

Integrated Regulation for Energy-Efficient Digital Circuits

Elad Alon, *Member, IEEE*, and Mark Horowitz, *Fellow, IEEE*

Abstract—Despite their use in analog or mixed-signal applications, the high power overheads of traditional linear regulators (both series and shunt) have precluded their successful adoption in regulating the supply of energy-efficient digital circuits. In this paper, we show that linear regulation can in fact reduce the effective supply impedance of digital circuits without increasing their total power dissipation. Achieving this goal requires minimizing the static power dissipation of the regulator, leading to a push-pull topology (similar to the regulators demonstrated by Wu and Sanders, 2001, Poon *et al.*, 1999, and Intersil, 1998) with comparator-based feedback and a switched source-follower output stage. Measured results from a regulator implemented in a 65 nm SOI test-chip verify that by using these techniques, regulation reduces the effective supply noise by $\sim 30\%$ while also enabling a slight decrease (1.4%) in total power dissipation.

Index Terms—Digital integrated circuits, power supply distribution, regulators.

I. INTRODUCTION

WITH THE roughly 1 m Ω impedance required by today's 1 V, 100 A microprocessors, concerns about the effects of supply noise on the robustness, performance, and power dissipation of modern chips have been steadily increasing. Supply noise has always been a major concern for sensitive analog or mixed-signal circuits, but it can have significant impacts on digital circuits as well. For example, noise on the supply voltage directly modulates the delay of digital gates. Thus, supply variations force increased buffering to avoid hold-time violations on fast paths, and decreased nominal delay (i.e., extra timing margin) on slow paths in order to meet performance. Furthermore, overshoots in the supply voltage can accelerate long-term degradation in the transistor characteristics.

Regulation is often used to combat the issue of supply noise in harsh mixed-signal environments, particularly to reduce the power supply noise seen by sensitive circuits such as VCOs [1], [2] or high-speed transceivers [3]. More recently, in order to alleviate the previously mentioned robustness concerns, techniques such as switched decoupling capacitors proposed by Ang [4], active decoupling capacitors proposed by Gu [5], and a noise canceller proposed by Nakamura [26] were explored for

active on-chip noise reduction in digital chips. However, these approaches are often suitable only for resonant load current profiles [4], [5] or isolated current steps due to turn-on transients [26]. Even more importantly, they have so far all increased the total power dissipation of the chip. Thus, despite the potential robustness benefits of such regulation, its associated power overhead has been a significant barrier to widespread adoption in today's power-limited chips.

To improve chip power by reducing the DC losses of the supply distribution (and to enable dynamic voltage scaling), works such as [6] and [7] have explored integrated switching DC-DC converters. However, at the > 1 W/mm² power density required in modern processors, these converters have not achieved the efficiencies needed to actually reduce total chip power. Specifically, Lee *et al.* in [6] projected a peak efficiency of less than 70% with a power density of approximately 70 mW/mm² for a fully integrated buck converter in a 0.13 μ m technology. Furthermore, the high frequency output impedance of such a switching converter is limited by its switching frequency—which for optimal efficiency will be relatively low [6] compared to what would be required to counter broadband load current noise. Thus, in this paper we will focus on integrated linear regulator topologies that efficiently reduce the supply impedance. Even if technology changes such as the addition of magnetic materials [8] enable improved switching converter efficiencies, such an integrated linear regulator could further improve the overall efficiency of these converters by reducing their output impedance.

Since a linear regulator must spend power to reduce noise on the supply voltage, at first glance it may seem impossible for the regulator not to increase total power. However, variations in the supply voltage (caused by variations in the load current) are themselves a direct cause of additional power dissipation. This is due to the fact that in synchronous digital systems, the operating frequency is set by the minimum average supply voltage over a clock cycle. Hence, to maintain the same performance, noise-induced drops in the minimum supply voltage necessitate an increase in the nominal voltage. Therefore, if a regulator consumes less power countering noise than the power saved by reducing the nominal voltage, it will not significantly add to the chip's total power—and in fact may even reduce it.

To achieve this power-neutral objective, the static power dissipation of the regulator (both in the output stage and in the feedback circuitry) must be minimized. Unfortunately, as we will show in further detail in Section II, single-supply linear regulators (both series and shunt) intrinsically burn significant static power in their output stages—making them unsuitable for this application. However, by adding a second, higher-than-nominal

Manuscript received December 14, 2007; revised February 25, 2008. Published July 23, 2008 (projected). This work was supported in part by the MARCO Focus Center for Circuit and System Solutions (C2S2, www.c2s2.org).

E. Alon is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: elad@berkeley.edu).

M. Horowitz is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA.

Digital Object Identifier 10.1109/JSSC.2008.925403

supply voltage, a push-pull shunt regulator (similar to the “active clamps” developed for board-level power distribution applications [9]–[11]) with zero nominal output current can be built.

By spending output power only when necessary to counter noise, this regulator has the potential to achieve no net increase in the chip’s total power consumption. Since such a regulator must make use of a second power supply, in Section III we examine some of the important considerations in the design of the additional shunt supply network. We then describe in Section IV circuit techniques for ensuring that the regulator burns minimal static power. Finally, in Section V we present measured results from a regulator implemented in a 65 nm SOI technology that meets the goal of reducing noise without increasing total chip power.

II. SUPPLY NOISE AND DIGITAL CIRCUIT EFFICIENCY

Since it sets the amount of power a regulator would be allowed to spend to reduce noise on the supply, we must first evaluate the power loss caused by supply variations. We will calculate this power loss by assuming that as the supply noise increases, the digital circuit’s performance—and hence its minimum supply voltage—must be kept constant. Therefore, as the noise on the supply increases, the nominal supply voltage must increase. An additional factor to consider is that since the dynamic current (and hence current variation) of a digital circuit is proportional to $C_{sw} \cdot f_{clk} \cdot V_{dd}$ (where C_{sw} is effective switched capacitance per cycle and f_{clk} is the operating frequency), the effective noise current increases along with the nominal supply voltage.

For this analysis, we define $V_{dd,min}$ as the minimum supply voltage required to achieve a given target operating frequency,¹ $V_{dd,nom}$ as the nominal supply voltage, and $I_{dd,nom}$ as the nominal load current. We will next assume that the variations in load current are zero-mean and equal in positive and negative magnitudes (i.e., the peak increase in load current is equal to the peak decrease in load current). Making the peak-to-peak load current variations equal to $k_n I_{dd,nom}$, the minimum supply voltage will be

$$V_{dd,min} = V_{dd,nom} - \frac{1}{2} k_n I_{dd,nom} \|Z_{dist}\| \quad (1)$$

where $\|Z_{dist}\|$ is the effective impedance of the supply network (e.g., for a worst-case sinusoid, $\|Z_{dist}\|$ is the supply network’s maximum impedance). Defining $R_{load} = V_{dd,nom}/I_{dd,nom}$, we can further simplify $V_{dd,min}$ to

$$\begin{aligned} V_{dd,min} &= \left(1 - \frac{1}{2} k_n \frac{\|Z_{dist}\|}{R_{load}}\right) V_{dd,nom} \\ &= \left(1 - \frac{1}{2} \frac{k_n}{LR_{dist}}\right) V_{dd,nom} \end{aligned} \quad (2)$$

where we have defined the load rejection of the supply network as $LR_{dist} = R_{load}/\|Z_{dist}\|$. With these definitions, the peak-to-peak noise on the supply voltage divided by the nominal supply voltage (i.e., $\Delta V_{dd,p2p}/V_{dd}$) is exactly k_n/LR_{dist} .

With the simplifying assumption that the chip’s power consumption is purely dynamic (i.e., no leakage current), the power

¹For a synchronous digital circuit, $V_{dd,min}$ is set by the worst-case supply voltage averaged over one clock cycle.




Fig. 1. Series regulator and example waveforms on input supply (V_{dd}) and regulated supply (V_{reg}), highlighting the need to margin the dropout voltage of the regulator to include the variations on V_{dd} .

consumption of the chip is proportional to $V_{dd,nom}^2$. In this case, the efficiency of the chip η (relative to the case with no noise) is set by $V_{dd,min}^2/V_{dd,nom}^2$, and hence

$$\eta = \left(1 - \frac{1}{2} \frac{k_n}{LR_{dist}}\right)^2. \quad (3)$$

A. Efficiency With a Series Regulator

Fig. 1 shows a conceptualized series regulator, where we have drawn the output stage as a variable resistor instead of the transistor that implements this resistor. We have drawn the regulator this way because the major limitation on the efficiency of the regulator is essentially independent of the actual implementation, and is inherent to the series topology itself.

For a series regulator to maintain an output impedance that is decoupled from (and hence can be lower than) the impedance of the input supply V_{dd} , the input supply voltage must remain above the regulated output V_{reg} . The minimum allowable voltage drop from V_{dd} to V_{reg} is typically referred to as the dropout voltage. As Fig. 1 also shows, this dropout voltage must be maintained not only in a DC sense, but must also take into account the dynamic V_{dd} variations.

Unfortunately, a series regulator does not alter (relative to an unregulated chip) the impedance of the input supply, nor does it significantly alter the magnitude of the load current variations. Thus, the voltage margin required to achieve a certain minimum V_{reg} will be just as large as in the unregulated case. Even worse, to keep the same performance, V_{dd} in the regulated system also has to be increased by the regulator’s dropout. Thus, it is clear that a series regulator can not achieve the goal of reducing the net power consumption of the chip.²

Intuitively, it is not very surprising that a series regulator reduces the overall efficiency of a digital chip. The noise on the power supply is set by the supply impedance and by the variations in the load current, and the series regulator by definition adds series resistance to the impedance of the supply network. Although the regulator uses negative feedback to reduce the voltage noise seen by the load, this clearly does not improve the impedance of the input supply network.

Despite making a series regulator unsuitable for supply impedance reduction of digital circuits, the regulator’s series resistance is exactly what allows it to isolate the regulated

²Since the nominal voltage seen by the series regulated digital circuits will be lower than in the unregulated case, the noise current magnitude in the regulated chip will be smaller than in the unregulated chip. However, this reduction in noise current will typically be small; the reduction is only large enough for the series regulated chip’s power dissipation to approach that of an unregulated chip when the noise is very large ($\sim 50\%$ supply variation).




Fig. 2. Shunt regulator and example load current (I_{load}) and shunt current (I_{shunt}) waveforms, highlighting the need for the shunt current source to statically draw current proportional to the maximum noise current ($I_{noise,max}$).

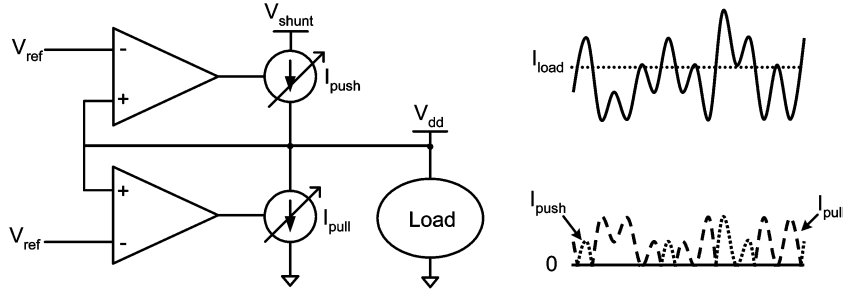


Fig. 3. Push-pull shunt regulator topology and example current waveforms.

supply from externally generated noise on the input supply. It is this isolation that makes the series topology effective in regulating the supply of sensitive analog or mixed-signal components.

B. Efficiency With a Shunt Regulator

Since the underlying cause of noise on the power supply of digital circuits is variation in the load current itself, a shunt regulator that directly reduces these current variations is much more suitable to this application than a series regulator. Fig. 2 shows such a shunt regulator, where we have once again depicted the output stage with a conceptual element (a current source) in order to highlight the fact that the efficiency limitation is inherent to the regulator's topology.

As also shown in Fig. 2, the instantaneous load current draw can be larger or smaller than the average current—in other words, the current variations are bidirectional. However, without an additional power supply or energy storage element, the shunt regulator's current source must be unidirectional, and can only pull current out of the power supply. Therefore, the shunt regulator must statically pull an output current that is proportional to the worst-case excess load current, so that it can then dynamically reduce this current when necessary.

Clearly, statically dissipating current proportional to the worst-case load current variations is very inefficient—especially if the load current variations have large peaks but on average are relatively small. The fundamental cause of this need to statically dissipate output current in a single-supply shunt regulator is that it cannot deliver energy when the load transiently demands more power. Thus, the regulator is forced to burn excess static power in its output stage, and then modulates this power dissipation to compensate for the load variations.

C. Efficiency With a Push-Pull Shunt Regulator

To enable the regulator to transiently deliver energy without static power overhead, we can introduce a second, higher-than-nominal supply voltage and build a push-pull shunt regulator [9]–[11] (Fig. 3). With the use of a second power supply, the nominal output current of the regulator can be set to zero, and the appropriate current source turned on only when necessary. Therefore, as also shown in Fig. 3, the average current through each of the regulator's output current sources is set by the average of one side of the noise current. Clearly, if the peaks of the noise current are significantly larger than the average current deviation, the average current flowing through a push-pull shunt will be significantly lower than that of a single-supply shunt.

In order to highlight the potential of this topology to meet the efficiency requirements of digital regulation, we will briefly describe a model for the overall efficiency of a digital chip using this regulator. Defining LR_{reg} as the additional rejection contributed by the regulator, the minimum regulated supply voltage $V_{dd,min,reg}$ will be

$$V_{dd,min,reg} = \left(1 - \frac{1}{2} \frac{k_n}{LR_{dist} LR_{reg}}\right) V_{dd,nom}. \quad (4)$$

To achieve this additional rejection, the push and pull sides of the regulator will each supply an average current of

$$I_{reg,avg} = \frac{1}{2} \frac{k_{Idiv}}{k_{pk/avg}} (1 - LR_{reg}^{-1}) k_n I_{dd,nom} \quad (5)$$

where k_{Idiv} accounts for the percentage of current the shunt regulator must handle for a specific noise excitation,³ and $k_{pk/avg}$

³The percentage of the noise current that the regulator's output stage must handle will vary with the spectral content of the noise, and must be calculated for a given noise current excitation. With the assumption of an ideal amplifier, for sinusoidal noise current $k_{Idiv} = 1$. For broadband (white), random (but bounded) current noise, k_{Idiv} is typically $\sim 1/3$.

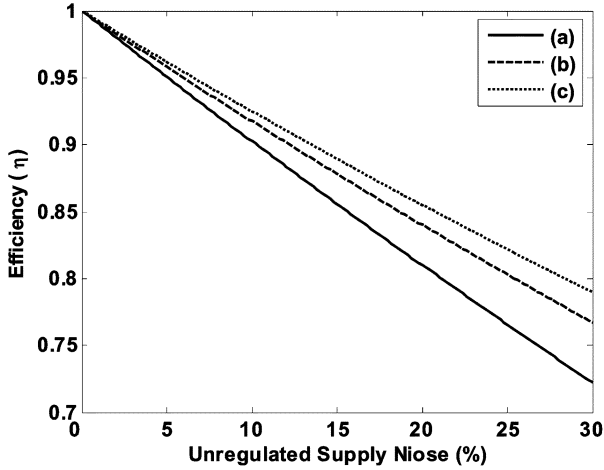


Fig. 4. Digital chip efficiency versus unregulated supply noise for (a) an unregulated chip, (b) a push-pull shunt regulated chip with $LR_{reg} = 1.5$, and (c) a push-pull shunt regulated chip with $LR_{reg} = 2$. In all cases, $k_{pk/avg} = 3$, $k_{div} = 1/3$, and $LR_{dist} = 5$.

is the peak current deviation divided by the mean current deviation. For simplicity, if we assume that the second power supply $V_{shunt} = 2 * V_{dd,nom}$, the total power dissipated by pushing and pulling this current is

$$P_{I_{reg}} = \left[\frac{k_{div}}{k_{pk/avg}} (1 - LR_{reg}^{-1}) k_n \right] \frac{V_{dd,nom}^2}{R_{load}}. \quad (6)$$

If we assume ideal feedback amplifiers with infinite bandwidth and no power consumption, the total power dissipated by the chip will be

$$P_{chip,reg} = \frac{V_{dd,nom}^2}{R_{load}} + P_{I_{reg}} \quad (7)$$

and therefore the overall efficiency is

$$\eta = \frac{V_{dd,min-reg}^2 / R_{load}}{P_{chip,reg}} = \frac{\left(1 - \frac{1}{2} \frac{k_n}{(LR_{dist} \cdot LR_{reg})}\right)^2}{1 + \left(k_{div}/k_{pk/avg}\right) (1 - LR_{reg}^{-1}) k_n}. \quad (8)$$

Using this simplified model and (1), Fig. 4 shows a comparison between the overall efficiency of regulated and unregulated digital chips. Even at a relatively moderate $k_{pk/avg}$ of 3, the regulator achieves a higher overall efficiency (i.e., a reduction in total power) while reducing the effective supply impedance.

III. SHUNT SUPPLY NETWORK DESIGN

Since modern chips already dedicate most of the available pad, pin, and metal resources to the supply distribution network, integrating a push-pull shunt regulator requires allocating these resources between the main power supply (V_{dd}) and the shunt supply (V_{shunt}). Clearly, taking resources away from the main power supply increases its loss (since its series resistance will increase), but allocating too few resources to the shunt supply makes it too lossy to be effective.

Ideally, V_{shunt} would be set to the minimum voltage required for the push-side of the shunt regulator to operate properly. In other words, as fewer metal resources are allocated to V_{shunt} ,




Fig. 5. Simplified model of the supply impedances in a push-pull shunt regulated system, highlighting the on-chip and off-chip return paths for I_{push} .

the nominal value of V_{shunt} would be raised in order to compensate for the increased IR drop. In this case, the resources for the two supplies should be allocated according to their relative contributions to the total loss, leading to

$$p_{V_{shunt}} = \frac{I_{push,rms}}{(I_{push,rms} + I_{load,rms})} \quad (9)$$

$$p_{V_{dd}} = \frac{I_{load,rms}}{(I_{push,rms} + I_{load,rms})}$$

where $p_{V_{shunt}}$ is the percentage of the total resources dedicated to the shunt supply, and $p_{V_{dd}}$ is the percentage dedicated to the main power supply.⁴ As long as the regulator delivers only transient currents with reasonable peak-to-average ratio, this allocation results in low $p_{V_{shunt}}$ (typically $< 5\%$), and hence the impact on the resistive losses of the main supply will be relatively minor.

In addition to properly allocating resources to minimize the resistive losses of the supply network, steps must be taken to ensure that I_{push} returns only through the on-chip path (highlighted in Fig. 5). If V_{shunt} is supplied by an external voltage source, the push side return current will flow almost entirely through the on-chip path when $\|Z_{V_{shunt,dist}}(j\omega)\|$ is significantly larger than $1/\omega C_{V_{shunt}}$ at the frequencies of interest. Thus, part of V_{dd} 's decoupling capacitance must be used for V_{shunt} . Of course, removing decoupling capacitance from V_{dd} increases the noise on the main supply. Therefore, to keep the capacitance reallocated to V_{shunt} low ($\sim 10\%$ of the total available), it is typically necessary to intentionally increase $\|Z_{V_{shunt,dist}}(j\omega)\|$ as well. To avoid sacrificing the DC resistive losses of the shunt supply, this increase in $\|Z_{V_{shunt,dist}}(j\omega)\|$ is best achieved by increasing the shunt supply's series inductance.⁵

Even though the average current through V_{shunt} will be relatively small, the transient currents are set by the full load current variations. Since the impedance of the shunt supply will be relatively high, these large transient currents make V_{shunt}

⁴If for simplicity V_{shunt} is set to a fixed voltage (independent of the resource allocation), then V_{shunt} should be allocated just enough metal tracks to maintain the minimum required voltage at the regulator.

⁵One interesting possible alternative that keeps the return current flowing on the die is to generate V_{shunt} directly from V_{dd} by using an integrated boost converter or charge pump. However, any current that flows from V_{shunt} would actually originate from on-chip V_{dd} . Thus, in order to avoid interfering with the operation of the push-pull shunt regulator at the frequencies of interest, the series input impedance of the converter would have to be significantly higher than the impedance in parallel with the converter's output (i.e., $1/j\omega C_{V_{shunt}}$).




Fig. 6. Push-pull shunt regulator with local reference generation, comparator-based feedback, and switched source-follower output stage to minimize static power consumption.

very noisy. Therefore, to guarantee that the regulator's push-side output device does not fall out of saturation, appropriate margin (typically 200–300 mV) must be added to V_{shunt} 's nominal value.

IV. REGULATOR DESIGN FOR MINIMAL STATIC POWER

In Fig. 3 we presented a conceptual schematic of the push-pull shunt regulator topology to highlight its potential for the regulation of digital chips. However, as we will describe in this section, several additional steps must be taken to ensure that the regulator burns minimal static power in both its output stages and its feedback circuitry.

A. Minimizing Static Output Power

Given the low resistance of the main supply distribution network, attempting to counter slowly varying IR drops by statically pushing or pulling current with the regulator would clearly be a waste of power. In fact, in modern chips it is particularly important that the regulator operate only on voltage transients since functional units or even the entire chip may be shut down (or placed into lower performance modes) by the operating system for extended periods of time [14].

To avoid burning static current due to such long-term load current variations, the regulator's reference is generated by locally RC filtering V_{dd} [9], [11], as shown in Fig. 6. Thus, any supply variations well below the bandwidth of the RC filter (which is typically set in the low-MHz range) will be ignored by the regulator. Since IR drop varies spatially within the die, a local filter for each regulator location⁶ (versus one global filter) is necessary.

Although this RC reference prevents the regulator from spending output current on static IR drops, the output devices could still burn significant current due to their nominal biasing.

⁶The impact of high-frequency current noise on the supply network is localized, and therefore a full chip would likely have many distributed regulator sites. This paper focuses on a single isolated regulator site, which can be used to project the impact of many such sites on an entire chip. An important topic for further investigation is the density of regulator sites required to avoid site-to-site interactions without sacrificing noise performance.

Minimizing this bias current rules out class-A operation of the regulator's power devices, necessitating instead a mode of operation where the quiescent current of the output stage is set by the power device's bias voltage, IV characteristics, and gate voltage swing.

In a regulator with linear, voltage-mode feedback amplifiers, the power device gate voltage swing is not an independent design parameter—it is set by the magnitude of the supply voltage noise and by the gain of the amplifier. The gain of the amplifier is itself tied to the effective transconductance of the output stage, as described in the Appendix. Therefore, it is not straightforward to guarantee gate voltage swings large enough to mitigate any errors (systematic or random) in the biasing of the output device with such a design—potentially leading to large static output current and poor efficiency. The active clamp described by Wu in [9] avoided some of these issues by feeding the output of a transconductance feedback amplifier into a cascade of multiplying current mirrors. This essentially creates an amplifier that is nonlinear in voltage, but linear in overall transconductance. However, in this application the degradation in feedback bandwidth caused by these current mirrors is highly undesirable.

In order to completely eliminate the regulator's static output power in a robust manner, as shown in Fig. 6, we employed comparator-based feedback. To avoid the limit cycles that can potentially arise in such a comparator-based system [15], the thresholds of the comparators are offset to create a symmetric dead-band.⁷

In addition, the offsets of both comparators and the preamp are tunable to compensate for mismatch and zero-center the dead-band. While the use of comparators increases the feedback gain required over a linear scheme (since the comparator outputs must swing full-rail independent of the input voltage), as shown in Fig. 7(b), it allows the use of CMOS inverters as gain stages in the comparator and as buffers to drive the output stages. As discussed by Hazucha *et al.* in [17], the efficiency of CMOS buffers in modern processes allows the power consumption of this approach to be the same or even less than a comparable design with class-A feedback amplifiers.

As shown in Fig. 7, the tunable offset and dead-band can be implemented with digitally programmable (5-bit with a thermometer decoder) current sources tied to the outputs of the pre-amp and comparator. For the test-chip described in Section V, this implementation was chosen for its simplicity, although other potentially more efficient schemes such as a pair of skewed-sizing differential pairs with programmable tail currents [18] could be adopted. In order to allow both the push side and pull side comparators to easily accept a single preamp output (as opposed to two preamps with potentially different offsets, gains, and bandwidths), common-mode feedback was

⁷In addition to limit cycle considerations, it is important to note that it can be beneficial from an efficiency standpoint to tune the dead-band along with the magnitude of the supply noise. In essence, since digital performance is set by the minimum voltage, the dead-band can be used to reduce the amount of current the regulator consumes without significantly impacting the minimum voltage. The exact relationship between the dead-band and noise magnitudes is strongly dependent upon the current noise distribution, but in general the dead-band width should be set at $\sim 30\%$ of the σ of the noise. Further analysis of appropriate dead-band settings can be found in [16, Appendix D].

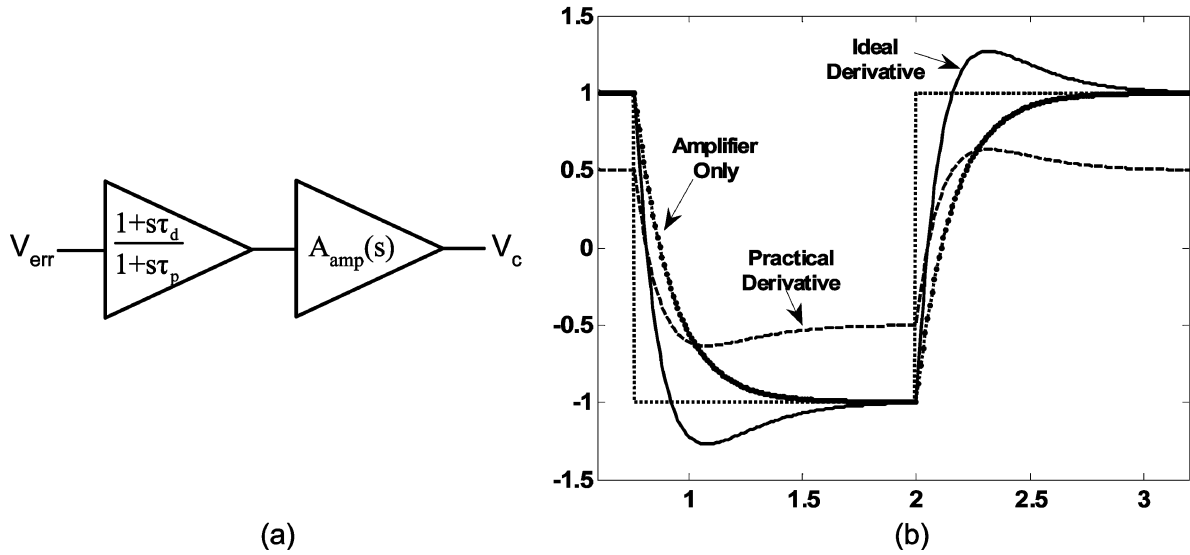


Fig. 8. (a) Additive derivative term at the input of a linear feedback amplifier. (b) Example output waveforms for an amplifier in isolation, an amplifier with an ideal derivative filter, and an amplifier with a practical (DC gain reduced) derivative filter.




Fig. 9. Preamp input stage modified to implement the derivative filter, where the filter time constants are programmable.

low-frequency gain. Therefore, with a properly chosen derivative time constant (typically, $1/\tau_d$ set to roughly half of the feedback bandwidth) the dominant impact of this scheme is to reduce the effective delay of the feedback action.⁸

In order to implement a derivative filter for the regulator described here, the input stage of the preamp can be modified as shown in Fig. 9. The derivative behavior is accomplished by subtracting an RC -filtered version of $V_{dd}(V_{hf})$ from V_{dd} . Thus, the effective input to each differential pair segment whose reference is connected to V_{hf} is

$$V_{dd} - V_{hf} = \left(1 - \frac{1}{1 + s\tau_p}\right) V_{dd} = \frac{s\tau_p}{1 + s\tau_p} V_{dd} \quad (10)$$

where τ_p is the time-constant of the RC filter—which is significantly smaller than that of the filter generating V_{ref} . By summing the V_{hf} segments' outputs with the V_{ref} segments' outputs

⁸The use of a $1 + s\tau_d$ feedback filter with comparator-based feedback is very similar to the sliding mode controllers used in switching DC-to-DC converters [22]–[24], although for efficiency the regulator described here rarely operates inside of the sliding region [16].

(which provide low frequency gain), the transfer function implemented by this circuit is

$$\begin{aligned} H_d(s) &= (1 - k_d) + k_d \frac{s\tau_p}{1 + s\tau_p} \\ &= (1 - k_d) \frac{1 + s\tau_p/(1 - k_d)}{1 + s\tau_p} \end{aligned} \quad (11)$$

where k_d is the number of segments tied to V_{hf} divided by the total number of segments.

Applying the positive phase shift of the filter at the appropriate frequency requires the derivative time constant (which in this implementation is $\tau_p/(1 - k_d)$) to be set as a relative fraction of the effective bandwidth of the feedback path. In addition, the bandwidth of the filter generating V_{hf} should roughly match the feedback bandwidth. For example, if the filter bandwidth is significantly greater than the feedback bandwidth, k_d would need to be increased to set τ_d . This would further lower the DC feedback gain and require additional gain margin from the comparators. Alternatively, if τ_p is too large it will reduce the filter's positive phase shift at the feedback bandwidth, which is where this phase shift has the largest impact on effective delay.

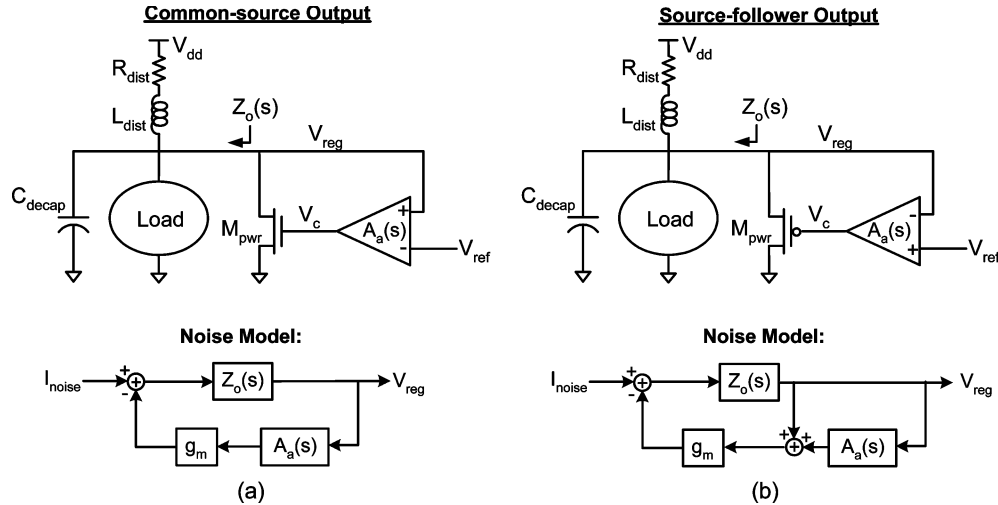


Fig. 10. Shunt regulators with (a) common-source and (b) source-follower output stages and their associated linearized noise models.

Since k_d can be relatively well-controlled through transistor matching and the derivative time constant scales with τ_p , the most important matching in this design is between the bandwidth of V_{hf} 's RC filter and the bandwidth of the feedback path. However, the feedback path bandwidth is set by several transistor parameters which can vary independently of the resistance (typically polysilicon) and capacitance (either metal-to-metal or gate) which set τ_p . Thus, τ_p should be made programmable in order to compensate for these relative variations.

Another very effective method to reduce the power required from the feedback circuitry is to employ a source-follower-based output stage. To understand the reasons for this, it is instructive to examine the small-signal noise models for shunt regulators (with linear feedback amplifiers) based on common-source and source-follower output stages (Fig. 10).

As described in the Appendix, most of the loop gain in these regulators will typically be contributed by the g_m of the output stage. In fact, the voltage gain of the amplifier itself is often one or less. Thus, in the common-source regulator shown in Fig. 10(a), the amplifier essentially acts only as a power-consuming wire routing the error signal to M_{pwr} 's gate. In contrast, with the source-follower design of Fig. 10(b), the feedback signal is directly applied to M_{pwr} through a high-bandwidth wire that has no power dissipation. For a given output impedance, this intrinsic feedback from the source-follower significantly reduces the requirements on the amplifier's GBW, hence reducing the amplifier's power. At low GBW, one can eliminate the amplifier entirely [25], while at higher gains it can decrease the required GBW by over 40%.

Although a regulator for digital circuits will most likely be based on comparators (for the previously described reasons), such a regulator similarly benefits from a source-follower-based output stage.⁹ For this application, the two principal requirements on the design of a source-follower-based output stage are

⁹As with a series regulator such as [25], the push-side source-follower output stage requires a higher voltage (versus a common-source) at M_{pwr} 's gate in order to accommodate for its V_{th} . For simplicity, in this implementation V_{shunt} was raised by V_{th} , but the regulator's efficiency could be further improved with a separate feedback path supply.

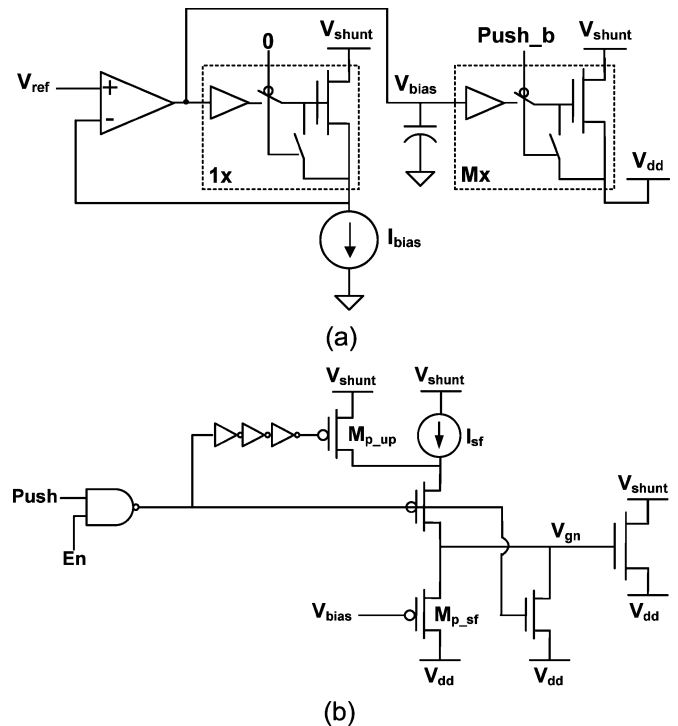


Fig. 11. (a) Replica-bias generation for a source-follower output driver. (b) Output driver implementation, highlighting the switched source-follower buffer isolating V_{bias} from V_{gn} . This figure shows the push side—the pull side implementation is once again complementary.

achieving low turn-on delay (since feedback delay is analogous to bandwidth in the linear loop), and maintaining predictable output current with a noisy V_{shunt} .

To achieve these goals, as shown in Fig. 11(a), we can employ a replica-biasing scheme where a bias voltage is driven onto the gate of the power device (through a buffer). For efficiency, the power consumption of the amplifier in the replica-bias loop should be minimized, making its output impedance relatively high. This means that directly switching V_{bias} onto the gate of the power device (which is nominally discharged to V_{dd}) would create disturbances on V_{bias} . While adding capacitance to V_{bias} can reduce the disturbance, this would exacerbate the

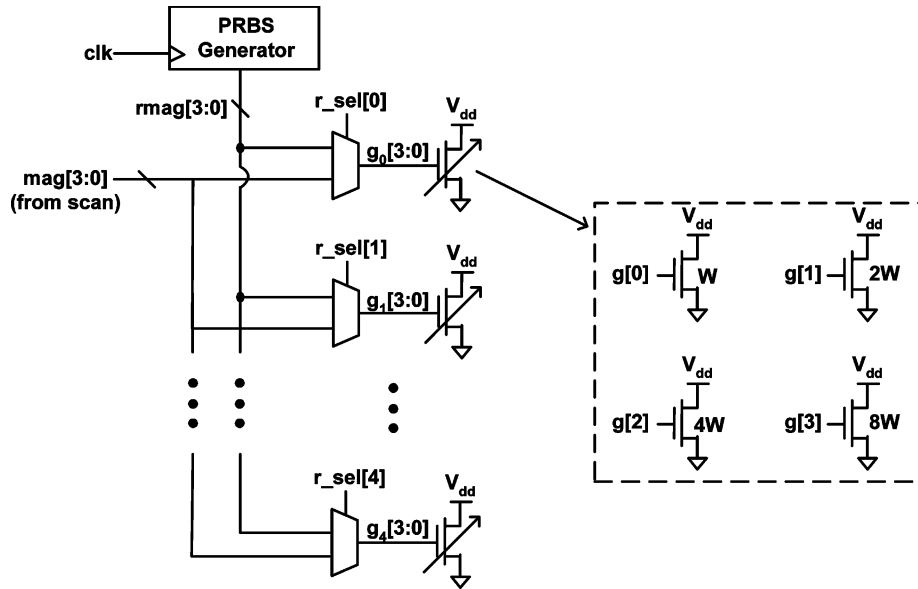


Fig. 12. On-chip programmable broadband current noise generator.




Fig. 13. On-chip performance monitor circuit.

dependence of I_{push} on the history of the comparator outputs. Hence, it is advantageous to include a buffer to isolate V_{gn} from the high-impedance V_{bias} . This additional buffer can be implemented as another source follower, as shown in Fig. 11(b). To eliminate its static power, the buffer is turned on only when current is being pushed. Since the buffer current flows into V_{dd} (or out of V_{dd} for the pull side) and only flows when the main output current does, I_{sf} can simply be treated as a portion of the regulator's total output current. However, waiting for I_{sf} to charge V_{gn} up to its final value can cause significant delay, so during the turn-on transition M_{pup} is left on to provide a large transient current. For implementation simplicity, the width and magnitude of the pre-emphasis current pulse may not be well-controlled, leading to some transient imprecision in the output current (preferably an overshoot). Fortunately, as long as the transient is short its effects will be minor, while the improvement in the feedback delay directly impacts the regulator's performance.

V. EXPERIMENTAL VERIFICATION

In order to experimentally verify the benefits of integrated regulation for digital logic, in collaboration with AMD, a push-pull shunt regulator was included in a 65 nm SOI test-chip used for technology bring-up. The die area of the regulator test-site was $\sim 400 \mu\text{m}$ by $400 \mu\text{m}$. The supply distribution network was designed with a methodology similar to the rest of the processor

test-chip, but with the previously described reallocation of resources for an externally generated V_{shunt} . Specifically, 10% of the total decoupling capacitance and $\sim 4\%$ of the on-chip power wires were reallocated from V_{dd} to V_{shunt} .¹⁰ To characterize the impacts of regulation, the test-chip included noise generators and a performance monitor circuit.

The programmable on-chip noise generator shown in Fig. 12 was used to inject different levels of broadband current noise onto the supply. As shown in the figure, the generator has five banks, where each bank consists of a binary-weighted, 4-bit array of nMOS devices tied between V_{dd} and V_{ss} . Each individual bank can have its magnitude set by a static control signal $\text{mag}[3:0]$ (such that the bank draws only DC current), or by the lower 4 bits of a $2^{15}-1$ PRBS generator (such that the bank draws a pseudo-random amount of current each clock cycle). The current noise generated by the PRBS is essentially white up to roughly half of the PRBS generator's clock frequency. The relative magnitude of the current noise is set by the number of banks whose inputs are controlled by the PRBS generator, i.e., the relative current noise can range from roughly 0 to 5.

¹⁰Due to the limited area of the test-site, the C4 bump and package pin allocation could not track the ideal ratios as closely: a single C4 bump was used for V_{shunt} as opposed to four for V_{dd} , and nine package pins were allocated to V_{dd} and one to V_{shunt} . However, from a total IR drop standpoint, these non-optimal ratios of C4 bumps and package pins were compensated by the package and PCB routing: a partial plane was used for V_{dd} while V_{shunt} was routed through a single trace (as well as a discrete external inductor).




Fig. 14. Measured noise reduction provided by the push-pull shunt regulator.

To measure the supply variations as they would impact the performance of the digital circuits themselves, the test-chip also included the performance monitor shown in Fig. 13. The monitor operates by generating a signal that transitions every clock cycle (*cdiv*), and then checking on every cycle that this signal was correctly captured by a flip-flop after N stages (typically 16–24) of fanout-of-four inverters. If on any cycle an error is detected (i.e., $dval \neq tval$), the *err* signal will remain high until the circuit is externally reset. In this manner, the monitor indicates whether the input clock frequency can be maintained at the worst-case supply voltage.

The measurement procedure began with picking a supply voltage $V_{dd,set}$, and with the regulator disabled, setting the noise generator to statically draw roughly half of its maximum current (i.e., no noise). The performance monitor was then used to determine the maximum frequency of operation (f_{op}). Leaving this frequency constant, measurements were then taken at 4 magnitudes of current noise. Throughout this characterization, all voltage and power measurements were performed by a single digital multimeter with a relative accuracy of roughly 0.03%.

At each level of current noise, three different measurements were taken. First, with the regulator still disabled and the supply voltage at $V_{dd,set}$, the average power consumption of the chip (at the current noise setting) was recorded. Since the devices within the noise generator banks may not match perfectly, this measurement provides the baseline power (P_{min}) against which to evaluate the chip's efficiency (for that level of noise). Next, with the regulator still off, the supply voltage was raised to the minimum level needed to operate at f_{op} . This voltage ($V_{dd,unreg}$) and the chip's power consumption (P_{unreg}) were once again recorded. Finally, the regulator was turned on, and its parameters were tuned along with the supply voltage (calling the final voltage $V_{dd,reg}$) to achieve minimum power consumption (P_{reg}) at the same f_{op} . The peak-to-peak voltage noise for each current noise setting was calculated as

$$V_{n,\{unreg,reg\}} = 2(V_{dd,\{unreg,reg\}} - V_{dd,set}) \quad (12)$$

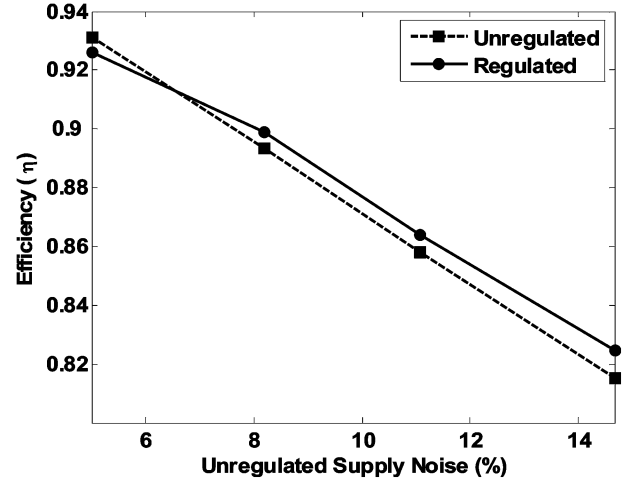


Fig. 15. Measured efficiencies of unregulated and push-pull shunt regulated systems versus magnitude of the supply noise without regulation.

and the efficiency as

$$\eta_{\{unreg,reg\}} = \frac{P_{min}}{P_{\{unreg,reg\}}} \quad (13)$$

Figs. 14 and 15 show the results measured from this chip using this procedure¹¹ with $V_{dd,set} = 1.188$ V and $V_{shunt} = 2.14$ V. At this $V_{dd,set}$, f_{op} was 1.67 GHz and P_{min} (at zero noise) was ~ 105 mW. The static power of the regulator's biasing and feedback circuitry was less than 1.5 mW. Based on the measured supply current and calculated peak-to-peak voltage noise, the effective (broadband) unregulated supply impedance was ~ 869 m Ω . For a 12 mm by 12 mm chip operating at the same power density (i.e., 94.5 W total chip power), this corresponds to an impedance of ~ 966 $\mu\Omega$.

In the range of expected supply noise ($\sim 8\%$ peak-to-peak and above), the regulator reduces the effective noise by $\sim 30\%$,¹² while successfully maintaining no net increase (in fact, a drop of up to $\sim 1.4\%$) in total chip power. Although these results already confirm the applicability of regulation to digital circuits, the f_t of the transistors in this developmental process was significantly lower than that of a production process.

With a higher f_t , the benefits of the regulator would be even larger. Fig. 16(a) shows a simulation of the regulated supply

¹¹Since (12) is based on measurements from the performance monitor, the measured noise voltage should be interpreted as the worst-case average voltage over one cycle. Note that the clock buffering was minimal (three stages) on this test-chip in order to minimize artifacts from supply-induced clock jitter. It should also be noted that these performance monitor-based measurements do not directly verify the operation of the pull side of the regulator. However, the pull-side circuits are complementary to the push-side circuits, and were verified to function before the noise characterization procedure was run by skewing the pull-side comparator until it triggered and caused additional V_{dd} current. Furthermore, the magnitude of the pull current and the pull-side dead-band were adjusted to match those of the push-side to achieve a symmetric response to positive and negative supply voltage variations.

¹²The external power supply used to generate V_{dd} for the test chip had a voltage resolution of ~ 5 mV. Hence, the measurements had to be snapped to this relatively coarse grid. This coarse snapping is the most likely cause of the low noise reduction at low unregulated noise (since the unregulated supply droops are only ~ 30 mV in this case), and of the non-monotonicity of the noise reduction curve.




Fig. 16. Simulated effective impedance of (a) regulated supply with production transistors, (b) regulated supply with test-chip transistors, and (c) unregulated supply. The simulations use a block-level model similar to Fig. 12(b), but with comparator-based feedback. The effective impedance of the regulator (which is a nonlinear system) was calculated with $\sim 14\%$ peak-to-peak unregulated noise.

impedance with the target f_t for the production process; with this f_t , the regulator's feedback loop meets its target delay of ~ 190 ps and the noise is reduced by $\sim 50\%$. Fig. 16(b) shows the same simulation, but with the transistor speed scaled according to the measured performance monitor frequency. Since the measured noise reduction of 30% closely matches the simulation of Fig. 16(b), in a production process we anticipate the regulator to achieve the simulated 50% reduction in noise. With this higher level of noise reduction, the expected reduction in total power increases to $\sim 4\%$.

VI. CONCLUSION

In order for on-chip regulation of digital power supplies to be applied to modern, power-limited chips, this regulation must not adversely impact the chip's total power dissipation. Fortunately, since variations on the supply voltage force the nominal supply to be increased, the net power overhead of a regulator can in fact be negative. However, single-supply regulator topologies cannot achieve a net reduction in power because of the significant static power they dissipate in their output stages, either because of the required drop-out voltage in a series regulator, or the static current draw in a shunt design.

To overcome this limitation, we can build a push-pull shunt regulator that makes use of a second power supply, enabling it to spend output power only when necessary to counter noise. To distribute this second power supply for the regulator, resources from the main power supply must be allocated to the additional supply. Thus, we derived a strategy for this resource allocation that minimizes the total loss. In addition, we showed that keeping the regulator's return current flowing through the on-chip path requires careful design of the second supply's impedance—including the allocation of decoupling capacitance to store energy for this supply.

The design of the regulator itself must be optimized to minimize all sources of static power consumption. The use of a local RC filter-based reference guarantees that the regulator will not inefficiently attempt to correct for quasi-static

shifts in the supply voltage. Comparator-based feedback with a dead-band robustly eliminates quiescent current from the output devices, and an input peaking filter exploits the excess gain of the comparators to reduce the effective feedback delay without increasing power consumption. Finally, a switched source-follower output stage making use of replica biasing and transient current peaking can be used to achieve low turn-on delay with relatively predictable output current.

Measurement results from a fabricated design confirm that regulation can significantly ($\sim 30\%$) reduce the effective noise on a digital chip's power supply while actually resulting in a slight improvement ($\sim 1.4\%$) in the total power dissipation. While issues related to distributing many regulator sites on a die and adapting the parameters of each of these regulators for maximum efficiency remain, the results presented here show that integrated regulation is a promising approach to improving the robustness of energy-efficient digital circuits.

APPENDIX

In this Appendix, we analyze the noise performance of regulators with linear feedback paths as a function of the feedback amplifier's design parameters, leading to design equations that maximize noise rejection for a given amplifier gain-bandwidth (GBW). Even though an integrated regulator for digital circuits will most likely use nonlinear comparator-based feedback, the design rules developed here can be extended to this case by using describing function [15] techniques to quasi-linearize the comparator into an equivalent gain.

The noise performance of both series and shunt regulators with common-source output devices can be analyzed using the model shown in Fig. 10(a) by treating $Z_o(s)$ as the intrinsic impedance of the regulated supply node (i.e., for the series regulator, $Z_o(s) = R_{load} || r_o || 1/(sC_d)$). While this model assumes that undesired coupling of the power device's gate (V_c) to the regulated supply (V_{reg}) is negligible, including this coupling would not change the nature of the tradeoffs shown by the model. Finally, since with proper layout on-chip decoupling capacitors can have very low equivalent series resistance (ESR), we will also assume negligible capacitor ESR.

Using a first-order feedback amplifier with a DC gain of A_a and bandwidth of ω_a , the regulated supply impedance is

$$Z_{reg}(s) = \frac{Z_o(s) \left(1 + \frac{s}{\omega_a}\right)}{\left(1 + \frac{s}{\omega_a}\right) + g_m A_a Z_o(s)}. \quad (14)$$

The amplifier's limited bandwidth causes a zero in Z_{reg} , such that the dynamic noise performance of the regulator suffers if the amplifier's bandwidth is too low. Of course, excessively high amplifier bandwidth with low gain will also result in poor performance. Therefore, we next analyze optimizing the allocation between amplifier gain and bandwidth in order to maximize the regulator's performance. Since the regulator's response depends upon the type of noise excitation, we will perform the analysis with two models of noise: worst-case sinusoidal, and random, white noise. As we will see, the optimal allocation for these two disparate types of noise is very similar, such that a




Fig. 17. Magnitude plot of $Z_{\text{reg}}(s)$ showing that both amplifier gain and bandwidth must be increased by 2x to decrease impedance by 2x.

fixed gain/bandwidth allocation remains nearly optimal across a broad variety of noise excitations. Because it is a lower-order system, we will only describe in detail here the analysis of a series regulator. However, particularly in the case of worst-case sinusoidal noise, the isolating effect of the inductance L_{dist} (mostly due to the package and external distribution network) makes the results for shunt regulators essentially identical.

To minimize the regulator's peak-to-peak response to load current steps, it has been shown that the impedance of the regulator should exhibit no peaking [12], [13]—this is also known as voltage positioning. As shown in Fig. 17, because of the first-order roll-off of $Z_o(s)$, maintaining this voltage-positioned response while reducing the regulator's impedance requires both the gain and the bandwidth of the amplifier to be increased. This means that the gain–bandwidth of the amplifier must scale with the desired noise rejection squared. Indeed, for sinusoidal noise at the worst-case frequency, minimizing the maximum value of (14) leads to

$$\text{GBW} \approx 2 \cdot \text{LR}_{\text{min}}^2 \cdot \frac{\omega_o}{A_o} \quad (15)$$

where we have defined the power device's gain as $A_o = g_m \cdot (R_{\text{load}} || r_o)$, $Z_o(s)$'s pole as $\omega_o = 1/(R_{\text{load}} || r_o \cdot C_{\text{decap}})$, and the minimum feedback-contributed load rejection as $\text{LR}_{\text{min}} = (R_{\text{load}} || r_o) / ||Z_{\text{reg,max}}||$. Correspondingly, the best allocation between gain and bandwidth increases both with the square-root of GBW:

$$\omega_a = \sqrt{\frac{3}{2} \text{GBW} \cdot \left(\frac{g_m}{C_{\text{decap}}}\right)} \quad A_a = \sqrt{\frac{2}{3} \frac{\text{GBW}}{\left(\frac{g_m}{C_{\text{decap}}}\right)}} \quad (16)$$

The first point to notice with the allocation of (16) is that because of the g_m provided by the output device itself, most of the GBW will be allocated to amplifier bandwidth. For example, with a GBW of $10/(R_{\text{load}} C_{\text{decap}})$, a g_m of $6/R_{\text{load}}$ will make $A_a \approx 1$. In addition, as GBW is increased, the allocation given by (16) will cause the regulator's damping ζ to approach a constant value of ~ 0.61 , leaving a small amount of peaking in the regulator's response (as seen in Fig. 17). This slight discrepancy from traditional voltage positioning results from the lack

of capacitor ESR, and from the fact that the response to sinusoidal current (instead of square current steps) was optimized. However, this does not mean that the ESR of C_{decap} should be intentionally increased, since this would significantly degrade high-frequency noise sensitivity while minimally reducing the peak-to-peak ripple from a current step.

While load currents are unlikely to display truly sinusoidal behavior, minimizing $R_{\text{reg,max}}$ nearly minimizes the sensitivity of the regulator to a variety of deterministic or repetitive current variations. Furthermore, optimizing the response of the regulator to random, white load current noise (i.e., maximizing $\sigma_{\text{Inoise}}/\sigma_{V_{\text{reg}}}$) leads to a very similar allocation of

$$\omega_a = \sqrt{\text{GBW} \left(\frac{g_m}{C_{\text{decap}}}\right)} \quad A_a = \sqrt{\frac{\text{GBW}}{\left(\frac{g_m}{C_{\text{decap}}}\right)}} \quad (17)$$

The slight reduction in bandwidth is explained by the fact that for white, random load current, the noise energy is not concentrated at a single frequency. Therefore, additional peaking is actually desirable to minimize the total voltage noise energy. Correspondingly, the regulator's ζ approaches a constant value of ~ 0.5 at high A_{GBW} .

ACKNOWLEDGMENT

The authors would like to thank AMD's Sunnyvale Design Center, particularly S. Sethi, A. Thiruvengadam, M.-L. Chiang, A. Daga, Y. Wang, and Z. Shen.

REFERENCES

- [1] V. von Kaenel, D. Aebischer, C. Piguat, and E. Dijkstra, "A 320 MHz, 1.5 mW@1.35 V CMOS PLL for microprocessor clock generation," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1715–1722, Nov. 1996.
- [2] S. Sidiropoulos, D. Liu, J. Kim, G. Wei, and M. Horowitz, "Adaptive bandwidth DLLs and PLLs using regulated supply CMOS buffers," in *IEEE Symp. VLSI Circuits 2000 Dig. Tech. Papers*, Jun. 2000, pp. 124–127.
- [3] H. Werker *et al.*, "A 10-GB/s SONET-compliant CMOS transceiver with low crosstalk and intrinsic jitter," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2349–2358, Dec. 2004.
- [4] M. Ang, R. Salem, and A. Taylor, "An on-chip voltage regulator using switched-decoupling capacitors," in *IEEE Int. Solid-State Circuits Conf. (ISSCC 2000) Dig. Tech. Papers*, Feb. 2000, pp. 438–439.
- [5] J. Gu, R. Harjani, and C. Kim, "Distributed active decoupling capacitors for on-chip supply noise cancellation in digital VLSI circuits," in *Symp. VLSI Circuits 2006 Dig. Tech. Papers*, Jun. 2006, pp. 216–217.
- [6] J. Lee, G. Hatcher, L. Vandenberg, and C.-K. K. Yang, "Evaluation of fully-integrated switching regulators for CMOS process technologies," in *Proc. Int. Symp. System-on-Chip*, Nov. 2003, pp. 155–158.
- [7] P. Hazucha *et al.*, "A 233-MHz 80%–87% efficient four-phase DC-DC converter utilizing air-core inductors on package," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 838–845, Apr. 2005.
- [8] C. R. Sullivan and S. R. Sanders, "Design of microfabricated transformers and inductors for high-frequency power conversion," *IEEE Trans. Power Electron.*, vol. 11, no. 2, pp. 228–238, Mar. 1996.
- [9] A. M. Wu and S. R. Sanders, "An active clamp circuit for voltage regulation module (VRM) applications," *IEEE Trans. Power Electron.*, vol. 16, no. 5, pp. 623–634, Sep. 2001.
- [10] F. N. K. Poon, C. K. Tse, and J. C. P. Liu, "Very fast transient voltage regulators based on load correction," in *Proc. IEEE Power Electronics Specialists Conf. (PESC'99)*, Jun. 1999, pp. 66–71.
- [11] Intersil Corp., Data Sheet for HIP6200 and HIP6201, 1998.
- [12] R. Redl, B. P. Erismann, and Z. Zansky, "Optimizing the load transient response of the buck converter," in *Proc. Applied Power Electronics Conf. Expo. (APEC'98)*, Feb. 1998, pp. 170–176.
- [13] A. Waizman and C.-Y. Chung, "Extended adaptive voltage positioning (EAVP)," in *Proc. IEEE Conf. Electrical Performance of Electronic Packaging*, Oct. 2000, pp. 65–68.

- [14] Advanced Configuration and Power Interface Specification, Hewlett-Packard Corp., Intel Corp., Microsoft Corp., Phoenix Technologies Ltd., and Toshiba Corp., 2004 [Online]. Available: <http://www.acpi.info/DOWNLOADS/ACPISpec30.pdf>
- [15] S. S. Sastry, *Nonlinear Systems: Analysis, Stability, and Control*. New York: Springer-Verlag, 1999.
- [16] E. Alon, "Measurement and regulation of on-chip power supply noise," Ph.D. dissertation, Stanford Univ., Stanford, CA, 2006.
- [17] P. Hazucha *et al.*, "A linear regulator with fast digital control for biasing integrated DC-DC converters," in *IEEE Int. Solid-State Circuits Conf. (ISSCC'06) Dig. Tech. Papers*, Feb. 2006, pp. 2180–2189.
- [18] B. Casper *et al.*, "An 8-Gbps/simultaneous bidirectional link with on-die waveform capture," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2111–2120, Dec. 2003.
- [19] R. Farjad-Rad *et al.*, "0.622-8.0 Gbps 150 mW serial IO macrocell with fully flexible preemphasis and equalization," in *IEEE Symp. VLSI Circuits*, Jun. 2003, pp. 64–66.
- [20] Y. Tomita *et al.*, "A 10 Gb/s receiver with equalizer and on-chip ISI monitor in 0.11 μm CMOS," in *Symp. VLSI Circuits 2004 Dig. Tech. Papers*, Jun. 2004, pp. 202–205.
- [21] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [22] R. Venkataramanan, A. Sabanovic, and S. Čuk, "Sliding mode control of DC-to-DC converters," in *Proc. IEEE Conf. Industrial Electronics and Control Instrumentation (IECON)*, San Francisco, CA, Nov. 1985, pp. 251–258.
- [23] J. Kim and M. Horowitz, "An efficient digital sliding controller for adaptive power-supply regulation," *IEEE J. Solid-State Circuits*, vol. 37, no. 5, pp. 639–647, May 2002.
- [24] S.-C. Tan, Y. M. Lai, M. K. H. Cheung, and C. K. Tse, "On the practical design of a sliding mode voltage controlled buck converter," *IEEE Trans. Power Electron.*, vol. 20, no. 2, pp. 425–437, Mar. 2005.
- [25] G. W. den Besten and B. Nauta, "Embedded 5 V-to-3.3 V voltage regulator for supplying digital IC's in 3.3 V CMOS technology," *IEEE J. Solid-State Circuits*, vol. 33, no. 7, pp. 956–962, Jul. 1998.
- [26] Y. Nakamura, M. Takamiya, and T. Sakurai, "An on-chip noise canceller with high voltage supply lines for nanosecond-range power supply noise," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2007, pp. 124–125.



Elad Alon (S'02–M'06) received the B.S., M.S., and Ph.D. degrees from Stanford University, Stanford, CA, in 2001, 2002, and 2006, respectively, all in electrical engineering.

In January 2007, he joined the University of California, Berkeley, as an Assistant Professor of electrical engineering and computer sciences, where he is now a codirector of the Berkeley Wireless Research Center (BWRC). He has also held visiting positions at Intel, AMD, Rambus, Hewlett-Packard, and IBM Research, where he worked on integrated circuits for a variety of applications using bulk and CMOS processes from 130 nm down to 45 nm. His research focuses on the design and implementation of energy-efficient integrated systems and the circuits and technologies that comprise them.



Mark Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founders Professor of the School of Engineering at Stanford University. In 1990, he took leave from Stanford to help start Rambus Inc., Mountain View, CA, a company designing high-bandwidth memory interface technology.

His current research includes microprocessor design, low-power circuits, high-speed links, and new graphical interfaces.

Dr. Horowitz is a Fellow of the ACM and a member of the National Academy of Engineering. He was the recipient of the 1985 Presidential Young Investigator Award, the 1993 ISSCC Best Paper Award, the ISCA 2004 Most Influential Paper of 1989, and the 2006 winner of the IEEE Donald Pederson Award in Solid-State Circuits.