July 2021

Power Supply Integrity for Onboard Optical Modules Using Remote Voltage Sensing

N/A

Power Supply Integrity for Onboard Optical Modules Using Remote Voltage Sensing

ABSTRACT

Optical modules on a PCB are typically physically spread out and share the same onboard power supply. To ensure the quality of the transmitted optical signal, the optical modules have tight requirements on the integrity of their power supply, including a maximum permissible DC voltage drop from the power supply output to the module. Modules that are close to the power supply experience smaller voltage drops and better power-supply integrity. Modules that are further away can experience voltage drops large enough to violate specification. This disclosure describes a voltage sensing-line design that leverages the operating configurations of the optical-module array on a PCB to achieve distributed power integrity using a low number of sensing lines.

KEYWORDS

- Voltage sensing
- Local sensing
- Remote sensing
- Power supply
- Optical module
- Wavelength division multiplexing (WDM)
- Serializer/Deserializer (SerDes)
- SerDes link
- Distributed sensing
- High performance computing (HPC)

BACKGROUND

Large-scale, high-performance computing (HPC) units comprise hundreds to thousands of application-specific integrated circuits (ASICs) working in a supercomputer configuration and are interconnected with high-speed interfaces, e.g., NRZ, PAM4 SerDes links, etc. Two physically proximate chips can be connected using onboard routing or copper cables; chips that are far away from each other are optimally connected using optical cables, as copper is too lossy for long-reach connections.

To drive signals through long-reach optical cables, electrical Serializer/Deserializer (SerDes) signals are routed out from each ASIC to onboard optical modules, transformed into optical signals, and transmitted to the receiver optical module at the other end. Depending on the interconnect technology and the number of onboard ASICs, a baseboard can host 30-100 such optical modules.

Onboard optical modules are typically physically spread out along the edge of the printed circuit board (PCB) to facilitate connection and cable management. They consume substantial power and share the same onboard power supply. To ensure the quality of the transmitted optical signal, the optical modules have tight requirements on the integrity of their power supply, including a maximum permissible DC voltage drop (the IR drop) from the power supply output to the module. Modules that are close to the power supply experience smaller voltage drops and better power-supply integrity; modules that are further away can experience voltage drops large enough to violate specification.

To ensure power-supply integrity to a module, a voltage-sensing line can be incorporated at the site of the load (remote sensing at a single location). The voltage sensed at the load is used to compensate for the DC voltage drop by elevating the output voltage at the power supply

output by an equivalent amount. However, in the present case of multiple optical modules, the load locations are distributed. A voltage-sensing line placed close to the power supply results in under-compensation at far-away modules. Conversely, a voltage-sensing line placed at far-away modules results in over-compensation for close-by modules. The situation is further complicated by the fact that the modules may be on and off independently and at random times. If the sensing line is connected to a module that is off (or not populated), compensation won't be effective, and the other modules receive incorrect voltages.

Local sensing, which monitors the voltage at the output of the power supply, does not compensate for the voltage drop at a distant module. Remote sensing at *all* locations followed by shorting the sensing lines works only when the number of modules is relatively small. Also, when *all* locations are remotely sensed, the number of sensing lines required is so large that routing them becomes a problem.

DESCRIPTION

This disclosure describes a sensing-line design that leverages the operating configurations of the optical-module array to achieve distributed power integrity using a low number of sensing lines.

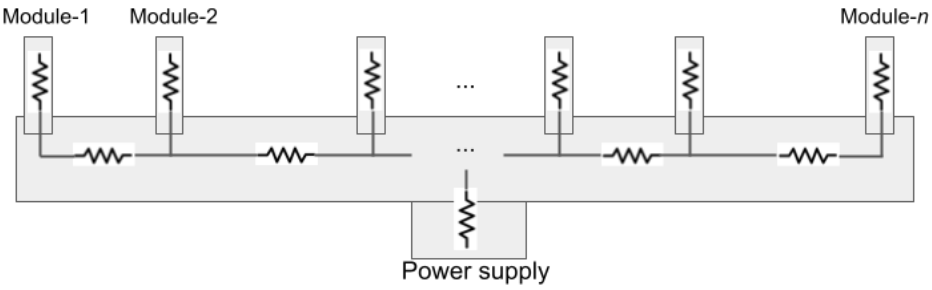

Fig. 1: Resistance network extraction

Per the techniques, illustrated in Fig. 1, the DC resistance network of the power

supply and the *n* modules are extracted into a tree-shaped netlist. The resistance of the main power delivery plane shared by all the modules is extracted in segments. The resistance of the branches of the planes leading to each individual module is calculated. By abstracting the power-supply network into a resistance-only circuit (ignoring the inductive and the capacitive components of the network, which have minimal effect on the DC voltage drop), circuit simulations are greatly accelerated.

	Module-1	**Module-2**	**Module-3**	**...**	**Module -(n-1)**	**Module-n**
Config-1	✓		✓			✓
Config-2	✓	✓			✓	
...						
Config-m	✓	✓				✓

Table 1: Table of configurations

Table 1 illustrates a table of configurations for the optical modules. For example, in config-1, modules 1, 3, and *n* are active; in config-2, modules 1, 2, and $n-1$ are active; etc. In contrast to a network switchboard, in an HPC baseboard, as illustrated in Table 1, not all of the optical modules are turned on or populated. Depending on the location of the HPC baseboard in the overall interconnected network, a particular board has only a limited number of operating configurations. In practice, it is not atypical for a 16-module baseboard with $2^{16} \approx 65{,}000$ theoretically possible configurations to have only 16 operating configurations.

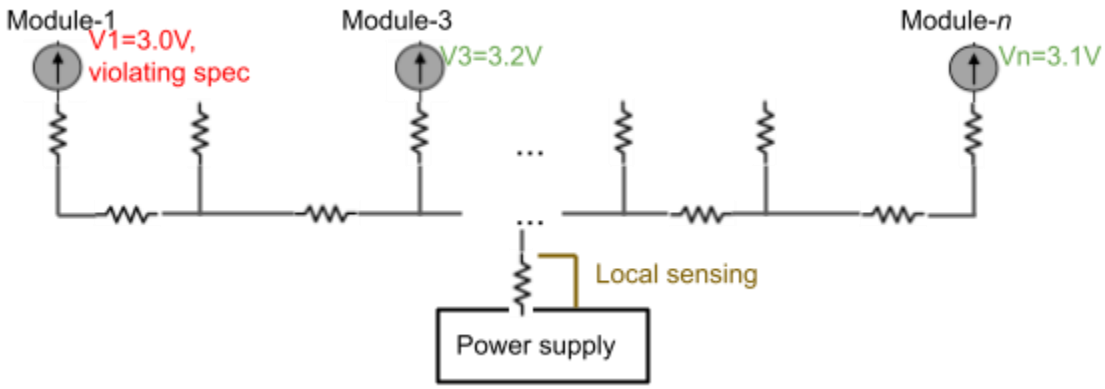

Fig. 2: Simulation of power drop under local sensing

For each board configuration listed in the table of configurations, the populated modules are modeled as DC current sources, as illustrated in Fig. 2. The resulting network, comprising the DC current sources, the power supply, and the extracted resistance network, is simulated using a circuit simulator, e.g., any SPICE-like simulator. The voltage drop at each populated module is determined.

The first round of simulations assumes local sensing, e.g., sensing close to the power supply. The module voltages are compared to the specification. If all module voltages meet their specifications, then local sensing is deemed sufficient. Otherwise, remote sensing schemes are investigated, as explained below. In the example of Fig. 2, module-1 violates the specification, while modules 3 and *n* meet the specification. Therefore local sensing does not work for the example of Fig. 2.

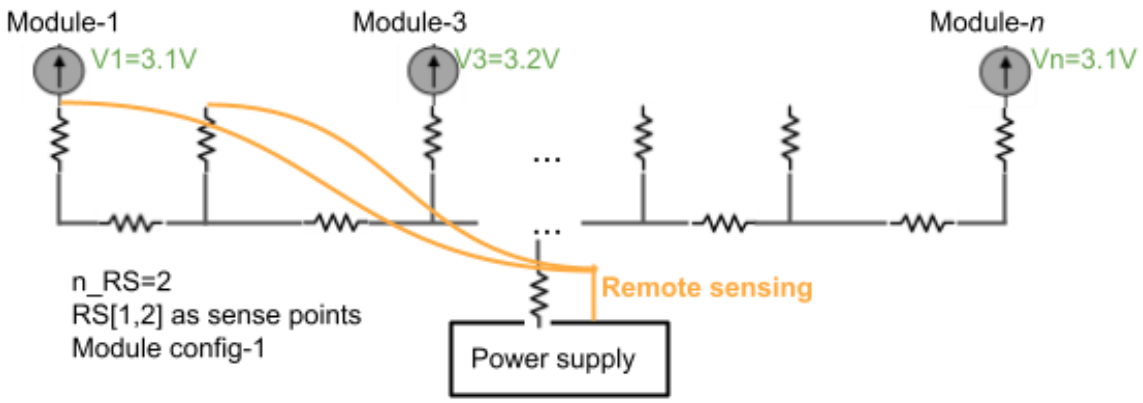

Fig. 3: Simulation of power drop under remote sensing

Starting from one remote sense (RS) line tapped at module-*i* (*i* is the module number, going from 1 to *n*), simulations are run to determine the voltages of all populated modules for every operating configuration. If no satisfying solution is found, one more RS line is added, as illustrated in Fig. 3 (orange lines), and all possible connection points at (module-*i*, module-*j*), $i, j \in [1, n]$ are simulated. With more than one RS line, the power supply sees an averaged out voltage feedback from the lines, and the voltage compensation process is handled in the simulation. In the example of Fig. 3, all populated modules meet their specification; thus, at least for the configuration of Fig. 3, two RS lines work (the two-RS line solution of Fig. 3 still has to be verified to work for all configurations in the table of configurations).

If no satisfying solution is found using two RS lines, the number of RS lines is incremented by one, and the simulation procedure is repeated until a set of $k \leq n$ RS lines and sensing points are found that enable all modules to meet specification under all operating configurations. The simulation being of a resistive network, the simulation time of each iteration is extremely fast. Sweeping the entire solution space is easily feasible. Fig. 4 illustrates pseudocode for determining an optimal RS connection scheme.

```
for n_RS from 1 to n:
     for RS[1..n_RS] as subset of [1..n]:
          Simulate V[1..n] for all m configurations, record
Vmin and Vmax among all populated modules
          if all Vmin and Vmax are within spec:
               Remote sensing RS[1..n_RS] is a viable
solution; break
          else:
               continue to next subset
```

Fig. 4: Pseudocode for determining an optimal RS scheme

In this manner, the described techniques determine the optimal RS connection scheme, including the least number of RS lines and the modules to connect them to. Although the optimal RS connection scheme is heavily dependent on the number of operating configurations, the resistive network, and the current load conditions of the modules, the techniques find an optimal RS connection scheme that satisfies the voltage drop requirements given the table of configurations. The techniques leverage the fact that an HPC baseboard has a relatively small number of operating configurations to efficiently search the space of RS connections. The RS connection scheme determined by the techniques is a single solution that fits all operating configurations: the DC voltage drops of all populated modules in all operating configurations meet their specifications.

CONCLUSION

This disclosure describes a voltage sensing-line design that leverages the operating configurations of the optical-module array on a PCB to achieve distributed power integrity using a low number of sensing lines.