

REMOTE

Site & Equipment Management

April/May 2006 Volume 6, Issue 2

www.RemoteMagazine.com

Optimizing Uninterruptible Power For Modern Data Processing Equipment

By Gary Mulcahy
Transistor Devices, Inc.

New generation data processing equipment provides ever-increasing speed and bandwidth utilization benefits, while Internet based communication and commerce activities continually fuel the need for more performance. One result of this self-reinforcing spiral is the amount and concentration of electrical energy being ported toward core data processing infrastructure. This has brought forth new challenges regarding energy and resources wasted on inefficient power systems.

Data processing equipment is generally segregated into channels built around microprocessor cores. In the past these microprocessors would typically utilize less than 100W of energy per channel. However, modern equipment quite often sees this power requirement exceed 200W per channel, with future equipment forecasted to go well in excess of this value. As these channels are paralleled together into clusters, the total energy being utilized presents unique challenges.

With current industry practices, AC power distribution systems are typically employed to power data processing equipment clusters. In addition, critical infrastructure installations generally employ power systems with an uninterruptible source of electrical power. Uninterruptible power is generally viewed at the facility level, as opposed to the equipment level. While this provides an easy demarcation line between the facility and data processing equipment, with each focusing on a different part of the problem, it also results in overall operating efficiency and total cost of ownership being difficult to ascertain and optimize.

Facility-level AC UPS systems are quite often constructed around model similar to the one portrayed in the following diagram.




Figure 1: Facility Level AC UPS

Figure 1: Facility Level AC UPS

The central AC UPS generally involves conversion from incoming raw AC power to DC, and then reconverting from DC back to AC. Moreover, since data processing equipment generally operates from 208/220VAC line voltage, facility-level entry power usually must be transformed from 480VAC down to 208VAC either before or after the UPS. A battery is then connected as an alternate input to the DC-AC converter so that if incoming raw AC power is interrupted, the system automatically switches over to battery power. Alternately, a rotating flywheel may be used to store energy and provide uninterruptible AC power. Uninterruptible AC power is then passed through an AC power distribution grid and fanned out to individual data processing apparatus.

Once inside DP equipment, AC power is once again converted into DC power. Off-line power supplies provide power factor correction as well as load isolation from the incoming power line.

This generally involves at least two stages of power con-

version. In addition, most modern microprocessors require very low voltages at fairly high currents, such as 1.1V at 100A.

Moreover, the precision of the voltage required is such that voltage regulation circuitry must be located directly next to the microprocessor. In order to effectively realize this circuitry, most processors require that an intermediate DC voltage, such as 12VDC, be delivered to the processor/local regulator combination. (In ATX compatible servers, off-line, “silver box” power supplies provide this directly to the motherboard. In Blade Servers, the off-line supply usually provides a higher voltage and an “intermediate bus converter” is included on the micro-controller board to provide 12V.) Thus, there can be up to six or more power conversion stages between facility power entry and the microprocessor.

1. 480VAC to 208VAC (Step Down Transformer)
2. 208VAC to 400VDC (ACUPS Front End)
3. 400VDC to 208VAC (ACUPS Back End)
4. 208VAC to 400VDC (Silver Box PS PFC Front End)
5. 400VDC to isolated 12VDC (Silver Box PS Back End)
6. 12VDC to 1.1VDC (DC-DC converter on mother board)

Limiting factors on power conversion efficiency are voltage ratings of semiconductors, along with their corresponding conduction losses, as well as power converter package size limitations, and economics. All of these combine to ultimately limit available efficiency improvements. Power conversion efficiency for best in class AC UPS equipment ranges as high as 92% for isolated, battery based systems. In the case where an ATX, computer grade AC-DC power supply is utilized, these generally do not provide efficiency higher than 75%. Non-isolated converters that provide final processor power conversion range as high as 88% efficiency. As detailed in Figure 2, in a typical ATX based solution, the net power efficiency from facility entry AC to the processor is 59.5%. For every watt of power utilized to process data, another 0.66W is required to support power conversion. (In the case of air conditioned facility cooling, up to another watt of power will be required for each watt utilized to cool the power conversion equipment.)

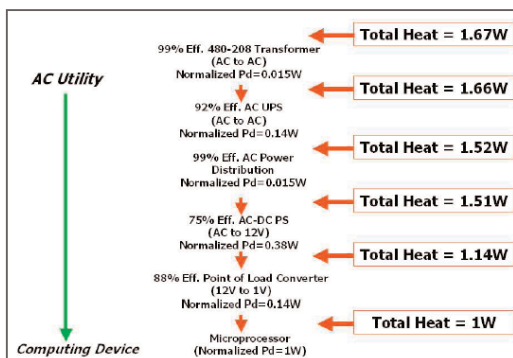


Figure 2: Normalized Heat loss for ATX Power Supply Model

will provide up to 94% efficiency. As shown in Figure 3, the net power efficiency from facility entry AC to the processor is 68.6%. For every watt of power utilized to process data, another 0.45W is required to support power conversion. (In the case of air conditioned facility cooling, up to another 0.68 watts of power will be required for each watt utilized to cool the power conversion equipment.)

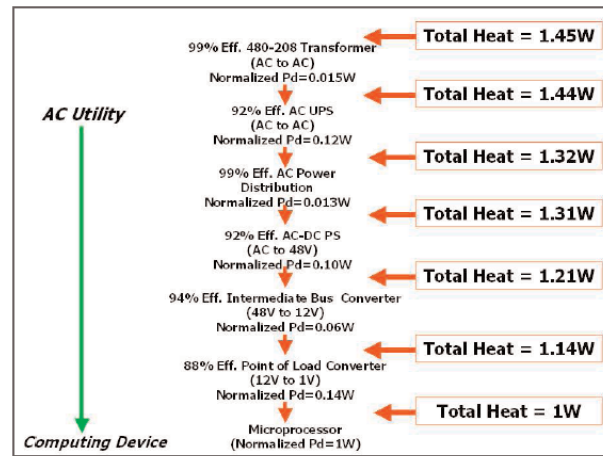


Figure 3: Normalized Heat Loss for Cabinet-level 48V Bus Model

While this is not too troubling at the individual microprocessor level, when overall data processing activities reach power usage levels on the order of 500kW, or more (not too uncommon in large internet hub installations), the impact of this level of inefficiency can be considerable.

Another area of complexity within large data processing installations is uninterrupted AC power distribution. Facility-level AC UPS systems generally have a centralized circuit breaker panel with power monitoring and bypass provisions for UPS and battery servicing. Quite often these are located a substantial distance away from the equipment being powered, leading to possible confusion regarding which equipment is fed by what breaker and raising the possibility of inadvertent DP equipment shut down due to operator confusion. In addition, AC power distribution must be installed in anticipation of eventual DP equipment usage, often requiring upfront capital expenditures larger than necessary for planned DP equipment phase in. In some cases the nature of the UPS being utilized is very sensitive to load current harmonics, resulting in a great level of care being required to assure data processing equipment and UPS compatibility.

Facility-level AC UPS can also represent exposure to a single point of failure for large amounts of data processing infrastructure.

Centralized battery plants utilized in facility level UPS systems may be a compromise between what is required for successful system realization and component limitations. Many times the actual battery run time required for acceptable system operation is a fraction of a minute, representing the time required to switch between alternate utility power feeds or the time to bring an auxiliary source of power on line, such as a motor generator. When sizing batteries for centralized AC UPS systems the nature of the voltages required or battery type chosen can often result in hold up times well in excess of what is required, resulting in wasted energy capability.

Likewise, the lower efficiency presented by the off-line power supplies utilized by DP equipment produces an extra load on batteries that only goes into producing heat for power conversion purposes.

Distributed DC UPS Opportunities for optimization are presented by employing the concept of DC uninterruptible power systems distributed throughout the data processing facility.

Unlike traditional central office telecommunications installations that have one large DC power plant feeding an entire facility, these systems are comprised of smaller power plants spread throughout the facility. They are sized to power clusters of end use equipment, with optimal effect being achieved up to 200kW per cluster. As an abundance of equipment exists from telecommunications technology built around -54.5VDC, this becomes a logical distribution voltage. At 200kW, the equivalent current for this voltage is approximately 3670A.

A distributed DC UPS approach is depicted in Figure 4.

In this system, each data processing cluster has a dedicated DC UPS system feeding it.

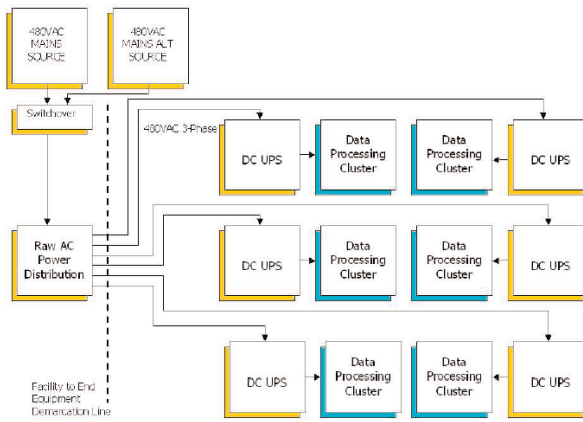


Figure 4: Distributed DC UPS




Figure 5: 200kW DC UPS for Distributed DC UPS Architecture

Figure 5 presents a 200kW DC power plant suitable for Distributed DC UPS applications in Data Centers. This power system is energized via a 3-phase, 480VAC power line, with a nominal AC Mains Current of up to 267A per phase. The power plant contains internal battery backup to supply 200kW for more than 15 seconds, which allows adequate time to start and switch in an alternate power source in the case of utility AC power failure. (Battery holdup time will increase as power draw is decreased.)

Figure 6 depicts the 200kW power plant in a massed server cluster application. Each 20kW Server Cabinet is fed by two 54.5V/180A feeds. Each feed is serviced by a 200A circuit breaker located in the DC UPS cabinet. With cables set at AWG-0, overall cable dissipation is less than 1600 watts for 200kW of delivered power, yielding DC Distribution Efficiency (=power consumed by distribution wires / delivered

power) better than 99%, which is similar to what is achieved in AC distribution systems. (Note that overall volume and weight of copper on the DC distribution system pictured is approximately twice that of an equivalent AC system with a service length of fifty feet and equal to that of an AC distribution system of one hundred feet.)

Distributed DC UPS technology can also be effectively deployed at the cabinet level, as depicted in Figure 7.

The DC UPS power plant depicted in Figure 6 provides enough battery backed up power to support 20kW of DC powered data processing equipment for approximately 15 seconds at full power.




Figure 6: Distributed DC UPS Physical Realization - 200kW Cluster




Figure 7: Cluster Server Cabinet and 20kW d²ups Uninterruptible DC Power Plant

Summary of Advantages of Distributed DC UPS (d2ups)

Upon examination of the Distributed DC UPS approach, many benefits become apparent.

1. Higher overall operating efficiency.

Best in class DC UPS equipment runs at greater than 91% power conversion efficiency. In addition, the traditional AC-DC “silver box” power supply is replaced with a DC-DC unit designed to operate from -54.5V. Highly efficient semiconductors available at this voltage provide an efficiency improvement to greater than 94%. Overall power conversion efficiency then increases to 75%, or a 50% improvement in power lost to power conversion over the central UPS/ATX case. This reduces the amount of energy wasted in power conversion by half (0.33W of power conversion wasted energy for every watt consumed, as opposed to 0.67W for the central AC UPS and ATX power supply case).

2. Better distribution of power conversion heat dissipation.

The majority of power converter heat dissipation is removed from data processing equipment. In the centralized AC UPS/ATX case, the biggest hitter on power conversion efficiency is the “silver box” power supply that is located within

the data processing equipment. With a Distributed DC UPS approach, this power supply is replaced by a high efficiency DC-DC converter that provides significant heat reduction within each data processing channel. For example, at 250W delivered power, a 75% efficient "silver box" power supply dissipates an additional 83W, or 285 BTU's per hour of heat into each processing channel. A well-designed DC-DC card operating at 94% efficiency dissipates 16W for the same delivered power, reducing the heat load per channel by 80% to 55 BTU's per hour.

3. Improved system reliability

The total number of power conversion steps is reduced from six to four (480VAC to 400VDC, 400VDC to -54.5VDC, -54.5VDC to 12VDC, 12VDC to 1.1VDC), thereby reducing the total number of electronic components and circuit complexity. Dividing up UPS functionality to individual clusters will also reduce the overall system's vulnerability to a single point of failure.

The DC-DC converter employed in individual servers is built around low voltage technology, which is fundamentally more reliable than higher voltage components utilized in off-line equipment. It is not unusual to have a Mean Time Between Failures (MTBF) of less than 100,000 hours on typical computer grade AC-DC power supplies. The DC-DC converter considered in this paper benefits from lower overall parts count, low voltage components with inherently higher reliability than higher voltage parts, and a well controlled manufacturing process that includes Highly Accelerated Stress Screening on each individual unit produced. Putting all these together results in a power converter that yields over 1 million hours MTBF. Due to the large numbers in which they are deployed, system down time and service requirements will be significantly improved. Reduced heat load within individual servers will also lead to improved reliability of data processing equipment.

4. Simple DC power distribution.

400A per cabinet distribution is easily handled via WG-0 cables. These provide low distribution losses, while also being easily adaptable to numerous installation scenarios (both above and below the computer room floor). Each data processing cabinet has a pair of 200A circuit breakers, providing branch circuit protection. Overall layout of the DC distribution bus is economical, safe and simple to understand. Other DC distribution arrangements are also available, such as DC plates (or "bus ducts") that can be arranged either above or below server cabinets.

5. Simplified AC power distribution.

Distributed DC UPS equipment can operate from 480VAC-three phase power lines, thereby simplifying facility AC wiring and eliminating bulky 480 to 208VAC step down transformers. At the beginning of a project AC power can be ported to the intended operating locations in a grid that anticipates ultimate installation size. Distributed DC UPS equipment does not have to be installed until that sector's data processing cluster is due to be installed.

6. Simplified control and improved fault containment.

As each data processing cluster has its own UPS, confusion regarding which equipment is being fed by what UPS breaker is eliminated. Likewise, maintenance or equipment upgrade activity

is localized to individual clusters, reducing the downtime risk to other clusters during these events.

7. Opportunity to optimize battery size.

A DC distribution bus built around 3700A of DC current at -54.5VDC provides the opportunity to utilize new technology batteries that are designed to provide very high currents for short periods of time. Thus, the battery pack to support a 3700A cluster occupies approximately a 10 cubic foot area and provides approximately 0.5 minutes of holdup time at full load.

8. Reduced capital, installation and operating costs.

DC distribution system components have been developed around telecommunications technology where a large number of cost competitive, mature alternatives are available. Analysis of industry standard centralized AC UPS costing versus the Distributed DC UPS technology described in this paper indicate capital and installation cost savings of up to 25%. Likewise, improved operating efficiency can provide savings in electrical power usage as high as 20%.

9. Other opportunities for optimization.

Localizing the bulk of power conversion heat load into a single cabinet provides other opportunities for cost management. Cooling moving air via HVAC equipment is generally a cost inefficient method. With most of the power losses in one cabinet, liquid cooling becomes a viable option. Likewise, as the ultimate challenge of cooling the data processing equipment is addressed, power conversion losses can be easily configured to latch on to similar cooling schemes (ducted air, circulating liquid, air-to-liquid heat exchange, heat pipes, etc.).

Conclusion

Ever increasing data processing power demands require a fresh look at how things are done and achieved total cost of ownership. DC UPS technology developed in the telecommunications industry provides a highly reliable, cost effective alternative to traditional AC UPS systems, especially when applied in a non-traditional, distributed manner.

Gary Mulcahy is Executive Vice President of the Commercial Products Division of Transistor Devices, Inc. (TDI). He received his BE-EE from New York University followed by graduate study at the Polytechnic Institute of New York. Mr. Mulcahy is a recognized authority in power conversion technology and the design, development and production of power systems for maximum performance and reliability with minimal life cycle cost of ownership. Contact TDI at www.tdi-power.com

