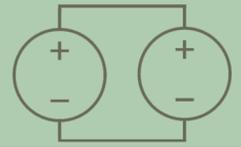




Load Sharing Considerations: Why Two is Not Always Better Than One



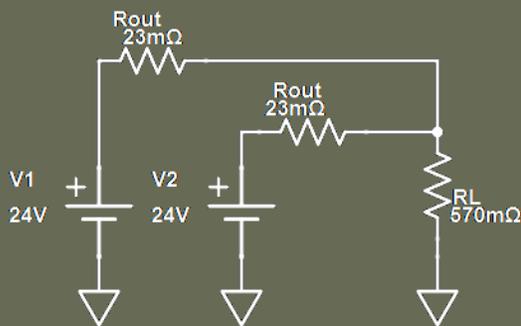
Introduction

From time to time, we get asked about the current sharing capabilities of our power supplies, and whether or not they can feed a load in parallel. This usually comes up when discussing modular system applications where the desire is for independent supplies to be able to be added and removed from a central power bus, or occasionally when our highest-power offering still isn't quite big enough for the application at hand. More often than not, the simple answer to the question is no, but to those who prefer to understand why, this article is for you.

In addition to analyzing the exact physical reasons why simple paralleling does not usually work as one might assume, this article will highlight the abundant constraints of one of the most commonly referenced, though perhaps most limited load sharing tactics, and will discuss what mechanisms must exist to facilitate proper load sharing.

The Problem with Paralleling

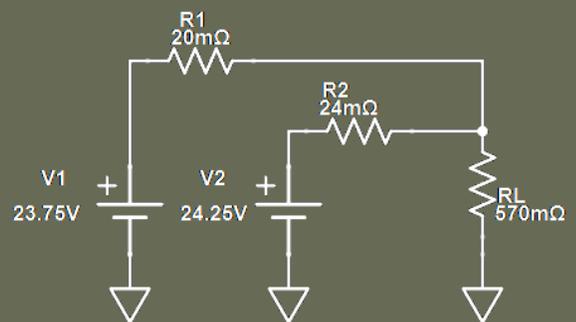
It may not be immediately obvious why feeding a load with two independent power sources in parallel simply does not work. If we first look at the problem from a simple circuit analysis standpoint, the flaw is not necessarily intuitive. Consider the example below, where two identical 24V, 500W power sources, with output impedance R_{out} work together to feed a 570m Ω resistive load. R_{out} can be determined by examining the load regulation of the supplies. If we assume the load regulation is $\pm 1\%$, then at full load (20.83A) there will be a voltage drop of 480mV which gives $R_{out} = \Delta V / \Delta I = 23\text{m}\Omega$.



A quick DC analysis reveals that in the proposed circuit, the load current is approximately 41.27A and that exactly half of that current comes from each of the independent power sources. That is, the sources are effectively sharing the load and neither source exceeds its 500W rated limit. Why then, do we insist that power supplies cannot simply be connected in parallel? The flaw in the example lies entirely in the use of the word *identical*.

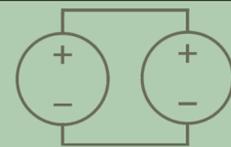
This idealistic load sharing example falls apart **very quickly** if we begin to consider that not all power supplies produced in a production environment are identical. Even with very tight tolerances on load regulation and on voltage set point, the slightest differences between the two paralleled supplies can cause extreme amounts of current to flow from one supply into the other.

Many power supply datasheets will specify the devices voltage set-point accuracy. This figure represents the acceptable range within which the output voltage can be tuned to on the production floor. A very common figure for setpoint accuracy is somewhere around $\pm 2\%$. This means the nominal voltage of our 24V supplies in the example above can actually be anywhere between 23.52V and 24.48V. Now assume that the same circuit in the example above was built with two *real*, non-identical power supplies, and assume that one is found to have a nominal set point of 23.75V and the other is found to have a nominal set point of 24.25V. Further, recognize that voltage regulation is typically specified as worst case, and is also subject to component tolerances, such that R_{out} will not always be exactly the same. Lets assume that supply one has a slightly lower output resistance of 20m Ω and that supply two has a slightly higher output resistance of 24m Ω . Take a closer look at the DC analysis for this real world configuration, and the problem with paralleling becomes immediately clear. Notice that in this very realistic example, there is a 500mV potential across a mere 44m Ω . A simple Ohms law calculation equates that more than 11A flowing from one supply to the other. One interesting note, is that the analysis tells us that the load voltage is largely unchanged (in theory, if the supplies did not shut down) which means we still see 41.27A flowing through the load. It is clear that even the slightest non-idealities can cause one of the two supplies to be severely overdrawn. In fact even just a few mV difference will have drastic effects. The next section introduces a generalized mathematical model for the system at hand, and illustrates just how quickly this phenomena can spiral out of control.



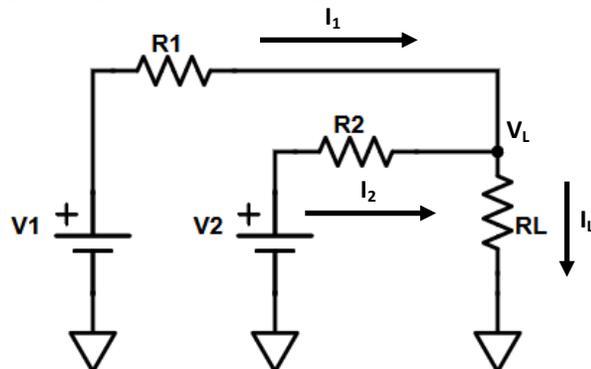


Load Sharing Considerations: Why Two is Not Always Better Than One



A Generalized Mathematical Model for Parallel Power Sources

The figure to the right shows a generalized model of the now familiar parallel power source configuration discussed on the preceding page. V_1 and V_2 are the nominal voltages of two independent power supplies, R_1 and R_2 are their output resistances respectively, and R_L is the load resistance. A DC analysis performed through the use of the superposition principle provides the following mathematical models for the branch currents and node voltages in the model:



$$I_1 = \frac{V_1}{R_1 + \frac{R_2 R_L}{R_2 + R_L}} - \frac{R_1 V_2}{(R_1 + R_L) \left(R_2 + \frac{R_1 R_L}{R_1 + R_L} \right)}$$

$$I_2 = \frac{V_2}{R_2 + \frac{R_1 R_L}{R_1 + R_L}} - \frac{R_1 V_1}{(R_2 + R_L) \left(R_1 + \frac{R_2 R_L}{R_2 + R_L} \right)}$$

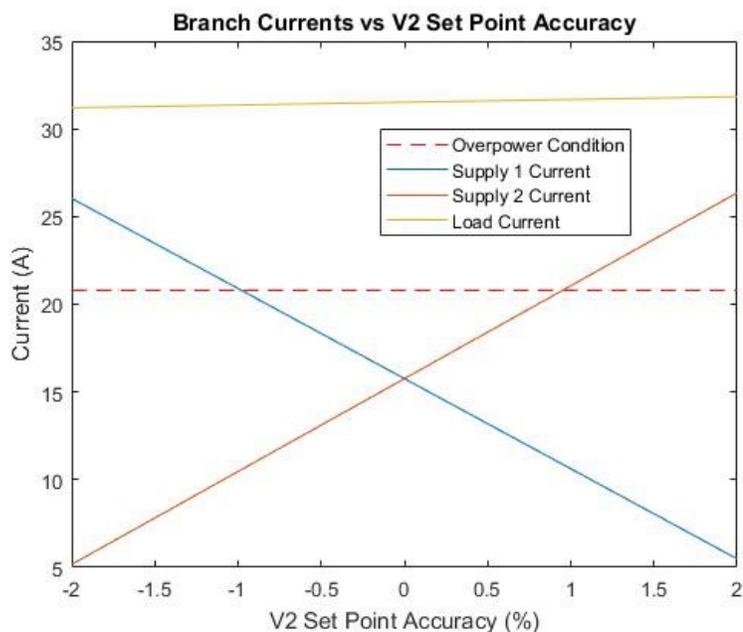
$$I_L = \frac{V_1}{R_1 + \frac{R_2 R_L}{R_2 + R_L}} - \frac{R_1 V_2}{(R_1 + R_L) \left(R_2 + \frac{R_1 R_L}{R_1 + R_L} \right)} + \frac{V_2}{R_2 + \frac{R_1 R_L}{R_1 + R_L}} - \frac{R_1 V_1}{(R_2 + R_L) \left(R_1 + \frac{R_2 R_L}{R_2 + R_L} \right)}$$

$$V_L = \frac{R_1 R_L V_2}{(R_1 + R_L) \left(R_2 + \frac{R_1 R_L}{R_1 + R_L} \right)} + \frac{R_2 R_L V_1}{(R_2 + R_L) \left(R_1 + \frac{R_2 R_L}{R_2 + R_L} \right)}$$

Using this model, and continuing in part from the example on the first page, we can graphically examine the impact for the non-idealities in question. Lets first start by assuming that each supply has exactly the same output resistance of 23mΩ. Lets also assume that the design was constructed with some reasonable headroom for both of the two supplies, with $R_L = 750m\Omega$, such that in an ideal world, each 500W supply would be operating at approximately 75% of its rated output. Say that V_1 happens to have been produced at the exact nominal voltage of 24VDC (a set point accuracy of ± 0%). Examine the plot below to see how the current balance between the two supplies changes quickly as the nominal value of the second supply is varied between a set point accuracy deviation of -2% and 2%.

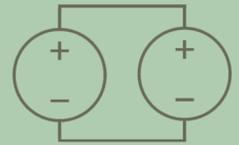
Notice how even with one ideally produced power supply, matching output resistances, and 25% headroom designed in for each supply, it only takes a set point deviation of ±1% to cause an over current event in one of the power sources. Notice also that the load current stays theoretically relatively constant in a parallel supply configuration.

In examining the branch current equations above, one will also notice that for some combinations of R_1 , R_2 , R_L , V_1 , and V_2 either I_1 or I_2 can actually go negative. This is indicative of current flowing directly into the output of one of the power supplies. Some topologies would be tolerant of such a condition, but many others would not. This will be discussed in further detail in later section detailing commonly referenced load sharing strategies.





Load Sharing Considerations: Why Two is Not Always Better Than One



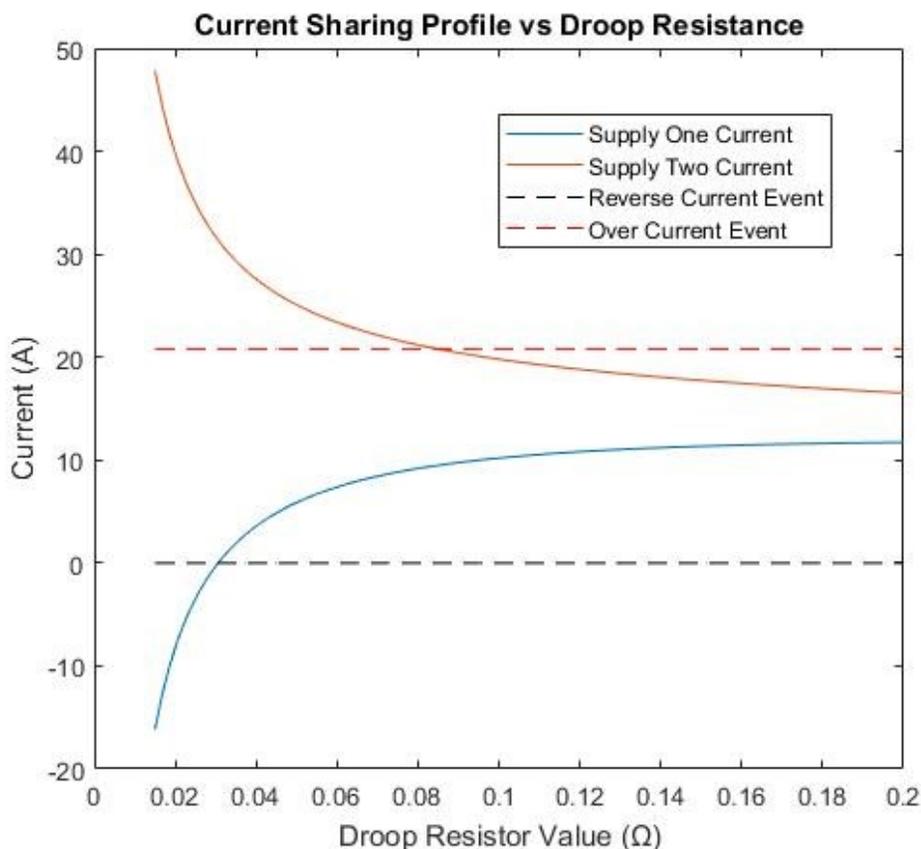
A Simple, but Limited Solution

In examining the mathematical model in the previous section, one will notice that the magnitudes of R_1 and R_2 play a significant role in countering the effects of a voltage differential between V_1 and V_2 . The larger the output resistances get, the lesser the effects are of set point discrepancies. This does not support the effort of implementing no-thought-needed load sharing between power supplies, as most designers strive to minimize the magnitude of load regulation, not maximize it. A high quality power source will have as small a voltage regulation as it possibly can, and smaller voltage regulation means smaller output resistance. This brings to life the concept of droop sharing.

Droop sharing is by no means a new idea, but it has an extremely limited range of applicability. Almost always, the cons will outweigh the pros, especially when dealing with higher power applications, like the example system that has been referenced throughout this article. The basic idea is to intentionally degrade the load regulation of each power supply, by inserting a resistor in series with the output of each supply, such that $R_1 = R_{out} + R_{droop}$. One can rather easily identify a resistor value that ensures that no two supplies, built to specification, will be overdrawn in a simple load sharing configuration.

Returning to our ongoing example, we found that setpoint deviations of magnitude greater than 1% would cause an overcurrent event in the proposed design. This was no good, as it is known that the set-point tolerance on the supplies is $\pm 2\%$. If we assume a worst-case situation whereby two supplies to be paralleled happen to be built at opposite ends of this set point tolerance, 23.52V and 24.48V respectively, we can identify a value for R_{droop} that will degrade the output regulation to such a degree, that the load is always shared.

Examine the plot below, which shows how the load in the previous example would be shared between two such supplies as the value of R_{droop} is increased. As values for voltage regulation are often specified for worst case tolerances, the safest approach is to assume $R_{out}=0$, such that $R_1 = R_2 = R_{droop}$. Notice that to prevent the possibility of a reverse current event on supply two, the droop resistor must be at least 30m Ω , but such a resistor would not prevent the possibility of an overcurrent even on supply one. To prevent either a reverse current *or* overcurrent event from occurring, the droop resistance must be at least 88m Ω . Even this does not facilitate true “sharing” in a worst case configuration, but should at least prevent damage to either supply.

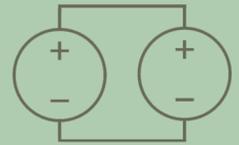


So what is the catch? Consider the voltage drop at the load, and the associated power losses in the droop resistors. If we set R_{droop} equal to 88m Ω , the voltage at the load becomes 22.67V at best, and more than 45W are dissipated as heat, just to prevent the two failure conditions.

To further complicate matters, chances are you're not going through all of this effort just to power up a 750m Ω resistor. Your load is likely not static, and the challenges are largely augmented as the load is allowed to vary. Consider for the sake of this example, that the load is reduced to 5A during certain operational modes. If you re-run the calculations, you will find that the droop resistance needed to prevent a reverse current event has shifted. Examine the load sharing profile on the next page for the same configuration, but with a load resistor of $R_L = 5\Omega$, rather than 750m Ω (~5A load).



Load Sharing Considerations: Why Two is Not Always Better Than One



A Simple, but Limited Solution (Continued)

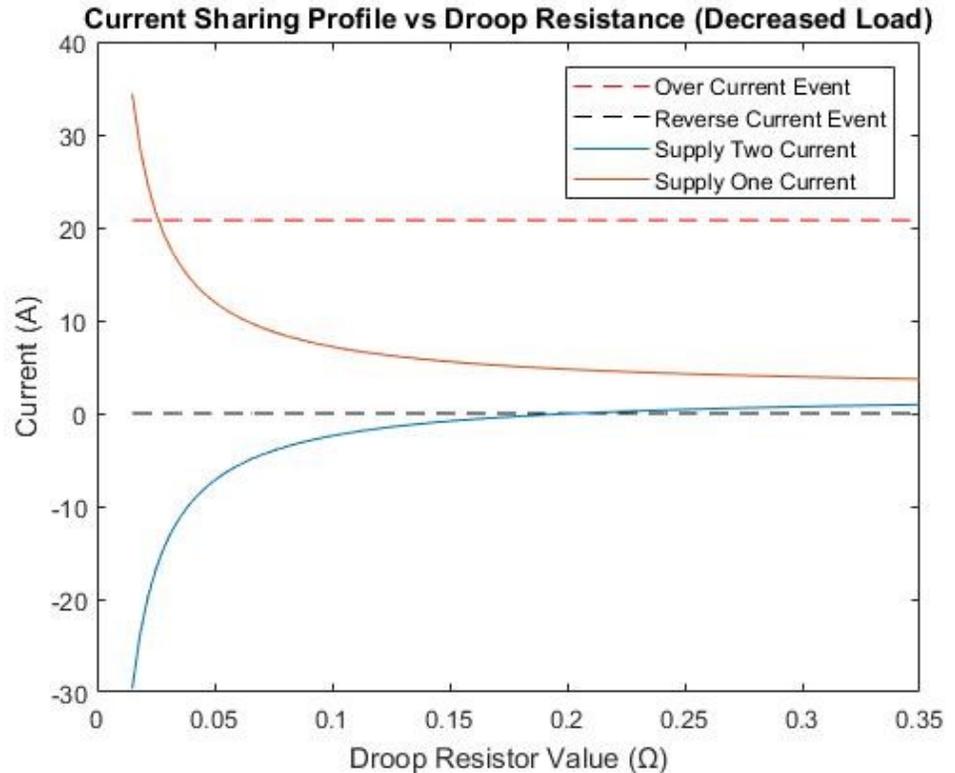
In the reduced load current plot to the right, notice that the required value of R_{droop} to prevent a reverse current event has increased drastically to 205m Ω . Now consider the associated voltage drop and heat dissipation when the operational mode returns to full load (31.25A) condition. Now an absolutely unmanageable 83.48W are dissipated as heat during full load operation, leaving your load with a 21.114V DC supply, and the 25% headroom that was designed in has been reduced to less than 20%.

The negative effects of droop sharing are obviously profound, but these examples have acted to identify the parameters to reflect on when considering droop sharing. We have shown how the effectiveness of droop sharing is affected by overall system power, voltage set point accuracy, load regulation, and variance in load during operation. If one is considering a droop sharing strategy, they should be certain that:

- the set point accuracy of their chosen supply is very narrow
- the load is relatively static during all operational modes of the device
- overall currents within the system are relatively low.

If all three of these are the case for your application, then droop sharing may actually work for you. An expanded droop sharing topology implements a series diode in each supply line in addition to the droop resistor to eliminate the threat of reverse current events, essentially widening the acceptable load variance range of the system, but bear in mind that these diodes contribute to losses and the generation of heat.

Some semiconductor manufacturers have designed external load sharing circuits that replace the droop resistors with MOSFETs. These circuits monitor the current through each parallel branch and adjust the MOSFET gate voltage to balance the current contributions from each supply. Such devices greatly decrease inefficiencies, and widen the range of applicability of droop sharing methods by adding some intelligence to the equation, but they are by no means lossless. These devices still operate on the principle of burning off excess energy in the form of heat.

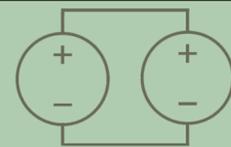


Ideal Load Sharing

All droop-sharing strategies introduce some amount of inefficiency and voltage drop into the power system. The only way to implement true, lossless load sharing is via access to the power supplies' feedback loop. The caveat here, is that the power supply must be designed with the intention of facilitating such a feature and the vast majority of off-the-shelf power supplies simply are not, as a function of limited market demand. The next section will take a closer look at how feedback based load sharing eliminates the many hassles of droop based sharing.

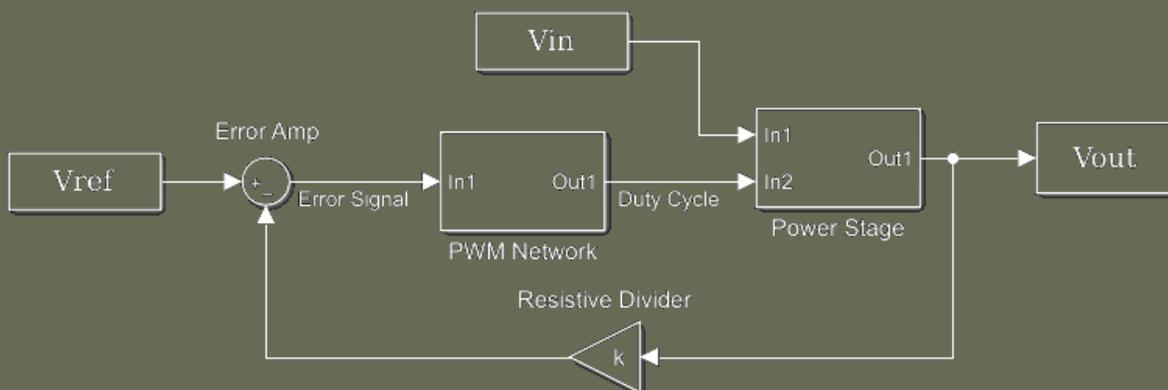


Load Sharing Considerations: Why Two is Not Always Better Than One



Basic Voltage Regulation Principles

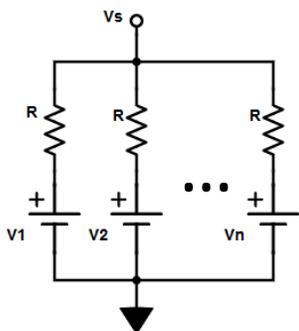
To understand how feedback based load sharing strategies can be implemented, it is helpful to first recapitulate the basic idea of how a power supply actually regulates its output voltage in the first place. The most typical regulation strategy is to use a resistive network to divide the power supplies' output voltage down to a nominal reference value, this reference value is usually between 0.6 and 2.5V and is typically established using a predictable and repeatable voltage drop, such as that which occurs across a forward biased diode. An error amplifier then compares the divided output to the true reference and generates an error signal, which is fed back to the pulse width modulation network, and is used to adjust the duty cycle of the main power switch, and subsequently the output voltage of the regulator. A systems level model of such a regulation network is shown in the figure below.



The basic voltage regulation network diagram illuminates the key to ideal load sharing. Notice that the power supplies output is really only a function of two primary signals, the input voltage, and the duty cycle of the control signal generated by the PWM network. If we consider the fact that in almost any load-sharing configuration, the conversion engines will all be sourcing their power from the same AC line, the output voltage regulation, as it would pertain to a load-sharing application, is really only a function of the duty cycle of the control signal sent to the power stage from the PWM network. The problem with this basic feedback loop, however, is that the error signal that is used to generate the desired power stage control signal, is itself only a function of the present output, and a reference voltage. That is, the regulation network is operating in an isolated environment with no "knowledge" of the world outside. To force a number of parallel power supplies to share a load current, the error signal must also contain a component proportional to the difference between a single devices own output current, and the output currents of other devices within the power system. Enter, the share bus.

The Share Bus

A share bus is a single node that is available to all individual power supplies within a parallel network that facilitates a sharing method known as Single-Wire-Sharing. The share bus carries a signal that is proportional the instantaneous algebraic average of the currents generated by each supply in a parallel network. Such a signal is incredibly simple to generate. Consider the diagram at left, where n independent voltage sources are connected through a resistor to a common node. If we assume that each resistor is exactly the same value, then R falls out of the nodal equations, and V_s can be described by the equation given below to the right of the averager diagram.

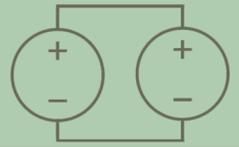


$$V_s = \sum_{k=1}^n \frac{V_k}{n}$$

Now suppose that each of these voltage sources ($V_1 \sim V_n$) is proportional to the instantaneous current through the output of one of a network of power supplies ($I_{PS1} \sim I_{PSn}$). This too would be easy to accomplish through the use of a small sense resistor in series with the main output rails, or even a Hall sensor. This simple architecture creates a bus, common to all parallel supplies in the power system which carries a voltage proportional to the instantaneous algebraic average of the currents generated by each supply. If, on each individual power supply, this share bus voltage is compared to the voltage generated by its own current sense element, the resulting error signal is indicative of that supplies relative contribution to the overall load current. This provides the individual supply with "sight" into the world around it, which we now know is the missing piece of the puzzle.

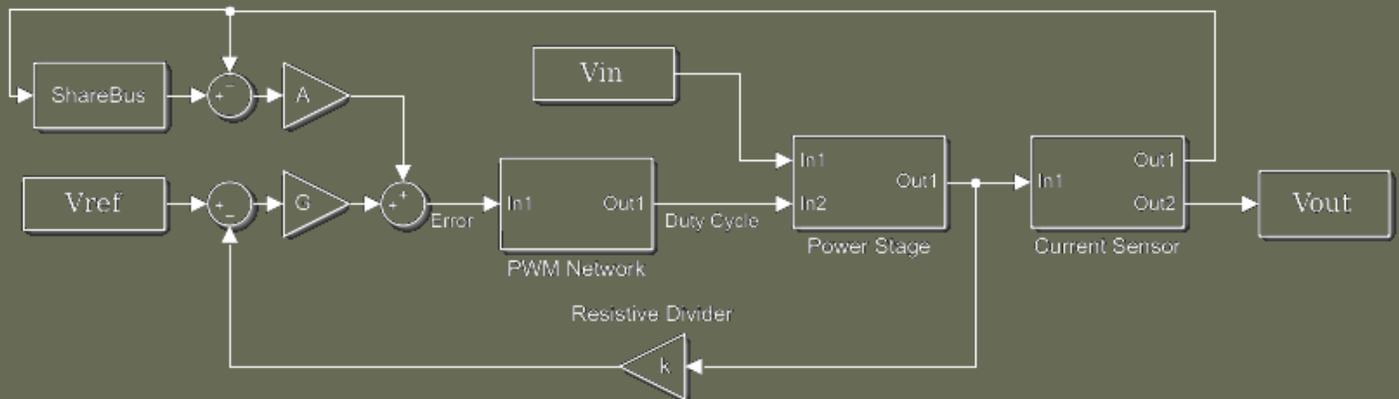


Load Sharing Considerations: Why Two is Not Always Better Than One



Single-Wiring –Sharing

Now that we have determined what signals must be introduced to a power supplies feedback network to facilitate proper load sharing, and have demonstrated simple tactics for generating those signals, lets put all of the pieces together, and revisit the voltage regulation network, this time with some minor modifications.



In this redesigned output regulation network, the error signal driven into the pulse width modulation network is a weighted sum of both the output voltage error, relative to a known reference voltage **and** the output current error, relative to the average current supplied by each supply in the power system. Walking through the signal paths, one can see that if the individual power supplies output current is greater than the average current supplied by other power sources in the system, the duty cycle will be driven down, subsequently reducing the energy transferred by this individual power source, and balancing current contributions of all the power supplies connected to the share bus. The original voltage regulation feedback loop is still intact, but is now weighted by gain G . Gain stages G and A are in place to ensure that the contributions to the error signal by the two independent error sources are appropriately balanced for optimum regulation performance, as the transfer function of the current sense elements is not likely to be naturally dimensionally equivalent to the output of the voltage error amplifier.

Why Two is Not Always Better Than One

In some applications, load sharing provides a number of benefits such as improved efficiency via reduction of current in individual switching semiconductors, and increased reliability via redundancy. To truly take advantage of either of these two benefits, one must use a power supply with a feedback network specifically designed to accommodate additional load sharing error signals such as a share bus. Other load sharing strategies that do not require special feedback signals can be applied to any power source through the use of some external circuitry, but as we have shown, these solutions have extremely limited ranges of practical applicability and often can cause more harm than good.

In the majority of applications, there is no substitute for a single, appropriately selected power solution. Single supply solutions have higher overall power densities, significantly reduced cost, and require less design resources to successfully implement. Consider the fact that undesirable power supply “side-effects” such as radiated and conducted emissions and leakage currents are all additive phenomena. Here at Power Partners, our sales and technical support staff are experts in the field of appropriately matching one of our products to the unique requirements of our customers applications. If the application still demands a current-share solution, our applications engineering team is here to help determine the feasibility of such a solution.

The load sharing strategies discussed in this article by no means encompass all of the possible load sharing solutions for all applications. Innumerable solutions have been designed or theorized, each tailored for use in one of the many different power conversion topologies. While every problem has a solution, it is important to recognize that when it comes to parallel power supplies, one plus one almost never equals two!

Dylan Howes

Applications Engineer

(978)567-9600 x230

