ISSUE: February 2021

A PSU Analytical Power Loss Model For Optimizing The Server Power Delivery Architecture

by Viktor Vogman, Power Conversion Consulting, Olympia, Wash.

In an effort to promote energy efficiency, the voluntary certification program 80Plus was established to certify computer and server system power supply units (PSUs) that have more than 80% energy efficiency at certain specified percentages of rated loads.[1] Because they reduce data center electricity costs considerably versus less efficient power supplies, 80Plus-certified PSUs have become the market (and industry) standards. Moreover, the 80Plus certifications are now being widely used as the benchmarks for power supply efficiency.

Nevertheless, even with the availability of these more-efficient power supplies, there are still opportunities for cost and energy savings. Specifically, the optimization of the sizes and ratings of 80Plus PSUs for the application could further reduce the total cost of ownership (TCO)[2] for server platforms and make them the most attractive option among market competition. Such optimization could be provided very effectively if a PSU power-performance analytical model were available for power architects.

Such a model would enable the most efficient and cost-effective overall system power delivery solutions. The analytical PSU model would also help system architects to guide PSU designers in providing PSU characteristics meeting targeted efficiency- and cost-optimized power delivery architecture needs.

Power loss is one of the major factors determining the overall market performance of highly efficient PSUs. Besides its 80Plus certification level, power loss also affects consumer energy bills, component heatsink sizes, the amount of required air cooling, and, ultimately, PSU size and cost. For PSUs with efficiency exceeding 90%, power loss is much more sensitive to efficiency improvements, than the efficiency rating itself. For example, a 1% efficiency increase over the 96% baseline (80Plus Titanium level) results in over 25% drop in P_{loss}. The higher the baseline efficiency, the more its improvement reduces the power loss.

Many of the 80Plus-certified power supplies also have so-called peak power requirements, often significantly exceeding PSU continuous ratings. This makes it relevant that the sought-after power loss model would be suitable for projecting the energy efficiency of power delivery at peak power operation as well. This model can help server system power architects make informed decisions on key tradeoffs so as to adopt the 80Plus requirements and properly specify continuous and peak power ratings.

This article presents an analytical PSU power loss model that provides a means to assess tradeoffs in continuous vs. peak power ratings of PSUs. This model also can be used for characterizing PSU dynamic efficiency and as a tool for optimization of the system power delivery spec.

Creating A PSU Power Loss Model

PSU power losses can be divided into two categories: load-dependent and load-independent (fixed) losses. The first category includes copper and active component conduction losses, and switching losses, while the second category consists of magnetizing power loss, and control and housekeeping power consumption. Copper and conduction losses are associated with heat produced by electrical currents in the conductors, power connectors, transformer and inductor windings, as well as in active devices in their on-states, such as switching and providing dc redundancy isolating ("ORing") MOSFETs.

Copper and conduction losses can be characterized by an equivalent resistance *r*, as referred to the secondary side (PSU output), which represents a passive component conducting the PSU load current. Power dissipated in it is proportional to the output current squared, or—assuming that output voltage is normally held constant by the PSU control—to the output power squared.

Switching power losses occur during active device transitions between on- or off-states. Such transitions occur not instantly but over some time intervals, over which the device current and voltage across it overlap, producing power dissipation. Normally, the current magnitude of active devices is proportional to the output current, so switching power loss increases linearly with output power. The relative significance of this power loss component depends on the fraction of power converted in hard-switching operating mode. (For PFCs and dc-dc converters, this usually occurs when they lose ZVS (zero-voltage switching) at light loads.)

Based on these considerations we can build an equivalent electrical network representing each power loss component. This network is shown in Fig. 1.

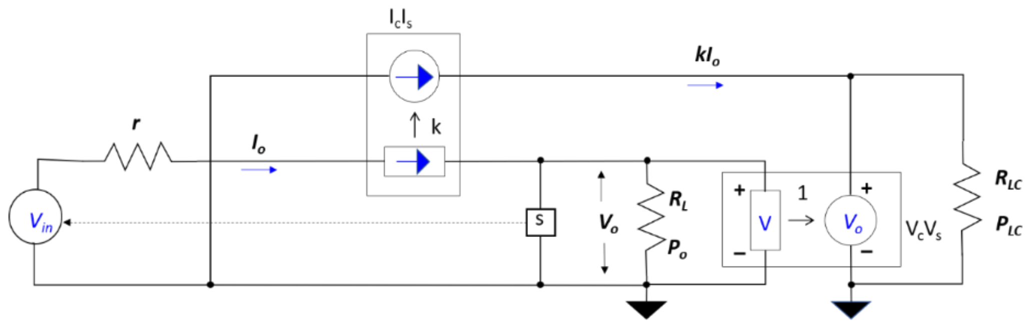

Fig. 1. Electrical interpretation of PSU power loss components. Load-independent (fixed) losses are dissipated in resistor R_{LC}, conduction/copper losses are dissipated in resistor r, and switching losses are absorbed by ideal voltage source V_o. Resistor R_L dissipates the useful output power P_o of power supply load.

In this network a voltage source V_{in} acts for the PSU with equivalent internal resistance *r*, which includes contributions from the combined resistance of power components and conductors referred to the PSU output. Resistor R_L dissipates P_o, the useful output power of the power supply load, which varies with load current.

The voltage across R_L is sensed by sensor *S*, which controls the PSU output so that the voltage across R_L remains constant, matching typical application requirements. The network also includes two controlled sources $I_C I_S$ and $V_C V_S$, which have gains of *k* and *1,* respectively.

The controlled current source $I_C I_S$ passes on the switching power loss, which is proportional to P_o, to the ideal voltage source V_o, which absorbs this power. Its gain *k* characterizes the weight of this power loss component. The resistor R_{LC} attached to the constant voltage output (V_o) dissipates the load-independent (fixed) power loss component P_{LC}.

Defining the power dissipated in resistor r as $P_r = (V_o/R_L)^2 \cdot r$, and multiplying the numerator and denominator in this expression by the same factor $V_o^2 P_{o.max}^2$ we can write the following equation for the total power loss in this network as a function of output power P_o:

$$P_{loss} = \frac{r}{R_{L.min}} \cdot \frac{P_o^2}{P_{o.max}} + k \cdot P_o + P_{LC} \qquad (1)$$

where $R_{L.min}$ is the minimum load resistance corresponding to the maximum continuous output power $P_{o.max}$ case: $R_{L.min} = V_o^2/P_{o.max}$. Normalizing (1) by $P_{o.max}$ we get:

$$\hat{P}_{loss} = \frac{r}{R_{L.min}} \cdot \hat{P}_o{}^2 + k \cdot \hat{P}_o + \hat{P}_{LC}.$$

The $\hat{P}_{loss.i}$ values at the output power levels $\hat{P}_{oi}$ required for certification can be determined from the 80Plus efficiency requirements for the standard 20%, 50% and 100% load levels using the equation:

$$\hat{P}_{loss.i} = \hat{P}_{oi}\left(\frac{1}{Eff_i} - 1\right) \qquad (2)$$

where Eff_i are PSU efficiencies at the certification power levels $\hat{P}_{oi}$.

Equating the right-hand sides of equations (1) and (2) we can determine, for given $\hat{P}_{oi}$ power levels, the unknown parameters $r/R_{L.min}$, k and P_{LC} in equation (1) as functions of the corresponding 80Plus efficiency requirements. This determination will require solving a system of three linear equations representing three main 80Plus certification power levels, $\hat{P}_{oi}$ and efficiency Eff_i levels (i = 1,2,3):

$$\hat{P}_{oi}\left(\frac{1}{Eff_i}-1\right)=\frac{r}{R_{L.min}}\cdot\hat{P}_{oi}^2+k\cdot\hat{P}_{oi}+\hat{P}_{LC}$$

$\hat{P}_{oi}$ and Eff_i data, along with computed power losses, for the most commonly used 80Plus Gold, Platinum and Titanium efficiency level certifications are provided below in Table 1.

Table 1. Efficiency values and normalized power losses for 80Plus power supply guidelines.

Output Power $\hat{P}_{oi}$	80Plus Gold		80Plus Platinum		80Plus Titanium	
	Eff_i	$\hat{P}_{loss.i}$	Eff_i	$\hat{P}_{loss.i}$	Eff_i	$\hat{P}_{loss.i}$
0.1	-	-	-	-	0.9	0.0111
0.2	0.88	0.0273	0.9	0.0222	0.94	0.0128
0.5	0.92	0.0434	0.94	0.0319	0.96	0.0208
1.0	0.88	0.1363	0.91	0.0989	0.91	0.0989

The solution of the system of equations for 80Plus efficiency PSUs yields negative k, which basically indicates that without either major reduction in fixed losses or significant increase of efficiency margin at heavy loads, i.e. without PSU overdesign, it would be really challenging to meet both light load and heavy load certification efficiency requirements.

Since power supply overdesign would highly impact its cost, in practice, this obstacle is overcome by implementing a load-adaptive PSU architecture, i.e. by adjusting PSU operating mode at some power level $\hat{P}_{oA}$. Such adjustment may include disabling of the PFC or one of its interleaved stages, lowering its output voltage setpoint, ac cycle skipping at light loads, and other techniques.[3-5]

Light And Heavy Load Domains

In the load-adaptive PSU architecture, the two load domains ($\hat{P}_o \le \hat{P}_{oA}$, $\hat{P}_o > \hat{P}_{oA}$) can be characterized by two separate analytical models: a light load—$\hat{P}_{loss1}$, representing the operating range where switching losses dominate, and a heavy load—$\hat{P}_{loss2}$, representing the operating range where copper and conduction power losses dominate:

$$\hat{P}_{loss1}=k\cdot\hat{P}_o+\hat{P}_{LC}\quad(\hat{P}_o\le\hat{P}_{oA})\tag{3}$$

$$\hat{P}_{loss2}=\frac{r}{R_{L.min}}\cdot\hat{P}_o^{\,2}+\hat{P}_{LC}\quad(\hat{P}_o>\hat{P}_{oA})\tag{4}$$

where $\hat{P}_{oA}$ is the power level corresponding to the PSU transition into a power saving mode at light loads. Typically, the transition point $\hat{P}_{oA}$ is selected around $\hat{P}_o$ = 0.3. This power-saving-mode transition is illustrated by the $\hat{P}_{loss}(\hat{P}_o)$ graph shown in Fig. 2.

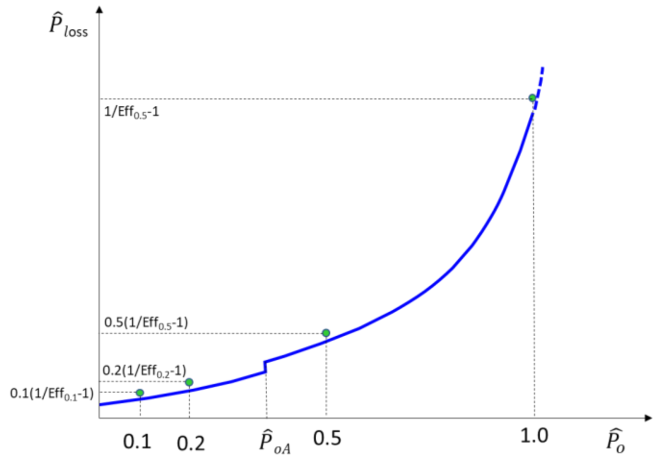

Fig. 2. At power level $\hat{P}_{oA} \approx 0.3$, a PSU with load-adaptive control transitions into a power-saving mode to reduce the switching power loss component and meet 80Plus efficiency requirements at light loads. At lighter loads ($\hat{P}_o \leq \hat{P}_{oA}$) power loss is represented by Eq. 3, while at heavier loads ($\hat{P}_o > \hat{P}_{oA}$) Eq. 4 represents the loss.

For dynamic peak power PSU operation, we will focus on the high-power range: ($\hat{P}_o > \hat{P}_{oA}$). In this range we have two 80Plus efficiency requirements, corresponding to 50% ($\hat{P}_o = 0.5$) and 100% ($\hat{P}_o = 1.0$) load. Based on equations (2) and (4), for these load conditions the following expressions are valid:

$$\begin{gathered} \frac{0.25r}{R_{L.min}} + \hat{P}_{LC} = 0.5 \cdot \left(\frac{1}{Eff_{0.5}} - 1\right) \\ \frac{r}{R_{L.min}} + \hat{P}_{LC} = \frac{1}{Eff_{1.0}} - 1 \end{gathered} \qquad (5)$$

Solving system (5) for $\hat{P}_{LC}$ we find:

$$\hat{P}_{LC} = \frac{2}{3 \cdot Eff_{0.5}} - \frac{1}{3 \cdot Eff_{1.0}} - \frac{1}{3} = \frac{1}{3}\left(\frac{2}{Eff_{0.5}} - \frac{1}{Eff_{1.0}} - 1\right)$$

Substituting into this equation $Eff_{0.5}$ and $Eff_{1.0}$ efficiency data from Table 1 we can determine the following $\hat{P}_{LC}$ levels supporting the efficiencies required for 80Plus certification:

Table 2. Load-independent power loss $\hat{P}_{LC}$ requirements for 80Plus certification.

80Plus cert	Gold	Platinum
$\hat{P}_{LC}$	0.012513	0.009583

Since in practice $\hat{P}_{LC}$ cannot be made significantly lower than 1%, then based on the selected model, 80Plus Gold and Platinum PSU heavy load efficiency requirements can be met without considerable overdesign, i.e. without excessive margin at the 100% load point. However, for the 80Plus Titanium efficiency certification system equation (5) yields a negative $\hat{P}_{LC}$ value (-0.00519). This result indicates that in Titanium-certified PSUs exactly matching both 50% and 100% load efficiency levels is not feasible.

Selecting a realizable $\hat{P}_{LC} = 0.01$ indicates that the 50% load point equation in system (5) yields an $r/R_{L.min}$ considerably lower than that required to match the 100% load Titanium level efficiency. This means that for 80Plus Titanium PSUs meeting the efficiency requirement at a 50% load results in a significantly better efficiency at 100% load, than required for certification, which agrees well with multiple efficiency data points from 80Plus Titanium PSU manufacturers.[1]

Let's find suitable *r/RL.min* and $\hat{P}_{LC}$ values for each of the 80Plus requirements by deriving the maximum possible efficiency at the 50% load point. The expression for PSU efficiency can be defined by a ratio of output power to the sum of output power and power loss, i.e. using equation (4) we can write:

$$Eff(\hat{P}_o) = \frac{\hat{P}_o}{\hat{P}_o + \hat{P}_{loss2}} = \frac{\hat{P}_o}{\hat{P}_o + \frac{r}{R_{L.min}} \cdot \hat{P}_o^{\,2} + \hat{P}_{LC}} \qquad (6)$$

The derivative of this function changes sign such that it is positive at light loads and negative at heavy loads, which indicates that function $Eff(\hat{P}_o)$ is not monotonic and has a maximum in the load range $\hat{P}_{oA} \le \hat{P}_o \le 1$, which is also supported by multiple sets of experimental data. Taking the first derivative with respect to the variable $\hat{P}_o$ and equating it to zero yields:

$$\hat{P}_{o.opt} = \sqrt{\frac{\hat{P}_{LC}}{r/R_{L.min}}}$$

This equation indicates that the abscissa $\hat{P}_{o.opt}$ of the efficiency function maximum can be moved towards higher or lower power levels by changing parameters *r/RL.min* and $\hat{P}_{LC}$. By selecting $\hat{P}_{o.opt}$ at 50%, corresponding to the maximum 80Plus efficiency point, we can optimize the PSU design for meeting the 80Plus requirements and establish a relationship between *r/RL.min* and $\hat{P}_{LC}$. The result yields the peak efficiency at this load level:

$$r/R_{L.min} = 4\hat{P}_{LC} \qquad (7)$$

This approach could also be used for finding the position of the peak-efficiency point in other switching power stages having the same power loss components.[7] It should be noted that the $\hat{P}_o$ position of the efficiency peak for lower (Silver/Bronze) 80Plus-certified PSUs, for which switching loss dominates over a wider load range, can be described by the same equation. That's because adding a switching loss component $k\hat{P}_o$ to the denominator of equation (6) does not change the abscissa $\hat{P}_{o.opt}$.

By solving equation (7) together with the first equation of system (5), representing the 50% load point power loss, we can determine the *r/RL.min* and $\hat{P}_{LC}$ parameters that we recommend for each of the 80Plus certification levels:

$$\hat{P}_{LC} = 0.25\left(\frac{1}{Eff_{0.5}} - 1\right)$$

$$\frac{r}{R_{L.min}} = \frac{1}{Eff_{0.5}} - 1$$

The values for these terms for the three common 80Plus certification levels are provided in Table 3.

Table 3. Recommended *r/RL.min* and $\hat{P}_{LC}$ values for 80Plus-certified PSUs.

80Plus cert	Gold	Platinum	Titanium
$r/R_{L.min}$	0.086	0.0638	0.0416
$\hat{P}_{LC}$	0.0215	0.01595	0.0104

The PSU efficiency and power loss data computed with these recommended *r/RL.min* and $\hat{P}_{LC}$ values are shown in Fig. 3.

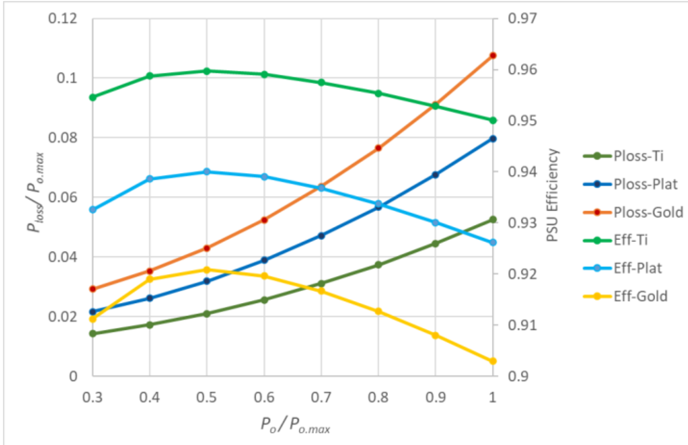

Fig. 3. Normalized $\widehat{P_{loss}}$ and efficiency vs. normalized output power $\hat{P}_o$ for typical 80Plus certified PSUs. The graphs were generated with the proposed model at heavy loads that exceed the transition level $\hat{P}_o \geq \hat{P}_{oA} = 0.3$. The model yields an efficiency higher than the efficiency level required for 80Plus certification at the 100% load point, which agrees well with actual PSU measurements.

As shown in Fig. 3, due to the dominating 50%-load requirement criteria, the model yields a significant efficiency margin at the 100% load point, which agrees within a few tenths of 1% with multiple actual PSU efficiency measurements provided on the 80Plus website.[1] This margin is greater for higher 80Plus certification levels.

The Impact Of Output Voltage Setpoint

The recommended $r/R_{L.min}$ and $\hat{P}_{LC}$ parameters for each 80Plus certification level can be used as reference points for actual PSU design. However, in some cases it becomes really challenging to provide an efficiency margin sufficient for mass production, especially at the 50% load point. Sometimes providing such a margin even requires making major PSU design modifications in the late stages of development.

Let's use the model to determine an alternative resource for increasing the PSU efficiency by a few decimal points. Equation (6) indicates that PSU efficiency depends not only on r but also on minimum load resistance $R_{L.min}$: the larger $R_{L.min}$ the higher the efficiency. Let's multiply the numerator and denominator of the $r/R_{L.min}$ ratio by V_o^2. In this case for 50% load point efficiency we will get:

$$Eff_{0.5} = \frac{1}{1 + \frac{r}{R_{L.min}}} = \frac{1}{1 + \frac{r}{V_o^2} \cdot P_{o.max}}$$

For constant power type loads, such as voltage regulators supplying power to CPUs, memory, etc., this expression indicates that increasing the PSU output dc voltage setpoint V_o can result in a reduction of $r/R_{L.min}$ and an efficiency boost. A similar efficiency rise will be provided at the 100% load point. In many cases such a "natural" $r/R_{L.min}$ reduction may represent the easiest and fastest way for satisfying a required efficiency margin.

For example, if in the 80Plus Titanium efficiency certification case, the actual $r/R_{L.min}$ ratio appeared to be higher than the 0.0416 recommended in Table 3, say 0.042, this would result in a 0.04% efficiency shortage at the 50% load point, i.e. 95.96% vs. 96% required by the 80Plus guideline. However, increasing V_o by 2.5%, which typically allows it to to keep the supplied voltage within its regulation range, boosts the 50% load point efficiency to 96.2%. This provides a suitable PSU efficiency margin for production without making any major design modifications.

Impact On Peak Power Operation

Normally, server power supplies have peak power ratings significantly exceeding their continuous ratings. This requirement is driven by various turbo modes accelerating processor and graphics performance for peak loads, automatically allowing processor cores to run faster and draw more power.[6] On the other hand, the power supply and voltage regulators' size and cost constraints push system architects toward making continuous power ratings lower. Selecting optimal continuous and peak power ratings plays a major role in minimizing the total cost of ownership (TCO) of power delivery.

In general, for an arbitrary power waveform (load profile), the PSU's average power must not exceed its maximum continuous rating. But even when this condition is met, the PSU power loss in dynamic mode must not exceed a so-called thermal design power $\hat{P}_{TDP}$, which is the maximum amount of heat generated in the PSU that it is designed to dissipate. In effect, PSU operation in the dynamic mode must satisfy the following two conditions:

$$\hat{P}_{o.avg} = \frac{1}{T_c}\int_0^{Tc} \hat{P}_o(t)dt \le 1$$

$$\hat{P}_{loss.dyn.avg} = \frac{1}{T_c}\int_0^{Tc} \left(\frac{r}{R_{L.min}} \cdot \widehat{P_o(t)}^2 + \hat{P}_{LC}\right) dt \le \hat{P}_{TDP}$$

where T_c is the peak power cycle.

In the dynamic/peak power operating mode, even if the standard requirement of keeping average power equal to maximum continuous rating ($P_{o.avg} = P_{o.max}$) is met, due to the quadratic nature of the power loss function, the actual PSU efficiency is always lower than the PSU efficiency at the 100% load point $Eff_{1.0}$, and average power loss $P_{loss.dyn.avg}$ is always greater than the loss $P_{loss.1.0}$ at the rated continuous power level:

$$P_{loss.dyn.avg} > P_{loss.1.0} = P_{o.max}\left(\frac{1}{Eff_{1.0}} - 1\right)$$

Thermal Design Power Considerations And Impact On PSU Efficiency

To illustrate these concepts, let's consider the most commonly specified load condition, which matches the situation we are describing. This load would be a square wave with peaks exceeding PSU continuous rating $P_{o.max}$. Let's also assume that an average power is maintained that is equal to the maximum continuous rating of the PSU and then quantify the impact that peak operating mode may have on the PSU thermal performance and thermal design power.

The square wave case with peaks exceeding the PSU continuous rating $P_{o.max}$ and average power $P_{o.avg} = P_{o.max}$ is illustrated in Fig. 4. The green-colored waveform represents the instantaneous power $P_o(t)$ drawn from the PSU. The red-colored waveform represents instantaneous PSU power loss $P_{loss.pk}(t)$ and coincides with the drawn instantaneous power waveform $P_o(t)$.

To evaluate the impact of peak power operation on average power loss we consider a ratio of average power dissipation in the peak power operating mode $P_{loss.dyn.avg}$ to the baseline-level power loss in the maximum continuous power operating mode $P_{loss}(P_{o.max})$.

At peak power operation, the fixed power loss is negligible compared to the power loss at maximum continuous and peak power points. Using equation (4) for power loss and the designations in Fig. 4 we can write the following expression for the power loss ratio:

$$\frac{P_{loss.dyn.avg}}{P_{loss}(P_{o.max})} \approx \frac{P_{o.pk}^2 \cdot D + P_{o.min}^2 \cdot (1-D)}{P_{o.max}^2} \qquad (8)$$

where D is the duty ratio of the peak power pulses.

If the peak power is applied for a thermally significant time duration, the D in equation (8) is effectively unity, which results in an increase in actual P_{loss} by a factor equal to the peak-to-average power ratio squared. However, in server applications that use various power management algorithms $P_{o.pk}$ levels are applied for thermally insignificant time intervals and are outbalanced by lower power-level conditions, which keeps average load power equal to $P_{o.max}$. The minimum power level $P_{o.min}$ supporting this condition can be derived by equalizing the shaded watt-second (power x time) areas in Fig. 4.

$$P_{o.min} = \frac{P_{o.max} - P_{o.pk} \cdot D}{(1-D)} \tag{9}$$

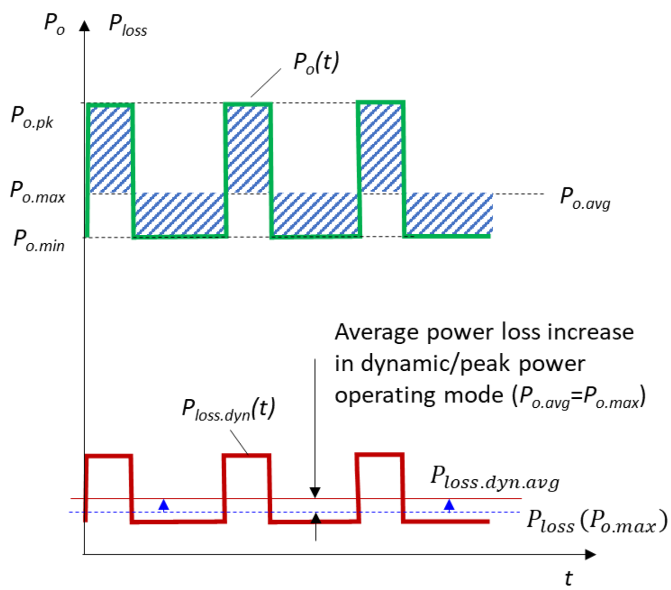

Fig. 4. PSU power loss increase in dynamic/peak power operating mode with average consumed power held equal to PSU continuous rating: $P_{o.avg} = P_{o.max}$. Due to the quadratic nature of the power loss function at heavy loads, and despite the fact that the same average power is drawn from the PSU ($P_{o.max}$), the average power loss in dynamic/peak power operating mode is always larger than in the continuous $P_o(t) = P_{o.max}$ = constant condition.

Substituting equation (9) into (8) we obtain the following expression for the power loss ratio:

$$\frac{P_{loss.dyn.avg}}{P_{loss}(P_{o.max})} = \frac{DA^2+1-2DA}{1-D} = 1 + \frac{D(A-1)^2}{1-D} \tag{10}$$

where A is the peak power factor, A = $P_{o.pk}/P_{o.max}$.

This equation shows that even though the same average power is drawn from the PSU, the PSU average power loss in peak power operating mode (A>1) is larger than in the continuous $P_o(t) = P_{o.max}$ = *constant* condition (Fig. 4). This power loss increase may have significant impact on system availability and needs to be taken into account either by providing a thermal design power (TDP) design margin of $D(A\text{-}1)^2/(1\text{-}D)$ in the PSU, or by system power management compensating for this increase (and throttling back on power consumption).

Equation 10 also demonstrates that at large peak power factors the impact on PSU power loss could be major. For example, at A = 2 and D = 0.5, equation 10 indicates that power losses in the PSU double as compared to the maximum continuous power operating mode. Along with characterizing the power loss increase, this equation allows us to quantify the projected impact on efficiency at different peak power conditions.

Graphs illustrating the increase in P_{loss} at different peak power factors and duty ratios are shown in Fig. 5a. Graphs in Fig. 5b show the projected impact on efficiency for an 80Plus Platinum PSU having the required 91% efficiency at 100% load.

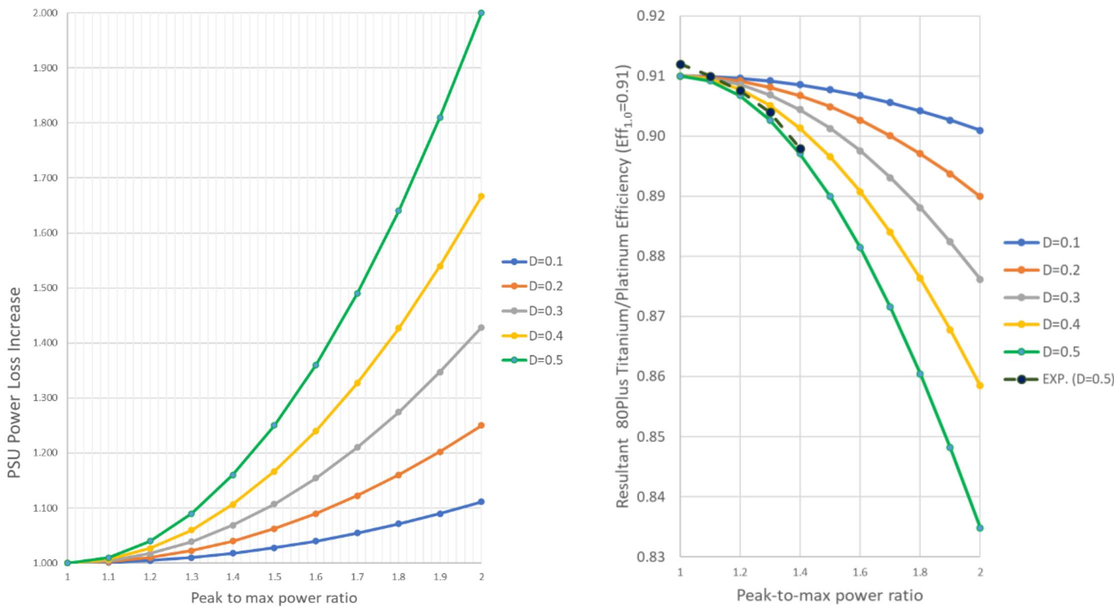

(a) *(b)*

Fig. 5. PSU power loss increase in peak power mode operation (a) and an example of the PSU efficiency drop in the 80Plus Platinum certification case (b). Experimental results obtained with a 1200-W Platinum efficiency PSU for D = 0.5, and A = 1.0 to 1.4 are given for comparison with theoretical predictions.

The P_{loss} increase graphs allow us to quantify TDP margins that are needed in the PSU to support the peak power operating mode continuously. In some cases, such support may require providing more airflow or the use of system-level power management to counterbalance the TDP increase.

As shown in Fig. 5, having peak power levels up to 1.5 of the max continuous rating at D < 30% would result in less than 10% power loss increase and less than 1% efficiency drop. This could make it a reasonable tradeoff for achieving minimal TCO if the specified TDP can accommodate the 10% power loss increase. In other words, the model represents a useful tool for power delivery spec optimization and tradeoff decision making.

For example, a reduction of power delivery efficiency that is associated with PSU peak load operation can be offset by a reduction of PSU cost and size when selecting a power supply with a lower continuous rating. On the other hand, at higher peak power levels and larger duty cycles, the greater energy cost can offset the cost savings associated with a reduction in the PSU continuous rating. The tradeoff decision in each case needs to be made based on the continuous power rating per-watt costs, electricity costs for a given application and device lifespan (product lifetime).

Experimental results obtained with a 1200-W CRPS (common redundant power supply) Platinum efficiency PSU for D = 0.5 and A = 1.0 to 1.4 (the peak factor was limited at 1.4 to prevent exceeding the actual PSU peak rating) are given in Fig. 5 for comparison with theoretical predictions. The efficiency in dynamic mode was evaluated by using input and output energy measurements taken over one-minute time intervals with a WT-210 Yokogawa digital power meter. The comparison demonstrates that the model effectively quantities the impact of PSU peak power mode operation on PSU efficiency.

The TDP impact always needs to be taken into account because the system can in fact operate in a dynamic/peak power mode for a thermally significant time period. It is also important to note that the impact of

PSU efficiency associated with peak power operation is the most critical for systems that are fully populated to their projected capacity with their PSUs operating in nonredundant mode. In frequently idling systems that operate with redundant PSUs, the efficiency drop associated with peak power operation will not be noticeable and selecting a power supply with the lowest possible continuous rating can be advantageous.

Conclusions

The proposed PSU power loss model provides sufficient accuracy for predicting efficiency and power loss. It can be used as a tool for optimization of the power delivery spec and cost by facilitating quantitative metrics for tradeoffs in PSU continuous vs. peak power ratings. Factors that need to be considered while making such tradeoffs include the system cost reduction when selecting a PSU with lower continuous power rating and the increase in electricity cost associated with the efficiency drop over the product lifetime.

The model supports multiple experimental measurement results indicating that PSU efficiency is not a monotonic function and has a maximum in a specified load range. The maximum position on the efficiency curve depends on values of the equivalent internal resistance and fixed losses.

In many cases, when the PSU load operates in a constant-power mode, which is characteristic of voltage regulator applications, increasing the PSU's output-voltage setpoint helps increase efficiency and achieve the needed 80Plus certification efficiency margins without major design changes.

The comparison with experimental results demonstrates the model's effectiveness in quantifying the impacts of PSU peak power mode efficiency and thermal design power.

Future Work

Future work could be focused on making the model an integral part of a tool for automatic computation of power delivery efficiency, adopting it for system power budget projection and generating TCO-optimized PSU spec requirements. Using the model for characterizing the energy efficiency of specific workloads, and applying it to the light load operating range could improve evaluation of the power and performance characteristics of single servers and multi-node servers running certification benchmarks.

References

1. 80 PLUS Certified Power Supplies and Manufacturers.
2. "Energy-Based Efficiency Metric Helps To Optimize Server Power Delivery For Dynamic Workloads" by Viktor Vogman, How2Power Today, October 2018.
3. Two New Control Methods to Achieve Both PFC and High Efficiency, Rohm's TechWeb, Nov. 22, 2018.
4. "AC cycle skipping improves PFC light-load efficiency" by Bosheng Sun, Texas Instruments. Analog Applications Journal, 2014. pp. 26-29.
5. "Improving Light-Load Efficiency of AC/DC Boost PFC Converters by Nonlinear Control Scheme" by K. Karthikeyan and I. Sayed Mohammed, International Journal Of Scientific Progress And Research (IJSPR), Volume-07, Number-01, 2015.
6. Intel Turbo Boost Technology 2.0.
7. "Peak-Efficiency Detection and Peak-Efficiency Tracking Algorithm for Switched-Mode DC–DC Power Converters," by V. Michal, IEEE Transactions on Power Electronics, vol. 29, no. 12, pp. 6555-6568, Dec. 2014.

About The Author

Viktor Vogman currently works at [Power Conversion Consulting] as an analog design engineer, specializing in the design of various power test tools for ac and dc power delivery applications. Prior to this, he spent over 20 years at Intel, focused on hardware engineering and power delivery architectures. Viktor obtained an MS degree in Radio Communication, Television and Multimedia Technology and a PhD in Power Electronics from the Saint Petersburg University of Telecommunications, Russia. Vogman holds over 50 U.S. and foreign patents and has authored over 20 articles on various aspects of power delivery and analog design.

For articles relating to powering servers, see How2Power's Design Guide, locate the "Application" category and select "Data Centers".