HOW2POWER TODAY
Your Power Design Newsletter

Exclusive Technology Feature

ISSUE: May 2019

Optimizing Transfer Switch N+1 Redundant Power Architectures

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

Automatic transfer switches (ATSs) have gained popularity in data center power distribution networks due to several advantages. They provide fail-safe ac power redundancy, achieve highly efficient power distribution in server racks, enable redundant power feeds for single ac power cord arrangements, and reduce UPS power conversion losses. In terms of ac redundancy, if the primary power feed becomes unavailable, the rack ATS will supply power from the secondary feed without interrupting server operation.

At the same time transfer switch arrangements can advance traditional dc redundant architectures supporting seamless dc power flow to the server when one of the power modules fails. In a conventional 1+1 (ac+dc) redundant server power subsystem without ATS each of the modules receives power from separate ac power lines so that when one ac line (or module) fails the subsystem remains active, receiving power from the second ac line and second power module (Fig. 1a).

Using an N+N dc redundant configuration (Fig. 1b) where each of the groups of N power supply units provides full system power has certain reliability advantages over 1+1 because such a configuration tolerates failures of N modules, while both ac lines are active. However, in many cases both of these arrangements can become cost prohibitive due to the high total installed power capacity (the sum of the power ratings of the power modules used in the power subsystem) or due to an excessive number of installed modules in the N+N case.

Another option, N+1 dc redundancy (Fig. 1c), for N>1, has noticeable cost advantages over 1+1 and N+N dc redundancy configurations due to the reduction of installed power capacity. However, providing ac redundancy in this case requires using transfer switch techniques widely employed in utility interactive reconfigurable microgrids.[1] Such ATS architectures are often not cost-optimized because in conventional applications they necessitate switching full system power, which significantly impacts the ATS size and ratings.

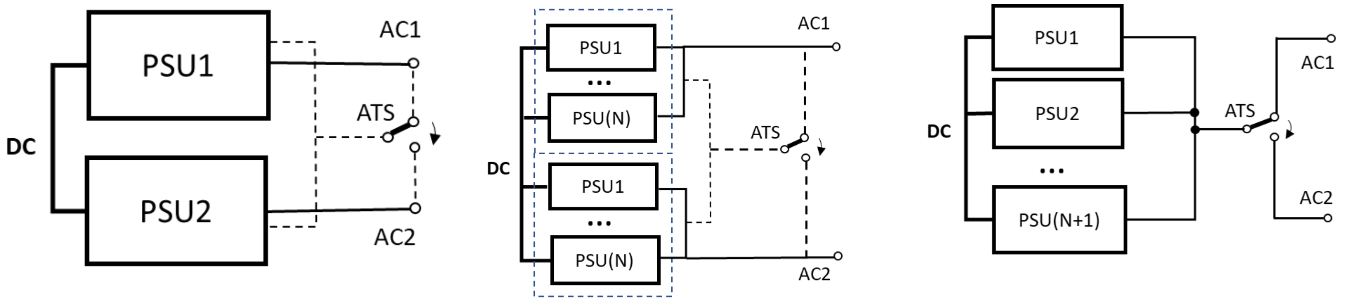

a. 1+1 (ac+dc) power redundancy arrangement. System availability can be increased by adding a transfer switch maintaining dc redundancy during ac fault time (dashed lines).

b. N+N (ac+dc) power redundancy arrangement. System availability can be increased by adding a transfer switch maintaining dc redundancy during ac fault time (dashed lines).

c. N+1 (ac+dc) power redundancy arrangement. Transfer switch maintains dc redundancy during ac fault time (dashed lines) and tolerates one PSU module failure.

Fig. 1. Block diagrams of redundant power subsystems. N+1 dc redundancy, for N>1 (c), has noticeable cost advantages over 1+1 and N+N dc redundancy configurations (a and b) due to the reduction of installed power capacity. However, providing ac redundancy in this case is not cost-optimized because it necessitates switching full system power, similar to 1+1 and N+N cases.

This article examines opportunities for adoption and optimizing of ATS techniques, reducing the cost and size of server power subsystems that provide both ac and dc redundancy. Specifically, separation of the safety isolation and power transfer functions in the ATS and reduction of the switched power level are discussed. Examples of optimized transfer switch implementations for 2+1 and 3+1 redundant configurations are presented.

Then, a look at the power levels transferred during the different stages of operation in a 2+1 configuration reveals how the configuration determines power capacity and requirements for the ATS, and these requirements are compared across 2+1, 3+1, and 4+1 configurations. Finally, some experimental results are presented, which confirm the predicted switching operation and robustness of a 2+1 redundant power system with ATS.

Advantages Of N+1 DC Redundant Configurations

If a server system requires ac power P, a 1+1 redundant power subsystem must have two power supply units (PSUs), each capable of drawing ac power level P. This brings the total installed ac power capacity to 2P, which is also the power capacity required for N+N redundant configurations, where each of the groups must deliver full system power.

For the N+1 redundancy case, the power subsystem must also tolerate one PSU failure, such that the total installed power capacity can be defined as: $P_{\Sigma} = P \cdot (N+1)/N$. Fig. 2 provides a bar graph showing the total installed power capacity as a function of the total number of PSUs in an N+1 redundant power subsystem.

Fig. 2 shows that the total installed power capacity can be significantly reduced by increasing the number of redundant PSU modules used, which in turn allows for a significant reduction in the power subsystem cost. (Of course, when adding PSU modules, you do reach a point of diminishing returns, probably above four to five modules, where power distribution infrastructure complexity and cost become dominating factors.)

Transfer switches, similar to those in utility interactive reconfigurable microgrids, can facilitate ac redundancy in such power subsystems.

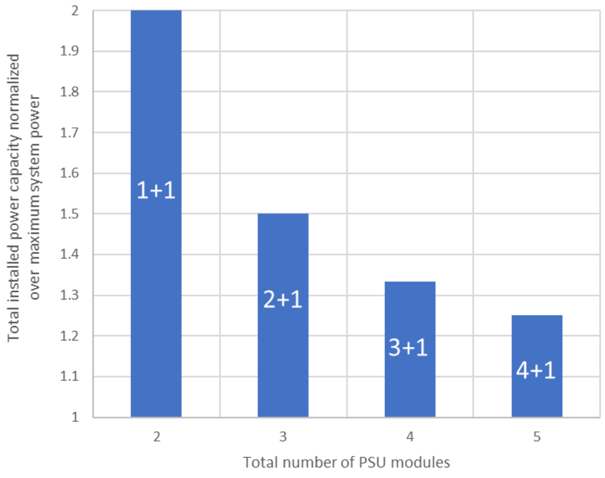

Fig. 2. Total installed power capacity as a function of the total number of power modules in an N+1 redundant power subsystem. Total installed power capacity can be significantly reduced by increasing the number of redundant modules.

A conventional transfer switch senses ac power from the ac source that supplies power, and automatically switches full system power to the redundant source at the time of the main power dropout. These switches are also required to provide a minimum clearance (4 mm) for galvanic isolation in primary circuits, as stipulated by the UL safety standard, UL 60950.[2]

Both of these constraints—the amount of power to be switched and the clearance for isolation—in turn impact the power subsystem holdup time requirement. Specifically, the greater the power to be switched and the greater the clearance, the longer the transfer switch's transition time and the greater the holdup time that is required. (This will be explained further in the next section.) These relationships represent an obstacle in redundant server power subsystems. However, a few techniques can be used to minimize the transfer switch size and shorten its operating time.

Providing Safety Clearance Without Impacting Transition Time

Typically, a mechanical relay is used as the active transfer switch component. The operating time of this switch is the transition time from the moment the control voltage is applied to the relay coil to the moment the contacts of the initially open switch close. The longer the distance between those contacts and the larger their size, the longer the operating time.

To shorten the distance between the contacts two features of the transfer switch—safety isolation and the power transition function—can be assigned to different devices (relays). To do so, two relays can be used—one providing the required safety clearance and utility protection, the other providing power transfer. A diagram illustrating this technique is shown in Fig. 3.

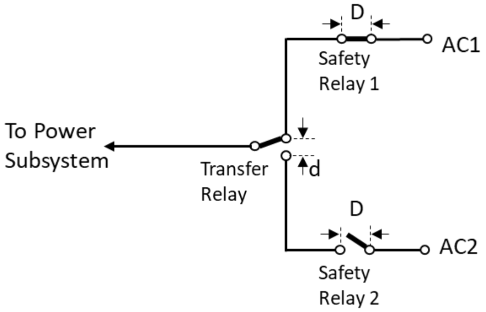

Fig. 3. To shorten the operating time of a transfer relay, safety and power transition features can be separated and assigned to different switches. While "slow" safety relay (SR) contacts travel a longer distance (gap) D, "fast" transfer relay (TR) contacts travel a shorter distance d. The safety relay contacts open whenever the ac cord is removed or the related ac line is not active. This ATS arrangement provides AC1-AC2 safety clearance equal to D+d with short TR operating time.

While the "slow" safety relay (SR) contacts travel a relatively long distance (gap) D, the faster transfer relay (TR) contacts travel the shorter distance d. Under steady-state conditions, the safety relay contacts remain open whenever the power supply's ac cord is removed or the associated ac line is not active. This transfer switch arrangement provides safety clearance equal D+d.

The time associated with ac cord removal and ac contacts exposure is always greater than practically any safety relay operating time, which is why the SR operating time can be selected as any value in the tens of milliseconds range without any safety impact.

Besides meeting the safety requirement, splitting power transfer and safety features between different switches facilitates a shorter transition time than the conventional PSU holdup time, which is typically equal to half of an ac cycle. In some cases, even a solid-state device designed for break-before-make operation can be used for this purpose. In a mechanical relay, the transition time would be minimized when the distance between transfer relay contacts is the shortest and the power it needs to switch is at its lowest level.

Minimizing The Switched Power Level

In a dc redundant power subsystem there is no need to switch all operating power modules from one ac line to another. Actually, the number of the switched PSUs must be minimized to keep the inrush current magnitude and the switched ac power level at a minimum, and therefore to enable usage of the smallest and fastest transfer relay.

The minimal number can be achieved if two modules are permanently attached to the two sources. In other words, to select the smallest and fastest transfer relay the number of the switched PSUs can be reduced to:

(N+1)-2 = N-1. For example, a 2+1 configuration would need to switch one PSU, a 3+1 configuration would need two PSUs, etc.

Examples of optimized transfer switch implementations for 2+1 and 3+1 redundant configurations are shown in Fig. 4a and 4b, respectively. When one ac line fails N PSUs remain active in both cases, which ensures the same level of redundancy as in the conventional (no-transfer switch) 1+1 case. Reducing the number of switched PSUs to N-1 minimizes the transfer switch stress, increases ATS reliability and enables selection of the smallest and fastest relay. For example, many miniature power relays, switching 1 kW of power have an operating time not exceeding 3 to 5 ms, which is well below the standard PSU holdup time requirement.

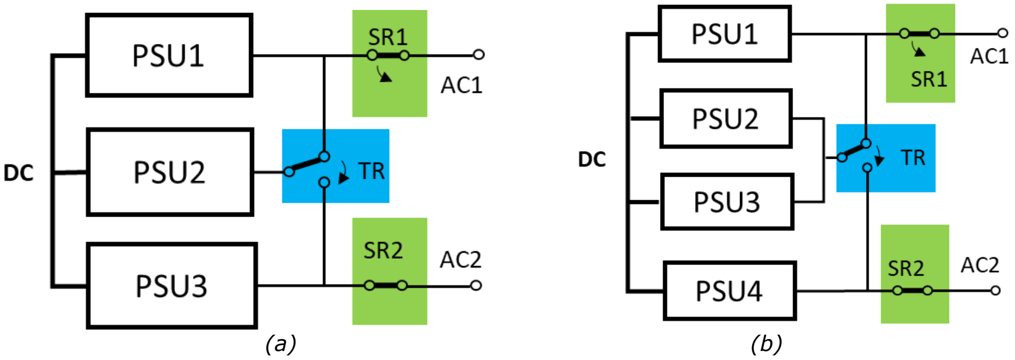

Fig. 4. Examples of the transfer switch implementations for 2+1 (a) and 3+1 (b) redundant configurations. When one ac line fails N PSUs remain active, which maintains the same level of redundancy as in the conventional (no-transfer switch) case. However, reducing the number of switched PSUs to N-1 allows minimization of the transfer relay stress.

Power Capacity Of Transfer Switch Architecture

The bar chart in Fig. 2 shows that the 2+1 configuration provides the largest reduction in installed power capacity as compared to the conventional (no ATS) 1+1 case. Let's examine the processes in a 2+1 arrangement shown in Fig. 4a and determine subsystem power capacity at each stage of its operation. A practical experimental case will be examined later.

For the sake of simplicity, let's begin by assuming that the transfer relay operating time is zero, such that the PSUs activate instantly and share power equally without delays. The timing diagram illustrating the power transition process is shown in Fig. 5.

The process starts at the initial state (t = 0), when both ac lines are active. Since PSU1 and PSU2 modules receive power from line AC1 (Fig. 4a) power supplied to these modules (P_1) and averaged over one ac cycle equals to 2/3 of the power consumed by the system (P). Power P_2 supplied to PSU3 from line AC2 equals to remaining: $P_2 = P-2P/3 = P/3$.

Once AC1 line dropout occurs (time t_1), the transfer relay switches PSU2 module to line AC2 and this line supplies full system power ($P_2 = P$) via two modules (PSU2, PSU3). Once AC1 recovers at time t_2 PSU1 becomes active and the power consumed by the system splits between AC1 ($P_1 = P/3$) and AC2 ($P_2 = 2P/3$).

When AC2 line dropout occurs (time t_3), the transfer relay switches PSU2 module back to line AC1 and this line now supplies full system power ($P_1 = P$) via two modules (PSU1, PSU2). Once AC2 recovers (time t_4) the system returns to its initial state: $P_1 = 2P/3$, $P_2 = P/3$.

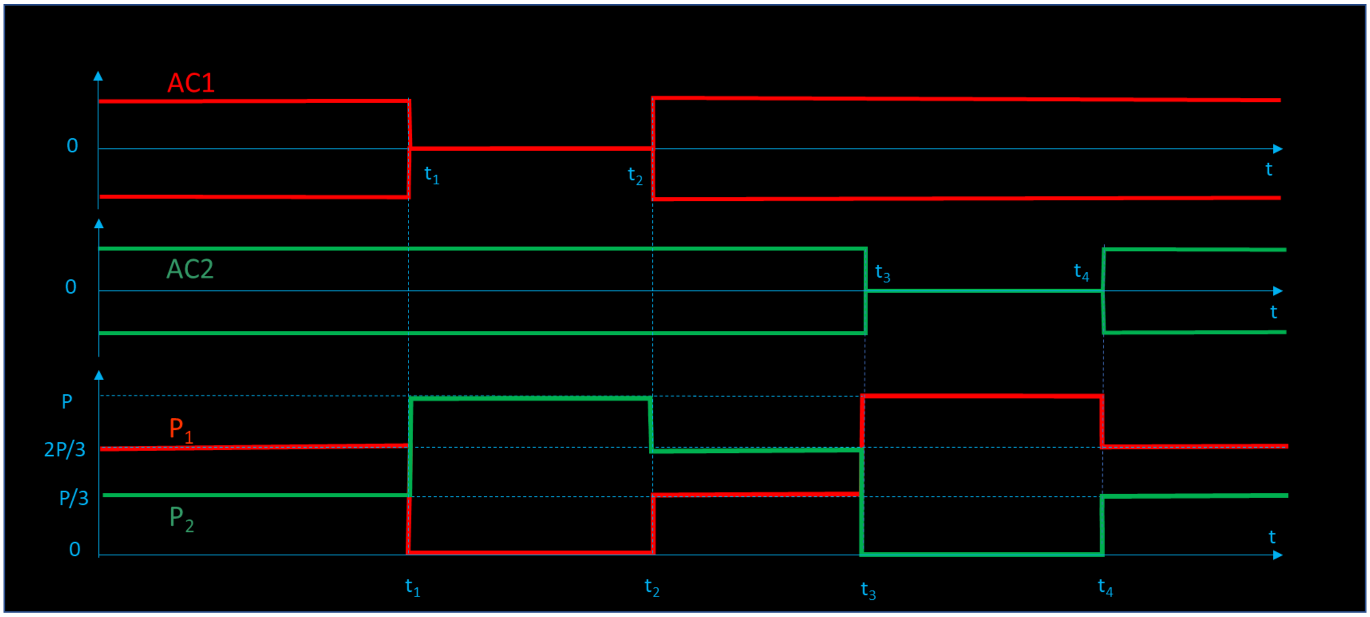

Fig. 5. The timing diagram illustrating the power transition process in a 2+1 power subsystem with a transfer switch. While the total power (averaged over one ac cycle) available for system operation always equals P, the module power reaches its maximum (P/2) during ac line outages, which means that the total subsystem installed power capacity does not need to exceed P/2 x 3=1.5 P.

The timing diagrams in Fig. 5 show that with system power level P and ideal power sharing, ac power supplied to each module reaches its maximum rating (P/2) during ac line outages, i.e. total subsystem installed power capacity does not exceed P/2 x 3 = 1.5P, which matches the corresponding level shown in Fig. 2 for 2+1 case. When both ac lines are active, failure of any of the installed modules does not cause the system to fail, as in the conventional (1+1 or 2+2) no-ATS cases.

Note that if the full system ac power equals P, the power level that needs to be switched by ATS is $P_{sw(N+1)} = P\cdot(N-1)/(N+1)$, e.g. in the 2+1 case $P_{sw(2+1)} = 0.33P$ and the PSU holdup time, which is inversely proportional to power, becomes noticeably longer than nominal. This means that in some cases safety relays SR1 and SR2 in Fig. 4 may not be required and the ATS can be implemented with a single switch.

If the PSU holdup time at nominal power is T_H and the total ac system power P is shared equally between active PSU modules, the following equations can be used to define the basic requirements for the N+1 redundant power subsystem with ATS:

Module ac power rating

$$P_{PSU} = \frac{P}{N}$$

AC power switched by ATS

$$P_{SW(N+1)} = \frac{P \cdot (N-1)}{(N+1)}$$

Installed power savings over conventional 1+1 case

$$P_S = \frac{P \cdot (N-1)}{N}$$

ATS operating time

$$T_O \le \frac{T_H \cdot (N-1) \cdot P_{PSU}}{P_{SW(N+1)}}$$

Based on these relations and assuming that the PSU holdup time at maximum power is one half of the ac cycle, we can tabulate the ATS power subsystem requirements and installed power savings for the most popular N+1 redundant configurations. The table shows that despite the reduction of installed power, the operating time and especially switched power level requirements become stricter as the number of redundant modules used increases, which needs to be considered for actual redundant architecture selection.

Table. ATS power subsystem requirements and installed power savings for redundant power systems.

Redundant Configuration	Module Power rating*	ATS operating time** must be less than:	Switched ac power*	Installed power savings*
2+1	0.50	0.758	0.33	0.50
3+1	0.33	0.666	0.50	0.67
4+1	0.25	0.625	0.60	0.75

* Normalized to max system power; ** Normalized to the ac cycle time.

Experimental Results

An experiment was performed to verify the robustness of the transfer switch architecture by performing repetitive voltage dropouts of sources AC1 and AC2 under different load and line conditions. A total of 28,000 dropouts were generated by two ac inrush current testers described in reference [3] while output dc voltages always stayed within their regulation limits.

Typical 2+1 power subsystem waveforms are shown in Fig. 6. If we discount the instantaneous signal nature, the experimental waveforms in Fig. 6 follow the reference power waveforms presented in Fig. 5: $P_1 = 2P/3$; $P_2 = P/3$ (0 to t_1 time interval), $P_1 = 0$; $P_2 = P$ (t_2 to t_3 time interval), $P_1 = P/3$; $P_2 = 2P/3$ ($t > t_4$). The real-time waveforms in Fig. 6 differ from the ideal averaged power waveforms in Fig. 5 due to the ordinary short inrush/rerush current spikes and transients associated with non-zero TR operating time (t_1 to t_2 time interval), PSU non-zero start time (t_3 to t_4 time interval), and non-zero current share settling time, which is natural for any redundant power subsystem.

The transition time of an active power supply from one ac line to another (time interval t_1 to t_2 in Fig. 6) must be shorter than the holdup time at a given switched power level. This requirement is usually easy to meet in the 2+1 transfer switch arrangement, because in the worst case its operating power level (P/3) is significantly (33%) lower than rated (P/2). The time lag associated with activating a previously deenergized PSU (t_3 to t_4 time interval in Fig. 6) is not critical as it just delays full (ac + dc) redundancy status by a few tens of milliseconds, which is also similar to the conventional case (a redundant configuration with no transfer switch). During the experiment all the dc outputs remained within spec limits.

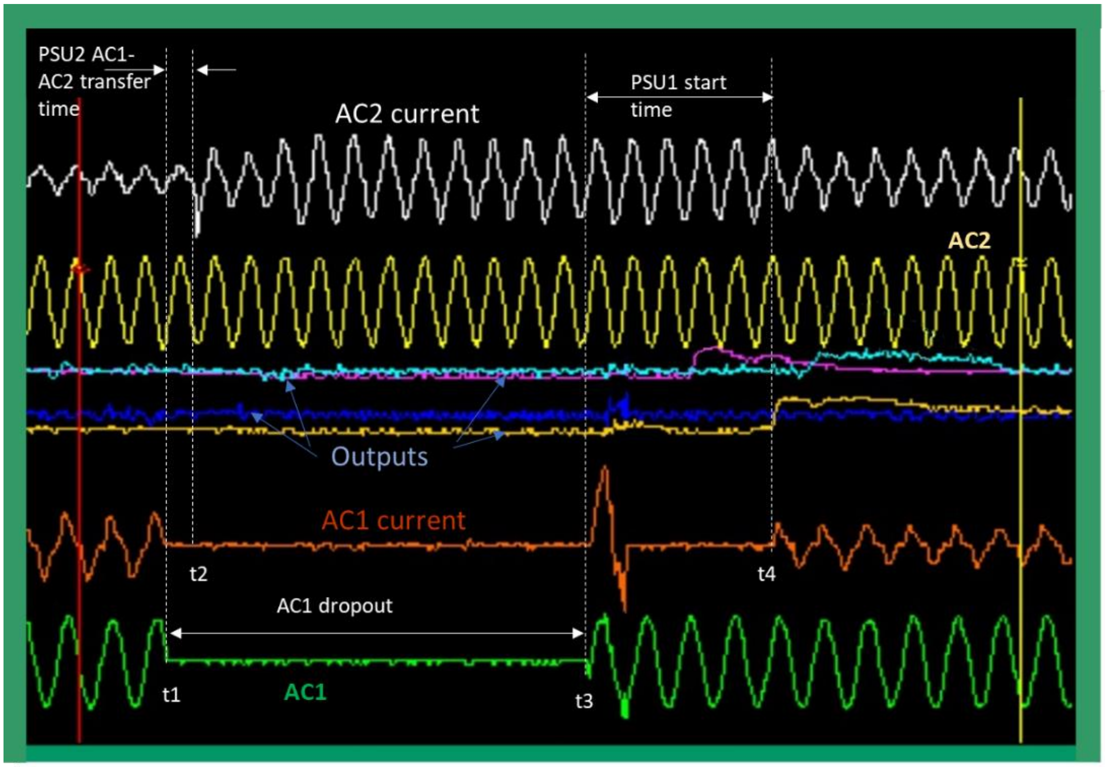

Fig. 6. Actual 2+1 power subsystem voltage and current waveforms captured by an eight-channel digital oscilloscope and displayed with Labview. The experimental waveforms follow the reference averaged power waveforms presented in Fig. 5 for AC1 dropout time interval: $P_1 = 2P/3$; $P2 = P/3$ (0 to t_1 time interval), $P_1 = 0$; $P_2 = P$ (t_2 to t_3 time interval), $P_1 = P/3$; $P_2 = 2P/3$ ($t > t_4$). They differ from the reference ones only due to short inrush spikes and transients associated with non-zero transfer time (t_1 to t_2 time interval), PSU non-zero start time (t_3 to t_4 time interval), and non-zero current share settling time.

Conclusions

Splitting transfer and safety features between different transfer switch relays results in a significant speed-up of the transfer process and makes the transition time shorter than the PSU holdup time.

The experimental results follow predictions represented by reference timing diagrams with sufficient accuracy for practical purposes consistent with actual PSU power sharing accuracy.

The proposed transfer switch architecture in N+1 (N>1) redundant configurations simplifies the PSU holdup time requirement and permits using a smaller and faster ATS, as compared to the conventional full power switching cases. This results in server power subsystem cost and size reduction. Despite the reduction of installed power, the operating time and the switched power level requirements become stricter as the number of modules in the power subsystem (N+1) increases, which needs to be considered when selecting a feasible redundant architecture.

References

1. "Control of Autonomous Single Phase Utility Interactive Reconfigurable Microgrid," by Shailendra Kumar, 2018 IEEE Energy Conversion Congress and Exposition (ECCE), p.31-36.
2. UL Standard for Safety for Information Technology Equipment, UL 60950, Third Edition, Underwriters Laboratories (UL), ISBN 0-7629-0470-4.
3. "Inrush Current Tester Enables More Thorough Validation Of Server Power Supplies" by Viktor Vogman, How2Power Today, April 2018.

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 more information on current sensing techniques, see How2Power's Design Guide, and do a keyword search on "current sensing." Also locate the Design Area category and select Test and Measurement.