, they don't reveal exactly how many servers they run, but I would estimate they have at least 36 data centers with 150 racks in each, this equates to 5400 racks and 216,000 servers. Even if half of those servers were Platinum, Google would save $37,800,000 in electricity costs over 5 years. Take away the cost of Platinum at $16,200,000 and Google would still have saved $21,600,000. In reality they probably have many more servers than this, so now tell me it's not worth it!!

The future of efficiency in the server and computer space still has some way to go. The current drive identified the AC DC power supply as a main offender, but we have to find a way to measure or define motherboard efficiency in a way that is simple, effective and repeatable. Only then can the discussions begin around overall system efficiency, to that aim being able to calculate the cost of a given task. With the platinum efficiency target set, further increases in AC DC efficiency are likely to be final or non-existent. Therefore further system efficiency improvements may well be restricted to system level hardware, DC converter efficiency, software changes and DC infrastructure implementation.

The future holds many possibilities, from changes in data center infrastructure to high power AC or DC supplies, entire server system architectural changes relative to power, all the way to changes at the rack level. The EPA tier 2 efficiency requirement states EPA intends to explore a *Net Power Loss* approach for Computer Server power supplies, this would be expected to become a requirement in October 2010. This approach would aim to specify a maximum allowed power loss through the power supply at actual operating conditions of the Computer Server (e.g., Idle and full load power). If a Net Power Loss approach is not developed, EPA plans to re-evaluate both Multi-Output and Single-Output power supply efficiency and power factor levels. Keep your fingers crossed they don't choose the Platinum levels!!

Net Power Loss (NPL) is being considered to prevent over sizing of power supplies to meet Energy Star, but the industry is concerned that this new requirement may penalize systems that reduce server idle power to very low levels. For example, assume a system has a maximum operating load of 850W; idle power (OS running but idle waiting for user input) may be around the 250W level. The system may have hardware and software techniques to reduce this idle power consumption further, to less than 80W. However, the current EPA guidelines may mean that a system with 250W idle can achieve Energy Star compliance, but the system with lower idle cannot. There is some confusion on the value add of NPL, it seems to add unnecessary complications to the current tier 1 specification as all systems will differ in their requirements, with system power consumption increasing or decreasing based on processors, DIMMS, hard drives and PCI cards.

Many question the value add of Climate Savers and the EPA direction relative to power. Climate Savers were formed to reduce power consumption, and the data speaks for itself. OEM's, end users and power design manufacturers have efficiency at the top of their list, with many end users more interested in power consumption than server feature sets. Climate Savers was also ensuring that nothing was asked of the industry that was not possible, and this is still true today. The EPA aside from sensor system reporting and adding a 10% load point has followed Climate Savers directives. Although the Platinum range is difficult, it is a target to aim for, there are no penalties for not meeting it, but, by being able to, and to ship products using Platinum supplies this becomes not only a cost saver to the end user, but produces a massively improved power supply in reliability terms than the same could be said if the initiative had not taken place at all. Based on this it is clear that CSCI and EPA directives have been successful.

Power supply manufacturers and design engineers deserve congratulating on the improvements they have made in efficiency performance over the last 3 years. For some who are not expert in power design, have seen this as an almost pointless task, and one that has taken far too long. They failed to realize that whilst design concepts can be discussed and jotted down on the back of notepads, that is a million miles away from producing a stable, reliable and safe power supply ready for mass production. Those individuals accept the complexities of motherboard design and architecture, but fail to realize the modern power supply is a complex component with safety issues surrounding its use.

Think for a moment about a Platinum power supply with Interleaved CCM, Bridgeless Rectification, LLC resonant output stages, Integrated or Planar magnetics and full digital DSP control. Add optimization of synchronous rectification stages, timing, bulk voltage control, phase shedding, burst or pulse skipping techniques. Now add accurate DSP controlled Vin, Iin, Pin, Iout, Vout and Pout with DSP controlled cooling, and it becomes apparent we have changed the look of power forever. Gone are the days of simple topologies to achieve power conversion with low 70% - 80% efficiencies. We now have complex designs, with specific differences between manufacturers, each finding their own best way to achieve the same end goal, and we are beginning to re-evaluate some of the most basic electrical laws as we uncover previously unseen effects as we work through these challenges.

The arrival of Renewable Energy with all that entails, and the real possibility of DC power driven data centers, will bring new problems for engineers to solve. For a period of time, old technology will have to integrate with new technology, this will be demanding and will require individuals and companies not adverse to change to make the transition successful. Much good work has already been done, and the world of power delivery, conversion, consumption and efficiency will be a very different one to that we know today by 2015.

Lee Harrison Peritus Power
http://www.perituspower.com
leeharrison@perituspower.com