Power Density v. Power Efficiency

By Paul Lee for Mouser Electronics

Power conversion efficiency is a headline metric, with module manufacturers vying with each other to show decimal-point improvements in 95%+ figures under carefully selected conditions. Ever-more complex conversion topologies are being utilized to achieve these figures such as 'Phase Shifted Full Bridges' and 'LLC Converters'. Diodes are being replaced by MOSFETs for lower losses wherever possible and 'Wide Band Gap' devices are hailed as the semiconductor of choice for the future with their spectacular switching speeds.

End-users however, look at the bigger picture and care more about the efficiency of their whole system or process in its ability to make profits while complying with environmental obligations. They understand that concentrating on incrementally reducing losses in one small element of the power conversion process does not necessarily lead to significant overall cost savings or environmental benefit, when all lifetime costs are factored-in. On the other hand, packing more power conversion equipment into a smaller volume – increasing its 'power density' – can utilize factory or data center floor space more efficiently and produce more output with existing overhead costs.

This article examines the real costs of chasing percentage points of power conversion efficiency in energy saved, acquisition/disposal costs and cabinet/floor space utilization compared with increasing power density and the system efficiency improvements that can follow.

Industry 0% or 1000% efficient?

In the power electronics world, efficiency is a term which is easy to conceptualize – 100% equals good, 0% equals bad, right? But you have to carefully set your reference; a data center is close to 0% electrically efficient overall – just about all power it draws from the grid is converted into heat in server blades, their power supplies and the electronics in cooling systems. It might then be 1000% efficient converting dollar value of electricity into dollar revenues and the same is true of most industries. You wouldn't expect otherwise, so if you want to save costs and the planet while making money, the real issue is how you minimize the total power draw while maximizing productivity.
Data center managers know this and face daily pressure to increase data processing capacity and speed while keeping the electricity bill as low as possible and getting payback from capital investment. They have little choice but to add servers in increments of many kilowatts of dissipation but can calculate the monetary value added to capacity and offset that against the extra energy and capital costs. In industry, if

another 100kW motor is needed, it is to produce more saleable output and the motor drive and its power supply is the unavoidable overhead. In all industry, power supplies are a necessary evil that add no commercial value in themselves, so every operating expense and watt dissipated in them is seen as reducing the bottom line. The spotlight therefore naturally turns on the power electronics manufacturers with pressure to reduce losses by increasing electrical efficiency.

Loss means more than efficiency

Power conversion efficiency seems easy to define - we can all quote the formula ‘power out divided by power in, as a percentage’ with the difference between the two dissipated as heat in the power converter. The problem is that efficiency is meaningless without quoting power levels and how they vary with operating and environmental conditions, leaving ‘efficiency’ as only a comparative measure between converters. It is then open to ‘creative’ specifications, picking out the sweet spots that show the equipment in the best light. Few converters are operated near their maximum power ratings so efficiency is normally designed to peak at around 50-75% of maximum rated load with some curve which must fall off to zero efficiency at zero load. At light load there can be huge variability between converter designs so under idling conditions, one power supply might dissipate several times that of another (**Figure 1**). At 5% load, the converter represented by the orange line is dissipating more than three times the one for the blue line. Light load losses therefore make a significant difference to total energy draw.

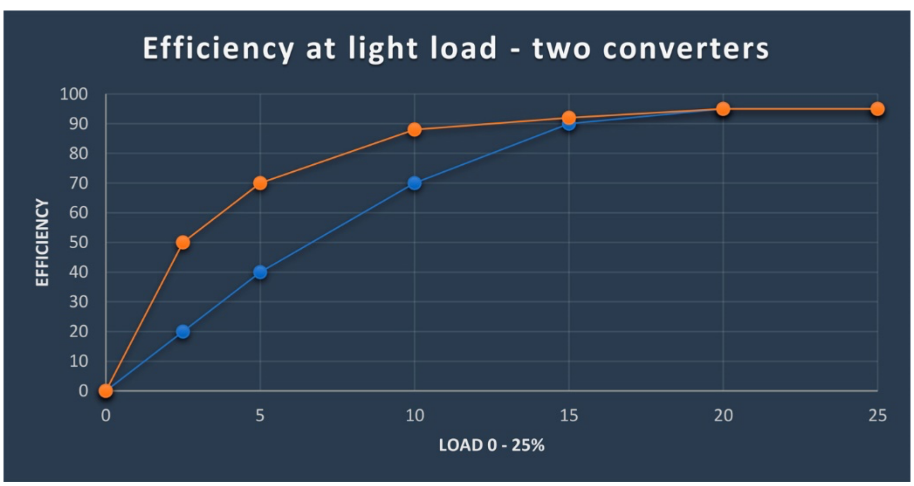

Figure 1: Efficiency at light load can vary widely between otherwise similar power converters

Fortunately, there are standards that set the shape of the efficiency curve, such as the ‘80 PLUS initiative’ with its various levels. ‘Titanium’ is the highest, demanding minimum 94% efficiency at 50% load and 90% at 10% load. These are for 115V systems, the figures for 230V are 96% and 90% respectively. (**Figure 2**).

80 PLUS Certification	115V Internal Non-redundant				115V Industrial			
% of Rated Load	10%	20%	50%	100%	10%	20%	50%	100%
80 PLUS	---	80%	80%	80%/PFC 0.9	---			
80 PLUS Bronze	----	82%	85%/PFC 0.9	82%	---			
80 PLUS Silver	---	85%	88%/PFC 0.9	85%	80%	85%/PFC 0.9	88%	85%
80 PLUS Gold	---	87%	90%/PFC 0.9	87%	82%	87%/PFC 0.9	90%	87%
80 PLUS Platinum	---	90%	92%/PFC 0.95	89%	85%	90%/PFC0.95	92%	90%
80 PLUS Titanium	90%	92%/PFC 0.95	94%	90%	---			

Figure 2: 80-PLUS initiative targets - 115V systems

These limits are quite tough to achieve – when the 80 PLUS certification scheme was conceived in 2004, the lowest level of 80% efficiency at 50% load was hard enough, but achieving the Titanium level at 94% means reducing losses in the power supply by three quarters. A mere 14% increase in efficiency but a kilowatt-rated supply has to reduce losses from 250W to 64W. This is not achieved by fine-tuning existing designs and has necessitated a radical re-think of converter topologies. Diodes are dropped in favor of synchronously driven MOSFETs, Phase Shifted Full Bridge (PSFB) and LLC resonant topologies are used to limit dissipation during switching transitions and new semiconductor technologies have arrived such as Silicon Carbide and Gallium Nitride for faster switching without a dissipation penalty. Even the humble bridge rectifier off the mains has morphed into a hybrid arrangement of MOSFETs that also form part of the necessary power factor correction circuitry. All this does not come cheap and without the 'risk of the new'. Still, customers and power supply manufacturers are in a spiral of demand and supply for yet higher efficiency figures, pushing towards 99% and beyond.

The cost of 1%

As power conversion efficiencies approach 100%, difficulty increases exponentially. From 97 to 98% means decreasing losses by a third. 98 to 99% means decreasing losses by a *further* half. Cutting losses by 50% in any converter design might force a complete re-start from scratch with the only route to use yet more complex techniques and more expensive components, often at the expense of size. A one-kilowatt supply is only dissipating 20.4W at 98% efficiency. How much is the huge effort worth to hit 99% and 10.1W loss? Think about the load taking 1kW – you would get 10.1W saving by reducing it by 1%. How much design effort would that take?

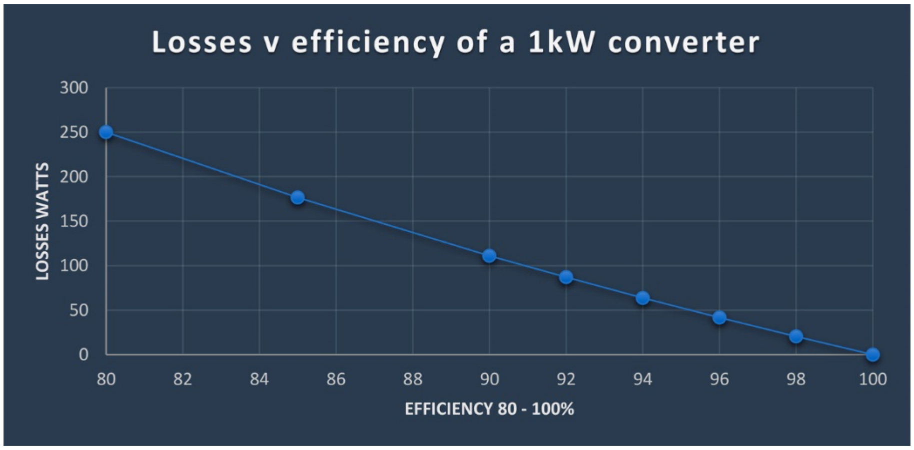

Figure 3: Losses versus efficiency in a 1kW power converter

Of course, all energy savings are worth having, but you need to look at the bigger picture. The average price paid for industrial electricity in the US is about 7 cents per kilowatt-hour [1]. If the 1kW power supply lifetime is say, five years or about 44,000 hours at 100% uptime, a reduction of 10.1W saves about $31 while the load power is costing over $3,100. Changing out the power supply has an acquisition cost, purchasing and qualification overhead, installation cost and a carbon footprint associated with typically hundreds of components, packaging and transport. Then there are disposal costs for the old equipment and the functionality risks with new cutting-edge products. It's difficult to see how this offsets the $31 saving compared with keeping previous generations of power supplies in place, assuming reliability is still adequate. Pursuit of high efficiency for its own sake can be an expensive business.

Leveraging size off efficiency gives better power density

Perhaps it's worth improving power converter efficiency to reduce internal temperatures and improve calculated life/reliability, but this only works if the case and cooling remain the same. There is an old rule-of-thumb that lifetime of electronics reduces by a factor of two for each 10-degree Celsius rise and according to reliability handbooks, semiconductor failure rate increases by about 25% and capacitors by about 50% for a 10°C rise. However, modern electronics is extremely reliable and durable, so these are percentage changes from a very long lifetime and high reliability figures anyway. Power electronics cooling has been historically set to maintain ideally around 21°C inlet temperature in data centers for example, but research by Intel and others has shown that this can be increased without significant effects on system reliability. A report by APC [4] quoting the American Society of Heating and Air-Conditioning Engineers (ASHRAE) predicts just 1.5 x increase in overall equipment failure rate for an inlet air temperature rise of 20 to 32°C (68 to 90°F), **Figure 4**. Each degree Fahrenheit increase in temperature in data centers is said to reduce associated cooling costs by about 4% so reducing case size, allowing equipment including power supplies to run hotter can make real savings, while freeing up rack space.

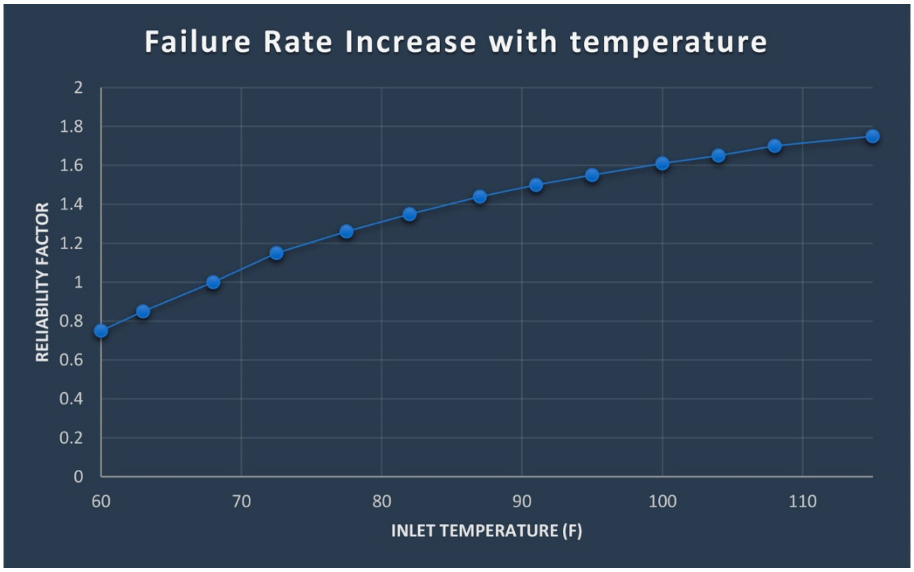

Figure 4: Equipment reliability with air inlet temperature – source ASHRAE

Another enabler for smaller power supplies running hotter is the use of Wide-Band-Gap (WBG) semiconductors fabricated in Silicon Carbide (SiC) or Gallium Nitride (GaN) materials. These devices have much higher operating temperature ratings than silicon types, particularly SiC with an allowed die temperature up to several hundred °C.

Power density is increasingly important

Competing suppliers of power conversion equipment into industry may vie with each other for claimed efficiencies under very specific conditions, but what matters to the end-user is productivity and profitability of their process. For sure, saving a few dollars by consuming less energy is a good thing but the dollars earned by increasing the density of equipment in a cabinet or rack and improving productivity per cubit foot may be more attractive. Floor space in data centers and manufacturing industry has a “dollar density”, a monetary value which it has to achieve to contribute to revenue, measured in thousands of dollars/square foot, so down-sizing the electronics to give more productive space is a real gain. If it means putting off provision of a complete extra cabinet when expansion is needed, even more dollars are saved short- and long-term.

Figure 5: Factory floor space has a $ value

Achieving higher density of electronics with associated power converters is driving system architects to think of 'power density' as an increasingly important metric. However, unlike end-to-end electrical efficiency, power density of a complete system is not easy to compare, and what do you include? In a typical industrial cabinet there may be switchgear, connectors, chassis-mounted EMI filters, an AC-DC converter generating an intermediate voltage, high current bus bars, DC-DC converters locally at the loads, fans and their own power supplies and mounting hardware. You might even include air conditioning units. In a controls cabinet, the loads might be external, perhaps motors. In this case, the volume of power conversion equipment is a significant proportion of the overall space and any savings in size allows more control electronics to be included. Returns diminish though, as more power would anyway be needed for the extra equipment added. Controls cabinets might also be limited by the requirement to use standardized hardware such as DIN-rails for equipment mounting, with suppliers launching ever-narrower products and the practicality of input/output connector size often defining the minimum. 30W AC-DCs are now down to about 21mm width while 480W parts can be around 48mm wide x 124mm high. Cooling, if any, in cabinets may just consist of fans with inlet temperatures ill-defined so power converters tend to be rated for just operation in high temperature airflow with no chassis heatsinking. This results in a relatively low value for power conversion density, perhaps 10-20 W per cubic inch.

Data center power converters are heated by their load

In data centers, the architecture of power provision affects power density strongly; latest trends are towards a 48V backplane bus with Point-of-Load (POL) converters on each server blade reducing the voltage down to IC levels, often sub-1V. Taken in isolation, the POLs can have dramatic power density – over 1kW per cubic inch, but need significant heatsinking or airflow to survive. The 48V bus can be derived from a rack AC-DC converter which might have a power density of only around 20W per cubic inch. Alternatively, 380VDC may be provided from an external central source with a conversion to 48V in the rack. With a DC supply and no losses from AC rectification and power factor correction circuitry, this converter can be very efficient and again have high power density of over 1kW per cubic inch (with adequate cooling). An additional advantage is that energy storage for supply loss or brown-outs can be centralized, unlike with AC-DCs in each rack, which have the overhead of large internal reservoir capacitors for ride-through, taking valuable space.

Unlike in industrial manufacturing cabinets, data center loads are the server blades themselves so each rack can be dissipating 10kW+ internally. This mandates active cooling with tightly controlled, high speed airflow at low inlet temperatures which is good news for the power converters which, with their high efficiency, are anyway only dissipating a fraction of the power of the blades. This allows the use of POLs and bus converters with minimal, if any, external heatsinking, keeping the overall power density high. In reality, a major consideration is to keep heat generated by the blades away from the power converters.

New WBG technology offers even higher power density

Power converter designers always have the option to increase efficiency by slowing switching speeds but this results in larger passive components and a consequently larger case size. Complex resonant converter topologies have allowed higher frequency operation with low losses but the arrival of SiC and GaN semiconductors have changed the game again with their combination of speed and low losses. Their ability to operate at higher temperatures reliably allows converter package sizes to reduce still further, pushing power density figures higher.

Chasing value

Chasing power conversion efficiency percentage points is a game of diminishing returns unless the improvement results in smaller products, leaving space for the equipment that directly adds to the bottom line. Power density is a good metric for the converters but should be compared carefully to include all elements in the system and can be expected to vary hugely between manufacturing industry cabinets and data center server racks.

References

[1] https://www.rockymountainpower.net/about/rar/ipc.html

[2] https://www.nuclear-power.net/nuclear-engineering/thermodynamics/laws-of-thermodynamics/thermal-efficiency/thermal-efficiency-of-nuclear-power-plants/