Understanding heat, it's common sources and common effects on a battery supported power system

Introduction

When it comes to UPS systems, heat is a topic that is rarely considered by newcomers to the field. For domestic or household applications, its relevancy is hard to see, but having an understanding of what it does to your power electronics can help you achieve a better outcome.

Heat is a hidden killer of batteries, and therefore the systems they support. It is one of the major reasons professional race cars have the battery relocated outside of the engine bay. Engine bay temperatures can easily reach 120 degrees when running flat out, and heat of that level will destroy a battery far more rapidly than you expect.

Part of the issue is that many people forget that the electronics in a power system represents only about half of the power solution. The other half is the energy storage in the form of a chemical reaction – and a battery is not just a box of volts.

This whitepaper looks at the issue of heat a little more holistically, in an effort to help non engineers understand the compromises and risks inherent in every battery backed up power system, and help find ways to minimise them.

Heat Sources.

Unless you're privy to the very latest superconductor technology, nothing is 100% efficient. Never is this truer than when you're working with electricity. Every time you draw power to a device, heat is generated by every part of it. Even the cables in the walls generate heat. Proper design and sizing of parts is undertaken to ensure the heat generated is absolutely minimised, but it **is** there. That's why there is a specialist field for power analysis called thermal imaging. Even the circuit boards in your mobile phone give off heat. It's smart to assume every electrical component gives off heat, because 99.99% of the time you'll be right.

In a Double Conversion UPS, power is converted from AC to DC, then back to AC again. Each of those conversion steps releases heat. The major heat losses occur from the power conversion steps, but there are minor heat sources as well. Fan bearings release heat from friction, battery chargers, switchgear, Integrated Circuits (IC's), anything that has a 'resistance' regardless of how small, releases waste heat.

How much waste heat can vary. Most UPS manufacturers laud the efficiency of their designs, usually using comparisons to 10 year old designs to make the results look significantly improved. What they fail to tell you is that the efficiency of a UPS system is a curve, and they are only particularly efficient in a small part of the operational loading curve.

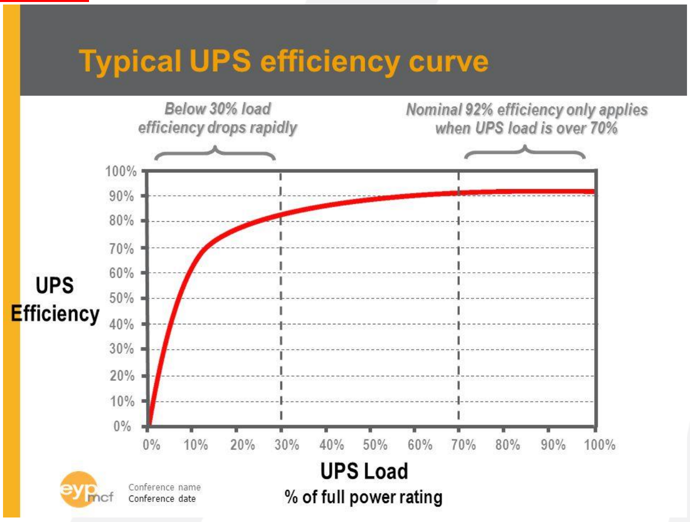

At idle (unloaded), all the energy drawn by a UPS is lost as heat, noise or earth leakage. For example, a 20kW UPS, unloaded, may draw 800 Watts of power, while out putting zero watts. That's not particularly efficient. The efficiency comes in when the UPS is running at something like 75% load, where it's putting out 17.5kW of energy, and is probably only dumping 1800 watts of waste energy as heat and noise.

However, that means your UPS is still dropping maybe 1600 watts of thermal waste into the environment around it.

If it's recharging the batteries, there will be a little more lost as the chemical reaction in the batteries is not 100% efficient either. You may find another 100 watts of waste heat is radiated from the batteries.

While the UPS is supporting the load, and recharging the batteries, there will be other things going on as well. For example, the Power Factor Correction (PFC) subsystem will be in full swing, working (sometimes very energetically) to correct power factor issues from the load. The amount of heat generated by this depends on the **load** performance and not the UPS, and sometimes it can be significant. While there are standards that govern most of the power factor requirements, there are some special instances where those standards are waived – and the heat generated from exotic equipment's poor power factor or harmonic disturbance can be extreme - extreme enough to turn a UPS into charcoal.

This photo was taken from a UPS at a university that was exposed to extremely large reflected harmonics from specialised laboratory equipment, for an extended period of time.
As a ‘Double Conversion’ UPS provides galvanic isolation between output and input (no direct power flow through), the UPS had to absorb and dissipate these harmonics. A UPS is not intended to do this at this level - it’s designed to deal with harmonics to the limit of what is allowable on the power grid (8% harmonic distortion). Instead it was exposed to 50% harmonic distortion.

This image is of a waveform analyser attached to a replacement UPS for the burned out one, showing the harmonic distortion at the UPS. The top section of the graph represents the input power – and what the grid sees. The bottom is the output – or what the UPS sees.

Short term ramifications of waste heat

In the short term, assuming everything in the UPS is close to new, the effects may not even be immediately noticeable. If you look at the battery system, a moderate (sub 10 degree) increase in battery temperature may actually result in better battery performance. It's a truism that the hotter a chemical reaction, the more 'energetic' it is. The down side to this is that the more energetic a reaction, the more uncontrolled it is, and the greater the risk of unexpected results. A UPS battery is a reversible reaction, but it is reversible **only** if it reacts as intended. If it reacts in an unexpected manner, and creates unexpected by products, then it will not be 100% reversible.

The moral of this story is that a temperature accelerated reaction in a UPS battery may form insoluble salts, reducing its capacity. The greater the difference between battery temperature and the manufacturers recommended limit, the greater the rate of insoluble salt production. The longer the increased heating exists, the more salts produced.

There is also an inherent risk of higher operating temperatures leading to a catastrophic event called thermal runaway. This is an event that occurs when the rate of heat generation in the battery exceeds its ability to shed that heat. There are a few reasons for this to occur, but they typically involve batteries that have increased internal resistance, an internal short circuit, or have had their ability to cool limited or significantly reduced. More info on the mechanics of thermal runaway can be found here - http://www.dtic.mil/dtic/tr/fulltext/u2/a477925.pdf

Another major short term effect of this heating, although indirect, is absolutely critical to the system's short term operation. If the total heat generated by the system exceeds the ability of the environment to shed that heat, then the system's internal temperature will climb. Obvious enough right? What's not so obvious is if the **rate** of temperature increase exceeds the operational temperature limits before the system's energy storage is depleted, is that the excess energy storage will never be used. If you have a 2 hour battery bank, and the system heats to temperature cut off point (usually 40 Degrees C) in half an hour because the aircon is no longer running, you will never be able to use 75% of your battery capacity.

Long term ramifications of waste heat

On batteries, the long term effects of higher temperatures are significant. Ordinary Valve Regulated Lead Acid VRLA batteries as used in UPS systems will exhibit significantly reduced lifespan due to increased average temperatures. The rule of thumb is that a 5 degree increase in average ambient temperature over the nominal 25 degrees C, that continues for 1 year, can result in a 50% service life degradation.

Degradation to the electronics can also be seen. Electrolytic capacitors will dry out more rapidly, shortening electronics' service life and the risk of component fault. Thermal transfer compounds will dry out more rapidly, electrical insulation components and even printed circuit board material will degrade. Most importantly, the failure rate of semiconductors increases as temperature increases (Arrhenius Law) – the rate of component failure is between 10 and 30 times **greater** at 60 to 80 degrees C, than it is at 30-40 degrees C.
http://www.apisteglobal.com/enc/technology_enc/detail/id=1262

Impact on UPS system design and decision making

One of the foundation principles that underpins my work as a system designer, is knowing the design principles applied to the products I use.

For example – one of the more common buzzwords / phrases in the IT industry is 'energy density'. Spruiking energy density is common, as it implies you can get more bang for the buck out of smaller spaces. Reduced infrastructure footprint means more space for customers or salespeople to utilise. What 'Energy Density' means to some of us engineers, is less available space, and therefore potentially less room for heat sinks, possibly poorer airflow, leading to higher average component temperatures – all of which add to greater thermal stress and more rapid 'aging' of components.

Good design can mitigate some of these issues, but not all of them. Good design is also often put aside for economic reasons – usually not to the point of 'bad design', but packaging and cost constraints are often allowed to interfere with 'ideal' design arrangements.

A good example of this effect is comparing a UPS design from 20 years ago with today. 20 years ago designs were far less efficient, heavier and larger, and the power quality output was much poorer. However they were also significantly more robust. Suffice it to say that the UPS designs used on oil Rigs and other ultra harsh, ultra critical locations are based on 20 year old designs and not the latest and greatest technology.

These design differences – Good Design, Economical design and so on do significantly contribute to how heat is dealt with in a UPS system, and therfor how rapidly a UPS system ages. If you compare 2 UPS systems made from **identical components,** a system that is more 'energy dense', packing more hardware into a smaller volume, will have greater difficulty in achieving the same degree of

1300 662 435
poweronaustralia.com.au
Power On Australia Pty Ltd - ABN 48 110 752 442
Unit 20, 256-258 Musgrave Road
Coopers Plains Qld 4108 (Head Office)

component cooling as in a larger, less dense system, and therefore may well have a shorter lifespan despite using identical reliability components.

Application and design considerations for UPS construction and installation

There is no one size fits all solution that does not have significant inherent compromises. An up to date data centre with its clean room environment will fare well with the latest and greatest or highes energy density technologies. They will be able to take advantage of the strong monitoring suites available and the near daily maintenance regimes that are standard operating procedures in information storage facilities. Whereas a fabricator with a UPS backed up CNC Mill, an industrial baker who needs ride through for his conveyor systems, or the mine site communications room will all have completely different requirements.

Manufacturers spend a lot of time and effort on thermal considerations in their designs, doing their best to hit the design targets they have. Those targets can include;

- Convection cooled systems so fan failures are not possible
- Redundant fans so fan failures are of limited effect
- Modular capacitor banks for easy and fast replacement at end of life
- Input air filters to minimize thermal degradation of systems
- Air ways and air paths design for most effective use of cooling air
- Component sizing to operate in the most thermal efficient operational bands
- Higher internal operating voltages, to lessen heating due to current considerations
- Operating frequency of the internal clock – lower is cooler, faster is higher efficiency

This information is not new to the UPS world, and there are ways of measuring the effect of systems designs, and how quickly they age.

All of these design considerations – and many more – are often rolled into a technical value called Mean Time Between Failures(MTBF). MTBF values are usually a calculated value, using a method specified by the US military, as the performance of many UPS systems (and other systems as well) exceeds easily monitored values. Only units that have been operational for many years can have MTBF values based on field proving – and by that time the design has changed.

MTBF values can be very telling in terms of equipment longevity. They can also be misleading because they don’t necessarily cover all parts of a UPS system. Often there will be subclauses written into the MTBF documentation stating things like ‘excludes consumable components’, which means fans, capacitors and other limited service life products are taken out of consideration to improve the MTBF result.

It’s typical for a 6kVA UPS to have an MTBF of 50,000 operational hours, a 40kVA UPS to have 100,000 hours, and 300kVA UPS to have 150,000 operational hours and a true industrial UPS (6 Pulse / 12 Pulse) to have 250,000 operational hours. All of these fade before telecommunications grade systems which routinely exceed 400,000 operational hours MTBF.

Batteries are usually 5 year design life or 10 year design life, with special application batteries like NiCad’s reaching 25 years design life.

Guess how the costs scale ?

Conclusion

The best advice is to know your application well. Look at the environment, the characteristics of the load that will be supported, and the performance required. In my experience, assumpotions in this area result in solutions that prove unreliable, have shortened lifespans, greater and earlier maintence requirements, as well as presenting unneccesary risks and hazards.

Here are some broad headings that may be useful when thinking about your application. These are relatively common hurdles power systems specialist have to jump.

- Assess your operating environment, temperature and humidity.
- Assess the nature of the load, especially inrush currents
- Assess your cooling system and how it's going to work during all conceivable UPS operational modes, especially in battery mode.
- Talk to your product expert, get some opinions on what technology (not what brand) will work best in your application
- Specify that technology when talking to your vendors / suppliers. If you let the vendor change the technology on offer, then make sure you understand the ramifications this will present.

All design is compromise, and no plan survives contact with the enemy . The best you can do is to manage your expectations, and know your required deliverables up front.

The UPS world is diverse, and there are thousands of voices clamouring for your attention. With such a critical piece of infrastructure in play, it's important that you get what you need.

The trick is knowing what you need.