A Study of Load-line Effect on Power Supplies of Digital Networking Processors

Tingdong Zhou, and James S. Golab
Digital Networking Group, NXP Semiconductor, Austin, TX
Email: tingdong.zhou@nxp.com, jim.golab@nxp.com

***Abstract*—Digital networking processor power supplies typically have voltage sense points at die level so that the die level voltages are regulated and kept at constant DC levels. In this paper, we studied the load-line effect on the power supplies by simulations and lab measurements. We found that the total dissipated power for digital networking processors can be reduced by the introduction of the load-line effect in the regulation loop. The thermal design point of the system could be relaxed by introducing load-line operation without compromising the performance of these processors.**

Keywords-Digital Networking Processors, load-line, Power Noise, Thermal Design Point.

I. INTRODUCTION

As the input/output (I/O) bandwidth and throughput of modern networking processors go up, more and more power is consumed by these processors. Due to switching activities designed into the processors, power noise becomes more and more an issue for them. Figure 1. is the current signature of

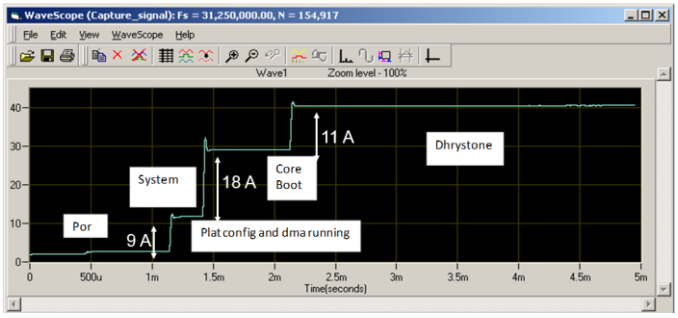

Fig.1, The current signature during power on, boot and application running on a high-end network processor.

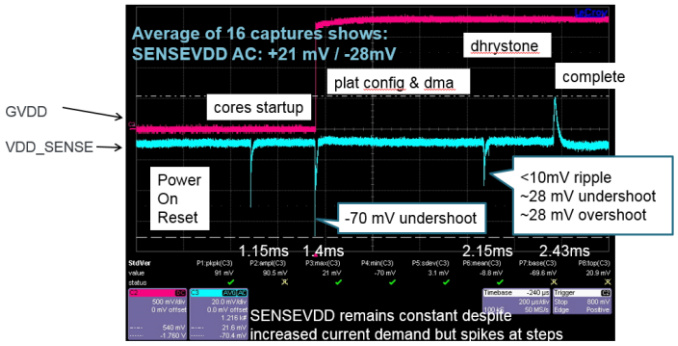

Fig. 2, VDD power noise measured during boot process.

our high end processor from power on to system running, core booting, and then an application program (in this case, the Dhrystone [1] benchmark) running. There are significant switching currents for these stages. It has also been observed that there is significant power noise corresponding to these current steps, as shown in figure 2. It also shows that at the die level, we wanted to keep the DC voltage constant – this is a legacy approach that has been taken for most the network processor chips' regulation scheme. The voltage and ground sense pins are routed from die level, e.g., Back-End-of-Line (BEOL) metals to the package Ball-Grid-Array (BGA) interface. The sense lines are routed by customers from package BGA to a board level filter and then the filter voltage sense lines are routed to the point of load power supply module to control the regulation loop. In this case, the power supply to the processors has a very flat V-I curve [1, 2]. This is a typical design for low power network chips, and is a different approach as compared to a power supply scheme with loadline that has been proposed [3, 4] for high power server or personal computer processor.

In this paper, we would like to investigate the impact of the loadline effect on the power supply of network processors. In section II, we will first study the impact using simulation with a power supply model and package board power delivery network models. In section III, we will examine hardware data collected using a power scheme that includes loadline resistance at the point of load power supply module. Section IV concludes the paper.

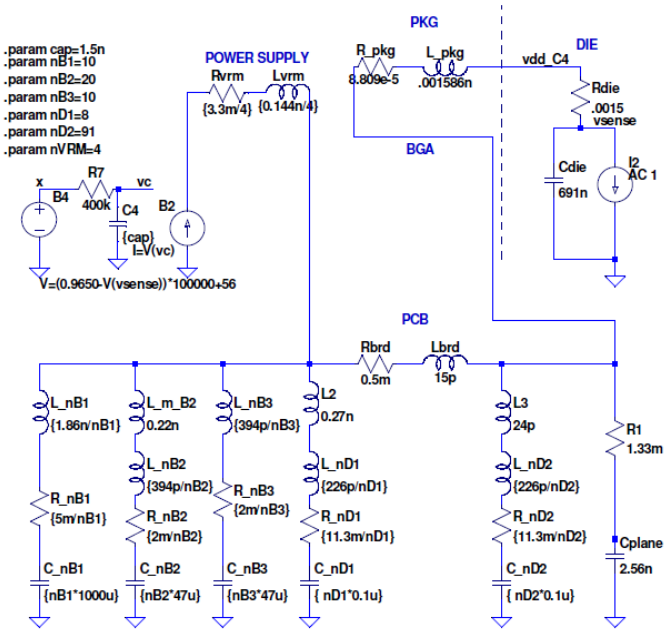

Fig. 3, The schematic for load-line embedded power delivery simulation.

II. POWER SUPPLY SIMULATION WITH LOAD-LINE

For the exact current signature shown in Fig. 1, we ran a simulation to understand how a load-line would affect the power noise waveform. Fig. 3 shows the schematic used in the simulation. The power supply is modeled as a voltage controlled current source. The power delivery network is modeled with lumped-element resistors, inductors, and capacitors as discussed in [1, 2]. Typically, for the power supplies of digital networking processors, there is no active load-line feature used in the point of load power supply. Therefore, we have a very flat V-I current for the power supply. Board level decoupling capacitors, including both the low frequency and high frequency capacitors, are modeled with lumped-element circuit models. The package is modeled with a simple R/L circuit. The die is modeled with a piecewise linear (PWL) current signature, with on-die resistance and capacitance.

For the legacy design, the voltage sense is routed to the die level. This corresponds to Fig.3, where the voltage sense node voltage is used as the set point for the regulator power supply. Another scenario is that we use the BGA node voltage as the set point so that the Rdie and R_pkg are built-in load-line resistances. Fig. 4 shows the die level voltage waveforms. The

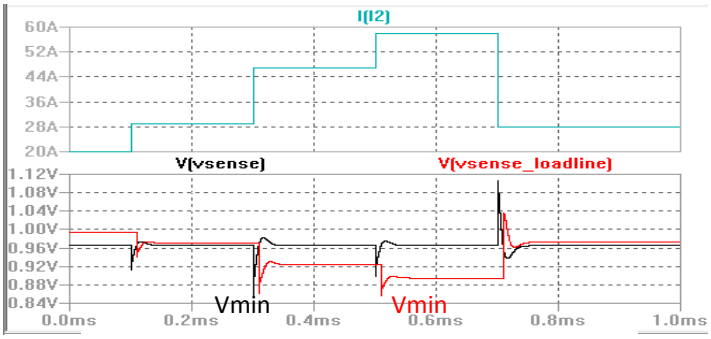

Fig. 4, Die level voltage waveforms during boot process without and with load-line.

DC level is adjusted for the load-line case so that the minimum voltage (Vmin) for two cases are matched at 857mv. It is well known that the performance is related to the Vmin, due to the voltage droop being much slower than the time scale of critical timing paths in modern processors. By introducing the load-line, we can see that the peak to peak noise is reduced from 249mv to 180mv. This allows tighter power tolerance, as AC

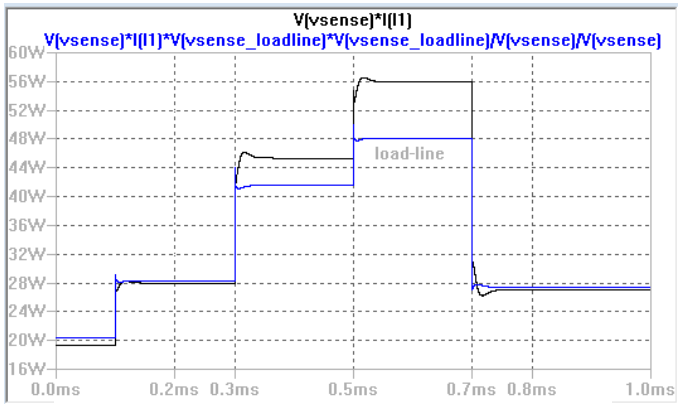

Fig. 5, Power consumption during boot.

power is proportional to V^2. Fig. 5 plots the power consumed during the boot process with and without load-line. Load-line enables power saving while meeting the same Vmin.

III. LOAD-LINE HARDWARE EXPERIMENTS

In order to fully understand the impact of the load-line and to confirm what the simulation shows, we have done hardware experiments with T2080 networking processor modules. For the experiment, we implemented the load-line at the point of load regulation module. The cases studied are 1) No load-line, 2) 1mΩ, and 3) 1.8mΩ load-line resistance. The maximum frequency (Fmax) grader application code, CPATTX8 code, was run to find the Fmax of T2080 parts while the part is cooled to a constant die temperature of 105C through thermal diode sensing. Fig. 6 shows the typical collected data where the

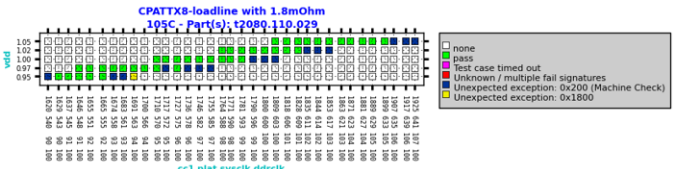

Fig. 6, Fmax collection for T2080 at different voltage levels.

x-axis represents the clock frequencies and the y-axis shows

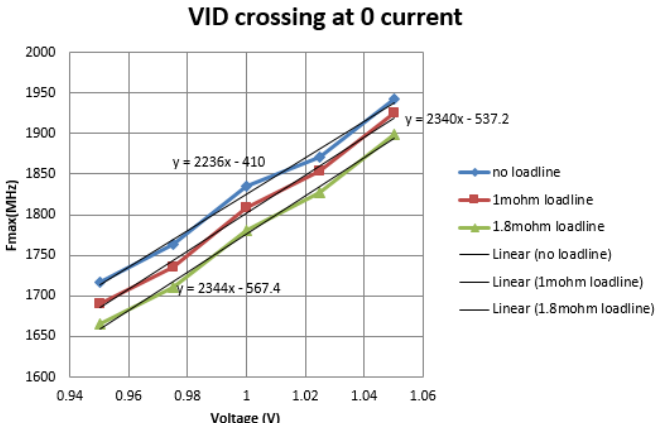

Fig. 7, Fmax data with three different load-line settings.

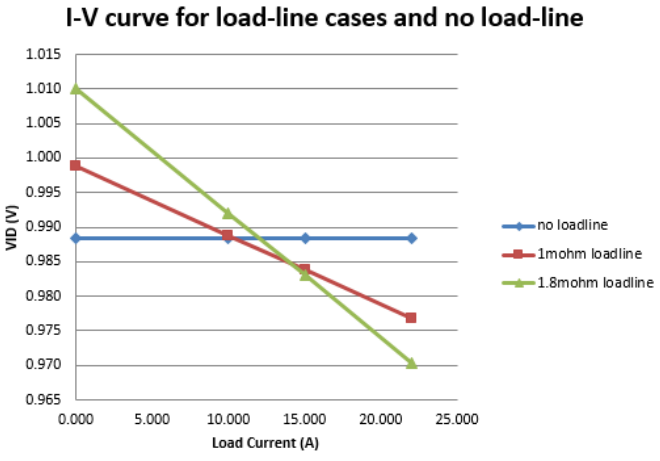

Fig. 8, Different load-line regulation settings to yield the same Fmax performance.

the different voltage levels that part has been tested to with the CPATTX8 code. The part is run with a faster and faster clock frequency until "Machine Check" failures occur, meaning that it has reached Fmax. Fig. 7 plots the Fmax data we collected. Linear regressions are applied to find the relation of Fmax and voltage applied for three load-line settings.

Based on the linear regression data, we find that the three different V-I regulation schemes (Fig. 8) yield the same Fmax performance (1800MHz). When the load current is low, the voltage at the die will be higher for steep load-line case (1.8mΩ). However, as the load current increases, the die voltage gets smaller and eventually lower than the no load-line case. This enables power saving during the time that system needs to deliver peak performance. For system cooling design, we have a smaller thermal design point (the product of voltage V and current I). With the same thermal solution, power regulation with load-line would yield lower junction temperature and avoid over-heating as peak performance is delivered. For the 1mΩ case, the VID is reduced by 11mV at 22A load current comparing to no load-line case. That corresponds to about a 2% power reduction. For 1.8mΩ case, the VID is reduced by 18mV at 22A, compared to the no load-line case, or about a 4% power saving.

We also recorded voltage regulator output current as shown in Fig. 9 for the three load-line cases when the processors were

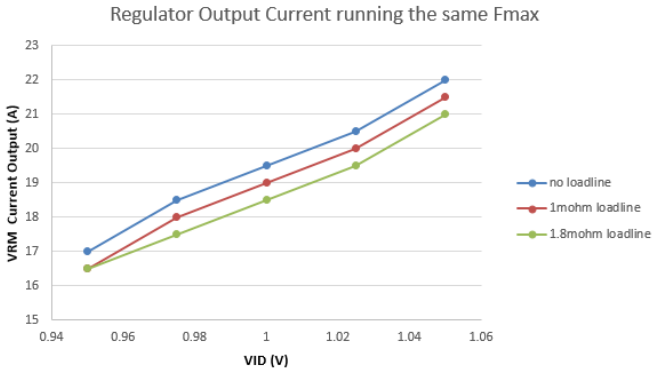

Fig. 9, Voltage regulator output current as processor running the same core frequency using different load-line settings.

running the same core frequencies. Just to note that in this experiment, the VID, meaning open output voltage, is set to be constant. The output current gets smaller for steeper V-I curves.

IV. Conclusions

We conducted a study of voltage regulation with a load-line scheme. Simulation data indicates that load-line would reduce the thermal design point, and the peak-to-peak noise amplitude. We also collected hardware Fmax data under different VID using Fmax grader code, and the data indicated that load-line would help reduce the die voltage level while still enabling the same Fmax performance. It also confirms that the thermal design point could be optimized by the introduction of load-line into the regulation loop.

V. Acknowledgment

The authors would like to thank our colleague Timothy Fike for collecting the hardware data in the lab, and Kelly Folts for the strong management support of this work.

References

[1] C. Corley, "Power supply network design for 3% voltage margin," in Proc. Freescale Technology Forum-NET-F0038, June 2012.

[2] C. Corley, Mohit Kedia, "Design to tight power supply requirements," in Proc. FTF-NET-F0036, APR. 2014.

[3] Intel Corporation, "Voltage Regulator Module (VRM) and Enterprise Voltage Regulator-Down (EVRD) 11.1-Design guidelines," Sep. 2009.

[4] G. Chinn, S. Desai, E. DiStefano, K. Ravichandran, S. Thakkar, "Mobile PC platforms enabled with Intel Centrino mobile technology," Intel Technology Journal, Vol. 7, Issue 2, May 2003.